



Munich Personal RePEc Archive

Hypothetical Bargaining and the Equilibrium Selection Problem in Non-Cooperative Games

Radzvilas, Mantas

Department of Philosophy, Logic and Scientific Method, London
School of Economics and Political Science

24 March 2016

Online at <https://mpra.ub.uni-muenchen.de/70248/>
MPRA Paper No. 70248, posted 24 Mar 2016 13:32 UTC

Hypothetical Bargaining and the Equilibrium Selection Problem in Non-Cooperative Games*

Mantas Radzvilas[†]

March 24, 2016

Abstract

Orthodox game theory is often criticized for its inability to single out intuitively compelling Nash equilibria in non-cooperative games. The theory of virtual bargaining, developed by Misyak and Chater (2014) suggests that players resolve non-cooperative games by making their strategy choices on the basis of what they would agree to play if they could openly bargain. The proposed formal model of bargaining, however, has limited applicability in non-cooperative games due to its reliance on the existence of a unique non-agreement point – a condition which is not satisfied by games with multiple Nash equilibria. In this paper, I propose a model of ordinal hypothetical bargaining, called the Benefit-Equilibration Reasoning, which does not rely on the existence of a unique reference point, and offers a solution to the equilibrium selection problem in a broad class of non-cooperative games. I provide a formal characterization of the solution, and discuss the theoretical predictions of the suggested model in several experimentally relevant games.

1 Introduction

A central solution concept of the classical game theory is Nash equilibrium – a strategy profile which is such that no rational player is motivated to unilaterally deviate from it by playing another available strategy. However, at least intuitively, not all Nash equilibria are equally convincing as outcomes of a rational gameplay: even the simplest games have Nash equilibria that involve strategies which seem unlikely to be chosen by players who understand the structure of the game.

One of the canonical examples of such a "puzzle" game is the Hi-Lo game (figure 1)¹

*This is a work-in-progress paper. Do not cite without permission.

[†]London School of Economics, Department of Philosophy, Logic and Scientific Method (e-mail: m.radzvilas@lse.ac.uk).

¹Unless it is stated otherwise, the payoff numbers in the matrices are Von Neumann and Morgenstern utilities. The payoffs are assumed to represent all the relevant motivations of

	<i>Hi</i>	<i>Lo</i>
<i>Hi</i>	2, 2	0, 0
<i>Lo</i>	0, 0	1, 1

Figure 1: The Hi-Lo game

At least intuitively, the outcome (Hi, Hi) stands out as the obvious solution because Hi-Lo is a common interest game: each player who knows the payoff structure of the game should also know that strategy profile (Hi, Hi) is the best outcome for both players, and that there is no conflict of players' interests in this game. The perfect alignment of player's interests in this game is the reason why we have a 'high-quality intuition' that strategy *Hi* is a much more likely choice than strategy *Lo* (Bacharach 2005: 35-68). Experimental results reveal that over 90% of the time people opt for Hi in this game².

From the perspective of classical game theory, however, this game has three rational solutions – pure strategy Nash equilibria (Hi, Hi) and (Lo, Lo) , and a mixed strategy Nash equilibrium $(\frac{1}{3}, \frac{2}{3}; \frac{1}{3}, \frac{2}{3})$. The theory does not single out Nash equilibrium (Hi, Hi) as a more likely or compelling solution than other Nash equilibria of this game. The reason of why standard game theoretic analysis leads to such conclusion becomes clear when we look into the model of reasoning which underpins it. In standard game theoretic analysis of complete information games, players' rationality and the payoff structure of the game is assumed to be common knowledge. A rational player is assumed to be a best-response reasoner – a player who always chooses a strategy that, given player's beliefs about the opponents' strategy choices, maximizes his or her expected payoff. Common knowledge of rationality implies that every player knows that none of the opponents' will choose their strictly dominated strategies – strategies which never are best-responses to any possible probabilistic beliefs that a rational player might hold about his or her opponents' strategy choices. If the payoff structure of the game is also common known, each player can iteratively eliminate the strictly dominated strategies of the game, thus leaving them with a set of strategies which are rationalizable³.

In non-cooperative games with multiple Nash equilibria at least one of the

players, including pro-social preferences, such as inequity aversion, altruism, sensitivity to social norms, and so on.

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.

³A rationalizable strategy is a strategy which is a best-response to some possible conjecture a player may have about his or her opponents. Bernheim (1984) and Pearce (1984) have shown that a set of rationalizable strategies of the game can be obtained via iterative elimination of strictly dominated strategies. For a detailed discussion and proofs, see Bernheim 1984: 1007-1028, and Pearce 1984: 1029-1050.

players have multiple rationalizable strategies. In case of the Hi-Lo game, it is easy to check that both pure strategies are rationalizable: if the player believes that the probability that the opponent is going to play strategy *Hi* is less than $\frac{1}{3}$, then his or her expected utility maximizing action is to play *Lo*, and if the probability of opponent playing *Hi* is more than $\frac{1}{3}$, then player's best-response strategy is *Hi*. The problem is that in non-cooperative games where players have multiple rationalizable strategies, best-response reasoners face a situation of strategic uncertainty, as common knowledge of rationality and of the payoff structure of the game gives them no further indication on what rationalizable strategies their opponents are going to choose. The classical game theory does not offer an explanation of how rational players are supposed to respond to such situations of strategic uncertainty, 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, is a likely outcome of the game (Olcina and Urbano 1994: 183-206).

Another game where the classical game theory solution seems to contradict our intuitions is the following "Impure Coordination game" (figure 2):

	<i>l</i>	<i>c</i>	<i>r</i>
<i>u</i>	6, 3	0, 0	0, 0
<i>m</i>	0, 0	3, 6	0, 0
<i>d</i>	0, 0	0, 0	5, 5

Figure 2: Impure Coordination game

In this game, players also face a situation of strategic uncertainty: there are three pure-strategy Nash equilibria (u, l) , (m, c) and (d, r) , and four Nash equilibria in mixed strategies $(\frac{2}{3}, \frac{1}{3}, 0; \frac{1}{3}, \frac{2}{3}, 0)$, $(\frac{10}{21}, \frac{5}{21}, \frac{2}{7}; \frac{5}{21}, \frac{10}{21}, \frac{2}{7})$, $(\frac{5}{8}, 0, \frac{3}{8}; \frac{5}{11}, 0, \frac{6}{11})$, $(0, \frac{5}{11}, \frac{6}{11}; 0, \frac{5}{8}, \frac{3}{8})$. Unlike in the Hi-Lo game, players face a conflict of interests

⁴Aumann and Brandenburger (1995) have established the epistemic conditions of Nash equilibrium for two player and n -player games. In a two player game, if rationality of the players and each player's conjecture (that is, a belief about what the other player is going to do) are mutually known (common knowledge of rationality is not one of the required conditions for the Nash equilibrium to obtain), 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. In standard game theory models, players' conjectures are not assumed to be commonly or mutually known, nor these private conjectures are assumed to be correlated. 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: 1161-1180 and Perea 2012: 1-67

in this game: the best personal payoff for the row player is associated with Nash equilibrium (u, l) , while the best personal payoff for the column player is associated with Nash equilibrium (m, c) . However, unshophisticated selfish behaviour is risky: if both players were to pursue their most preferred options, they would end up playing strategy profile (u, c) which, in this game, is associated with the worst possible personal outcome for both players. Each of the mixed strategy Nash equilibria is also a possible resolution of this coordination problem, yet it would be compelling only if both players believed that his or her opponent will respond to strategic uncertainty by playing a randomized strategy constituting a particular mixed Nash equilibrium of the game. A belief that an opponent will respond to strategic uncertainty by playing his or her end of a particular mixed Nash equilibrium does not follow from the common knowledge of rationality.

At least intuitively, Nash equilibrium (d, r) seems to be the most compelling solution of this game: it seems reasonable that players would respond to strategic uncertainty “cooperatively”– by choosing strategies that would ensure both players the second-best personal payoff, rather than by guessing the opponent’s actions and facing the risk of getting the worst possible payoff. In this particular version of the Impure Coordination game, this intuition is further supported by the fact that the payoff associated with Nash equilibrium (d, r) is, for both players, higher than payoffs associated with each mixed Nash equilibrium of this game. However, standard game theory analysis does not single out Nash equilibrium (d, r) as the unique solution of this game.

One of the more recent theories which offers some support to the aforementioned intuitions about real-world strategic reasoning is the theory of Virtual Bargaining, developed by Jennifer Misyak and Nick Chater (2014). Arguably the main conceptual innovation of this theory is the idea that real-world player’s are not best-response reasoners, but virtual bargainers. When the player reasons as virtual bargainer, s/he focuses on the question "what would we agree to do in this game if we could openly bargain? The virtual bargainer then searches for the best feasible agreement, and, if such an agreement exists, plays his or her part in realizing it.

The theory of virtual bargaining shares some conceptual similarities with hypothetical bargaining, which has received considerable attention in economics, political philosophy and legal theory. Roughly, it is the idea that solutions to interdependent decision problems can be found by considering a counterfactual situation where rational and self-interested agents openly negotiate a solution to the problem. An agreement that could be reached in such hypothetical negotiations is taken to be a rational (sometimes fair) solution of the problem.

The theory of virtual bargaining, however, differs from the aforementioned approach, since it is proposed as a psychological theory of reasoning. In other words, the model of virtual bargaining is offered as an approximation to the actual process of reasoning by which people arrive at their choices, rather than as a merely useful conceptual tool for finding reasonable solutions to various decision problems⁵ to non-cooperative games.

⁵The BER model suggested in this paper aims to represent the behaviour of rational

The idea that people aim to resolve non-cooperative games by identifying the mutually beneficial outcomes is intuitively compelling, and is fairly well supported by experimental observations⁶. In addition, there are at least two reasons of why the theory of virtual bargaining is conceptually appealing.

First, the axiomatic bargaining theory is, essentially, a correlated equilibrium selection theory. In bargaining games where agreements are not externally enforceable, the set of feasible basic agreements is the set of correlated equilibria. Each axiomatic bargaining solution of such game is a correlated equilibrium (or a set of correlated equilibria) which satisfies a certain set of intuitively desirable properties. These bargaining solutions can be interpreted as reasonable predictions of the result of the open negotiations⁷. It stands to reason to assume that the desirable properties of bargaining agreements may also be relevant for finding reasonable solutions to equilibrium selection problems in other types of games, such as matching problems and non-cooperative games.

Second, bargaining theory is a branch of non-cooperative games, and bargaining solutions are based on the concept that underly all the areas of non-cooperative game theory, including non-cooperative one-shot games (Myerson 1991: 370). This means that bargaining solutions are consistent with the concept of player's rationality employed in the analysis of non-cooperative one-shot games.

However, the theory of virtual bargaining is relatively new, and still has substantial conceptual limitations. At its current state⁸, the theory of virtual bargaining is, essentially, the Nash bargaining solution applied to non-cooperative games: the set of mutually optimal bargains of the non-cooperative game is taken to be the set of outcomes that maximize the product of players' payoff gains.

I will argue that such a direct application of the Nash bargaining solution to non-cooperative games is problematic for two reasons. First, all standard bargaining solutions, including the Nash bargaining solution, rely on the existence of a unique non-agreement point – an outcome which obtains in case the negotiators fail to reach an agreement. In games with multiple Nash equilibria, however, there is no such unique reference point, meaning that virtual bargainers would face a problem of identifying the common non-agreement point. I also show that, in some games, players who selected different non-agreement points would identify different Nash-optimal agreements, and their actions would not

individuals, rather than the actual process of reasoning by which players arrive at their choices. That is, it suggests that players' resolve the non-cooperative games as if they have adopted hypothetical bargaining approach to determine the solution of the problem.

⁶See, for example, Coleman and Stirk (1998) who report experimental results from several different mixed motive games, including the Chicken game, the Stag Hunt Game, the Battle of the Sexes and the Leader game, as well as participants' justifications of their choices. A substantial proportion of participants have justified their choices in one-shot games by appealing to some notion of mutual benefit ("most points for both", "mutual benefit", etc.).

⁷For an extensive discussion, see Myerson 1991: 370-416.

⁸Misyak and Chater explicitly state that Nash bargaining theory is merely a "*useful starting point for the analysis of virtual bargaining*". They never claim that the Nash bargaining theory is the best possible approximation to the process of reasoning by which players identify the optimal and feasible agreements. See Misyak and Chater 2014: 1-9.

lead to mutually beneficial outcomes. Second, in non-cooperative games, there may be multiple Nash optimal agreements, and each Nash-optimal agreement may be associated with a different allocation of players' personal payoff gains. I will argue that Nash bargaining theory does not offer an explanation of how self-interested players would resolve such benefit allocation problems, and therefore a search for a different approach to virtual bargaining is a warranted endeavour.

In this paper, I propose an alternative model of reasoning, called the Benefit Equilibration Reasoning (later abbreviated as BER) which offers an explanation of how the players may use the commonly known information about the payoff structure of the game in identifying the mutually agreeable outcomes in non-cooperative one-shot games. It is a model of ordinal bargaining, which does not require a non-agreement point, and does not rely on players' ability to make interpersonal comparisons of payoffs⁹.

The rest of the paper is structured as follows. In section 2 I discuss the virtual bargaining theory, and offer some reasons of why the standard bargaining solutions, such as Nash bargaining solution, cannot be used to represent the outcomes of hypothetical negotiations in non-cooperative games. In section 3 I present an alternative benefit-equilibratin (BE) ordinal bargaining solution and its axiomatic characterization. In section 4 I illustrate the theoretical predictions of the BER model using a number of experimentally relevant examples. In section 5 I discuss the empirical relevance of the BER model. With section 6 I conclude and discuss some limitations of the model.

2 Virtual Bargaining

The theory virtual bargaining suggests that people resolve the non-cooperative games by identifying those strategy profiles that they would agree on playing if they were engaged in open bargaining – real negotiations in which each player communicates his or her offers to other players, and receive their counter-offers.

To see the intuition behind this model of reasoning, consider the Hi-Lo game. This game has multiple Nash equilibria, so the players 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 on playing (Hi, Hi) , as it is the best outcome for both players. The joint action plan to realize strategy profile (Hi, Hi) is an enforceable agreement: each player knows that, if he or she plays strategy Hi , the co-player will not have an incentive to defect by playing Lo (Misyak and Chater 2014: 1-9).

