Virtual Bargaining
A new Disagreement Point and Interpretation

Matthias Neudert, 15.09.2020
E-Mail: md.neudert@gmx.de

Abstract

Virtual bargaining tries to explain why some actions suggested in games by best response reasoning appear unreasonable and why real players often succeed easily in some coordination problems, although orthodox game theory is unable to resolve them. Yet the original account of virtual bargaining was lacking a proper mathematical formalism, a convincing motivation for players to follow the reasoning as well as a deliberate notion of the feasible agreements and the disagreement point. However, the ideas of two recent papers on virtual bargaining, one of them yet unpublished, can overcome some of the initial problems, while also raising new issues.

I will argue here that the latest account rests on an implausible disagreement point, an in general unintuitive assumption of non-spiteful best responses and an inconsistent definition of the worst payoff function. Yet I propose a new disagreement point, which I argue to be convincing and which justifies a slightly weaker assumption than non-spitefulness for an arbitrary number of players such that virtual bargaining accounts for the phenomena, it tries to explain. Moreover, I will adequately generalize the worst payoff function and the virtual bargaining equilibrium (VBE) for n-player games, while also outlining the epistemic conditions for a VBE to be chosen.

Acknowledgements

I'm very grateful to Dr. Jurgis Karpus for supervising this thesis.

Contents

Chapter 1

Introduction

Although orthodox game theory is accepted by many philosophers and economists, there are still puzzling phenomena. For example some actions suggested in games by best response reasoning seem unreasonable and real players easily manage certain coordination problems, while the current theory can't resolve them. In sequential-move games not all players choose a subgame perfect solution, but precommit to strategies with certain principles, even though they can receive lower payoffs by doing so. Some deviations from orthodox game theory seem reasonable and suggest that something is missing.

To account for the above phenomena in game solutions, numerous alternative reasonings for players' choices were proposed until now, one of which is *virtual bargaining*. It's founded on the idea that players can imagine a bargaining situation to solve interactions, even when they're unable to communicate. By considering possible strategies of the other players in their mind they search for agreements called *virtual bargains*, on which they would coordinate.([8],pp.1ff)

In the past virtual bargaining was able to resolve certain coordination problems for mutual advantage, while maintaining the players as individual reasoners. Yet it fails to account for precommitment and reasonable deviation from subgame perfect strategies in sequential-move games as well as to solve exploitation dynamics in favor of mutual advantage. The reason has been that the definition of feasible agreements didn't fully represent players' intuitions. The account was further lacking a universal notion for the disagreement point of the virtual bargaining process, a well-considered mathematical framework and sufficient motivation to follow the reasoning. Recent papers on virtual bargaining ([10, 11]), one of them yet unpublished, were able to extend the account and provide suggestions for some deficits.

I will argue in this paper that the new theory has the potential to account for the above phenomena, but that many of the underlying assumptions and definitions are implausible. To resolve the remaining problems, I will propose a new definition for the disagreement point, which motivates best response reasoners to adopt virtual bargaining and justifies a slightly weaker version of non-spitefulness for an arbitrary number of players. Moreover, I will consistently generalize the worst payoff function and the virtual bargaining equilibrium for n-player games, while outlining the epistemic conditions for a VBE to be chosen. This modified virtual bargaining theory is cleared of many challenges and describes the above

phenomena.
The guideline for the paper is the following. In chapter 2 I introduce the current account of virtual bargaining, explain its mathematical definitions and demonstrate the related reasoning in game examples. In chapter 3 I compare the existing notions of feasible agreements to show which one is suited best to describe the phenomena. Moreover, I will discuss the major problems of the account concerning the justification of the definitions and its accordance with best response reasoning. In chapter 4 I solve these problems by proposing a new disagreement point to motivate virtual bargaining and by adjusting non-spitefulness. Then I will generalize the worst payoff function and the virtual bargaining equilibrium (VBE) for n players, while outlining the epistemic conditions for a VBE. Chapter 5 serves as a summary and an outlook to future research.

Chapter 2

The Virtual Bargaining Account

The idea of virtual bargaining comes from an approach of cognitive science to enrich economics and game theory with more adequate views on agents' behavior in social interaction.[5] Yet phenomena like coordination between agents for mutual advantage remain not entirely understood in cognitive science as well as game theory. Virtual bargaining can help to understand and model many processes of social interaction, assuming that they fundamentally depend on *virtual bargains*.[9]

In general, a virtual bargain describes the idea of a focal point for interaction problems, which agents bear as a proposition in their mind to solve many challenges of day to day life without communication. Take for example two agents trading. If you want something I have and I want something you have, this is a bargaining problem for mutual advantage, where our interests are not aligned. In reality incomplete knowledge often raises the need for actual bargaining. But if both players are equally rational and the whole situation is clear, players often propose a solution right away — here exchanging goods of equal worth. The account describes such virtual bargains within the framework of game theory.

In this chapter I want to introduce the current state of the (formal) virtual bargaining account. The goal is to show the theory's definitions and assumptions as well as its predictions in game examples. However, the presentation will lack certain justifications and motivations because the authors of virtual bargaining provide insufficient answers so far. Instead, these questions and the arising problems are discussed in chapter 3.

Currently, the theory consists of two elements: the notion of feasible agreements and a solution concept to select one of these agreements, which is proposed to be the Nash bargaining solution. A third and rather understated element is the disagreement point, necessary to identify a virtual bargain.[8, 9]

2.1 An Example of Virtual Bargaining

Before I present the formal account, it's instructive to discuss a simple example to illustrate the need of virtual bargaining in game theory. Let's start by looking at the the famous "Prisoner's Dilemma", which is given for two players in normal-form by the payoff matrix in table 2.1.[1] In pure strategies it's known from many discussions of the game that for rational players choosing D guarantees them a higher payoff no matter what the other player chooses. As D therefore strictly dominates C, both players will choose D, which yields [1, 1] as the outcome of the game. It's the unique Nash equilibrium[2] (NE), although many players would prefer coordinating on (C,C) for mutual advantage. Yet in the Prisoner's Dilemma this seems unreasonable as long as the opponent rationally maximizes his expected payoff and plays a best response.([15], pp.20f/44)

	C	D
C	3, 3	0, 4
D	4, 0	1, 1

Table 2.1: A Prisoner's Dilemma

But now consider the "Boobytrap Game" (table 2.2), an extended Prisoner's Dilemma discussed in the virtual bargaining research.([8], p.3) In this game there's the additional strategy B, which allows to punish defection at the cost of one's own payoff. Yet here again the only rational solution by best response reasoning is the NE (D,D). Reason is that C strictly dominates B and if B is eliminated from the rationally permissible strategies, the remaining subgame is just a Prisoner's Dilemma. This procedure of "iterated deletion of strictly dominated strategies" is a common method of best response reasoning to determine rational game solutions.([16], p.58)

	C	D	B
C	30, 30	10, 40	30, 29
D	40, 10	20, 20	-100, 9
B	29, 30	9, -100	29, 29

Table 2.2: The Boobytrap Game

But both players would prefer (B,B) for mutually advantageous payoffs. Playing (C,C) is irrational because you have to expect that your opponent might exploit you. Yet (B,B) seems a feasible agreement, as no player can gain and make the other worse off. The possibility to change to C doesn't hurt anybody — if both try to best respond to (B,B), they even reach (C,C). It doesn't seem intuitive to believe that D is the only rational choice if both players could gain by coordinating on (B,B). In expression, it's rational to choose B or C if you believe the other player chooses B. Chater et al. conducted an experiment on real players' choices in the Boobytrap Game and found that less than 10% of the participants had chosen D.([8], pp.5f) If the rational suggestion of orthodox game theory doesn't describe human behavior, we need a different kind of game solution. Moreover, there are many phenomena of coordination problems, mutual advantage and precommitment in sequential move games, which can't be explained by best response reasoning, but which virtual bargaining can describe.

[1] The row player receives the left number and Column the right one, respectively.

[2] A Nash equilibrium is defined by best response as where no player can gain by a unilateral change of strategy.([13]; [16], p.58)

2.2 Current Definitions of feasible Agreements

As a first approach to the formal account let's focus on the definition of the "feasible agreements". This set of strategy profiles consists of all the possible candidates for a virtual bargain, on which the players might be able to agree. Every player on her own has to consider the available actions of all players, while asking herself: what would be feasible for everybody, including me? Since all players try to anticipate the other's situation as well, the individually attained set of feasible agreements is then supposed to be common knowledge. In the following I will introduce the two proposals of feasible agreements in the virtual bargaining research.

2.2.1 Not-exploitable Agreements

In the early developement of the theory any possible virtual bargain was required to be not-exploitable such that no player could gain on the expense of another player by varying her strategy. Such an agreement seemed sensible because no single player deviation could harm any other player and if the virtual bargain is appealing, then no player would want to best respond. The concept was defined for two players ([8], p.4):

Definition 2.2.1 (Not-exploitable Agreements)**.** Any (pure or mixed)[3] strategy profile is called *feasible* if no player can increase her payoff by a unilateral change of strategy, while reducing the other player's payoff.

Note that for $n > 2$ players there's a problem with this kind of feasibility because of the best responses available to the players; I will discuss this in section 3.3.
For an example look again at table 2.1. Since the game is popular for its exploiting betrayal, (D,D) with [1, 1] is the only not-exploitable strategy profile. Try for example (C,C), which would yield the Pareto-superior outcome [3, 3], but both players could instead get a payoff of 4 by changing to D, while the other stays at C. Varying strategy in this way exploits the other player and thus (C,C) isn't feasible.
Note that in the Prisoner's Dilemma the feasible agreement is also a NE. In general, it's true that every NE is already not-exploitable. If nobody would change strategy in a NE, as there's no way to gain, exploitation is impossible. On the other hand there are feasible agreements, which aren't NEs, so the not-exploitable (feasible) agreement is a weaker definition, yielding more strategy profiles.([8], p.4) To illustrate, the strategy profile (B,B) in the Boobytrap Game is a not-exploitable agreement and it's not a NE.

2.2.2 Worst-payoff Agreements

The other proposal of feasible agreements for virtual bargaining is to view players as less pessimistic, when thinking about a strategy profile. It's not necessary to exclude every possibility of exploitation, but it's enough to look out for your worst

[3] It means a single action or a probability distribution over several actions.

payoff. This means an agreement is feasible if no player can increase her worst payoff by varying her strategy unilaterally. The worst payoff for a player related to some strategy profile is given by the minimum payoff of two cases: either all other players stick to the agreement *or* (at least) one of them plays his best response to it, while the rest follows the agreement. The idea is that what counts in a feasible agreement isn't that nobody can exploit it. But a player needs to be willing to accept a lower payoff, as she gets exploited, because it's still better than any of her alternatives. A player always thinks of a possible agreement as a hypothetical case, where every player has chosen a certain strategy. Then every player compares her worst payoff in that strategy profile to her worst payoff, when she would change (only) her strategy. A strategy profile is feasible if nobody can increase her worst payoff by changing.
For more clarity of the reasoning let's review table 2.1 and examine if (C,C) is worst-payoff feasible. Row would then ask herself: "What scenario is worse for me? Column sticking to C, when I play C, or playing his best response? Because D is better for Column no matter what I choose, his best response will be D, while I play C. But then my payoff is 0, which is worse than 3." Indeed 0 is the worst payoff regarding (C,C) for both players and they search to increase it. If you examine (D,C), you recognize that for Column 0 is already the worst possible, but for Row the worst payoff is 1, when Column best responds with D — similarly for Column in (C,D). So for Row the worst payoff of 1 in (D,C) is larger than the worst payoff of 0 in (C,C), which she can't increase further, as there's no other strategy to change to; (C,D) for Column, respectively. Hence, (D,C) is feasible for Row, but not Column, and (C,D) is feasible for Column, but not Row. However, in (D,D) both players already receive their worst payoff, as sticking to the agreement and best responding to it are the same thing. Can they increase it? No because the worst payoff of (C,D) for Row is 0, which is lower than 1, and (D,C) with 0 is worse for Column than (D,D). So (D,D) is a feasible agreement.

The formal definition of this type of feasible agreements relies on a player's worst payoff function, which is defined for two players as ([11], pp.6f):

Definition 2.2.2 (Worst Payoff Function)**.** Let $\sigma_i \in \Sigma_i$, $i = 1, 2$ be player i's strategy and $u_i(\sigma_i, \sigma_{-i})$ player i's payoff function, where σ_{-i} denotes the strategy of i's opponent. The worst payoff of an agreement $\sigma^A = (\sigma_1^A, \sigma_2^A)$ for player i is defined as

$$w_i(\sigma_i^A, \sigma_{-i}^A) = \min \left\{ u_i(\sigma_i^A, \sigma_{-i}^A), \max_{\sigma_{-i} \in R_{-i}(\sigma^A)} u_i(\sigma_i^A, \sigma_{-i}) \right\}, \tag{2.1}$$

where $R_{-i}(\sigma^A)$ states the set of player $-i$'s best responses to the strategy profile σ^A.

This definition looks at a strategy profile (σ_1^A, σ_2^A) as a possible agreement. Then every player i defines the worst payoff related to σ^A as the minimum of either both players following the agreement or the case where i's opponent plays a best response to it. From all such best responses $-i$ chooses the one such that player i's payoff is largest (maximum). In other words, even though $-i$ might play what's best for him, i expects him to be nice and give her the maximum payoff

available, while doing so.([10], p.5602) Note that a worst payoff is always lower than or equal to a player's actual payoff.
For a simple example look at table 2.3 and the strategy profile (C,C). Here Column has two possible best responses, while Row plays C; Column can either play C or D. In both cases he gets 4, but if he plays D, Row gets 0. The worst payoff definition states that Row should expect Column to choose C here, because Column gets 4 anyway, but he can ensure the maximum payoff of 4 (instead of 0) for Row.

	C	D
C	4, 4	0, 4
D	4, 0	1, 1

Table 2.3: A Coordination Game

The authors interpret this as a player being a limited pessimist because she does believe that her opponent might play a best response, but she also believes that he gives her the best result with his available best responses. Column doesn't gain anything from playing D against C, he only hurts Row.
However, assuming non-spiteful best responses is not plausible in general (also for two players). With more than two players the assumption of "non-spitefulness" becomes highly problematic because the authors' suggested generalization to n players suffers from serious inconsistencies.([11], p.6) More about this follows in section 3.3 and I will give an adequate generalization of the worst payoff function for n players in section 4.2. Until then I will only consider examples of two player games.

The worst-payoff (feasible) agreements are then given as ([11], pp.7f):

Definition 2.2.3 (Worst-payoff Agreements)**.** Any agreement σ^F is called *feasible* if and only if

$$\sigma_i^F \in argmax_{\sigma_i \in \Sigma_i} w_i(\sigma_i, \sigma_{-i}) \text{ for all } i \in \{1, 2\}. \tag{2.2}$$

This means in any feasible agreement every player i chooses the strategy σ_i^F, which maximizes her worst payoff w_i. It's given by definition 2.2.2 as the minimum of two cases: the other player sticks to the agreement or he best responds to it, while giving i her maximum payoff in compliance with his best response.

For an example look again at table 2.3. The NEs in pure strategies are now (D,D) and (C,C). The latter is a weak one, which results in both players having two best responses, while the other plays C. In that case they don't care if they play C or D because their payoff is always 4. According to the previous definitions both NEs are feasible, whereas the (C,D) and (D,C) profiles are not. The reason is that, as argued before, both players can expect a worst payoff of 4 in the profile (C,C) because they believe their opponent doesn't try to hurt them. Can they maximize this worst payoff even further by unilaterally changing their strategy? No because for example in (D,C) Row has a worst payoff of 1 as Column will best respond with D. For Row that's worse than Column sticking to C, which gives Row 4. The worst payoff for Column in (D,C) on the other hand is 0. So both players receive a lower worst payoff in (D,C) than in (C,C). By symmetry the same applies to (C,D) and (C,C) is a feasible agreement. (D,D) is also feasible because sticking to the agreement or any of them best responding to it yield the same payoff in a NE, so 1 is already worst for both and they can't increase it.

Note that in general every NE in pure or mixed strategies is worst-payoff feasible because sticking to the NE and best responding to it are the same thing, so no player can increase his worst payoff. But any weak NE is only feasible because of the maximum operator in the worst payoff function. Yet there are worst-payoff agreements, which are not NEs, and some yield higher (expected) payoffs for everybody than any NE or any not-exploitable agreement. An example for that will be discussed in section 2.3.2.

2.3 The Virtual Bargaining Solutions

Now that I've explained the two notions of feasible agreements, it's still necessary to have a solution concept, which selects the right virtual bargain. For that Chater et al. use the Nash bargaining solution (see [8, 9, 10, 11]), invented by Nash to solve situations of actual bargaining, the concept of which I introduce first.[14]

2.3.1 The Nash Bargaining Solution and Disagreement Points

As for rationality assumptions, the Nash bargaining solution as well as best response reasoning with expected utility maximation require cardinal payoff functions[4] so that payoff differences have precise meaning.[14] Therefore, the same holds for virtual bargaining.
Based on that, the Nash bargaining solution rests on three axioms regarding the selection of the optimal bargain: (1) Pareto efficiency, (2) symmetric gains for symmetric payoff sets and (3) independence of irrelevant alternatives.[14]
Pareto efficiency means that any bargaining solution should be Pareto-optimal, i.e. there's no strategy set Pareto-superior to it.[5] Note that for virtual bargaining this assumption only regards the feasible agreements, not any strategy profile in the game. The second axiom states that, whenever the bargaining set is symmetric such that all players can gain equally with respect to certain disagreement payoffs, the solution should do so. At last, there's a consistency postulate: any solution to some domain of strategies should be the solution to a subset of those strategies if the solution is available in that subset.
Taken together it was proven that these axioms determine the bargaining solution as the strategy profiles, which maximize the product of players' payoff differences with respect to some disagreement point. Formally the Nash bargaining solution reads:

Definition 2.3.1 (Nash Bargaining Solution)**.** Let u_i, $i = 1, 2$, be player i's payoff function. The strategy profile σ^N is called the *Nash bargaining solution* if

[4]This means that a player's payoff function is invariant up to positive affine transformations.([16], pp.42f)

[5]In table 2.1 (C,C), (C,D) and (D,C) are the Pareto-optimal outcomes of the game.

and only if

$$\sigma^N \in argmax_{\sigma\in\Sigma} \prod_{i=1}^{2} (u_i(\sigma) - d_i), \tag{2.3}$$

where d_i are the payoffs related to the disagreement point.

In general, the Nash bargaining solution can select more than one strategy profile as the solution of a game. You might notice that the product of payoff differences varies when different disagreement points are inserted. In Nash's original bargaining game the disagreement point was a clearly defined outcome in the rules of the game. Generalizing his solution to other games thus requires a universal definition of a unique disagreement point.[14]
The choice of the disagreement point needs to make sense with the feasible agreements because together they should determine and motivate an intuitive, unique and general solution for virtual bargaining. For the older version of not-exploitable agreements Chater et al. didn't propose a disagreement point, while for the worst-payoff agreements they defined one, unambiguously.([8]; [11], p.8)
I will now introduce the disagreement point for the worst-payoff agreements and extend it to the not-exploitable agreements as well. The definition relies on the idea that the Nash bargaining solution is only supposed to look at the feasible agreements for the solution set of the game.([9], p.514) However, the authors never justified why the bargaining axioms make sense with the feasible agreements and why the disagreement point, determining the solution, should be among them. Thus it's difficult to explain their reasoning here. But it seems that in their picture the feasible agreements come first in motivating a player to follow virtual bargaining. If only feasible strategies are considered to be a played, the disagreement point should be among them. The virtual bargain is then picked by following Nash's bargaining axioms. I will argue in section 3.4 that the bargaining axioms make sense, but I will also argue in section 3.2 that the disagreement point should come first in the motivation of virtual bargaining. After all, the bargaining axioms are proposed on that basis.