The Virtual Bargaining procedure is as follows. Each virtual bargainer determines the set of agreements – a set of combinations of actions of the players, which includes all the combinations of players' randomized actions¹⁰. The preference relation of each virtual bargainer must be defined on the set of lotteries

⁹BER model is an ordinal solution concept, which relies on players' ability to make interpersonal comparisons of their ordinal preferences over outcomes.

¹⁰In cases where players' agreements are enforced by an external party, the feasibility set includes all the possible combinations of players' actions. If the external enforcer is not available, the set of feasible agreements is composed only of correlated equilibria – the self-

over agreements¹¹. A virtual bargainer then searches for the best agreement – a combination of (possibly randomized) players’ actions which maximizes the product of players’ personal payoff gains, relative to their non-agreement payoffs. If such an agreement exists, then each virtual bargainer plays his or her part in realizing it. (Misyak and Chater 2014: 1-9).

The idea that players search for mutually beneficial solutions to games seems to have intuitive appeal, and is fairly well supported by empirical evidence¹². An idea that players find the mutually advantageous solutions of games by looking for combinations of actions that openly negotiating players would agree on carrying out seems to be a psychologically plausible explanation of how people reason about mutual advantage in interdependent decision problems. However, the formal model of virtual bargaining, which is offered as an approximation to the actual process of reasoning by which people arrive at their choices, has substantial conceptual and empirical limitations.

The reliance of virtual bargaining on the Nash bargaining solution seems to be problematic for several reasons. First, standard bargaining solutions, such as Nash bargaining solution, had been developed for a particular class of interdependent decision problems, known as bargaining games. These solutions rely on the existence of a unique non-agreement point – an outcome that obtains when individuals fail to reach an agreement following a bargaining process. For example, in the Nash bargaining game this is assumed to be the outcome in which both players gain nothing. The existence of a unique non-agreement point in a bargaining game is important for two reasons. First, it is used to determine each player’s personal utility gains from each feasible agreement, and thus it plays a fundamental role in formal characterizations of bargaining solutions¹³. Second, the non-agreement point serves as a threat point in strategic (alternating offers) bargaining models, where threats are interpreted as players’ actions: 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 counter-offer by threatening the opponent to play his or her non-agreement strategy, thus harming the opponent by bringing him down to his or her personal disagreement payoff as well¹⁴.

enforcing agreements of the game. Misyak and Chater seem to suggest that virtual bargaining mimics the procedure of bargaining where external enforcer is not available. For an extensive discussion of the Nash bargaining problem, see Myerson 1991: 370-416.

¹¹The idea is that players’ preferences must be defined over the set of lotteries, where each lottery "prize" is a particular combination of players’ actions, since players’ preferences over agreements must capture their attitude to risk. For an extensive discussion of the role of players’ attitude to risk in bargaining problems, see Myerson 1991: 370-416

¹²See, for example, Coleman and Stirk (1998), Crawford et al 2008: 1443-1458, Bardsley et al. 2010: 40-79, Misyak and Chater 2014: 1-9, Faillo et al. 2016: 1-61

¹³That is, the reference point is used in axiomatic bargaining theory to identify the feasible agreements which satisfy certain desirable properties, such as Pareto efficiency, symmetry, independence of irrelevant alternatives, proportionality, etc. See, for example, Luce and Raiffa 1957: 114-154, Kalai and Smorodinsky 1975: 513-518, Kalai 1977: 1623-1630, Myerson 1977: 1631-1637, Roth 1979: 775-778.

¹⁴In 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: 80-109, Rubinstein 1982: 97-109 and Binmore et al. 1986: 176-188.

The problem with virtual bargaining is that the structural properties of most non-cooperative games differ significantly from the properties of bargaining games. Some non-cooperative games have unique inefficient solutions that can, in principle, serve as non-agreement points. For example, Misyak and Chater suggest that a unique Pareto inefficient Nash equilibrium of the Prisoner’s Dilemma game can be interpreted as the disagreement point (even though no mutually beneficial agreements are feasible in this game¹⁵). In games with multiple Nash equilibria, however, there is no such unique non-agreement point: every Nash equilibrium of the game is a potential non-agreement point, meaning that virtual bargainers face a non-agreement point selection problem, which is the equilibrium selection problem of the original game. The fact that each Nash equilibrium of the game is a potential non-agreement point makes the theory of virtual bargaining useless in offering a solution to the equilibrium selection problem, which is present in most games that are used to represent real-world social interactions.

To understand the limitations of the theory, consider the following version of an asymmetric Battle of the Sexes game (figure 3):

	<i>e</i>	<i>f</i>
<i>a</i>	6, 11	0, 0
<i>b</i>	4, 4	8, 7

Figure 3: The Asymmetric Battle of the Sexes game

This game has two pure strategy Nash equilibria – (a, e) and (b, f) , and a mixed Nash equilibrium $(\frac{3}{14}, \frac{11}{14}; \frac{4}{5}, \frac{1}{5})$. Each one of the Nash equilibria satisfies the feasibility criterion – none of the players can, by changing his or her strategy, gain advantage to the disadvantage of the other. Also, every Nash equilibrium is an enforceable non-agreement point: if one of the players reverts to playing his or her non-agreement strategy, the other player cannot do better than to revert to his or her non-agreement strategy as well.

The problem with using Nash equilibria as non-agreement point is that there is no obvious reason as to why one of the multiple equilibria should be selected as a non-agreement point – in terms of the formal properties that a non-agreement point has to satisfy, all Nash equilibria are equivalent¹⁶.

¹⁵Misyak and Chater discuss a version of the Prisoner’s Dilemma game where the cooperative outcome uniquely maximizes the product of players’ payoff gains. However, according to their theory, this outcome is not a feasible agreement: each player can gain a personal advantage by defecting from the agreement, to the disadvantage of the player who sticks to the agreement. For details, see Misyak and Chater 2014: 1-9.

¹⁶In this particular game, as well as in many other games with multiple Nash equilibria, there is no feasible agreement which, relative to any pure Nash equilibrium point, would be mutually beneficial for the players. For the feasible agreement to be mutually beneficial, the

Even in cases where there seems to be a clearly mutually beneficial agreement available to the players, such as the Pareto-efficient Nash equilibrium (Hi, Hi) of the Hi-Lo game (figure 1), the virtual bargainers would only identify the outcome (Hi, Hi) as the optimal bargain if they were to choose Nash equilibrium (Lo, Lo) as disagreement point. The theory does not provide any explanation of why players would choose an inefficient equilibrium as the non-agreement point.

Note that, in case of the Hi-Lo game, the players who chose different equilibria as non-agreement points would still end up playing Nash equilibrium (Hi, Hi) : the player who chose the Nash equilibrium (Hi, Hi) as a reference point would reach the conclusion that no mutually beneficial agreements are available, and would stick to playing his or her non-agreement strategy (Hi) . The player who chose Nash equilibrium (Lo, Lo) would identify Nash equilibrium (Hi, Hi) as Nash optimal agreement, and would therefore play strategy Hi as well. In other games, however, players' failure to choose the same non-agreement point may lead them to making choices that would leave the players with suboptimal outcomes. For example, consider this extended Battle of the Sexes game (figure 4):

	x	y	z
a	4, 10	1, 1	0, 1
b	4, 4	5, 6	0, 1
c	1, 0	1, 0	4, 7

Figure 4: The Extended Battle of the Sexes game

This game has three pure strategy Nash equilibria – (a, x) , (b, y) and (c, z) , and four mixed Nash equilibria¹⁷, and each of these can serve as virtual bargainers' non-agreement point.

The selection of the non-agreement point actually determines what combinations of actions the players would identify as optimal agreements. For example, consider the possibility that players would select one of the pure Nash equilibria as a non-agreement point. The Nash equilibrium (c, z) is a maximin non-agreement point – an outcome which would obtain if, in case of disagreement, the players were to revert to playing their maximin strategies¹⁸. If this Nash

product of players' payoff gains must be strictly positive. In the aforementioned Battle of the Sexes-Chicken hybrid game, if one of the Nash equilibria is chosen as non-agreement point, then no strategy profile has a positive product of players' payoff gains.

¹⁷The four mixed Nash equilibria of this game are: $(\frac{2}{11}, \frac{9}{11}, 0; 1, 0, 0)$, $(\frac{7}{61}, \frac{63}{122}, \frac{45}{122}; \frac{4}{7}, 0, \frac{3}{7})$, $(\frac{7}{16}, 0, \frac{9}{16}; \frac{4}{7}, 0, \frac{3}{7})$, and $(0, \frac{7}{12}, \frac{5}{12}; 0, \frac{1}{2}, \frac{1}{2})$.

¹⁸A maximin strategy is a strategy which maximizes player's minimum payoff. A maximin payoff (also known as security payoff) is the maximum payoff that a player can guarantee to himself or herself irrespective of the opponents' actions. For an extensive discussion, see Luce

equilibrium was selected as a non-agreement point, then no feasible agreement would be Nash optimal. If the outcome (b, y) was selected as a non-agreement point, then the Nash equilibrium (a, x) would be an optimal agreement. Finally, if players' were to choose profile (a, y) as a non-agreement point, then, again, they would find no optimal agreements in this game.

This observation points to another problem associated with the presence of multiple non-agreement points: the virtual bargainers who selected different Nash equilibria as non-agreement points (since virtual bargainers do not communicate, this scenario is entirely possible), might identify different solutions of the game. In such cases, each player's individual actions could lead to sub-optimal outcomes for both players. For example, if the column player were to select the Nash equilibrium (c, z) as a non-agreement point, then s/he would conclude that the game has no feasible bargains, and would play z . If the row player chose the Nash equilibrium (a, y) , s/he would identify outcome (a, x) as Nash optimal agreement, and would play strategy a . The players would thus end up playing strategy profile (a, z) , which is, for both players, associated with low personal payoffs.

Another problem besides the reference point selection problem is that non-cooperative games may have multiple Nash optimal agreements with different allocations of players' personal payoff gains, even relative to the same reference point. The Nash bargaining solution has been developed to resolve a specific type of game, known as the Nash bargaining game. 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 involving a set of feasible allocations of the good, and the Nash bargaining solution of the game is a unique distribution of the good, which satisfies a set of desirable properties: symmetry, invariance with respect to affine utility transformations, Pareto optimality, and independence of irrelevant alternatives¹⁹. In non-cooperative games, however, players' utility functions can represent any kind of personal motivations, and the game may have multiple feasible agreements which maximize the Nash product, yet each agreement may be associated with different allocations of players' personal payoff gains. This result seems to be problematic: since players are assumed to be self-interested individuals, 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 which of the agreements would be reached in open negotiations becomes a crucial one.

For example, consider this two player four strategy coordination game (figure 5):

This game has four Nash equilibria in pure strategies – (a, h) , (b, i) , (c, j) and (d, k) , and eleven Nash equilibria in mixed strategies. For the sake of simplicity, let us assume that the disagreement point is strategy profile (d, k) – a Nash equilibrium, which would obtain if, in case of disagreement, the players were

and Raiffa 1957: 278-326.

¹⁹For an extensive discussion of Nash bargaining theory, see Luce and Raiffa 1957: 114-154, and Myerson 1991: 370-416.

	<i>h</i>	<i>i</i>	<i>j</i>	<i>k</i>
<i>a</i>	17, 2	0, 0	0, 0	0, 1
<i>b</i>	0, 0	2, 17	0, 0	0, 1
<i>c</i>	0, 0	0, 0	5, 5	0, 1
<i>d</i>	1, 0	1, 0	1, 0	1, 1

Figure 5: Coordination game with an asymmetric allocation of payoff gains

to revert to playing their maximin strategies. Relative to non-agreement point (d, k) , all three pure strategy Nash equilibria are Nash optimal agreements: each outcome maximizes the product of players' payoff gains, relative to the maximin non-agreement point²⁰

Despite the fact that all three pure strategy Nash equilibria are Nash optimal agreements, they are associated with different allocations of players personal payoff gains. Since players are assumed to be self-interested, it seems unreasonable to assume that they would be indifferent between these outcomes. More specifically, it seems more likely that players would agree on jointly realizing outcome (c, j) rather than outcomes (a, h) and (b, i) . Given the fact that Nash equilibrium (c, j) is a feasible alternative agreement, at least intuitively it seems very likely that the disadvantaged party would reject the offers (a, h) (b, i) as unfair, since these agreements lead to intuitively unequal advancement of players' personal interests²¹. Also, note that the disadvantaged player can always threaten the other player to play his or her maximin strategy, thus bringing the other player down to his or her disagreement payoff. Since each player prefers the personal payoff associated with agreement (c, j) over his or her personal payoff associated with non-agreement point, it seems likely that both players would

²⁰It is easy to check that the product of players' payoff gains, relative to disagreement outcome (g, p) , is 16 for each pure strategy Nash equilibrium:

Equilibrium (a, h) : $[(17 - 1) \times (2 - 1)] = 16$

Equilibrium (b, i) : $[(2 - 1) \times (17 - 1)] = 16$

Equilibrium (c, j) : $[(5 - 1) \times (5 - 1)] = 16$

²¹Misyak and Chater discuss a version of the Battle of the Sexes game, where one Nash equilibrium is associated with what they refer to as "asymmetric payoffs" $(1, 11)$, while another Nash equilibrium is associated with "mutually good" payoffs $(10, 9)$. Misyak and Chater suggest that a Nash equilibrium with "mutually good" payoffs is a more likely outcome of the bargaining process than the Nash equilibrium with "asymmetric payoffs", as the disadvantaged players is likely to reject the latter offer. However, Misyak and Chater do not consider a situation where both outcomes maximize the Nash product (in their example, the "mutually good" outcome is the unique Nash-optimal outcome of the game), nor do they offer a theoretical explanation of how the two Nash equilibria can be compared without interpersonal comparisons of players' payoffs. For details, see Misyak and Chater 2014: 1-9.

agree on playing (c, j) rather than ending their negotiations with no agreement at all.