For the worst-payoff agreements the disagreement point is defined as ([11], p.8):

Definition 2.3.2 (The worst-payoff Disagreement Point)**.** Any tuple (d_1, d_2) is called the *disagreement point* if and only if for all $i = 1, 2$ every d_i is given by

$$d_i = w_i^m = \min_{\sigma^F \in F} w_i(\sigma^F), \tag{2.4}$$

where F denotes the set of feasible agreements and w_i^m is called *minimum feasible worst payoff.*

This means every player i searches all the worst-payoff agreements and picks her minimum payoff among them as d_i. Then all d_i are gathered in a tuple, which is inserted in formula (2.3). As opposed to Nash's fallback position this disagreement point doesn't have to pick the d_i to be all related to a joint strategy profile in the game. The question is if such a definition is compatible with the assumptions of the Nash bargaining solution and the idea of virtual bargaining. To be clear, the definition has no problems with uniqueness of the disagreement payoffs, but whenever the tuple (d_1, d_2) isn't related to a specific strategy profile in the game, this needs to be justified. I will address the matter in section 3.2. For an example consider again table 2.2. The worst-payoff agreements are (D,D), (B,B) and the mixed strategy ($\frac{129}{140}$D,$\frac{11}{140}$B; $\frac{129}{140}$D,$\frac{11}{140}$B) with an expected payoff of $\frac{74}{7} \approx 10.6$, which is also the expected worst payoff. So the disagreement point is ($\frac{74}{7}$, $\frac{74}{7}$) from a mixed strategy profile. I argue that no player would think on their own that such a specific mixed profile is the obvious disagreement point. The profile (D,D) would make more sense, as D is the only rational choice by best response reasoning.

Similarly, the disagreement point for the not-exploitable agreements follows as:

Definition 2.3.3 (The not-exploitable Disagreement Point)**.** Any vector (d_1, d_2) is called the *disagreement point* if and only if for all $i = 1, 2$ every d_i is given by player i's minimum payoff of all not-exploitable agreements.

	C	D
C	3, 2	1, 1
D	0, 0	2, 3

Table 2.4: A modified Battle of the Sexes Game

The idea is exactly the same as in definition 2.3.2, but here the disagreement payoffs are chosen from the not-exploitable agreements.
For an example look at table 2.4 as a version of the Battle of the Sexes, where definition 2.3.3 would pick (0, 0) as disagreement point.[19] However, it's more plausible that players would go for their maximin strategies (C,D) to secure [1, 1] for themselves if they fail to agree on a bargain. The maximin strategy for a player is when he maximizes his minimum payoff, i.e. secures himself the highest payoff possible. So why not take that as the disagreement point?
As there are counterintuitive examples of the minimum feasible worst payoffs for both notions of feasibility, I will discuss the problems of the disagreement point in section 3.2. Yet for this chapter I will use definition 2.3.2 and 2.3.3 to display the current state of virtual bargaining.

2.3.2 The Virtual Bargaining Equilibrium in normal- and extensive-form Games

Having defined the Nash bargaining solution with a unique disagreement point for the two notions of feasible agreements, I will now introduce the solution concepts of virtual bargaining.
The virtual bargaining equilibrium (VBE) for the worst-payoff agreements for normal-form games follows as ([11], p.8):

Definition 2.3.4 (Worst-payoff Virtual Bargaining Equilibrium)**.** Any strategy profile $\sigma^V = (\sigma_1^V, \sigma_2^V)$ for two players is called a *worst-payoff virtual bargaining equilibrium* (wpVBE) if and only if

$$\sigma^V \in argmax_{\sigma^F \in F} \prod_{i=1}^{2} (w_i(\sigma^F) - w_i^m). \tag{2.5}$$

A wpVBE thereby takes all minimum feasible worst payoffs as d_is and selects the feasible agreement with the Nash bargaining solution that maximizes the product of *worst payoff* differences related to definition 2.2.2.
Similarly, the VBE for the not-exploitable agreements in normal-form games can be defined as ([9], p.514):

Definition 2.3.5 (Not-exploitable Virtual Bargaining Equilibrium)**.** Any not-exploitable strategy profile for two players is called a *not-exploitable virtual bargaining equilibrium* (neVBE) if and only if it maximizes the product of all players' payoff differences with respect to the disagreement point 2.3.3.

	H	L
H	2, 2	0, 0
L	0, 0	1, 1

Table 2.5: The Hi-Lo Game

For now some examples should be discussed to illustrate the alignment and differences between definition 2.3.4 and 2.3.5 due to the different notion of feasibility. Let's start with the Hi-Lo Game (table 2.5), where the NEs are (H,H) and (L,L) as well as the mixed strategy profile ($\frac{1}{3}$H,$\frac{2}{3}$L; $\frac{1}{3}$H,$\frac{2}{3}$L) with an expected payoff of $\frac{2}{3}$. These are also the worst-payoff feasible agreements, since all payoffs in the game are already worst for each player according to definition 2.2.2. As all interests are perfectly aligned between the players, every strategy profile is a not-exploitable agreement. The neVBE of the game is (H,H) with the disagreement point (0, 0) and a maximum payoff product of $(2-0)(2-0) = 4$. (H,H) is also the wpVBE with the minimum feasible worst payoffs ($\frac{2}{3}$, $\frac{2}{3}$) and a maximum worst payoff product of $(2-\frac{2}{3})(2-\frac{2}{3}) = \frac{16}{9}$. Here both disagreement points are joint profiles in the game and both solution concepts resolve the coordination problem in favor of the Pareto-optimal NE. The solution makes sense intuitively and the disagreement payoffs are given by maximin payoffs related to pure or mixed strategies, respectively.

For an example, where the neVBE and the wpVBE are not NEs, but yield higher payoffs for the players, consider again table 2.2. The only NE is (D,D), yet the worst-payoff agreements are (D,D), (B,B) and the mixed strategy ($\frac{129}{140}$D,$\frac{11}{140}$B; $\frac{129}{140}$D,$\frac{11}{140}$B) with an expected payoff of $\frac{74}{7} \approx 10.6$. The not-exploitable agreements are (D,D), (B,B), (D,B), (B,D) and some mixed strategies, which are not important here. Now, the neVBE of the game is (B,B) with the disagreement point (-100, -100) and a maximum product payoff of 129^2. Similarly, the wpVBE is also (B,B) with the disagreement point ($\frac{74}{7}$, $\frac{74}{7}$) and a maximum worst payoff product of $\frac{129}{7}^2$. The solution seems intuitive because for example (C,B) is exploitable by Column and the worst payoff for Row is higher in (B,B). Although B is a strictly dominated strategy, (B,B) is mutually advantageous compared to (D,D) and can be played without harm as long as one believes the other will not just play D. Like in the Hi-Lo Game both solution concepts choose the same solution, but with a different disagreement point. Both are unintuitive, since why wouldn't one just choose (D,D)? Also this is an example where the disagreement point (-100, -100) isn't a joint profile of the game.

	H	L	N
H	60, 60	33, 62	-10, 61
L	62, 33	36, 36	-6, 40
N	61, -10	40, -6	0, 0

Table 2.6: An Exploitation Game

For an example, where the wpVBE is not a NE nor a neVBE, but yields higher payoffs for the players, consider table 2.6 for an Exploitation Game.([10], p.5603) Here the only NE is (N,N), which is also the only not-exploitable agreement and thus also the neVBE of the game. For the worst-payoff agreements there's also the mixed strategy profile ($\frac{6}{7}$H,$\frac{1}{7}$N; $\frac{6}{7}$H,$\frac{1}{7}$N), where the expected worst payoff according to definition 2.2.2 is 34. That's higher for both players than in any other strategy profile and ($\frac{6}{7}$H,$\frac{1}{7}$N; $\frac{6}{7}$H,$\frac{1}{7}$N) is thus worst-payoff feasible. It's also the wpVBE because it yields the maximum product of worst payoff differences with respect to the disagreement point (0, 0).
But does the wpVBE make sense? Best response reasoning says H is never a rational strategy to play because it's strictly dominated by L and N. Yet following virtual bargaining one receives an expected payoff of approximately 50.3, 34 if the other best responds non-spitefully with L instead of N. Here the assumption yields a feasible agreement, which isn't a weak NE, so non-spitefulness has implications beyond that and without it the wpVBE would be just (N,N). Therefore, I will discuss the assumption further in section 3.3.

Next up to define and illustrate, are the wpVBE and neVBE in extensive-form games.

Definition 2.3.6 (The extensive wpVBE)**.** Any strategy profile $\sigma^V = (\sigma_1^V, \sigma_2^V)$ for two players is called a *worst-payoff virtual bargaining equilibrium of an extensive-form game* Γ (or simply wpVBE) if σ^V is a wpVBE of the normal-form game corresponding to Γ. The normal-form game is given for two players by labeling the columns and rows with all possible tuples of the respective player's pure strategy options at each decision node of the game.([11], pp.11f)

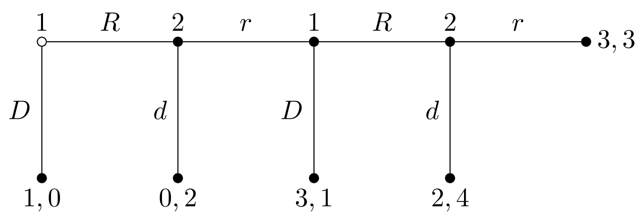

Figure 2.1: A Centipede Game

Take for example the Centipede Game with player 1 and 2 in figure 2.1.[19] If you write player 1's strategy choices at each decision node as a tuple and list all possibilities as the rows of the table, while player 2's choice tuples are listed as the columns, you receive table 2.7. In general, every extensive-form game can be transformed into a normal-form game using this method, though for more than two players this becomes hard to depict.

The definition of the neVBE of extensive-form games follows similarly:

Definition 2.3.7 (The extensive neVBE)**.** Any strategy profile for two players is called a *not-exploitable virtual bargaining equilibrium of an extensive-form game* Γ (or simply neVBE) if it's a neVBE of the normal-form game corresponding to Γ.