Despite the fact that Nash equilibrium (c, j) seems to be an intuitively appealing solution of this game, it is not a unique Nash bargaining solution of this game. The Nash bargaining theory does not offer an answer to the question of how negotiators would resolve the conflicts over alternative allocations of personal payoff gains, and this question seems to play a major role in real-world negotiations²². If the theory of virtual bargaining is supposed to be an approximation to the actual process of reasoning by which players arrive at their choices, it seems that equity considerations should be incorporated into the model.

There are several game theoretic models which purport to explain player's choices in experimental games by incorporating players' fairness considerations and other types of pro-social preferences into players' payoff functions, such as the well known inequity aversion theory suggested by Fehr and Schmidt (1999). These theories, however, have two conceptual limitations. First, although these theories are useful in explaining players' choices in games with material payoffs, they cannot be applied to games where players' payoffs are their Von Neumann Morgenstern utilities, which represent all their relevant motivations, including, among other things, players' pro-social preferences, such as inequity aversion, altruism, sensitivity to social norms, and so on. Second, the aforementioned models rely on the assumption that interpersonal comparisons of players' payoffs are meaningful, which goes beyond the principles of the orthodox expected utility theory.

In the following section, I will suggest an alternative bargaining-based explanation of how players resolve the benefit-allocation problems, which does not rely on players' ability to make interpersonal comparisons of payoffs, and does not require the existence of a unique non-agreement point. More specifically, I will argue that foregone opportunities play an important role in players' judgments of whether a solution of the game is mutually beneficial. I will suggest a formal Benefit-Equilibration Reasoning model, which offers an explanation of how such comparisons of foregone opportunities may determine players' choices in non-cooperative games.

3 The Benefit-Equilibrating Solution

3.1 The Intuition Behind the Benefit-Equilibrating Solution

To understand the intuition behind the Benefit-Equilibrating solution (later abbreviated as BE solution), consider the Impure coordination game (figure 6):

This game has three pure strategy Nash equilibria (u, l) and (m, c) and (d, r) ,

²²for an extensive discussion of the role of equity considerations in bargaining, see Myerson 1991: 370-416.

	l	c	r
u	6, 3	0, 0	0, 0
m	0, 0	3, 6	0, 0
d	0, 0	0, 0	5, 5

Figure 6: Impure Coordination game

and four mixed strategy equilibria²³.

Each player's payoff function represents his or her preferences over the strategy profiles (i.e. outcomes) of the game. Each outcome of the game can be interpreted as a possible state of the world that players can bring about via joint actions. Although players' payoffs are not interpersonally comparable, the common knowledge of payoffs implies that players know each other's preferential ordering of the outcomes of the game. In other words, each player who knows the payoff structure of the game can rank, for each player of the game, the outcomes of the game from that player's most preferred to his or her least preferred outcome. The row and the column players' preferential rankings of the pure strategy profiles of the three strategy coordination game (figure 6) are shown below:

$$\begin{array}{l}
 \text{Row player:} \\
 \text{Column player:}
 \end{array}
 \left[\begin{array}{c}
 (u, l) \\
 (d, r) \\
 (m, c) \\
 (u, c), (u, r), (m, l), (m, r), (d, l), (d, c) \\
 (m, c) \\
 (d, r) \\
 (u, l) \\
 (u, c), (u, r), (m, l), (m, r), (d, l), (d, c)
 \end{array} \right]$$

Imagine that player's preferential ranking of pure strategy profiles is a set of written "contracts", or possible agreements, that s/he can offer to the other player. The row player's most preferred contract is (u, l) , meaning that his or her personal interests would be maximally advanced if this contract was carried out. The column player's interests would be maximally advanced if players carried out contract (m, c) . Note that both player's personal interests cannot be maximally advanced with any of the available contracts. If players carried out contract (u, l) , the players would forego the opportunity to maximally advance

²³In this particular example, I consider only pure strategy profiles as possible agreements in order to make the intuition behind the suggested solution as clear as possible. The omission of mixed strategy profiles will not make a difference in this case, since the BER solution of the game is a pure strategy profile. It must be noted, however, that a formal BER solution will be defined over the set of mixed strategy profiles, and, in many games, the BER solution of the game is a mixed Nash equilibrium.

column player's personal interests by carrying out contract (m, c) . If, on the other hand, players carried out contract (m, c) , they would forego the opportunity to maximally advance the personal interests of the row player by carrying out contract (u, l) .

Since players' have partially conflicting preferences over possible agreements, they could only reach an agreement if at least one of the player's were to give up his or her preferred contract and accept the opponent's offer²⁴. However, notice that each agreement that could potentially be reached in open bargaining would be associated with a specific distribution of players' *foregone preferred alternatives*. For example, if contract (u, l) was carried out, the column player would forego two opportunities²⁵ to advance his or her personal interests with contracts (m, c) and (d, r) , while the row player would forego no opportunities to advance his or her personal interests at all. If, on the other hand, contract (m, c) was carried out, the column player would forego no opportunities to advance his or her personal interests at all, while the row player would forego two opportunities to advance his or her personal interests with contracts (u, l) and (d, r) . In other words, the implementation of either of the two contracts would mean that one of the players would forego more opportunities to advance his or her personal interests than the other player. An agreement which leads to an unequal distribution of players' foregone preferred alternatives would not be a symmetric, or *benefit-equilibrating*, resolution of the conflict. An asymmetric resolution of players' conflict seems to be problematic in the context of non-cooperative games where players' roles are symmetric²⁶.

Assuming that bargaining did not involve external enforcements (which seems to be a reasonable assumption since players in non-cooperative games make independent choices), the players would only consider the feasible agreements of the game, meaning that they would constrain their negotiations to a set of self-enforcing contracts of the game. In our example, such players would constrain their negotiations to a subset of contracts containing only the pure strategy Nash equilibria (u, l) , (d, r) and (m, c) ²⁷. However, this would not resolve the problem, since contracts (u, l) and (m, c) are self-enforcing, and thus

²⁴We can imagine this open bargaining as a process similar to the non-cooperative iterative vetoing procedure suggested by Anbarci (1993), where the contract is reached by players' iteratively vetoing their least preferred contract. For details, see Anbarci 1993: 245-258.

²⁵Note that, in some games, a player may be indifferent between several contracts (as in the Impure coordination game. These contracts will have the same rank value assigned to them in player's preferential ranking of contracts. The contracts with the same rank value will be counted as one opportunity to advance players' personal interests

²⁶For a discussion of the role of symmetry in bargaining problems, see Luce and Raiffa 1957: 114-154, and Myerson 1991: 370-416.

²⁷A self-enforcing contract is a possible agreement such that a player who believes that his or her opponent will carry out his or her part of the contract cannot gain personal advantage by deviating from it. In bargaining problems without external enforcement, the set of feasible agreements is the set of correlated equilibria of the game. This feasibility constrains reflects each player's belief that his or her opponent will always deviate from the agreement if such a deviation is personally beneficial. If both players express a common belief in rationality, they should expect each other to consider only the set of feasible agreements of the game. For details, see Myerson 1991: 370-416.

they would both remain in the set of feasible agreements of the game. The players could not settle on any of the two contracts, since, in the absence of an external enforcer, each player would respond to the opponent's contract offer by replacing it with his or her preferred contract, and this cycle would continue indefinitely.

A benefit-equilibrating agreement could, however, be reached if the players, in order to reach an agreement, would constrain their negotiations to a set of contracts that did not contain strategy profiles (u, l) and (m, c) ²⁸. In this constrained set of options, strategy profile (d, r) would be the best available contract for both players. If one of the players were to offer it, the other player would have no incentive to swap the offer with his or her most preferred contract in the constrained set. If one of the player's preferred contract from the constrained set were implemented, the other player could not object that the implementation of the opponent's offer prioritized the advancement of the opponent's personal interests over the advancement of his or her own personal interests²⁹.

Note that for both players to constrain the set of contracts in the aforementioned way, each of them would have to give up his or her most preferred contract. This means that each player would forego one opportunity to advance his or her personal interests. But it also means that no player could object that the constrained set of contracts prioritizes the advancement of his or her opponent's personal interests, since each player would forego *one* opportunity to advance his or her personal interests. In this sense, contract (d, r) would be an ordinally egalitarian solution of the Impure Coordination game³⁰.

3.2 A Formal Characterization of the BE Solution

The BE solution of the ordinally symmetric games³¹, such as Impure coordination game (figure 6), can be easily identified using a common numerical representation of players' ordinal preferences over the feasible agreements³². For

²⁸The reasons of why the players would be motivated to constrain their set of contracts will be discussed in subsection 3.3.

²⁹In a sense that an implementation of one player's preferred offer meant a foregone opportunity to advance the interests of the other player. For an extensive discussion of the relationship between equity and the structures of opportunity sets, see Kolm 2007: 69-162.

³⁰The principle of equal ordinal sacrifice of preferred allocations was introduced by Conley and Wilkie (2012) as a solution to the Pareto-optimal point selection problem for bargaining games with finite sets of Pareto-optimal alternatives. An agreement is ordinally egalitarian if both players give up equal numbers of *preferred alternatives* in order to reach it. This solution principle does rely on comparisons of players' ordinal preferences, but it does not require interpersonal comparisons of players' cardinal payoffs. The benefit-equilibrating solution differs from the aforementioned approach, since it applies to games with multiple ordinally egalitarian agreements, and offers a criterion of how players resolve the ordinally-egalitarian solution selection problem. For an extensive discussion of ordinally egalitarian solution, see Conley and Wilkie 2012: 23-42.

³¹Ordinally symmetric games are symmetric with respect to the ordinal structure of the payoffs. For the discussion of some of the problems with BE solutions of ordinally asymmetric games, see subsection 3.3.

³²To simplify the analysis of this example, mixed strategy profiles have been omitted.

example, each player's preferences over the feasible agreements of the Impure coordination game (figure 6) can be numerically represented by assigning a rank value of 1 to the most preferred feasible agreements, the second most preferred feasible agreement can be assigned a rank value of 2, and so on. The row and the column players' personal rankings of the feasible agreements are shown below:

$$\text{Row player: } \begin{bmatrix} (u, l) : 1 \\ (d, r) : 2 \\ (m, c) : 3 \end{bmatrix} \quad \text{Column player: } \begin{bmatrix} (m, c) : 1 \\ (d, r) : 2 \\ (u, l) : 3 \end{bmatrix}$$

There are 3 preferentially differentiable contracts³³ to advance players' personal interests. If the players were to implement one of these possible agreements, they would lose the opportunity to implement the remaining ones. The rank values associated with each contract, or possible agreement, represent each player's ordinal preferences over them. A contract with an assigned preferential rank value of 1 is one of the players' most preferred possible agreement available in the game. If this contract was carried out, the player who prefers it over all the remaining contracts would not forego any opportunities to advance his or her personal interests. However, for any agreement with a rank value lower than 1, each player has at least one opportunity to advance his or her personal interests. For example, in the aforementioned coordination game (figure 6), each player has one opportunity that advances his or her personal interests more than the contract with a rank value of 2, two opportunities that advance his or her preferences more than the contract with a rank value of 2, and so on. If any of the contracts with a rank value lower than 1 was implemented, the player would forego at least one opportunity to advance his or her personal interests. For example, if contract (m, c) with a rank value of 3 was carried out, the row player would forego two opportunities to carry out contracts that s/he personally prefers over the one that was carried out – contract (u, l) with a rank value of 1, and contract (d, r) one with a rank value of 2.

Each contract with different rank values in players' personal preferential rankings is the contract over which players have conflicting preferences. For example, contract (u, l) is the best available contract for the row player, yet the second-worst contract for the column player. The difference in the rank-value that a contract has in each player's preferential ranking of outcomes can be used to compare the foregone opportunities of the players to advance their personal interests. For example, if contract (u, l) was carried out, the row player would not forego any opportunities to advance his or her personal interests, while column player would forego two opportunities to carry out contracts which s/he personally prefers over contract (u, l) . This means that, if contract (u, l) was carried out, the column player could object that this resolution of the conflict of interests is asymmetric, since it is associated with an unequal distribution of players' foregone preferred alternative agreements. The contracts which have the same rank in players' preferential rankings are the ones that, if carried out,

³³The contracts are said to be preferentially differentiable if the player is not indifferent between them. The preferentially non-differentiable contracts will count as one opportunity to advance players' personal interests.

would leave the players with equal numbers of foregone opportunities to advance their personal interests. If contract (d, r) was carried out, both players would forego one opportunity each to advance their personal interests, and it is the best feasible contract which has this property³⁴. We will call the Nash equilibrium (d, r) a benefit-equilibrating solution of the Impure Coordination game³⁵.

Notice that if the same numerical representation were used to represent both players' ordinal preferences over the agreements, then a benefit-equilibrating solution of the game will be the agreement with the lowest common rank. However, the type of the ranking function used to represent players' preferences does not matter, as long as the same ranking function is used to represent every player's ordinal preferences.

A rank-based representation of the BE solution can be formally characterized in the following way.

Let $\Gamma (I, (S_i, u_i)_{i \in I})$ be a finite normal form game where $I = (1, \dots, m)$ is a set of players, S_i is a set of pure strategies of player $i \in I$, and $u_i : (\times_{i \in I} S_i) \rightarrow \mathbb{R}$ is the utility function of player $i \in I$. Let $\Delta (S_i)$ be a probability distribution over the set of pure strategies of player $i \in I$, generating a set Σ_i of mixed strategies of player $i \in I$. A mixed strategy of player $i \in I$ will be denoted as $\sigma_i \in \Sigma_i$. A mixed strategy profile is a vector $\theta = (\sigma_1, \dots, \sigma_n)$, where each component $\sigma_i \in \Sigma_i$ is a mixed strategy of player $i \in I$. Hence, a set $\Theta = (\times_{i \in I} \Sigma_i)$ of mixed strategy profiles of the game Γ is a Cartesian product of the sets of players' mixed strategies.

In the standard bargaining problems, players' utility functions are defined over $\mathcal{L}(\Theta)$, the set of lotteries over Θ . Player i 's preference relation over $\mathcal{L}(\Theta)$ will be represented by utility function $\mu_i : \mathcal{L}(\Theta) \rightarrow R$. However, the BE solution of the game will be defined over the set of feasible agreements of the game – strategy profiles that players can actually implement via joint actions³⁶. In addition, the BE bargaining solution can be identified using purely ordinal information about players' preferences over the feasible *basic* agreements (that is, preferences over the Nash equilibria of the game), and thus it can be applied to games where players' cardinal utility functions containing information about their attitudes to risk are not available.

Let \mathcal{F} be a set of possible agreements that players believe to be feasible, with $\alpha \in \mathcal{F}$ being the typical element of \mathcal{F} . In standard bargaining games, the set of self-enforcing agreements is taken to be the set of *correlated equilibria* of

³⁴To simplify the discussion, the mixed strategy Nash equilibria have been omitted. In this game, however, each mixed strategy Nash equilibrium gives both players lower personal payoffs than pure strategy Nash equilibrium (d, r) , meaning that BE solution is the aforementioned pure strategy equilibrium.

³⁵Conley and Wilkie (2012) refer to this solution as the ordinal egalitarian bargaining solution. For details, see Conley and Wilkie 2012: 23-42

³⁶Sakovics (2004) has argued that a realistic bargaining problem should assume that players only rank the set of possible agreements that they can actually implement, and not the set of all the possible utility allocations, some of which the players may find impossible to obtain. For example, in non-cooperative game, the players cannot resolve the game with an actual lottery over the combinations of their actions. Sakovics have shown that such a restriction on the set of agreements leads to consistent ordinal bargaining solutions, which are not affected by Shapley's impossibility result. For details, see Sakovics 2004:1-7.

the game:

$\mathcal{F} = \{(\mu_1(\theta_g), \dots, \mu_m(\theta_g)) \mid \theta_g \in \Theta^{\mathcal{CE}}\}$, where $\Theta^{\mathcal{CE}} \in \mathcal{P}(\Theta)$ is the set of correlated equilibria of Γ .

If the players were able to communicate, they could indeed implement any correlated equilibrium of the game. In the context of non-cooperative games, however, the players would not be able to agree to coordinate their actions by observing some correlation device, such as a toss of a fair coin, and therefore they would not be able to implement the correlated equilibria of the game. Therefore, it seems natural to assume that the set of feasible agreements of a non-cooperative game is the set of its Nash equilibria:

$\mathcal{F} = \{(\mu_1(\theta_k), \dots, \mu_m(\theta_k)) \mid \theta_k \in \Theta^{\mathcal{NE}}\}$, where $\Theta^{\mathcal{NE}} \in \mathcal{P}(\Theta)$ is the set of Nash equilibria of Γ .

Let $\mathcal{F} = (\alpha_1, \dots, \alpha_n)$ be a finite set of feasible agreements of the game Γ . Each player $i \in I$ has a complete and transitive ordinal preference ranking \succeq_i over the set \mathcal{F} of feasible basic alternatives. For each feasible agreement $\alpha_x \in \mathcal{F}$, we can define the *cardinality of the preferred set of alternatives* for each player $i \in I$:

$$\mathcal{C}_i(\alpha_x, \mathcal{F}) \equiv \{|\mathcal{T} \mid \alpha_h \in \mathcal{T} \text{ if and only if } \alpha_h \in \mathcal{F}, \text{ and } \alpha_h \succ_i \alpha_x\}$$

In other words, the cardinality of the preferred set of alternatives is the number of feasible agreements that player $i \in I$ prefers over the agreement $\alpha_x \in \mathcal{F}$.

Let $\wp_i := \{h \in Z^+ \mid h \leq n\}$, where n is the number of agreements in \mathcal{F} , be a ranking function of player $i \in I$. The rank of each feasible agreement $\alpha_x \in \mathcal{F}$ is defined in the following way:

$$\wp_i(\alpha_x, \mathcal{F}) = [\mathcal{C}_i(\alpha_x, \mathcal{F}) + 1]$$

Notice that each players' most preferred agreement will have a rank of 1, his or her second most preferred agreement will have a rank of 2, and so on.

The cardinality of the preferred set of alternatives $\mathcal{C}_i(\alpha_x, \mathcal{F})$ represents the number of opportunities to advance personal interests that player $i \in I$ would give up if the agreement $\alpha_x \in \mathcal{F}$ was carried out. For any two feasible contracts $\alpha^x \in \mathcal{F}$ and $\alpha^y \in \mathcal{F}$, player $i \in I$ always prefers contract $\alpha^x \in \mathcal{F}$ over contract $\alpha^y \in \mathcal{F}$ if and only if $\mathcal{C}_i(\alpha_x, \mathcal{F}) < \mathcal{C}_i(\alpha_y, \mathcal{F})$. Since it is the case that $\{\mathcal{C}_i(\alpha_x, \mathcal{F}) = [\wp_i(\alpha_x, \mathcal{F}) - 1]\}$, the aforementioned condition is equivalent to the following condition: $\alpha_x \succ_i \alpha_y$ if and only if $\wp_i(\alpha_x, \mathcal{F}) < \wp_i(\alpha_y, \mathcal{F})$.

A contract $\alpha_x \in \mathcal{F}$ is ordinally equitable (that is, achieves the equality of ordinal sacrifices of opportunities) if and only if it is the case that $\mathcal{C}_i(\alpha_x, \mathcal{F}) =$

$$\mathcal{C}_j(\alpha_x, \mathcal{F}) \text{ for every pair } \begin{pmatrix} i, j \\ j \neq i \end{pmatrix} \in I. \text{ Given that } \{\mathcal{C}_i(\alpha_x, \mathcal{F}) = [\wp_i(\alpha_x, \mathcal{F}) - 1]\},$$

it follows that every ordinally egalitarian contract $\alpha^{\mathcal{E}} \in \mathcal{F}$ is such that, for every

$$\text{pair of players } \begin{pmatrix} i, j \\ j \neq i \end{pmatrix} \in I, \text{ it is the case that } \wp_i(\alpha^{\mathcal{E}}, \mathcal{F}) = \wp_j(\alpha^{\mathcal{E}}, \mathcal{F}).$$

Let $\mathcal{E}(\mathcal{F}) \in \mathcal{P}(\mathcal{F})$ be a set of ordinally egalitarian agreements of Γ , where each agreement $\alpha^{\mathcal{E}} \in \mathcal{E}(\mathcal{F})$ is such that $\wp_i(\alpha^{\mathcal{E}}, \mathcal{F}) = \wp_j(\alpha^{\mathcal{E}}, \mathcal{F})$, for all $j \neq i$

$$\binom{i, j}{j \neq i} \in I.$$

THE BENEFIT-EQUILIBRATING SOLUTION FUNCTION $f^{BE} : \mathcal{P}(\mathcal{F}) \rightarrow \mathcal{P}(\mathcal{F})$ selects a subset $\mathcal{F}^{BE} \in \mathcal{P}(\mathcal{F})$ of the set \mathcal{F} of feasible contracts of the game Γ , where each contract $\alpha^{BE} \in \mathcal{F}^{BE}$ satisfies two conditions:

1. ORDINAL EQUITY CONDITION: $\alpha^{BE} \in \mathcal{E}(\mathcal{F})$
2. MAXIMAL PERSONAL ADVANTAGE CONDITION: $\alpha^{BE} \in \underset{\alpha \in \mathcal{E}(\mathcal{F})}{\operatorname{argmin}} [\varphi_i(\alpha, \mathcal{F})] := \{\alpha^{BE} \mid \forall \alpha^E \in \mathcal{E}(\mathcal{F}) : \varphi_i(\alpha^{BE}, \mathcal{F}) < \varphi_i(\alpha^E, \mathcal{F}), \forall i \in I\}$

3.3 The Role of Threat Strategies

In standard strategic models of bargaining, bargaining is modelled as a process of negotiations where players communicate their offers to the other players, and receive their counteroffers. An important component of this process is each player's ability to reject the opponents' offers, and force them to consider their counteroffers³⁷. For the player to be able to reject the proposed contract and force the opponent to consider his or her counteroffer, s/he must have a threat strategy – a strategy which, if played, would bring the opponents' down to their disagreement payoff³⁸. Players' ability to threaten his or her opponents is also important in BER models. Note that in the aforementioned Impure Coordination game (figure 6), both the row and the column player can force the opponent to consider a counteroffer because each player in this game have a threat strategy. The column player can reject the row player's preferred contract (u, l) by threatening him or her to play strategy d . If the column player chose strategy d , the row player could not get the payoff associated with contract (u, l) , no matter what s/he did. The row player could do no better than to play strategy r in order to minimize the payoff loss.

The row player could reject the column player's preferred contract (m, c) by threatening him or her to play strategy d . If the row player played this strategy, the column player could not get the payoff associated with contract (m, c) , no matter what s/he did. The column player could do no better than to choose strategy r to minimize the loss.

If both players were to end the negotiations and carry out their threats, they would end up with outcome (d, r) , which is the BE agreement of this game³⁹.

In games where players do not have credible threat strategies, BE solution does not look particularly convincing, since at least one of the players has no

³⁷For an extensive discussion of strategic bargaining models, see Binmore 1980: 80-109, Rubinstein 1982: 97-109 and Binmore et al. 1986: 176-188.

³⁸For an extensive discussion of the role of threat strategies, see Myerson 1991: 370-416.

³⁹In this particular game, players' threat strategies would leave them with the same personal payoff as the BE solution of the game. In many other games, however, the players who carried out their threats would end up with worse payoff than the one associated with BE agreements. For example, the players who carried out their threats in the Chicken game (figure 10) would both end up with the worst possible outcome of the game.

motivation to resolve the game in a mutually beneficial way. For example, consider the ordinally asymmetric two-player "Indifference game" (figure 7):

	d	e
a	7, 5	7, 6
b	7, 3	7, 1

Figure 7: The Indifference game

This game has two pure strategy Nash equilibria (a, e) and (b, d) , and two mixed strategy Nash equilibria $(\frac{3}{4}, \frac{1}{4}; 1, 0)$ and $(\frac{3}{4}, \frac{1}{4}; 0, 1)$. Players' rankings of the feasible agreements are shown below:

$$\text{Row: } \left[\quad (\mathbf{a}, \mathbf{e}), (b, d), (\frac{3}{4}, \frac{1}{4}; 1, 0), (\frac{3}{4}, \frac{1}{4}; 0, 1) : 1 \quad \right] \text{ Col: } \left[\begin{array}{l} (\mathbf{a}, \mathbf{e}) : 1 \\ (\frac{3}{4}, \frac{1}{4}; 1, 0), (\frac{3}{4}, \frac{1}{4}; 0, 1) : 2 \\ (b, d) : 3 \\ (u, l) : 3 \end{array} \right]$$

Although this game is ordinally asymmetric, this game has a BE solution, which is the Nash equilibrium (a, e) . However, unlike the column player who prefers Nash equilibrium (a, e) over all other feasible agreements, the row player is indifferent between all the Nash equilibria of this game. The column player does not have a threat strategy against the row player: no matter what strategy the column player chose to play, the row player would always get the same payoff. It means that, in open negotiations, the column player could not use a threat strategy to reject row player's offer (b, d) , and force the row player to consider counteroffer (b, d) . Therefore, in an asymmetric bargaining game where at least one of the players has no threat strategy, the BE solution does not seem to be particularly convincing.

3.4 BE Solution and Mutual Advantage

For the players to be motivated to resolve the game by implementing a feasible agreement, it must be *mutually beneficial*. In other words, every player of the game must see the feasible agreement as personally beneficial. In the standard bargaining problem, each agreement which gives, to every player, a payoff higher than his or her disagreement payoff is considered to be mutually beneficial. However, in non-cooperative game, the question of how the concept of mutual advantage ought to be defined becomes significantly more complicated, since, in many games, it is difficult to identify each player's disagreement payoff⁴⁰.

Some economics and philosophers have suggested that players' maximin payoffs should be considered as reference points for determining whether a strategy

⁴⁰For an extensive discussion, see section 2.

profile or, in this case, self-enforcing agreements are mutually beneficial⁴¹. This suggestion seems compelling, since the maximin payoff is the maximum payoff that a player can guarantee to himself or herself, no matter what the other players are going to do⁴². It stands to reason to assume that a player would not be motivated to participate in any joint actions which would not guarantee him or her an expected payoff which is at least as good as the personal maximin payoff.

Another possibility, is that, in case of disagreement, players would try to maximize the opponent's payoff loss, even at the expense of their own personal payoff. For example, in the aforementioned Impure Coordination game (figure 6), both players could deliberately play their threat strategies, thus ending up with the worst possible personal payoffs for both. This option is less appealing from the perspective of rational strategic reasoning, since, in case of disagreement, rational players should aim to maximize their minimum disagreement payoff, and then the worst possible disagreement outcome should be a strategy profile associated with player's maximin payoffs. However, it is not a completely unreasonable assumption about the real-world bargaining, with less-than-perfectly rational players⁴³.

The BE solution suggested in this paper will always be one of the Nash or, in case communication is possible, correlated equilibria of the game. Therefore, in all games, a BE solution is an agreement which, for every player of the game, is at least as good as his or her maximin payoff. Therefore, a BE solution of every game will invariably be mutually beneficial⁴⁴.

3.5 The Axiomatic Characterization of the BE Solution

A BE solution satisfies the following axioms:

- **ORDINAL SYMMETRY:** A BE solution to every ordinally symmetric choice set is ordinally symmetric. A choice set O is said to be ordinally symmetric if and only if, for all $x, y \in O$ such that $\varphi_i(x, O) = \varphi_j(y, O)$, it holds that $C_i(x, O) = C_j(y, O)$. A solution concept $f(O)$ satisfies the axiom of symmetry if and only if, for any ordinally symmetric set O^s , it is the case that

⁴¹For example, this idea has been suggested by Myerson 1991: 370-416, Gauthier 2013: 601-624, and Sugden 2015: 143-166

⁴²For an extensive discussion of the maximin concept, see Luce and Raiffa 1957.

⁴³For an extensive discussion on disagreement points and threat strategies, see Myerson 1991: 370-416.

⁴⁴Hargreaves-Heap and Varoufakis (1995) discuss games where player's payoff associated with a unique pure strategy Nash equilibrium of the game is equal to his or her maximin payoff, yet playing a Nash equilibrium strategy is extremely risky. Hargreaves-Heap and Varoufakis argue that it is unreasonable to expect that such a player would play his or her equilibrium strategy, and therefore Nash equilibrium is not a compelling solution of the game. Since these games have a unique Nash equilibrium, it is unreasonable to expect the players to apply BER reasoning to resolve the game. For details, see Hargreaves-Heap and Varoufakis 1995: 61-68.