These two solution concepts for virtual bargaining in extensive-form games rely on a kind of precommitment of the players to their strategy tuples for a virtual bargain. Players would thus simulate the game in their minds to find the feasible agreements before any sequential move has been made. After they found the wpVBEs or neVBEs of the game,[6] the players would precommit to the virtual bargain, even in case of mutually damaging consequences to punish deviation.([11], pp.12ff)
Chater et al. also proposed a version of the virtual bargaining equilibrium for extensive-form games, which doesn't require precommitment but only subsequently feasible strategies for all sequential moves of both players. It's called the *renegotiable virtual bargaining equilibrium* (RVBE).([11], p.12) So after every move the players can reconsider their initially agreed VBE by looking at the new subgame and they might adjust to the other player's chosen strategy. How they adjust, depends on the informational circumstances of the game. But if everything is open to all players, i.e. they know exactly where they are in the decision tree and they know the whole tree, then the sets of RVBE and subgame perfect equilibria coincide. That means backward induction paths lead to a RVBE and punishment isn't credible.
The advantage of the RVBE is that players know what to choose as virtual bargainers, even if equilibrium paths are left in an extensive-form game. Yet as virtual bargaining tries to explain phenomena, which best response reasoning can't, the RVBE is not a plausible solution concept. It just shows, under which conditions a virtual bargain is found by backward induction for complete information. The RVBE can't account for precommitment nor mutual advantage in many games, but these are phenomena, which happen with real players.[12] So the RVBE hurts the goals of the account by restricting the theory to subgame perfection.

	rr	rd	dr	dd
RR	(3, 3)	(2, 4)	(0, 2)	(0, 2)
RD	(3, 1)	(3, 1)	(0, 2)	(0, 2)
DR	(1, 0)	(1, 0)	(1, 0)	(1, 0)
DD	(1, 0)	(1, 0)	(1, 0)	(1, 0)

Table 2.7: Centipede Game in normal Form

Finally, the Centipede Game (table 2.7) is discussed to show a game where the wpVBE suggests cooperation for mutual advantage, but backward induction, any NE or neVBE don't. In the Centipede Game two players take turns, either to choose D to end the game or to choose R to continue. After every turn the payoffs increase, but the one who ends the game always has the advantage. Here if nobody ends it, both get a fair cut of the generated payoffs.

[6] Extensive games often yield more than one wpVBE or neVBE.

Best response and backward induction suggest that every player should end the game at each decision node, so he can't be exploited by the other player. This reasoning yields the subgame perfect NE (DD,dd) in favor of player 1 due to the game structure. The other NEs also lead to player 1 choosing D right away, but they consist of all possible variations of DR, DD, dr and dd as well as all kinds of mixed strategies between them, e.g. $(D\frac{1}{3}R,\ D\frac{2}{3}D;\ dr)$. These are also the neVBEs of the game.

Regarding the wpVBEs all NEs are worst-payoff feasible, too, but there is also the feasible agreement (RR,rr), since its worst payoff for both players is higher than that of any other strategy profile. Explicitly, for player 1 the worst payoff would be 2, when player 2 chooses her best response (rd) to the agreement, and for player 2 the worst payoff is 3, as player 1 would best respond non-spitefully by sticking to the agreement instead of playing RD. In comparison (RR,rd) would yield a worst payoff of 2 for player 1 and 1 for player 2, which is less for player two than in (RR,rr), so he would change his strategy to rr for a higher worst payoff. (RD,rr) yields also a lower worst payoff of 0 for player 1 and 1 for player 2 because of the other's best responses — the same with all other strategy profiles. So there is no way for any player to change their strategy unilaterally and increase his worst payoff. Therefore, (RR,rr) is worst-payoff feasible and the unique wpVBE of the game, maximizing the product of worst payoff differences from the disagreement point (1, 0).

The latter would be also the disagreement point of the neVBE, which seems intuitive, as the game is constructed to the advantage of player 1 and she has the power to end the game right away. Conclusively, the wpVBE suggests a strategy profile with greater payoff for both players than any NE or neVBE of the game. For more game examples of wpVBEs see [10] or [11].

Chapter 3

Critical Discussion of Virtual Bargaining

As I have explained the current account of virtual bargaining in the last chapter, let's analyze it in more detail now. First I will argue that virtual bargaining should be based on worst-payoff feasibility, provided one can justify and extend it to n-player games (section 3.1). Then I will discuss the most important issues of the previous definitions: the implausibility of the current disagreement point (3.2), the complications of the worst-payoff function due to non-spitefulness as well as the challenges of n-player feasibility (3.3) and the plausibility of Nash's bargaining axioms on the feasible agreements (3.4).

3.1 Relations between the feasible Agreements

In this section I want to show that the worst-payoff agreements are better fit than the not-exploitable ones to account for the phenomena which can't be explained by best response reasoning.
Let's start by summarizing what we've learned in chapter 2 about the feasible ageements. The examples showed that both definitions of feasibility imply every NE in pure or mixed strategies, weak or not, to be a feasible agreement. Therefore they exist in every finite game according to the (mixed) Nash equilibrium existence theorem.[13] Moreover, both definitions can denote strategy profiles as feasible, which are not NEs, but yield higher expected payoffs for all players than any NE (like in the Boobytrap Game). Although the Nash bargaining solution sometimes selects the same strategy profile as solution of the game for both definitions regardless the disagreement point, the sets of feasible agreements are not the same in general. They can overlap, but no definition includes the other. Yet it was indicated that the worst-payoff agreements can yield higher (expected) payoff for all players than the not-exploitable agreements (like in the Exploitation and Centipede Game) — there's no example known the other way round.

But to decide, with which definition of feasibility to settle, one has to justify why the related reasoning and assumptions are plausible and show that it accounts for the phenomena, which virtual bargaining tries to explain. The former I will

discuss in section 3.3, as it also depends on whether the disagreement point is intuitive; the latter on the other hand I argue for now.
More precisely, virtual bargaining is supposed to describe how and why real players are able to solve many coordination problems, though best response reasoning can't suggest a solution. Further, the theory should explain how mutual advantage can be achieved, while players are motivated from an individual perspective. At last, virtual bargaining is to account for precommitment in sequential-move games, which violates backward induction and subgame perfection (as in the Centipede Game).[11]
Both types of feasible agreements can resolve certain coordination problems like the Hi-Lo Game irrespective of the disagreement point. Yet with the not-exploitable agreements players can't accept possible exploitation in favor of a higher individual payoff, which is unintuitive. Instead, we'd expect a solution of mutual advantage in a game designed with multiple stages of exploitation like the Exploitation or Centipede Game.

Let's revisit the Centipede Game (figure 2.1 and table 2.7). Here almost no real player ends the game outright as suggested by backward induction with the subgame perfect NE. Instead, most players take a few turns (in a longer Centipede Game) and then somebody plays D/d before reaching the end.[12] Virtual bargaining suggests that it's more rational not to focus on being exploited if one doesn't play D/d oneself, but to maximize one's worst payoff by playing R/r all the way to the end. The player with the last turn can pursue the highest mutual advantage or get himself most, but even that is better for everyone than ending the game earlier.
Real players, who don't play R/r to the end, are then considered not entirely rational with respect to virtual bargaining, but they share the intuition that they can receive more payoff by taking a few turns. Yet the same can be said with respect to best response reasoning because players choose D/d at some point, but not at the start as suggested. Both reasonings seem close to the actual phenomenon and the belief that the opponent is boundedly rational justifies both views. There seems to be a conflict between best response and the virtual bargain with mutual advantage. Without examining the beliefs and motivations of real players further it's hard to argue, which reasoning is descriptively correct.
Nonetheless, if virtual bargaining should account for phenomena like in the Exploitation and Centipede Game, it needs to rationalize mutual advantage despite best response reasoning and precommitment not to defect earlier. The worst-payoff agreements try to do that, while the not-exploitable agreements clearly can't.

In conclusion there's good reason to settle with worst-payoff feasibility in virtual bargaining if the theory is supposed to explain coordination, mutual advantage and precommitment. Whether the definition indeed motivates such behavior, can only be answered after the disagreement point, non-spitefulness and n-player feasibility have been examined. In general, it might be possible to find an entirely new notion of feasible agreements, which is descriptively better suited than the existing versions. But such a notion is only needed when it were clear that the worst-payoff agreements and the subsequent proposals of this paper still don't account for real players' actions. So I won't consider new definitions here.

3.2 Critique on the Disagreement Point

Now I will show that the disagreement point introduced by Chater et al. is not convincing. However, before criticizing the current definition, I want to explain why the disagreement point is crucial to motivate virtual bargaining.

The authors assume the feasible agreements to be the core notion of describing and motivating a virtual bargaining process. Although this isn't wrong because the reasoning to select agreeable bargains needs to be convincing, it's more important to state how players come to follow such reasoning in the first place. Maximizing worst payoffs may seem reasonable, supposed that we accept all related assumptions, but it's not obvious why a best response reasoner would do that. Such a player maximizes his expected payoff given his beliefs about the other players' strategy choices.([21], sec.2.2) If best response reasoners are the standard of game theory, I argue that they need to be motivated to maximize their worst payoffs instead.
Right now virtual bargaining doesn't provide any reason to do so except that it's often more advantageous compared to best response solutions. Moreover, the authors didn't state if virtual bargaining can or should be motivated for a best response reasoner. They just seem to give an alternative to best response reasoning, which is still based on the same rationality axioms. But such an enterprise is ill justified, when the axioms are commonly understood as best response reasoning.
To fill the gap and motivate a best response reasoner to follow virtual bargaining, the disagreement point has to state intuitively what would happen without virtual bargaining, i.e. the fallback position. It also needs to point the direction to an advantageous bargain. If finding the disagreement point isn't understandable for a best response reasoner because he has to commit to virtual bargaining first, one loses the connecting motivation. So either the strategy profile related to the current definition is so obvious in games that it leads to worst payoff reasoning or the disagreement point, as it is, should be modified. According to these requirements I discuss the current disagreement point now in game examples.