$\mathcal{C}_i(f(O^s), O^s) = \mathcal{C}_j(f(O^s), O^s)$, for every pair $\binom{i, j}{j \neq i} \in I$. A BE solution is always, by definition, ordinally egalitarian, meaning that, for any O , it must be the case that $\mathcal{C}_i(f^{\text{BE}}(O), O) = \mathcal{C}_j(f^{\text{BE}}(O), O)$, for every pair $\binom{i, j}{j \neq i} \in I$. It follows that, for every ordinally symmetric set O^s , a BE solution must be such that $\mathcal{C}_i(f^{\text{BE}}(O^s), O^s) = \mathcal{C}_j(f^{\text{BE}}(O^s), O^s)$ for every pair $\binom{i, j}{j \neq i} \in I$. Therefore, a BE solution satisfies the axiom of ordinal symmetry.

- **INVARIANCE UNDER ORDER-PRESERVING SCALAR TRANSFORMATIONS OF PAYOFFS.** Pure Nash equilibria are invariant under positive affine transformations of players' payoffs. However, this does not apply to mixed Nash equilibria. Since BE solutions can be mixed Nash equilibria, it follows that BE solutions are not invariant under positive affine transformations of payoffs. However, every BE solution remains stable if one or multiple players' payoffs are multiplied by positive real numbers. In other words, BE solutions are invariant under order-preserving scalar payoff transformations of the form $u'_i = \lambda_i(u_i)$, where $\lambda_i > 0$ for every $i \in I$ (for proofs, see subsection 3.6).
- **INDEPENDENCE OF IRRELEVANT ALTERNATIVES:** If any game which has a BE solution is extended with any set of strictly dominated strategies, a BE solution of the extended game remains the same as the solution of the original game. This follows from the fact that adding strictly dominated strategies to any game does not change the set of its Nash equilibria. Since a BE solution of every game is, invariably, an equilibrium, it is always invariant under any additions of strictly dominated strategies.

3.6 Existence

In this part I provide a partial characterization of the existence of BE solutions in finite, normal form games.

A BE SOLUTION EXISTS IN EVERY CARDINALLY SYMMETRIC, FINITE, NON-COOPERATIVE GAME. To see this, notice that a BE solution of the finite non-cooperative game can only be a Nash equilibrium $\theta^{\mathcal{N}\mathcal{E}} \in \Theta^{\mathcal{N}\mathcal{E}}$ which is such that $\mathcal{C}_i(\theta^{\mathcal{N}\mathcal{E}}, \mathcal{F}) = \mathcal{C}_j(\theta^{\mathcal{N}\mathcal{E}}, \mathcal{F})$, for all $i \in I$. A normal form finite game

Γ is cardinally symmetric if and only if it is the case that the players have identical pure strategy spaces ($S_1 = S_2 = \dots S_I = S$), and it is the case that $u_i(s_i, s_{-i}) = u_j(s_j, s_{-j})$, for $s_i = s_j$ and $s_{-i} = s_{-j}$ for all $(i, j) \in I$, where s_{-i} is the combination of all the strategies in profile $\mathbf{s} \in \mathbf{S}$ except for $s_i \in S_i$.

Let $\Theta^{\mathcal{N}\mathcal{E}^s}$ be a set of Nash equilibria of the symmetric game Γ . Every set $\Theta^{\mathcal{N}\mathcal{E}^s}$ of the cardinally symmetric game is ordinally symmetric: for any two Nash

equilibria $\theta_x^{\mathcal{N}\mathcal{E}} \in \Theta^{\mathcal{N}\mathcal{E}s}$ and $\theta_y^{\mathcal{N}\mathcal{E}} \in \Theta^{\mathcal{N}\mathcal{E}s}$ which are such that $\wp_i(\theta_x^{\mathcal{N}\mathcal{E}}, \Theta^{\mathcal{N}\mathcal{E}s}) = \wp_j(\theta_y^{\mathcal{N}\mathcal{E}}, \Theta^{\mathcal{N}\mathcal{E}s})$, it holds that $\mathcal{C}_i(\theta_x^{\mathcal{N}\mathcal{E}}, \Theta^{\mathcal{N}\mathcal{E}s}) = \mathcal{C}_j(\theta_y^{\mathcal{N}\mathcal{E}}, \Theta^{\mathcal{N}\mathcal{E}s})$, for all $\binom{i, j}{j \neq i} \in I$.

As a special case of Nash's (1951) theorem, every finite symmetric game has a symmetric Nash equilibrium⁴⁵. A (mixed) Nash equilibrium $\theta^{\mathcal{N}\mathcal{E}} = (\sigma_1^*, \dots, \sigma_m^*)$ is symmetric if and only if $\binom{\sigma_i^* = \sigma_j^*}{j \neq i}$, for all $i \in I$. From the definition of a cardinally symmetric game it follows that $u_i(\sigma_i^*, \sigma_{-i}^*) = u_j(\sigma_j^*, \sigma_{-j}^*)$, for all $\binom{i, j}{j \neq i} \in I$. Since the set $\Theta^{\mathcal{N}\mathcal{E}s}$ is ordinally symmetric, it follows that, for every Nash equilibrium $\theta_x^{\mathcal{N}\mathcal{E}} \in \Theta^{\mathcal{N}\mathcal{E}s}$ which is such that $u_i(\theta_x^{\mathcal{N}\mathcal{E}}) = u_j(\theta_x^{\mathcal{N}\mathcal{E}})$ for all $\binom{i, j}{j \neq i} \in I$, it must be the case that $\mathcal{C}_i(\theta_x^{\mathcal{N}\mathcal{E}}) = \mathcal{C}_j(\theta_x^{\mathcal{N}\mathcal{E}})$, and so a BE solution of the game always exists.

A BE SOLUTION EXISTS IN EVERY ORDINALLY SYMMETRIC, FINITE, NORMAL FORM GAME WHICH CAN BE TRANSFORMED INTO A CARDINALLY SYMMETRIC GAME VIA ANY POSSIBLE COMBINATION OF ORDER-PRESERVING SCALAR TRANSFORMATIONS OF PLAYERS' PAYOFFS. A game is said to be ordinally symmetric if the ordinal ranking of one player's payoffs is equivalent to the ordinal ranking of the transpose of the other player's payoffs. Every cardinally symmetric game is also ordinally symmetric, and so an infinite number of ordinally symmetric games can be derived from each cardinally symmetric game via transformations of one or multiple players' personal payoffs. Therefore, it is important to identify the type of transformations under which the BE solutions of the cardinally symmetric games are invariant. The set of pure strategy Nash equilibria is invariant under positive affine transformations of payoffs, but this is not the case for mixed Nash equilibria. However, it turns out that both pure and mixed Nash equilibria are invariant under order-preserving scalar multiplications of one or multiple players' personal payoffs.

First, we need to show that a set of mixed Nash equilibria is invariant under order-preserving positive scalar transformations of players' payoff functions. Let $\Gamma(I, (S_i, u_i)_{i \in I})$ and $\Gamma'(I, (S_i, u'_i)_{i \in I})$ be any two finite, normal form, ordinally symmetric games that differ only by some combination of positive scalar transformations of players' payoff functions. That is, in game Γ' the payoff function of each player $i \in I$ is multiplied by some positive real number $\lambda_i > 0$ such that $u'_i(\mathbf{s}) = \lambda_i u_i(\mathbf{s})$, for all strategy profiles $\mathbf{s} \in \mathbf{S}$. It is relatively easy to check that the set of mixed Nash equilibria of Γ is the same as the set of Nash equilibria of Γ' .

Let $\beta_i^\Gamma : \Theta \rightarrow \Sigma_i$ be the mixed-strategy best-reply correspondence of player $i \in I$ in Γ , which maps each mixed strategy profile to the non-empty finite set $\beta_i^\Gamma(\sigma_{-i}) = \{\sigma_i^* \in \Sigma_i : u_i(\sigma_i^*, \sigma_{-i}) \geq u_i(\sigma_i, \sigma_{-i}) \forall \sigma_i \in \Sigma_i\}$.

⁴⁵For an extensive analysis and proofs, see Nash (1951), and Cheng et al. 2004: 71-78

Let $\Gamma' (I, (S_i, u'_i)_{i \in I})$ be a game which is such that $u'_i(\mathbf{s}) = \lambda_i u_i(\mathbf{s})$, for all $\mathbf{s} \in \mathbf{S}$. We need to show that $\beta_i^{\Gamma'}(\sigma_{-i}) = \beta_i^\Gamma(\sigma_{-i})$. Notice that this will be the case if $\beta_i^{\Gamma'}(\sigma_{-i})$ is such that

$$\beta_i^{\Gamma'}(\sigma_{-i}) = \{\sigma_i^* \in \Sigma_i : \lambda_i u_i(\sigma_i^*, \sigma_{-i}) \geq \lambda_i u_i(\sigma_i, \sigma_{-i}) \forall \sigma_i \in \Sigma_i\}$$

Let $\Gamma (I, (S_i, u_i)_{i \in I})$ be a symmetric, finite, normal form game. Let $\sigma_i \in \Sigma_i$ be a mixed strategy of player $i \in I$, which assigns some probability distribution $p(S_i)$ on the set S_i of pure strategies of $i \in I$. Let $\mathcal{B} = (s_1, \dots, s_k)$, where $2 \leq k \leq n$, be any k -tuple of pure strategies with strictly positive probabilities assigned by a mixed strategy $\sigma_i \in \Sigma_i$ of player $i \in I$. Let us assume that $\sigma_i \in \Sigma_i$ is an equilibrium strategy of $i \in I$, which means that $\sigma_i \in \beta_i(s_{-i}, u_i)$. A mixed strategy $\sigma_i \in \Sigma_i$ is a best-response strategy against the k -tuple \mathcal{B} if player $j \in I$ is indifferent between the pure strategies in the k -tuple $\mathcal{B} \in \mathcal{P}(S_i)$. It means that, for every pair of pure strategies $s_b, s_c \in \mathcal{B}$, it must be the case that $u_j(s_b, \sigma_i) = u_j(s_c, \sigma_i)$ for $j \in I$.

Let $\Gamma' (I, (S_i, u'_i)_{i \in I})$ be a finite, ordinally symmetric, normal form game where $u'_j = \lambda_j u_j(\mathbf{s})$ for all $\mathbf{s} \in \mathbf{S}$. Assume that $\sigma'_i \in \Sigma'_i$ is a Nash equilibrium strategy of player $i \in I$, which means that $\sigma'_i \in \beta'_i(s_{-i}, u'_i)$. It follows that $\sigma'_i \in \Sigma'_i$ is such that, for every pair of pure strategies $s_b, s_c \in \mathcal{B}$, it must be the case that $\lambda_j u_j(s_b, \sigma'_i) = \lambda_j u_j(s_c, \sigma'_i)$ for $j \in I$. Since $\lambda_j > 0$ is a constant for a particular player (order preserving scalar multiplier), it follows that $\sigma'_i = \sigma_i$, for every $\sigma'_i \in \beta'(s_{-i}, u'_i)$, and so $\Theta^{\mathcal{N}\mathcal{E}(\Gamma)} = \Theta^{\mathcal{N}\mathcal{E}(\Gamma')}$.

Suppose that $\Gamma (I, (S_i, u_i)_{i \in I})$ is a finite, cardinally symmetric, normal form n -player game. From Nash's theorem (1951) it follows that every such game has a symmetric Nash equilibrium. Let $\Gamma' (I, (S_i, u'_i)_{i \in I})$ be a finite, normal form ordinally symmetric game derived from the cardinally symmetric game Γ , where $u'_i = \lambda_i u_i$, and $\lambda_i > 0$. Each finite, symmetric, normal form game has a symmetric Nash equilibrium, which is invariant under order-preserving positive scalar transformations of players' payoffs. Therefore, a symmetric Nash equilibrium will exist in every ordinally symmetric game derived from the cardinally symmetric game via order-preserving positive scalar transformations of players' payoffs. Since every symmetric game has a BE solution, a BE solution will exist in every ordinally symmetric game derived via order-preserving positive scalar transformations of players' personal payoffs. It also follows that a BE solution will always exist in an ordinally symmetric game which can be transformed into a cardinally symmetric game via order-preserving positive scalar transformations of players' payoffs.

IF ANY NORMAL FORM TWO PLAYER GAME CAN BE EXTENDED WITH ANY POSSIBLE CORRELATING DEVICE, A BE SOLUTION ALWAYS EXISTS IN THE EXTENDED GAME. In a cooperative setting, the set of feasible agreements is the set of the correlated equilibria of the game. An important property of the correlated equilibrium is that it is always specific to a particular correlating device and a given set of players' information partitions. Assuming that players have

access to any possible correlating device, it is always possible to construct a “lottery” over the Nash equilibria of the game, which will be a BE solution of the game.

To see why this is the case, we need to look into the concept of correlated equilibrium. Let $(\Omega, \{H_i\}_{i \in I}, p)$ be a correlating device where Ω is a (finite) space of states corresponding to the outcomes of the device with a typical element $\omega \in \Omega$, p is a probability measure on Ω , and $\{H_i\}$ is a partition of Ω of player $i \in I$. An information partition H_i assigns an $h_i(\omega)$ to each $\omega \in \Omega$ in such a way that $\omega \in h_i(\omega)$, for all $\omega \in \Omega$.

A correlated strategy of player $i \in I$ is a function $f_i : \Omega \rightarrow S_i$, which is measurable relative to information partition H_i : if it is the case that $h_i(\omega) = h_i(\omega')$, then $f_i(\omega) = f_i(\omega')$. It follows that a set of possible payoff allocations attainable under the I -tuple (f^1, \dots, f^I) of correlated strategies is a convex hull of the payoff vectors given in the payoff matrix of the game.