	C	D
C	1, 1	0, 2
D	2, 0	-1, -1

Table 3.1: A Chicken Game

I mentioned in section 2.3.1 that the unique minimum feasible worst payoffs don't have to select a joint strategy profile of the game. Now look at table 3.1 for a Chicken Game, where the worst-payoff agreements are the NEs (C,D), (D,C) and ($\frac{1}{2}$C,$\frac{1}{2}$D; $\frac{1}{2}$C,$\frac{1}{2}$D) with an expected payoff of $\frac{1}{2}$. Here the minimum feasible worst payoffs are 0 for both players, but from different strategy profiles; for Row it's related to (C,D) and for Column to (D,C), respectively. If Row plays C for (C,D) and Column plays C for (D,C), then both get a payoff of 1 in (C,C), while the worst payoffs are still 0. If that were the fallback position related to the minimum feasible worst payoffs, the wpVBE from (C,C) would be the mixed NE ($\frac{1}{2}$C,$\frac{1}{2}$D; $\frac{1}{2}$C,$\frac{1}{2}$D), which yields lower payoffs but higher worst payoffs than (C,C). That's odd. First of all, (C,C) isn't worst-payoff feasible and thus not what definition

2.3.2 intended. Secondly, it doesn't motivate the mixed NE as the virtual bargain for a best response reasoner when he effectively loses payoff. So whenever the minimum feasible worst payoffs don't regard a joint strategy profile, they're incompatible with the idea of an actual fallback position. But what else do they mean then and is that a problem for the bargaining axioms?
One interpretation would be that players should imagine the game as a bargaining problem over the feasible agreements, where the disagreement point (0, 0) is a real fallback position, in order to apply Nash's axioms. But then it's unclear why a bargaining solution to that hypothetical case should be the virtual bargain of the actual game, as it's not the same game. This is like claiming that the bargaining axioms apply, although the disagreement point is just an idea. Yet why should all players gain equally from a disagreement point when there's only a certain chance to receive their disagreement payoffs? The expected payoff would normally be higher than the minimum feasible worst payoff. Also, the third bargaining axiom states that the solution should only change when it's not available in the players' strategies. Why shouldn't the virtual bargain change when the disagreement point is unavailable? Chater et al. didn't answer any of these questions and it seems unlikely that the bargaining axioms can be applied without further problems regarding their justification. Moreover, best response reasoners aren't motivated to pursue virtual bargaining when the disagreement payoff has no obvious meaning (without worst payoffs) and doesn't state what happens without coordination.
Conclusively, virtual bargaining would have to be restricted to games where the disagreement point is a joint strategy profile to justify the bargaining axioms and the players' motivation. That would be unsatisfying because this restriction is just due to an improper disagreement point and we'd expect a solution proposal for games like the Chicken as well. Besides, (C,C) *does* make sense as a fallback position because the players can choose it as their maximin strategy to avoid (D,D).

Furthermore, there are other games, where the current disagreement point is implausible and the reason is that it often selects a mixed strategy profile, when it chooses a joint fallback position. One has already seen that in the Boobytrap Game the minimum feasible worst payoffs were (10.6, 10.6) from a mixed strategy worst-payoff agreement. It's unclear why players wouldn't just fall back on (D,D) with the outcome [20, 20], as it is a NE and the only best response rational profile. Without committing to worst payoff reasoning a best response player wouldn't understand why ($\frac{129}{140}$D,$\frac{11}{140}$B; $\frac{129}{140}$D,$\frac{11}{140}$B) should make sense. In general a mixed strategy profile is a fallback position in the broader sense, but it would be more plausible if a disagreement point is chosen in pure strategies, when there is a convincing candidate. Conclusively, if virtual bargaining should solve problems like the Boobytrap Game, the disagreement point needs to be appealing and not choose a rather exceptional mixed strategy, just to fix the solution.

In summary there are games like the Chicken, where the current disagreement point isn't related to a specific strategy profile of the game. Secondly, even if the disagreement point is a definite fallback position, it's often implausible. As the standard of game theory is best response reasoning, players should be expected as such, choosing a more preferable virtual bargain under certain conditions.[1] The motivation to change or extend a player's reasoning is crucial and Chater et al. don't provide it. Therefore, I will propose a new definition of the disagreement point in section 4.1.

3.3 Non-Spitefulness and N-Player Feasibility

One has seen earlier in chapter 2 that many worst-payoff agreements (and not just weak Nash equilibria) rely on non-spitefulness. Thus I want to discuss the worst payoff function now in more detail to see when non-spitefulness can be plausible. Afterwards, I will show the difficulty to generalize the worst payoff function and non-spitefulness to games with $n > 2$ players. Moreover, one will see that not-exploitibility has the same issues.

Let's start with two players. I explained before that with non-spitefulness a player assumes her opponent to give her the maximum payoff possible, while he is best responding. Chater et al. called it limited pessimism regarding players' worst payoff evaluations, yet the authors gave no justification beyond that intuition.([11], p.6)
Why should Row believe in a game like 2.3 that, assuming Column betrays the agreement and best responds, he looks out for her if he can? According to rational expected payoff maximation Column is first of all indifferent between best responding such that Row gets her maximum payoff and best responding such that she gets less. Both C and D yield him the same payoff, given that Row chooses C, because they are *best* responses. Which of Column's beliefs would bring him to consider playing C? If he believes that Row will play C, there is no reason in Column's preferences to choose C over D or any other mixed strategy. In fact D weakly dominates C for both players, meaning that D is better against D and both C and D yield equal payoff against C. So Column as a best response reasoner might believe D is best for him, anyway. If Column *should* believe that Row chooses C only if he chooses C, too, this would require Column to acknowledge that Row as a pessimist doesn't choose C if she believes he could play D. Additionally, if Column doesn't also believe that Row is a virtual bargainer, but simply a best response reasoner, Column thinks that Row will choose D herself without sufficient reason to believe that Column chooses C.
But if one assumes now that Column believes Row to be a virtual bargainer, who maximizes her worst payoff function and believes Column to be nice to her if it doesn't matter for him, would he play C? In table 2.3 a virtual bargainer considers the same profiles as feasible/rational as would a best response reasoner unless C is never considered rational.[2] Row's tendency to coordinate on the most advantageous agreement suggests her to play C. In such a case Column would believe Row to choose C. If Column would believe that Row only does that, when he is nice, he should at least seem to be nice and play C as long as he prefers a payoff of 4. But in a one-shot game Column doesn't have to be nice and if he thinks Row to play C due to her disposition, he can again play whatever he

[1] This follows similar reasoning as the *team reasoning* paper [7].

[2] This is a rather difficult topic in epistemic game theory (see [18] or [20], sec.5.2).

wants. So additionally, for both of them to certainly play C, Column must also be motivated for mutual advantage. However, he doesn't have to share all the reasoning of Row and also not all the bargaining axioms. Yet it needs specific requirements for Column's beliefs in order to attain the weak NE as the virtual bargain.
The other side is whether Row as a virtual bargainer should believe Column to best respond nicely. It seems at least consistent that if she is motivated towards a bargain with more expected payoff for her than best response alone can suggest, she is convinced Column will see that in a symmetric game, too. Then it seems not unreasonable in her place to consider non-spitefulness, assuming that both players want to be sure to reach mutual advantage. But even though it might not be inconsistent, it's as if she assumes both players to agree to a proposition before making a move such that they can coordinate on (C,C). This is precisely the idea of a virtual bargain. So in order to reason like a virtual bargainer, Row is assumed to honor a virtual bargain. For her that might be coherent, though it seems circular, but to assume that for Column is only plausible if she believes him to virtually bargain as well.
Conclusively, in a case of codependent beliefs about the other player's disposition towards virtual bargaining or another motivation for mutual advantage the assumption can be plausible. Yet its motivation can not emerge from worst payoff reasoning, but either depends on a player's norms, convictions and beliefs about the other player. Or it must be genuinely motivated by the idea of virtual bargaining itself and *not* just assumed in the reasoning.

Now I want to explain why the authors' suggestion to generalize the worst payoff function and non-spitefulness to n-player games is inconsistent. Their paper ([11], p.6) assumes that any of i's opponents best respond independently, while believing that all but himself stick to the agreement. On the other hand it states that any subgroup of other players might best respond and that their *joint* best response would give i the maximum payoff available. But if every opponent independently chooses his non-spiteful best response, the resulting joint deviation for some subgroup of opponents doesn't have to maximize their payoffs individually, i.e. it may not be a best response.
To claim that some subgroup's deviation is a best response raises the question if their payoffs need to be maximized on the constraint that these players deviate together. But if that's true, then one needs a notion of this deviating subgroup's joint interest and how it determines the individual payoffs, i.e. some concept of collusion. The authors didn't provide any concept of the sort. Also to give the players, who stick to the agreement, their maximum payoff, any deviator needs to know, who else is deviating, and together they have to coordinate on some joint strategy such that it gives all deviators most — whatever that is. I argue that such a level of knowledge and collusion is completely implausible in a game without communication. Coordination between best response reasoners is what virtual bargaining should explain and thus can't just be assumed. Conclusively, there's a direct contradiction between opponents best responding individually and the claim that they best respond as a group. Even if that would make sense, it requires an unlikely level of kowledge and coordination between the deviators to best respond non-spitefully.
But can you expect every unilateral deviator to best respond non-spitefully? No,

because for more than two players a deviator might have to decide to whom of the non-deviating players he best responds non-spitefully. It's possible that the best response of player 3 is non-spiteful for player 1, but spiteful for player 2. A player can't expect to be the one, who receives her maximum payoff. So non-spitefulness is implausible for more than two players except for very specific games.
A similar objection can be made against not-exploitability. This notion of feasibility is supposed to yield a stable bargain because no player can make somebody else worse off, while best responding unilaterally. Yet for $n > 2$ players more than one might be able to best respond for a higher payoff in such a way and if they do so, the result might yield lesser payoff for everybody than the feasible agreement. This possibility isn't adressed by the definition nor the reasoning.[8] It's rather assumed that in a feasible agreement nobody wants to change strategy. But if there are two best response players, which can still increase their payoff, you can't just claim that they won't best respond. In other words, not-exploitability for n players is only unproblematic in a Nash equilibrium or when only one player can best respond to the agreement. Although worst payoff reasoning has similar challenges, it addresses them by considering the other players' best responses; not-exploitability is simply ignorant. Therefore, I propose to generalize the *worst payoff* function in a consistent way if virtual bargaining is to account for more strategy profiles than NEs in n-player games. Indeed it should because collusion in Bertrand competition is a real phenomenon of economics and it's not explained by Nash equilibrium play.([10], pp.5600f)

In summary, unless a different disagreement point can motivate non-spitefulness for two players or there are certain conditions matched about the players' beliefs, the assumption is implausible. I will return to this matter in section 4.2 after I've proposed the new disagreement point and I will show that a slightly weaker version of non-spitefulness can be justified for arbitrary players.

3.4 Motivation of the Bargaining Axioms

As mentioned in section 2.3.1, it needs to be justified why the axioms of the Nash bargaining solution should be required on the set of feasible agreements. I argue that the bargaining assumptions make sense if both the disagreement point and the feasible agreements are plausible.

Nash proposed that the ideal bargain given two rational players should have the earlier mentioned properties (see 2.3.1). Yet also in his game actual communication between the players wasn't necessary to choose the bargain. Both players had only to accept the bargaining axioms and the equivalent formula would select the optimal solution with respect to the disagreement point. The idea in virtual bargaining is to generalize this to arbitrary games for n players without communication; especially then the bargaining axioms should be plausible. As in Nash's original bargaining game all strategy profiles were part of the bargaining set, the feasible agreements need to transform a general game into a bargaining situation to apply his solution concept. Consequently, I will assume now that virtual bargaining bears a convincing notion of feasible agreements for n players and the disagreement point to establish the theory. Based on that I will now argue that Nash's axioms are plausible to identify a virtual bargain.

Given that virtual bargainers are also (worst) payoff maximizers, it's fairly intuitive that the solution should be Pareto-optimal among the feasible agreements.[3] By restricting the strategies to the feasible agreements the game has been limited enough, so players would want to use all available payoff ressources. They would like to get the most possible and the feasible agreements fix the options how the payoffs are reasonably distributed. So far, that doesn't determine if the solution yields a fair share or gives everything to one player.
Especially if players can't communicate, it's plausible that every player has to gain equally from the disagreement point whenever possible or he wouldn't agree to the virtual bargain. This also implies that every player gains strictly from her disagreement payoff. If the game is contructed to the advantage of one player, it would be wrong to generally claim equal payoffs for the virtual bargain because the advantaged player doesn't have to agree. Hence, it's important that a player's payoff is measured in relation to the disagreement point.
At last, for every subset of the feasible strategies available to the players the overall solution should be the same if it's available in that subset. This is as if one player didn't look at all mixed strategies and didn't find all feasible agreements, but the virtual bargain is attained by pure strategies, which the other players identified, too. Such a subset of feasible agreements considered by one player shouldn't change the virtual bargain. Although virtual bargaining requires rational players, this axiom is important for consistency.
In total, this determines the virtual bargain to be identified by the Nash bargaining solution on feasible agreements. The bargaining axioms seem intuitive, but there's a need for experiments and cognitive science to find out if real players actually refer to these assumptions or what other principles set up a plausible solution. Yet in a lot of situations in real life where people come to an agreement they often propose a solution with few to none communication — only by guessing, what's best for everybody.
Consider for example a stand on the road, where one can buy eggs. Often there's no seller to bargain with, but just the eggs and a box to leave money; the amount noted somewhere. It seems most people, who want eggs, throw in (at least some) money instead of simply taking eggs and walking away. If that weren't true, these stands wouldn't exist. Both parties could think of all kinds of prices for eggs as the bargaining set, but they coordinate on a price, which is mutually beneficial. In expression, they use the given ressources of eggs and money in one's pocket and try to gain equally from the disagreement point of no money versus no eggs. Most people give only less than the proposed amount of money if they have less in their pocket, as when the optimal solution is not available due to limited strategies. By not taking eggs if they have no intent or money to pay, they honor the obvious virtual bargain of trade.

[3] Pareto-efficiency in the whole game would be implausible in a game like the Prisoner's Dilemma.

Chapter 4

Solution to the Problems

Having understood the main problems of virtual bargaining to be an implausible disagreement point, a problematic assumption of non-spitefulness and incoherent reasoning for n players, I now propose solutions to these issues. Section 4.1 will therefore introduce a new definition of the disagreement point, which is motivated by best response and maximin reasoning. Subsequently, I will argue that a weaker version of non-spitefulness can be justified for an arbitrary number of players (section 4.2) and that worst payoff reasoning can be generalized for $n > 2$.

4.1 A new Disagreement Point

I argued before that the disagreement point should motivate a standard best response reasoner to adopt virtual bargaining for a more preferable solution. With respect to this idea the disagreement point has to be intuitive even without the feasible agreements and be represented by a strategy profile with definite meaning in the game. Moreover, it should only be a mixed strategy if no pure strategy profile is plausible. Given these requirements I will propose a new disagreement point.

	L	M	R
U	2, 2	2, 0	2, 0
M	1, 0	3, 1	1, 0
D	1, 0	1, 0	0, 0

Table 4.1: Another Coordination Game

To clarify why the following definition is chosen, consider table 4.1. In this game the strategy D is strictly dominated by U for Row and R is weakly dominated by L and M for Column, so best response would suggest not to play them. Thus the NEs are (U,L), (M,M) and the mixed strategy profile ($\frac{1}{3}$U, $\frac{2}{3}$M; $\frac{1}{2}$L, $\frac{1}{2}$M) with expected payoffs [2, $\frac{2}{3}$]. These are also the current worst-payoff agreements, which don't depend on non-spitefulness.
The game is interesting because the virtual bargain isn't obvious, as (U,L) and (M,M) are the Pareto-optimal outcomes, but the distribution of payoffs seems "fair" at the former, but favoring Row at the latter. Thus the plausibility of the

virtual bargain crucially depends on the disagreement point. What would the players do without a virtual bargain?
It seems very convincing that Row would play U if she had uncertain beliefs about Column's choice. For Column it's more difficult, as the game is constructed to the advantage of Row and Column needs to coordinate with Row or he gets nothing. Maybe he would anticipate Row's maximin strategy and consider Row to play U with high probability, so he would pick L for his maximum payoff. But Column shouldn't play L if he believes that Row might just play M to reach her maximum payoff of 3. Yet as D and R are unreasonable by best response, it makes sense that both players, if believed rational, will only choose within the best response rationalizable (BRR) left upper subgame. With this constraint Column could achieve an expected maximin payoff of $\frac{2}{3}$ through the mixed strategy ($\frac{1}{3}$L, $\frac{2}{3}$M) and Row would still have U for her maximin payoff of 2. Their joint strategy profile would be (U, $\frac{1}{3}$L, $\frac{2}{3}$M) with expected payoffs [2, $\frac{2}{3}$]. For standard best response reasoning that's the highest payoff both players can reach on their own without having specific beliefs about the other player's actions.
I argue this to be the adequate disagreement point of the game, not the mixed strategy NE with the same expected payoffs and not any other Pareto-inefficient strategy profile. For example no player wants to receive a payoff of 0, but (D,R) is not a convincing fallback position, as rational players wouldn't play dominated strategies. (U,M) is also intuitive, but it's simply the maximin payoff within best response rationalizable *pure* strategies. But why is the mixed NE not plausible? This strategy is only best response rational if one believes the other to play it, too. To demonstrate, if Row plays U and Column plays ($\frac{1}{2}$L, $\frac{1}{2}$M) from the mixed NE, he gets an expected payoff of 1, but if Row plays M, Column receives only $\frac{1}{2}$. Without a more specific belief about what another player will do and in these situations virtual bargaining is needed most, Column can increase her expected minimum payoff by playing (U, $\frac{1}{3}$L, $\frac{2}{3}$M). The latter yields her at least $\frac{2}{3}$ no matter Row's choice. The virtual bargain would then be (M,M) with [3, 1]. Although this solution isn't as obvious as (H,H) in the Hi-Lo Game, it recognizes the advantage of Row to receive 2 for sure. As Column can end up with nothing, it's still a mutual improvement to coordinate on (M,M). If this result would occur with real players, is yet another matter.
You might ask now: if (U, $\frac{1}{3}$L, $\frac{2}{3}$M) is a convincing disagreement point, probable to happen without virtual bargaining, why would a standard best response reasoner not just best respond to it for more payoff? If both players believe the other to go for (U, $\frac{1}{3}$L, $\frac{2}{3}$M), but Column best responds with L and Row best responds with M, they reach (M,L) with [1, 0]. So (U, $\frac{1}{3}$L, $\frac{2}{3}$M) is convincing because a player never receives less than the BRR maximin payoff if the other best responds to it. But the deterrent of both best responding makes the disagreement point a better option than to best respond to it — this carries on to n players.

Therefore, I propose the disagreement point to be:

Definition 4.1.1 (The Disagreement Point)**.** Any strategy profile $\sigma^d = (\sigma_1^d, \ldots, \sigma_n^d) \in \Sigma$ is called *disagreeable* if and only if for all $i = 1, \ldots, n$ the strategy $\sigma_i^d \in \Sigma_i$ is player i's maximin strategy on the subset of BRR strategy profiles $\Sigma^{br} \subseteq \Sigma$ or formally with the opponents' strategies $\sigma_{-i} = (\sigma_1, \ldots, \sigma_{i-1}, \sigma_{i+1}, \ldots, \sigma_n)$:

$$\sigma_i^d \in \Sigma_i^d \equiv argmax_{\sigma_i \in \Sigma_i^{br}} \{ \min_{\sigma_{-i} \in \Sigma_{-i}^{br}} u_i(\sigma) \}. \tag{4.1}$$

The *disagreement point* is given as $\sigma^D = \sigma^d \in \Sigma^d$ with the payoffs $u_i(\sigma^D)$ if σ^d is unique. Otherwise, $u_i(\sigma^D) = \max_{\sigma^d \in \Sigma^d} u_i(\sigma^d)$. The respective σ^D would be the disagreement point whenever it's unique; else, the disagreement point is given by the BRR maximin profile on the subgame of BRR maximin strategies of the original game.