A strategy profile (f^1, \dots, f^I) is a correlated equilibrium relative to a correlating device $(\Omega, \{H_i\}_{i \in I}, p)$ if and only if, for every $i \in I$ and every correlated strategy \tilde{f}_i , it is the case that

$$\sum_{\omega \in \Omega} u_i(f_i(\omega), f_{-i}(\omega)) p(\omega) \geq \sum_{\omega \in \Omega} u_i(\tilde{f}_i(\omega), f_{-i}(\omega)) p(\omega)$$

It follows that any randomization over the pure or mixed strategy Nash equilibria of the game is a correlated equilibrium, and the set $\Delta^{\mathcal{CE}}$ of correlated equilibria of Γ contains the convex hull $\Delta^{\mathcal{NE}} \subseteq \Delta^{\mathcal{CE}}$ of the set of (mixed) Nash equilibria of Γ . It means that the set of correlated equilibrium payoff profiles of any game Γ is *at least as large* as the convex hull $\Delta^{\mathcal{NE}} \subseteq \Delta^{\mathcal{CE}}$ of the set of Nash equilibrium payoff profiles of Γ ⁴⁶. Notice that any convex combination of correlated equilibrium payoff profiles of Γ is a correlated equilibrium payoff profile of Γ . Assuming that players are not constrained in their choice of the correlating device, they can implement any possible randomization over the Nash equilibria of Γ ⁴⁷.

Let $\Gamma(S_1, S_2, u_1, u_2)$ be any two-player game with a set $\Theta^{\mathcal{NE}}$ of Nash equilibria. Each Nash equilibrium $\theta_x^{\mathcal{NE}} \in \Delta^{\mathcal{NE}}$ can be interpreted as a payoff profile $\mathcal{U}_{12}(\theta_x^{\mathcal{NE}}) = [(u_1(\theta_x^{\mathcal{NE}}), u_2(\theta_x^{\mathcal{NE}}))]$, which is a point of the convex hull of the set of Nash equilibrium payoff profiles. In games where players’ preferences over the Nash equilibria coincide, a BE solution of the game will be players’ preferred Nash equilibrium. The relevant case is where players have conflicting preferences over the pure Nash equilibria of the game, meaning that $\mathcal{C}_1(\theta_x^{\mathcal{NE}}) \neq \mathcal{C}_2(\theta_x^{\mathcal{NE}})$, for all $\theta_x^{\mathcal{NE}} \in \Theta^{\mathcal{NE}}$.

Let $\theta_*^{\mathcal{NE}} \in \Delta^{\mathcal{NE}}$ and $\theta_{**}^{\mathcal{NE}} \in \Delta^{\mathcal{NE}}$ be two pure strategy Nash equilibria of the game Γ . Suppose that each player 1 and player 2 prefer different pure strategy Nash equilibria of Γ , and so $\mathcal{C}_1(\theta_*^{\mathcal{NE}}, \Theta^{\mathcal{NE}}) = \mathcal{C}_2(\theta_{**}^{\mathcal{NE}}, \Theta^{\mathcal{NE}}) = 0$. Given the ranking function defined in subsection 3.2, it follows that $\wp_1(\theta_*^{\mathcal{NE}}, \Theta^{\mathcal{NE}}) =$

⁴⁶for an extensive analysis and proofs, see Aumann 1987: 1-18, and Lehrer et al. 2011: 1-6.

⁴⁷Each correlated equilibrium of Γ can be implemented with a *universal mechanism*— a correlating device $\Omega^{\mathcal{U}}$ where the set of states is the set $\Theta = (\times_{i \in I} \Sigma_i)$ of (mixed) strategy profiles of Γ .

$\wp_2(\theta_{**}^{\mathcal{N}\mathcal{E}}, \Theta^{\mathcal{N}\mathcal{E}}) = 1$. Given that players have conflicting preferences over $\theta_*^{\mathcal{N}\mathcal{E}} \in \Theta^{\mathcal{N}\mathcal{E}}$ and $\theta_{**}^{\mathcal{N}\mathcal{E}} \in \Theta^{\mathcal{N}\mathcal{E}}$, the following set of inequalities must hold for $(i = 1, 2)$:

$$\begin{aligned} u_i(\theta_*^{\mathcal{N}\mathcal{E}}) &> u_i(\theta_{**}^{\mathcal{N}\mathcal{E}}) \\ u_j(\theta_*^{\mathcal{N}\mathcal{E}}) &< u_j(\theta_{**}^{\mathcal{N}\mathcal{E}}) \end{aligned}$$

Note that the maximum payoff attainable by each player from playing a pure strategy Nash equilibrium must occur at one of the vertices of the convex hull $\Delta^{\mathcal{N}\mathcal{E}} \subseteq \Delta^{\mathcal{C}\mathcal{E}}$ of the set of Nash equilibria of Γ , which means that $\theta_*^{\mathcal{N}\mathcal{E}}, \theta_{**}^{\mathcal{N}\mathcal{E}} \in bd(\Delta^{\mathcal{N}\mathcal{E}})$. The other vertices of $\Delta^{\mathcal{N}\mathcal{E}} \subseteq \Delta^{\mathcal{C}\mathcal{E}}$ are the mixed strategy Nash equilibria associated with players' lowest personal payoffs. It follows that payoff profiles $U_{12}(\theta_*^{\mathcal{N}\mathcal{E}})$ and $U_{12}(\theta_{**}^{\mathcal{N}\mathcal{E}})$ associated with pure strategy Nash equilibria $\theta_*^{\mathcal{N}\mathcal{E}} \in \Delta^{\mathcal{N}\mathcal{E}}$ and $\theta_{**}^{\mathcal{N}\mathcal{E}} \in \Delta^{\mathcal{N}\mathcal{E}}$ are the vertices of the convex hull of the set of Nash equilibrium payoff profiles of Γ .

Each correlated equilibrium of Γ is a Nash equilibrium of the extended game Γ^+ , which is the original game Γ augmented with some correlating device $(\Omega, \{H_i\}_i, p)$. Assuming that players are unconstrained in their choice of the correlating device⁴⁸, the players can implement any correlated equilibrium $\psi^* \in \Delta^{\mathcal{C}\mathcal{E}}$ of Γ .

Suppose that players implement a correlated equilibrium $\psi^* \in int(\Delta^{\mathcal{N}\mathcal{E}})$ which is in the interior of the convex hull of the set of Nash equilibria $\Delta^{\mathcal{N}\mathcal{E}} \subseteq \Delta^{\mathcal{C}\mathcal{E}}$. It follows that any correlated equilibrium $\psi^* \in int(\Delta^{\mathcal{N}\mathcal{E}})$ is such that the payoff profile $U_{12}(\psi^*)$ is in the interior of the convex hull of the set of Nash equilibrium payoff profiles. Since it is the case that $\theta_*^{\mathcal{N}\mathcal{E}}, \theta_{**}^{\mathcal{N}\mathcal{E}} \in bd(\Delta^{\mathcal{N}\mathcal{E}})$, it follows that the following set of inequalities must hold:

$$\begin{aligned} u_1(\psi^*) &< u_1(\theta_*^{\mathcal{N}\mathcal{E}}) \text{ and } u_1(\psi^*) < u_1(\theta_{**}^{\mathcal{N}\mathcal{E}}) \\ u_2(\psi^*) &< u_2(\theta_*^{\mathcal{N}\mathcal{E}}) \text{ and } u_2(\psi^*) < u_2(\theta_{**}^{\mathcal{N}\mathcal{E}}) \end{aligned}$$

It follows that, in the extended game Γ^+ augmented with a correlating device which implements any correlated equilibrium $\psi_* \in int(\Delta^{\mathcal{N}\mathcal{E}})$, the correlated equilibrium $\psi_* \in \Theta^{\mathcal{N}\mathcal{E}^+}$ is a Nash equilibrium such that $\mathcal{C}_1(\psi^*, \Theta^{\mathcal{N}\mathcal{E}^+}) = \mathcal{C}_2(\psi^*, \Theta^{\mathcal{N}\mathcal{E}^+}) = 1$, and so it is a BE solution of Γ^+ .

IF ANY NORMAL FORM GAME CAN BE EXTENDED WITH ANY POSSIBLE CORRELATING DEVICE, A BE SOLUTION ALWAYS EXISTS IN THE EXTENDED GAME. In any n-player game, the set of correlated equilibria is at least as large as the convex hull of the set of Nash equilibria of the game. Given that players can implement any correlated equilibrium, they can implement a correlated equilibrium which is in the interior of the convex hull of the set of Nash equilibria. Therefore, each n-player game which can be extended with any correlating device will always have a BE solution.

⁴⁸This condition can be interpreted as an assumption that players can use a universal mechanism $\Omega^{\mathcal{U}}$ with any chosen probability distribution over the set $\Theta = (\times_{i \in I} \Sigma_i)$.

4 BER and the Equilibrium Selection: Examples

4.1 Equilibrium selection in 2x2 Normal Form Games

In ordinally symmetric 2x2 games where one pure strategy Nash equilibrium Pareto-dominates all other Nash equilibria of the game, such as the Hi-Lo and the Stag Hunt games, the BE solution of the game will always be the Pareto-dominant pure strategy Nash equilibrium. To see why this will always be the case, consider the following Stag Hunt game (figure 8): This game has two pure

	<i>s</i>	<i>r</i>
<i>s</i>	4, 4	1, 3
<i>r</i>	3, 1	2, 2

Figure 8: The Stag Hunt game

strategy Nash equilibria (s, s) and (r, r) , and a mixed strategy Nash equilibrium $(\frac{1}{2}, \frac{1}{2}; \frac{1}{2}, \frac{1}{2})$. Players' preferential rankings of the three feasible agreements are shown below⁴⁹:

$$\text{Row : } \left[\begin{array}{l} (\mathbf{s}, \mathbf{s}) : 1 \\ (\frac{1}{2}, \frac{1}{2}; \frac{1}{2}, \frac{1}{2}) : 2 \\ (r, r) : 3 \end{array} \right] \quad \text{Column : } \left[\begin{array}{l} (\mathbf{s}, \mathbf{s}) : 1 \\ (\frac{1}{2}, \frac{1}{2}; \frac{1}{2}, \frac{1}{2}) : 2 \\ (r, r) : 3 \end{array} \right]$$

Notice that Nash equilibrium (s, s) , which Pareto-dominates all the remaining feasible agreements has the highest rank in both players' preferential rankings of feasible agreements, meaning that it is maximally personally advantageous agreement for both players. Also, since this agreement is the best feasible agreement for both players, none of the players foregoes any opportunities to advance his or her personal interests with alternative feasible agreements. Therefore, it satisfies the ordinal equity condition, meaning that Nash equilibrium (s, s) is a BE solution of the Stag Hunt game.

In ordinally symmetric 2x2 games where players' have conflicting preferences over the pure strategy Nash equilibria of the game, the pure strategy Nash equilibria, by definition, will not satisfy the ordinal equity condition: each pure strategy equilibrium will have a different rank in players' preferential rankings of feasible agreements, meaning that, for each player, the numbers of foregone opportunities associated with each pure strategy equilibrium are not equal. Therefore, a BE solution of this type of games can only be a mixed strategy Nash equilibrium. In payoff symmetric games, a BE solution of the

⁴⁹In BER model, pure strategies are treated as "extreme" mixed strategies where one of the pure strategies is played with probability 1, while other pure strategies are played with probability 0. These extreme probability distributions will be omitted in the tables representing players' preferential rankings of outcomes in order to make the argumentation easier to follow.

game will be a symmetric Nash equilibrium. In ordinally symmetric games with asymmetric payoffs, a BE solution will be an asymmetric mixed Nash equilibrium.

For an example of a payoff symmetric game, consider the following Battle of the Sexes game (figure 9):

	<i>o</i>	<i>b</i>
<i>o</i>	6, 3	0, 0
<i>b</i>	0, 0	3, 6

Figure 9: The Battle of the Sexes game

This game has two pure strategy Nash equilibria (o, o) and (b, b) , and a mixed strategy Nash equilibrium $(\frac{2}{3}, \frac{1}{3}; \frac{1}{3}, \frac{2}{3})$. A mixed strategy Nash equilibrium is, for both players, associated with the expected payoff of 2. The row and the column players' preferential rankings of feasible agreements are shown below:

$$\text{Row : } \left[\begin{array}{l} (o, o) : 1 \\ (b, b) : 2 \\ (\frac{2}{3}, \frac{1}{3}; \frac{1}{3}, \frac{2}{3}) : 3 \end{array} \right] \quad \text{Column : } \left[\begin{array}{l} (b, b) : 1 \\ (o, o) : 2 \\ (\frac{2}{3}, \frac{1}{3}; \frac{1}{3}, \frac{2}{3}) : 3 \end{array} \right]$$

Notice that if one of the pure strategy Nash equilibria were implemented, one of the players would forego no opportunities to advance his or her personal interests with an alternative agreement, while the other player would forego one opportunity to advance his or her personal interests. It means that pure strategy Nash equilibria (o, o) and (b, b) do not satisfy the ordinal equity condition, and therefore cannot be BE solutions of this game. A mixed strategy Nash equilibrium $(\frac{2}{3}, \frac{1}{3}; \frac{1}{3}, \frac{2}{3})$ is the unique feasible agreement which satisfies the ordinal equity condition: if this outcome were implemented, each player would forego two opportunities to advance his or her personal interests with alternative contracts. Therefore, a mixed strategy equilibrium is the BE solution of the Chicken game. Notice that this example shows that BE solution may not be Pareto optimal. The expected payoff associated with mixed strategy Nash equilibrium $(\frac{2}{3}, \frac{1}{3}; \frac{1}{3}, \frac{2}{3})$ is 2 for both players of this game. This means that both players would be better off if they were playing one of the pure strategy Nash equilibria of the game. However, the payoff associated with the mixed strategy Nash equilibrium is better for both players than the payoff that they would get from carrying out their threats. Notice that, in a situation of open strategic bargaining, the row player could reject the column player's offer (b, b) by threatening him or her to play strategy o . The column player could reject the row player's offer (o, o) by threatening him or her to play strategy b . If both players were to carry out their threats, they would end up playing strategy profile (o, b) , which gives a payoff of 0 for both players. Therefore, the players would

be motivated to continue their negotiations in order to reach an agreement⁵⁰.