The disagreeable payoffs $u_i(\sigma^d)$ are thus given by a player's maximin payoff, restricted to the subgame of strongly BRR strategy profiles. This means dominated strategies are not rationalizable no matter the belief about the other player and should be excluded. If you repeat this procedure, as shown for the Boobytrap Game in section 2.1, until there are no more dominated strategies to remove, the remaining strategies will be BRR. Sometimes the method of iterated deletion of dominated strategies yields different results depending on the order of the deleted strategies. However, the NEs of a game can also be identified without this method. Any strategy of a player is best response rationalizable if it takes part in some NE profile. Strongly BRR strategies would then exclude strategies related to *weak* NEs from the disagreement point, whenever they aren't safe for best response reasoners to play (like discussed for table 2.3). This is consistent with the idea of secure strategies that yield at least the maximin payoff in case of uncertainty.
You could argue that weakly dominated strategies like R for Column in 4.1 are rationally admissible, but it will never reduce your BRR maximin payoff to delete them and it secures uniqueness of the disagreement point in many cases. As Column still wouldn't assign positive probability to R for his mixed BRR maximin strategy, I propose to delete them.[1] But the final decision on this matter would depend on the explanatory range that virtual bargaining is supposed to have, which is yet to be fully determined.

	B	C	D
B	0, 0	1, 0	0, 0
C	0, 1	1, 1	0, 2
D	0, 0	2, 0	-1, -1

Table 4.2: Extended Chicken Game

	B	C
B	0, 0	1, 0
C	0, 1	0, 0

Table 4.3: A BRR maximin Subgame

Note that by mathematical definition of the maximum and minimum operator the disagreeable payoffs are unique in every game and $u_i(\sigma^d)$ is greater than or equal to the BRR maximin payoff for all σ^d. Yet there can exist more than one strategy σ_i^d related to a player's BRR maximin payoff, so σ^d isn't always unique and $u_i(\sigma^D)$ wouldn't be fully determined. For that look at table 4.2, as for this extended Chicken Game there are infinite BRR maximin strategies, which guarantee a minimum payoff of 0 for both players, namely B, C and all mixtures

[1]Nonetheless, this subset of BRR strategies will always include at least one NE.

between them. Yet the NEs and worst-payoff agreements are the same as in table 3.1. It matters which profile is picked as the disagreement point because the virtual bargain from (B,B) with [0, 0] would be the mixed NE ($\frac{1}{2}$C,$\frac{1}{2}$D; $\frac{1}{2}$C,$\frac{1}{2}$D), whereas e.g. with (C,C) there's no more beneficial solution. However, for table 3.1 the disagreement point would be (C,C). But in the subgame of 4.2 given by B and C for both players all strategy profiles would be BRR maximin profiles and NEs, so best response reasoning can't give any solution and the disagreement point would be unclear. Yet I argue most players would simply choose (C,C) in such a subgame for their highest payoffs. Thus I proposed that every player should regard her disagreement payoff $u_i(\sigma^D)$ as the maximum of all the disagreeable profiles and the respective σ^D as the disagreement point. However, if the subgame of BRR maximin strategies looks like table 4.3, the personally maximum disagreeable profiles (B,C) and (C,B) don't coincide. In that case I proposed the disagreement point as the BRR maximin profile on the subgame of BRR maximin strategies of the original game 4.2, i.e. (B,B) with an expected payoff of 0.
Whenever the disagreement point given by the successive method of definition 4.1.1 is unique, it describes the fallback position for virtual bargaining. Should there be a case where it's still not unique, the maximum disagreeable payoffs still hold and the players can virtually bargain with them for a worst-payoff agreement, which yields higher worst payoffs for all players. If no beneficial bargain exists and the fallback position is unclear, then virtual bargaining is ambiguous. Note that in such a case best response reasoning can't suggest a solution either, which means such games are generally unsolved. Nonetheless, the set of σ^ds wouldn't be empty or without meaning, they're just not unique. Yet most of the time $\sigma^d = \sigma^D$ is unique and yields intuitive disagreement points for all game examples of this paper.

In summary definition 4.1.1 regards the individual perspective of a rational best response reasoner. I propose it as the best response solution of a general game for uncertain beliefs about another player's actions, when σ^D is uniquely given. Otherwise, $u_i(\sigma^D)$ is a player's threshold any virtual bargain would have to exceed. I will further show in the next section that the definition is closely related to the worst payoff function.

4.2 Non-spitefulness and generalized Worst Payoff Reasoning

Based on the disagreement point 4.1.1 I will now motivate a weaker version of non-spitefulness for an arbitrary number of players and show that worst payoff reasoning is plausible. For this I will specifiy simple assumptions about the players' beliefs in a virtual bargaining situation and show that a player bears inconsistent beliefs if he best responds spitefully. I further generalize the worst payoff function and the VBE for n-player games. To clarify the results, I will state the conditions of n players' beliefs to optain a virtual bargain.

Imagine a general game where the disagreement point was identified by BRR maximin strategies. The players estimate this to be the highest expected payoff they can secure on their own.[2] Because of the rationality axioms a player recognizes, of course, when there exist other strategy profiles, which would yield him higher expected payoff. I argue that it's common belief among best response rational players that every player would prefer more payoff to less. This means every player would like to gain with respect to the disagreement point, but they know that they can't reach higher payoffs on their own. To resolve the situation, I propose that players should imagine to bargain.

Any best response reasoner would then believe that a player will only contribute to a virtual bargain, if he gains payoff with respect to the disagreement point.[3] Further, a player only has reason to deviate from a virtual bargain, if he (strictly) gains payoff by doing so. If I deviate from the bargain without incentive, then I might get the same payoff, but the rest may not. The idea to virtually bargain tells me that the other players wouldn't help to achieve a bargain if they believed me to deviate in such a way, yielding them lower payoffs than the disagreement case. This simply follows from worst payoffs and the common belief that everybody wants to gain above the disagreement point. But then I have to believe that I reduce the probability of the virtual bargain to happen if I deviate without incentive because other players would want to fall back to the disagreement point. Such an action is inconsistent with my own interest to gain with respect to the disagreement point. Moreover, it's irrational regarding expected payoffs if I must expect to reduce my own payoff back to the disagreement case with such a best response.

For more clarity let's revisit table 2.3. Row should expect Column to be indifferent between an expected payoff of 1 in the strategy profile ($\frac{1}{4}$C, $\frac{3}{4}$D; C), where Row best responds in part spitefully to the agreement (C,C), and a payoff of 1 from the other feasible agreement (D,D). So if Row expects Column to believe that she chooses C with a probability greater than $\frac{1}{4}$, then she would expect him to accept (C,C) as the virtual bargain — similarly the other way round for Column. Conclusively, if both players believe the other to assign a probability greater than $\frac{1}{4}$ to playing C, the virtual bargain (C,C) would yield higher expected (worst) payoff for both players than the disagreement point.

This gives a precise threshold in my beliefs when the other player would prefer to stick to the disagreement point by best response *and* worst payoff reasoning if I deviate from the bargain without further gain. Hence, if a player considers a spiteful best response to the bargain, which assigns lesser probability to play the bargain than the other player can accept in compliance with the common interest to gain above the disagreement point, he chooses inconsistently. In expression, he believes that everybody wants to gain further, including himself, and he does believe that everybody is best response rational, but he nonetheless considers a strategy, which makes the other player prefer the disagreement point. Therefore, to avoid inconsistency in one's own beliefs and choices, a player should only best respond if he strictly gains payoff by doing so. In many games (like the ones in this paper) that's equivalent to non-spitefulness. For n players the argument holds still, but it doesn't result in non-spiteful best responses, but in the common

[2]Note that this refers to the maximin payoff. The actual disagreement payoff can be higher according to definition 4.1.1, but a player has no influence on receiving more than the maximin payoff.

[3]That's simply the second bargaining axiom.

belief that every player only deviates from a bargain if he strictly gains payoff. Note that such a conclusion is only possible because the disagreement point is best response rational and motivates to virtually bargain. Non-spitefulness on the other hand comes out of nowhere and is ill justified.

Next I want to show why worst payoff reasoning makes sense to find a virtual bargain for a best response reasoner, who accepts the disagreement point. It's reasonable because it's a measure to tell me if other players would accept an agreement, when their payoff (despite exploitation) is still higher than in the disagreement point. Additionally, the players don't have to adopt a new unrelated reasoning because the best everybody can do on their own is a maximin payoff. In expression, maximin reasoning maximizes the minimum payoff given all other players' possible strategies and worst payoff reasoning maximizes the minimum payoff regarding other players' best responses or the agreement itself. Considering only best responses and not any deviation for the worst payoff is convincing because players try to find a solution additional to best response reasoning. So if the players are assumed to think about a virtual bargain, too, they will only deviate from a possible agreement if it benefits them. Thus it's plausible that in a feasible agreement all players' worst payoffs should be maximized so that everybody can accept the bargain.
In conclusion, the worst payoff function 2.2.2 for two players needs to be formally adjusted with respect to the assumption of best responses with strict gain (instead of non-spitefulness):

Definition 4.2.1 (Worst Payoff Function)**.** Let $\sigma_i \in \Sigma_i$, $i = 1, 2$ be player i's strategy and $u_i(\sigma_i, \sigma_{-i})$ player i's payoff function, where σ_{-i} denotes the strategy of i's opponent. The worst payoff of an agreement $\sigma^A = (\sigma_1^A, \sigma_2^A)$ for player i is defined as

$$w_i(\sigma_i^A, \sigma_{-i}^A) = \min_{\sigma_{-i} \in B_{-i}(\sigma^A)} \left\{ u_i(\sigma_i^A, \sigma_{-i}^A), u_i(\sigma_i^A, \sigma_{-i}) \right\}, \tag{4.2}$$

where $B_{-i}(\sigma^A)$ states the set of player $-i$'s best responses to the strategy profile σ^A with strict gain.

Effectively, the new worst payoff of 4.2.1 will equal the one of 2.2.2 with non-spitefulness in the maximum operator for all the games of this paper. However, there can be games where the definitions yield different payoffs.

Using the previous results, I will now consistently generalize this concept for n-player games. The arguments of section 3.3 and 4.2 showed so far that for $n > 2$ a player can't expect to receive her highest payoff, while an opponent individually best responds to the bargain (non-spitefulness). But she can generally expect that an opponent sticks to the bargain if he doesn't gain by a best response. That's important because it allows weak Nash equilibria into the feasible agreements for n players. Section 3.3 also showed that more than one player unilaterally best responding to an agreement can lead to suboptimal payoffs for everybody. Therefore, to generalize worst payoff reasoning, a player needs to consider all possible best responses of all other players and all combinations of them simultanously best responding to the agreement.[4]

[4] This may render many profiles unfeasible, which are not-exploitable by a single player, but anything else ignores the challenge of n-player dynamics.

The worst payoff function for more than two players is thus given as:

Definition 4.2.2 (Worst Payoff Function for n Players)**.** For $n > 2$ players let $\sigma_i \in \Sigma_i$, $i = 1, \dots, n$, be player i's strategy and $u_i(\sigma_i, \sigma_{-i})$ player i's payoff function, where $-i = \{1, \dots, i-1, i+1, \dots, n\}$.
The worst payoff of an agreement $\sigma^A = (\sigma_1^A, \dots, \sigma_n^A)$ for player i is defined as

$$w_i(\sigma_i^A, \sigma_{-i}^A) = \min_{\sigma_l \in B_l(\sigma^A),\, l \in L,\, L \in \mathcal{P}(-i)} \left\{ u_i(\sigma_i^A, \sigma_{-i}^A), u_i(\sigma_i^A, \sigma_{-i\setminus L}^A, \sigma_L) \right\}, \tag{4.3}$$

where $B_l(\sigma^A)$ states the set of player l's best responses to the strategy profile σ^A with strict gain, σ_L is the strategy tuple of opponents from the set $L \in \mathcal{P}(-i)$[5], who unilaterally best respond to σ^A, and $\sigma_{-i\setminus L}^A$ is the strategy tuple of opponents from the set $-i\setminus L$, who don't best respond to σ^A.

Now a player i looks at all possibilities of one opponent best responding to σ^A and all possibilities of two opponents best responding and so on until the case of every opponent best responding. If some opponent has no option to best respond to σ^A, this narrows down the cases. In general, the number of best response possibilities is bounded from above by the sum over the binomial coefficient of $k = 1, \dots, n-1$ opponents best responding to σ^A: $\sum_{k=1}^{n-1} \binom{n-1}{k} = 2^{n-1} - 1$.([11], p.7) Since non-spitefulness is implausible for n players, the maximum in 4.2.2 is dropped for an overall minimum. A player can't expect to receive her maximum payoff if only one opponent best responds and she can't anticipate how many other players might best respond at once. Thus the worst payoff for σ^A should be given by the worst best response scenario or the agreement itself.
Like in definition 2.2.3 a worst-payoff agreement for $n > 2$ players is given as where all players' worst payoff functions 4.2.2 are maximized. Finally, the virtual bargaining equilibrium is adjusted to:

Definition 4.2.3 (Virtual Bargaining Equilibrium)**.** Any strategy profile $\sigma^V = (\sigma_1^V, \dots, \sigma_n^V)$ for n players is called a VBE if

$$\sigma^V \in argmax_{\sigma^F \in F} \prod_{i=1}^{n} (w_i(\sigma^F) - u_i(\sigma^D)), \tag{4.4}$$

where $w_i(\sigma^F) \geq u_i(\sigma^D)$ for all $i = 1, \dots, n$. If no such σ^V exists, the VBE is given by σ^D.

With definition 4.1.1 i's disagreement payoff can now be higher than her worst payoff of some feasible agreement, so she wouldn't accept the latter as the virtual bargain. Therefore, definition 4.2.3 excludes all worst-payoff agreements from the bargaining set where at least one player receives lower worst payoff than in σ^D. Else, there would be problems of minus signs in the product. Hence, a worst-payoff agreement is only *feasible* if every player's worst payoff is higher than or equal to $u_i(\sigma^D)$. Note that $w_i(\sigma^F)$ is related to definition 4.2.1 for $n = 2$ and to definition 4.2.2 for $n > 2$. Whenever the set of $\sigma^V s$ is empty, players fall back on

[5] $\mathcal{P}$ denotes the power set, e.g. $\mathcal{P}(\{1,2\}) = \{\{1\},\{2\},\{1,2\}\}$.

the disagreement point. For extensive games the VBE is still given by the VBE of the corresponding normal form game.

At last, the necessary conditions of n players' beliefs for a virtual bargain are the following. If the players' payoff functions and best response rationality are mutually known and there's mutual belief in the disagreement point 4.1.1, the idea to virtually bargain, maximized worst payoff reasoning with respect to definition 4.2.1 and 4.2.2 and the bargaining axioms on feasible agreements, the resulting strategy profile chosen by the players is a VBE if it's unique[6] or commonly known and no player chooses to best respond to it.[1, 2, 4, 21]
To clarify, this means that every player knows the structure of the game and correctly believes himself and every other player to be best response rational as to choose a possibly mixed BRR maximin strategy to solve a game for uncertain beliefs about the other players' actions. Further, all players need to believe that everybody imagines bargaining by maximizing worst payoff functions with the previous assumptions about other players' best responses such that the feasible agreements are commonly known. Then the virtual bargain is optained if all players believe in the bargaining axioms, coordinate on a specific one and don't best respond to it.
In summary, virtual bargaining as modified in this chapter is now able to account for coordination, mutual advantage and precommitment in the sense of the earlier discussed game examples. In comparison with other theories, which motivate mutual advantage such as *team reasoning* (e.g. [3, 6, 7, 17]) by requiring team identification between the players, virtual bargaining still regards individual reasoners and is based on weaker assumptions.

[6] In fact, the VBE uniquely exists except for generally unsolvable games, e.g. like the Battle of the Sexes.

Chapter 5

Summary and an Outlook

From the outline of the latest research I introduced the concepts and formal definitions of virtual bargaining. It became clear that the theory yields reasonable solutions in many games, which are not suggested by best response, but that it rested on certain problematic assumptions. In expression, the notion of the disagreement point was implausible in certain games, the idea of non-spiteful best responses was unjustified with respect to orthodox game theory and the worst payoff function was inconsistent for n-player games.
Settling with the worst-payoff feasible agreements, I addressed its problems by introducing an adequate disagreement point, which is intuitive for best response reasoners and gives a close connection to worst payoff reasoning. Moreover, I argued that a player will only deviate from a virtual bargain if it strictly increases his payoff, which is equivalent to non-spitefulness for two players in the games of this paper. Furthermore, I consistently generalized the worst payoff function and the VBE for n players. In the end this modified account is able to model and explain many phenomena that puzzled orthodox game theory such as coordination, mutual advantage and precommitment (as in [8, 10, 12]). It's able to do so by introducing the idea to virtually bargain, which is a weaker assumption than for example team identification and entirely motivated by individual reasoning.
Nonetheless, there is much work to be done in validating the underlying claims by real players' intuitions for interactions, which virtual bargaining is to describe. Further, it's unclear how the idea of virtual bargains is represented in an agent's mind and if or how such behavior is learned, reproduced or influenced by social norms. Moreover, it should be examined when real players' choice behavior is guided by the idea of virtual bargains instead of team interest. Another question to answer is how virtual bargaining as a complex reasoning mode can be simplified for boundedly rational agents to meet real players' abilities for advantageous coordination.

Bibliography

[1] Aumann, R., Brandenburger, A., "Epistemic Conditions for Nash Equilibrium", Econometrica, Vol. 63(5), 1161-1180 (1995)

[2] Bach, C., Tsakas, E., ´´Pairwise epistemic conditions for Nash equilibrium", Games and Economic Behavior, 85, 48–59 (2014)

[3] Bacharach M., Gold N., Sugden R., "Beyond Individual Choice: Teams and Frames in Game Theory", Princeton University Press, Princeton (2006)

[4] Barelli, P., "Consistency of beliefs and epistemic conditions for Nash and correlated equilibria", Games and Economic Behavior, 67, 363–375 (2009)

[5] Chater, N., "Can cognitive science create a cognitive economics?", Cognition, 135, 52–55 (2015)

[6] Colman, A., Gold, N., "Team reasoning: Solving the puzzle of coordination", Psychological Bulletin & Review, 25, 1770-1783 (2018)

[7] Karpus, J., Radzvilas, M., "Team Reasoning and a measure of mutual advantage in games", Economics and Philosophy, 34, 1-30 (2018)

[8] Misyak, J., Chater, N., "Virtual bargaining: a theory of social decision making", Phil. Trans. R. Soc. B, 369, 20130487 (2014)

[9] Misyak, J., Melkonyan, T., Zeitoun, H., Chater, N., "Unwritten rules: virtual bargaining underpins social interaction, culture, and society", Trends in Cognitive Sciences, 18(10), 512-519 (2014)

[10] Melkonyan, T., Zeitoun, H., Chater, N., "Collusion in Bertrand vs. Cournot Competition: A Virtual Bargaining Approach", Management Science, Informs, https://doi.org/10.1287/mnsc.2017.2878 (2018)

[11] Melkonyan, T., Zeitoun, H., Chater, N., "Virtual bargaining as a formal account of tacit agreements", http://dx.doi.org/10.2139/ssrn.3168669 (2018)

[12] Nagel, R., Tang, F., "Experimental Results on the Centipede Game in Normal Form: An Investigation on Learning", Journal of Mathematical Psychology, 42, 356-384 (1998)

[13] Nash, J., "Equilibrium points in n-person games", Proceedings of the National Academy of Sciences, 36(1), 48-49 (1950)

[14] Nash, J., "The bargaining problem", Econometrica 18, 155–162 (1950)

[15] Peterson, M., "The Prisoner's Dilemma", Cambridge University Press, Cambridge (2015)

[16] Reiss, J., "Philosophy of Economics: A Contemporary Introduction", Routledge, New York (2013)

[17] Sugden, R., "The logic of team reasoning", Philosophical Explorations 6(3), 165-181 (2003)

[18] Trost, M., "Epistemic characterizations of iterated deletion of inferior strategy profiles in preference-based type spaces", Int J Game Theory, 42, 755–776 (2013)

Internet Sources:

[19] http://www.gametheory.net/Dictionary/games/, lastly visited at the 01.06.2020

[20] https://plato.stanford.edu/entries/epistemic-game/, lastly visited at the 01.07.2020

[21] https://plato.stanford.edu/entries/game-theory/, lastly visited at the 01.07.2020