In most games where BE solution is a mixed strategy Nash equilibrium, the solution will not be efficient, since, in most games, the mixed strategy Nash equilibrium is inefficient. However, at least intuitively, the mixed strategy solution of this game is a fair resolution of the conflict, since, by independently randomizing their strategies, both players effectively participate in a lottery, and no player can be said to have a strategic advantage over the other player. In this sense, BE solution is symmetric⁵¹

For an example of a payoff asymmetric game, consider the following Chicken game (figure 10):

	<i>l</i>	<i>r</i>
<i>u</i>	6, 3	0, 0
<i>d</i>	5, 5	3, 8

Figure 10: The Chicken game

This game has two pure strategy Nash equilibria (u, l) and (d, r) , and an asymmetric mixed strategy asymmetric Nash equilibrium $(\frac{1}{2}, \frac{1}{2}; \frac{3}{4}, \frac{1}{4})$ ⁵². The row and the column players' preferential rankings of the feasible agreements are shown below:

$$\text{Row : } \left[\begin{array}{l} (u, l) : 1 \\ (\frac{1}{2}, \frac{1}{2}; \frac{3}{4}, \frac{1}{4}) : 2 \\ (d, r) : 3 \end{array} \right] \quad \text{Column : } \left[\begin{array}{l} (d, r) : 1 \\ (\frac{1}{2}, \frac{1}{2}; \frac{3}{4}, \frac{1}{4}) : 2 \\ (u, l) : 3 \end{array} \right]$$

A BE solution of this game is a mixed strategy asymmetric Nash equilibrium $(\frac{1}{2}, \frac{1}{2}; \frac{3}{4}, \frac{1}{4})$, which is not an efficient Nash equilibrium of this game. For each player of the game, the personal payoff associated with a mixed strategy Nash equilibrium is larger than his or her personal maximin payoff. This solution of the game is also weakly Pareto optimal⁵³. Notice that although mixed Nash

⁵⁰Notice that each player's maximin payoff of this game is 0, since it is the maximum payoff that each player could guarantee to himself or herself irrespective of the other player's choices. For an extensive discussion, see Luce and Raiffa 1957.

⁵¹In addition, the mixed strategy Nash equilibrium of the Battle of the Sexes game is a unique evolutionary stable equilibrium of this game. It means that a population where all players' resolved the Battle of the Sexes game by playing a mixed strategy Nash equilibrium could not, in evolutionary time, be invaded by mutants playing a different mixed or pure strategy. This offers, although in a very limited capacity, support to the idea that solution concepts, such as BE solution, which prescribe randomized strategies can be interpreted as evolved responses to strategic conflicts of individuals' personal interests. For an extensive discussion of evolutionary stability of mixed strategy profiles, see Hofbauer and Sigmund 1998.

⁵²The row player's expected payoff associated with mixed strategy Nash equilibrium is 4.5. The column player's expected payoff from playing mixed strategy Nash equilibrium is 4

⁵³An allocation of payoffs is said to be weakly Pareto optimal if there are no alternative allocations which would make each player strictly better off.

equilibrium has the same rank in players' preferential rankings of feasible agreements, rank-egalitarian solution does not imply that players' personal payoff gains from the agreement, relative to their maximin payoffs, will be equal. The row player's personal payoff gain, relative to his or her maximin payoff of 3, is 1.5. The column player's payoff gain, relative to his or her maximin payoff of 3, is 1.

4.2 Benefit-Equilibration and the Envy-Free Allocations of Resource

In bargaining games where players have to split a perfectly or imperfectly divisible resource, a BE solution function will always select an envy-free allocation of resource. An allocation of resource is said to be envy free if no player would swap the received amount of resource for that received by the other player (Foley 1967: 45-98). To understand this result, let us first consider the discrete Divide-the-Cake game with an even number of slices of cake, which is a particularly simple version of the well known Nash Bargaining game (figure 11):

	<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 11: The Divide-the-Cake game (even number of pieces)

In this game, two players are presented with a cake that is cut into four 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 asked for. If, on the other hand, the sum exceeds 4, they both get nothing. The game has six pure strategy Nash equilibria: (4, 0), (3, 1), (2, 2), (1, 3), (0, 4) and an inefficient (4, 4). The row and the column players' preferential rankings of outcomes are shown below⁵⁴:

⁵⁴To make the discussion easier to follow, the mixed strategy Nash equilibria will be omitted. This will not effect the result, since BE solution in this particular game is a pure Nash equilibrium

$$\text{Row player: } \left[\begin{array}{l} (4, 0) : 1 \\ (3, 1) : 2 \\ (2^*, 2^*) : 3 \\ (1, 3) : 4 \\ (0, 4), (4, 4) : 5 \end{array} \right] \quad \text{Column player: } \left[\begin{array}{l} (0, 4) : 1 \\ (1, 3) : 2 \\ (2^*, 2^*) : 3 \\ (3, 1) : 4 \\ (4, 0), (4, 4) : 5 \end{array} \right]$$

Strategy profile (2, 2) is the unique BE solution of this game. This result usually appeals to most decision-makers, and is supported by experimental results⁵⁵.

Notice that the BE solution of this particular Divide-the-Cake Game is in line with the Nash bargaining solution, as well as with the Kalai-Smorodinsky (1975) solution of this game. Assuming that the disagreement point in this game is (0, 0), the strategy profile which uniquely maximizes the product of players' personal payoff gains is (2, 2), and so it is a Nash bargaining solution of this particular game⁵⁶. Strategy profile (2, 2) is, also, a Kalai-Smorodinsky solution of the game: assuming that the disagreement point is (0, 0), it is a unique Pareto optimal profile which maintains the ratios of players' maximal, or ideal, payoff gains⁵⁷.

Let us now 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 pieces is shown in figure 12:

A BE solution of this game is a mixed strategy Nash equilibrium where each player demands 2 pieces of cake with a probability $\frac{2}{3}$, and 3 pieces of cake with probability $\frac{1}{3}$. Each player's expected number of pieces from playing this mixed strategy Nash equilibrium is 2. Unlike in the Divide-the-Cake game with an even number of slices of cake BE solution is not in line with the Nash bargaining solution, since it does not maximize the product of players' payoff gains.

The Nash bargaining solution of this game are strictly Pareto optimal pure strategy Nash equilibria (2, 3) and (3, 2). Notice that neither of these equilibria is envy-free, since the player who received 2 pieces of cake would prefer to swap his or her share for that received by the other player. A BE solution of this game is an envy-free allocation of cake, since none of the players would prefer to sap his or her expected share of cake for that expected by the other player⁵⁸.

BE solution is, in line with Kalai-Smorodinsky bargaining solution, in a sense that it maintains the ratio of players' ideal payoff gains⁵⁹. However, BE solution

⁵⁵See 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.

⁵⁶For a detailed discussion of the Nash bargaining solution, see Luce and Raiffa 1957, 124-134

⁵⁷. For a discussion of Kalai-Smorodinsky solution, see Kalai and Smorodinsky 1975: 513-518.

⁵⁸It is important to note that this solution only ensures an *ex ante* envy-free allocation of cake. The *ex post* allocation of cake may, in fact, not be envy-free.

⁵⁹Each player of this game can, in the idea case (say, when the other player deliberately chooses a strategy that maximizes his or her opponent's payoff) get a maximum payoff of 5.

	<i>0</i>	<i>1</i>	<i>2</i>	<i>3</i>	<i>4</i>	<i>5</i>
<i>0</i>	0, 0	0, 1	0, 2	0, 3	0, 4	0, 5
<i>1</i>	1, 0	1, 1	1, 2	1, 3	1, 4	0, 0
<i>2</i>	2, 0	2, 1	2, 2	2, 3	0, 0	0, 0
<i>3</i>	3, 0	3, 1	3, 2	0, 0	0, 0	0, 0
<i>4</i>	4, 0	0, 0	0, 0	0, 0	0, 0	0, 0
<i>5</i>	5, 0	0, 0	0, 0	0, 0	0, 0	0, 0

Figure 12: The Divide-the-Cake game (odd number of pieces)

is only weakly Pareto efficient⁶⁰.

Pareto efficiency is often taken to be one of the key desirable properties of the bargaining solution⁶¹. However, it is not clear whether Pareto optimality is as important factor in real-world allocation problems as it is in theoretical bargaining, especially in cases where the resource is not perfectly divisible⁶². Experimental evidence suggests that real-world decision makers tend to focus on fairness rather than efficiency considerations when dealing with resource allocation problems⁶³, which suggests that a weakly Pareto efficient allocation of an imperfectly divisible cake should not strike as an unreasonable solution of the problem.

Therefore, assuming that disagreement point is (0, 0) the ratio of players' maximum payoff gains in this game is $\frac{(5-0)}{(5-0)} = 1$. The Kalai-Smorodinsky solution of this game is a strategy profile (u_1, u_2) which is such that $u_1/u_2 = 1$, and so the BE solution satisfies this requirement. For details, see Kalai and Smorodinsky 1975: 513-518.

⁶⁰A resource allocation is said to be Pareto efficient in a weak sense if there are no feasible resource allocations which would strictly increase each player's personal payoff.

⁶¹See, for example, Luce and Raiffa 1957: 114-154, Kalai and Smorodinsky 1975: 513-518, Myerson 1977: 1631-1637, Kalai 1977: 1623-1630, Roth 1979: 775-778.

⁶²According to Weller's theorem, it is always possible to divide a perfectly divisible cake among n players with additive value functions in a way that is both Pareto efficient and envy-free. See Weller 1985: 5-17.

⁶³An experimental study conducted by Herreiner and Puppe suggests that people are willing to sacrifice Pareto efficiency in order to reach envy-free allocations of resource in fair division problems. See Herreiner and Puppe 2009: 65-100.

5 BER and Social Coordination

5.1 Social Interactions and the Set of Feasible Agreements

The formal BER model rests on the assumption that, in non-cooperative setting, the set of feasible agreements is the set of the Nash equilibria of the game. Theoretically, this assumption seems reasonable, since Nash equilibria are stable solutions of the game: if the players who expressed a common belief in rationality⁶⁴ were to agree on playing a Nash equilibrium, each player would have an assurance that his or her opponents will carry out their part of the agreement.

In real-world social interactions, however, the players may be using a different criterion of assurance in identifying the feasible agreements of the game. For example, the players may only be concerned about the personal payoff losses associated with opponents' deviations from the agreement, and therefore deem feasible all the agreements where opponents' deviations cannot lead to a loss of their personal payoff. If social agents used such a feasibility criterion, some of the non-equilibrium strategy profiles of the game would actually 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 (figure 13).

	<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 13: The Boobytrap Game

Like the Prisoner's Dilemma game, the Boobytrap game has a unique pure strategy Nash equilibrium (*defect, defect*). However, each player has an additional option of playing a *boobytrap* strategy. If one of the players chooses this strategy, the other player's best-response strategy is *cooperate*, while the choice of the strategy *defect* is associated with player's worst personal payoff in this game. If both players choose strategy *boobytrap*, each of them gets a payoff higher than the one associated with the Nash equilibrium (*defect, defect*).

⁶⁴A player expresses a common belief in rationality if s/he believes that every player is rational, believes that every player believes that every player is rational, believes that every player believes that every player believes that every player is rational, and so on *ad infinitum*. A common belief in rationality implies that neither of the players believes that his or her opponents will choose a strategy that is not a best response to his or her beliefs about the opponents' strategy choices. It also implies that neither of the players believes that his or her opponents will ever play a strictly dominated strategy. For a formal discussion of this concept, see Perea 2012: 68-124.

Strategy profile $(defect, defect)$ is not a Nash equilibrium, since both players have an incentive to deviate from this profile by choosing the strategy *cooperate*. However, any individually advantageous unilateral deviation from this strategy profile does not harm any of the players of this game: if one of the players chose to deviate from the agreement $(boobytrap, boobytrap)$, the player who plays the strategy *boobytrap* gets the same payoff of 29, while the deviating player gets a slightly better payoff of 30.

The experimental results seem to suggest that players identify the strategy profile $(boobytrap, boobytrap)$ as a feasible agreement of the game: the majority of participants ended up playing strategy profiles $(boobytrap, cooperate)$, $(cooperate, boobytrap)$ and $(cooperate, cooperate)$, and such behaviour is only consistent in the light of a belief that the opponent will choose a *boobytrap* strategy⁶⁵.

Misyak and Chater’s suggestion has two important implications. First, if social agents really use a less restrictive criterion of feasibility then, in some games, the set of feasible agreements becomes broader, and the BE solution of the game may be a non-equilibrium strategy profile of the game. For example, consider the game suggested by Hargreaves-Heap and Varoufakis (figure 14):

	<i>C1</i>	<i>C2</i>	<i>C3</i>
<i>R1</i>	1, 1	100, 0	−100, 1
<i>R2</i>	0, 100	1, 1	100, 1
<i>R3</i>	1, −100	1, 100	1, 1

Figure 14: A Game with a Risky Nash Equilibrium

This game has a unique pure strategy Nash equilibrium $(R1, C1)$. Note, however, that the row player can secure, with absolute certainty, the same payoff by playing his or her maximin strategy *R3*, and the column player can secure the same payoff by playing maximin strategy *C3*. Playing a Nash equilibrium

⁶⁵The experimental results reveal that more than 48% of participants ended up playing strategy profile $(cooperate, cooperate)$, 30% of participants ended up playing strategy profiles $(boobytrap, cooperate)$ and $(cooperate, boobytrap)$ and 4% ended up playing strategy profile $(boobytrap, boobytrap)$. There were no players who ended up playing the Nash equilibrium $(defect, defect)$. Choosing a strategy *cooperate* is a best-response to a belief that the opponent will play strategy *boobytrap*, but this choice is not rational in the light of a belief that the opponent will play an equilibrium strategy $(defect)$. Some of the players’ decision to stick to the strategy *boobytrap* can be explained by a bit more sophisticated considerations of strategic risk: if the opponent predicts that the player will deviate from the profile $(boobytrap, boobytrap)$, s/he can best respond by playing strategy $(defect)$ (this happened in 13% of the observed cases). A player who expects the opponent to expect his or her deviation is better off by playing the strategy *boobytrap*. For a detailed discussion of the experimental results, see Misyak and Chater 2014: 1-9

strategy is risky: if the opponent plays his or her maximin strategy, the player who plays a Nash equilibrium strategy will end up with the worst possible personal payoff. Therefore, it seems that Nash equilibrium $(R1, C1)$ is not a particularly compelling solution of this game, and it seems reasonable to expect that rational players will end up playing their maximin strategies⁶⁶.

Notice that maximin strategy profile $(R3, C3)$ satisfies the feasibility criterion suggested by Misyak and Chater. If one of the players chose his or her maximin strategy, the other player would have an incentive to defect. However, any personally advantageous deviation from the maximin strategy profile would not affect the payoff of the player playing the maximin strategy. If both players were to deviate from the agreement, they would end up playing strategy profile $(R2, C2)$, which, for each player, is associated with the same payoff as the one associated with the maximin strategy profile $(R3, C3)$.

If the players looking for a BE solution of the game were to adopt the aforementioned feasibility criterion⁶⁷, they would identify both the Nash equilibrium $(R1, C1)$ and the maximin strategy profile $(R3, C3)$ as BE solutions of this game, and would therefore face a second-order coordination problem. In a situation of such strategic uncertainty, the players should assign a positive probability to the event where the players fail to coordinate their strategy choices (that is, the events where players end up playing either strategy profile $(R1, C3)$ or profile $(R3, C1)$), and so they should play the maximin strategy profile $(R3, C3)$, rather than Nash equilibrium $(R1, C1)$.

Second, the BE solution may be relevant in explaining social agents' out-of-equilibrium choices in games with inefficient Nash equilibria. The Boobytrap game is an example of a game with a unique inefficient Nash equilibrium. If players looking for a BE solution of the game were to adopt the feasibility criterion suggested by Misyak and Chater, they would identify a more efficient feasible solution of the game – strategy profile $(boobytrap, boobytrap)$.

5.2 The Role of BER in Cooperative and Repeated Social Interactions

So far, the discussion of the BE solution was based on the assumption that players operate in non-cooperative context where communication is not possible. In such non-cooperative contexts, the set of feasible agreements is limited to the set of Nash equilibria of the original game. In most games where players have conflicting preferences over the pure strategy Nash equilibria, the only feasible

⁶⁶for an extensive discussion of this problem, see Hargreaves-Heap and Varoufakis 1995: 61-68.

⁶⁷A thorough analysis of the reasons of why social agents would adopt the BER approach in games with unique Nash equilibria fall outside of the scope of this paper. However, experimental evidence suggests that people seem to be able to compare the Nash equilibria of the game in terms of their efficiency, or in terms of risk associated with playing a Nash equilibrium strategy (see, for example, Colman and Stirk 1998: 279-293). It seems reasonable to assume that players could identify the inefficiency of a unique Nash equilibrium by comparing it with other strategy profiles of the game, and then search for a more efficient or less risky feasible solutions.

BE agreement is a mixed strategy Nash equilibrium. In general, mixed strategy Nash equilibria tend to be inefficient, and so the benefit-equilibrating solution of such games is, in general, inefficient.

There are few real world interactions where individuals face one-shot non-cooperative games in game theoretic sense. In most social interactions, players have some means of communicating with each other, and/or they interact on a more or less regular basis. Even a relatively limited level of communication enables the players to implement the correlated equilibria of the game, thus making coordination much more efficient.

For example, consider the aforementioned Battle of the Sexes game (figure 9). In non-cooperative environment, the only feasible BE solution of the game is a mixed Nash equilibrium. This solution is clearly Pareto inefficient: for both players, the expected payoff from playing a mixed Nash equilibrium is 2 – lower than the expected payoff that they could get from playing either of the pure strategy equilibria.

In cooperative setting, however, players could implement a more efficient BE solution by agreeing to play each pure strategy Nash equilibrium with probability $\frac{1}{2}$. The expected payoff of each player from implementing this agreement would be 4.5, and this BE solution would be weakly Pareto optimal. For example, two individuals could agree on a convention that, whenever they face a Battle of the sexes game, they will toss a fair coin, and follow a pre-agreed coordination plan. The players could also *roughly approximate* the BE correlated equilibrium by agreeing to coordinate their actions by date (with even days representing one pure strategy Nash equilibrium, and odd days representing another) or even by clock (with even hours representing one pure strategy Nash equilibrium, and odd numbers representing another). Given the context where players can communicate and establish a coordination plan, the list of correlating devices that players could use for implementing a correlated equilibrium is virtually endless.

Given this possibility, it is natural to ask whether benefit-equilibration plays any role in real-world social interactions, where players can communicate and/or engage in repeated interactions with each other.

A conclusive answer to this question cannot be given without an extensive empirical study, yet very basic empirical observations of social agents' behaviour suggest that benefit-equilibration does play a role in social interactions. People often face situations where their personal preferences over the feasible combinations of joint actions do not perfectly coincide. For example, friends and couples repeatedly face situations which are structurally similar to the Battle of the Sexes game (figure 9): each individual wants to engage in joint activities with the other individual, yet at the same time individuals have conflicting preferences over the types of activities which are available to them. If individuals were only concerned with coordination, they could easily establish a rule that, in every Battle of the Sexes type of interaction, one of the individuals is a "dictator", whose preferences always take priority over the preferences of others. In other words, social agents could establish a convention that, in every Battle of the Sexes game, a pure strategy Nash equilibrium of the game preferred by

a particular individual or a particular type of individual (say, a person with a special social status) will always be played. Such convention would ensure a successful and strictly Pareto efficient coordination of players' actions.

There is no denying that such conventions do exist, especially when it comes to individuals who stand in a relationship that they deem asymmetric. For example, in many cultures, it is not uncommon that an individual considered to be the head of the family makes the final choice in situations where family members have conflicting preferences over alternative courses of joint actions.

On the other hand, in social interactions where players perceive their roles as symmetric, such conventions are far less prevalent. Friends or couples rarely follow a convention that one individual's preferences always have a decisive role in determining the course of joint actions. In a considerable number of cases, individuals resolve such conflicts of interests either by using an improvised "lottery", such as the Rock-Paper-Scissors game, or, in cases where such conflicts are expected to be recurring, by taking turns at playing the role of a "dictator" whose preferences over the alternative courses of action determine the joint actions of the group. For example, friends and couples often take turns in deciding, what kind of activities they will engage in together on their leisure time. Such randomized correlated equilibria are not strictly Pareto optimal (although they may be weakly Pareto optimal). The social agents who deem their relationship symmetric may be motivated to find a resolution of the conflict which would ensure a symmetric allocation of the benefits associated with the feasible joint actions. This may be the reason of why players choose to randomize between the pure strategy Nash equilibria, rather than play one of the pure strategy Nash equilibria of the game. In games where players cannot meaningfully compare their personal payoff gains, an ordinally egalitarian solution seems to be one of the simplest and intuitive symmetric solutions of players' conflicts of personal interests.

6 Conclusion

Misyak and Chater's Virtual Bargaining theory, which relies on Nash bargaining solution, is the first attempt to apply the hypothetical bargaining theory to a general class of non-cooperative games. In this paper I have proposed some arguments against the use of the standard Nash bargaining solution for modelling players' hypothetical reasoning in non-cooperative games. My arguments focused on the standard bargaining models' reliance on the existence of a unique reference point – a condition which is not satisfied by games with multiple Nash equilibria – as well as the failure of the Nash bargaining theory to differentiate potential agreement points in terms of players' personal benefit allocations associated with them. I attempted to fill this gap by proposing a model of Benefit-Equilibration Reasoning, which does not depend on the existence of a unique non-agreement point, and offers an explanation of how players can resolve the conflict of their personal interests in games where interpersonal comparisons of payoffs are not meaningful. I have shown that BE solution can be applied

to non-cooperative games where payoffs are assumed not to be interpersonally comparable.

While the proposed model seems to offer an intuitively compelling story of how social agents resolve non-cooperative games, further empirical tests will need to be constructed to test its empirical validity. In principle this task is possible, since BER model provides testable predictions in many experimentally relevant games.

The second area requiring further research concerns cases where the proposed model does not yield a unique solution. As an example, consider the Extended Hi-Lo game (figure 11):

	<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

Figure 15: The Extended Hi-Lo game 1

There are two BE solutions of this game – pure strategy Nash equilibria ($Hi1, Hi1$) and ($Hi2, Hi2$), and so the ordinal bargainers would face a second-order coordinations problem. The model proposed in this paper does not offer an answer of how players would coordinate their actions in such games.

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 the feasible and mutually beneficial equilibria in non-cooperative games. In games with a unique BE solution, the ordinal bargainers can coordinate their actions merely by identifying the BE solution of the game. In game with multiple BE solutions, however, the bargainers would have to use some additional decision rules to coordinate their actions. For example, they could choose their BE strategies randomly (that is, they could play their BE strategies with equal probabilities), or they could considering those ordinally egalitarian equilibria which are not maximally individually advantageous. For example, in the aforementioned extended Hi-Lo game, BE deliberators could choose the egalitarian Nash equilibrium (Lo, Lo) which, although less personally advantageous for both players than Nash equilibria ($Hi1, Hi1$) and ($Hi2, Hi2$), is a unique second-best option for both players. By playing their part in realizing a Nash equilibrium (Lo, Lo), both players could achieve a sure payoff of 9, which is higher than the expected payoff of 5 from choosing strategies $Hi1$ and $Hi2$ randomly.

The BER model suggested in this paper could, in principle, be modified to include uniqueness into the formal characterization of the BE solution. In that

case, the outcome (Lo, Lo) would be the unique BE solution of this game.

However, such *ad hoc* modifications would be conceptually problematic. First, the solution could not be taken to represent the outcome of the actual bargaining process, since, in open negotiations, players would definitely agree on playing either the Nash equilibrium $(Hi1, Hi1)$, or the Nash equilibrium $(Hi2, Hi2)$.

Second, such a modification of the suggested model would not resolve the problem in every possible scenario. For example, consider the following version of the Extended Hi-Lo game:

	<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	5, 5

Figure 16: The Extended Hi-Lo game 2

In this game, the players would get the same expected payoff from randomizing between strategies $(Hi1)$ and $(Hi2)$ as they would get from playing Nash equilibrium (Lo, Lo) . The question of what the players would choose to do in such situations of strategic uncertainty cannot be answered with the tools of the theory suggested in this paper, and further research into the psychological factors that influence players' belief formation process may be necessary to explain coordination in such games.

References

- Aumann, R. J. [1987]: 'Correlated Equilibrium as an Expression of Bayesian Rationality', *Econometrica*, pp.1-18.
- Aumann, R. J. and Brandenburger, A. [1995]: 'Epistemic Conditions for Nash Equilibrium', *Econometrica*, 63, pp. 1161-1180.
- Bacharach, M. [2006]: *Beyond Individual Choice: Teams and Frames in Game Theory*, Princeton: Princeton University 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 Behaviour', *Econometrica*, 52, pp. 1007-1028.

- Binmore, K. G. [1980]: 'Nash Bargaining Theory I', International Center for Economics and Related Disciplines, London School of Economics, Discussion Paper 80-09.
- Binmore, K. G., Rubinstein, A., Wolinsky, A. [1986]: 'The Nash Bargaining Solution in Economic Modelling', *RAND Journal of Economics*, pp. 176-188.
- Cheng, S. Reeves, D. Vorobeychik, Y. Wellman, P. [2004]: 'Notes on Equilibria in Symmetric Games', *Proceedings of the 6th International Workshop On Game Theoretic And Decision Theoretic Agents GTDT 2004*, pp. 71-78.
- Colman, A. M. and Stirk, J.A. [1998]: 'Stackelberg Reasoning in Mixed-Motive Games: An Experimental Investigation', *Journal of Economic Psychology*, 19, 279-293.
- Conley, J. P., Wilkie, S. [2012]: 'The Ordinal Egalitarian Bargaining Solution for Finite Choice Sets', *Social Choice and Welfare*, 38, pp. 23-42.
- Fehr, E. and Schmidt, K. M. [1999]: 'A Theory of Fairness, Competition, and Cooperation', *Quarterly Journal of Economics*, 114, pp. 817-868.
- Foley, D. [1967]: 'Resource Allocation and the Public Sector', *Yale Economic Essays*, 7, pp. 45-98.
- Herreiner, D. K. and Puppe, C. D. [2009]: 'Envy-Freeness in Experimental Fair Division Problems', *Theory and Decision*, 67, pp. 65-100.
- Hofbauer, J., Sigmund, K. [1998]: *Evolutionary Games and Population Dynamics*. Cambridge: Cambridge University Press.
- Kalai, E. and Smorodinsky, M. [1975]: 'Other Solutions to Nash's Bargaining Problem', *Econometrica*, 43, pp. 513-518.
- Kalai, E. [1977]: 'Proportional Solutions to Bargaining Situations: Interpersonal Utility Comparisons', *Econometrica*, 45, pp. 1623-1630.
- Lehrer, E. Solan, E., Viossat, Y. [2011]: 'Equilibrium Payoffs of Finite Games', *Journal of Mathematical Economics*, 2011, 47, pp. 48-53.
- Luce, R. D. and Raiffa, H. [1957]: *Games and Decisions: Introduction and Critical Survey*, New York: Dover Publications, Inc.
- Misyak, J. and Chater, N. [2014]: 'Virtual Bargaining: A Theory of Social Decision-Making', *Philosophical Transactions of the Royal Society B*, 369, pp. 1-9.
- Myerson, R. B. [1977]: 'Two-Person Bargaining Problems and Comparable Utility', *Econometrica*, 45, pp. 1631-1637.

- Nash, J. [1951]: 'Non-Cooperative Games', *Annals of Mathematics*, 54, pp. 286–295.
- 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.
- Pearce, D. [1984] 'Rationalizable Strategic Behaviour and the Problem of Perfection', *Econometrica*, 52, 1029-1050.
- Perea, A. [2012]: *Epistemic Game Theory: Reasoning and Choice*. Cambridge: Cambridge University Press.
- Roth, A. [1979]: 'Proportional Solutions to the Bargaining Problem', *Econometrica*, 47, pp. 775-778.
- Rubinstein, A. [1982]: 'Perfect Equilibrium in a Bargaining Model', *Econometrica*, 50, pp. 97-109.
- Weller, D. [1985]: 'Fair Division of a Measurable Space', *Journal of Mathematical Economics*, 14, pp. 5-17.