

Negotiation Strategies for Agents with Ordinal Preferences

Sefi Erlich¹, Noam Hazon¹, Sarit Kraus²

¹ Department of Computer Science, Ariel University

² Department of Computer Science, Bar-Ilan University

erlichsefi@gmail.com, noamh@ariel.ac.il, sarit@cs.biu.ac.il

Abstract

Negotiation is a very common interaction between automated agents. Many common negotiation protocols work with cardinal utilities, even though ordinal preferences, which only rank the outcomes, are easier to elicit from humans. In this work we concentrate on negotiation with ordinal preferences over a finite set of outcomes. We study an intuitive protocol for bilateral negotiation, where the two parties make offers alternately. We analyze the negotiation protocol under different settings. First, we assume that each party has full information about the other party's preference order. We provide elegant strategies that specify a sub-game perfect equilibrium for the agents. We further show how the studied negotiation protocol almost completely implements a known bargaining rule. Finally, we analyze the no information setting. We study several solution concepts that are distribution-free, and analyze both the case where neither party knows the preference order of the other party, and the case where only one party is uninformed.

1 Introduction

Negotiation is a dialogue between two or more parties over one or more issues, where each party has some preferences regarding the discussed issues, and the negotiation process aims to reach an agreement that would be beneficial to the parties. The basic automated negotiation protocol, which consists of two parties that alternate offers, was introduced by [Rubinstein, 1982]. Since then a lot of work has been done to develop other types of protocols, and to extend the basic bilateral negotiation protocol [Fatima *et al.*, 2014].

Many common negotiation protocols work with cardinal utilities, i.e., with utility functions that give different outcomes a specific numerical value, according to the agents' preferences. This representation requires the agents to specify the magnitude of how much they prefer one outcome over another. However, this specification is not always readily available. Moreover, in many cases the agents need to act on behalf of humans, and the use of cardinal utilities for representing human preferences has been widely criticized on

the grounds of cognitive complexity, difficulty of elicitation, and other factors (e.g., [Ali and Ronaldson, 2012]). On the other hand, ordinal preferences only rank the outcomes, so they reduce cognitive burden and are easier to elicit.

Indeed, there are some negotiation protocols that work with ordinal preferences. However, these protocols only start with the ordinal preferences, and they then convert them to a cardinal utility according to some assumptions [Nash Jr, 1950; Shapley, 1969]. Moreover, the traditional assumption in the negotiation theory is that there is a continuum of feasible outcomes. However, many real-life situations involve a finite number of outcomes, such as two managers choosing from among a few job candidates, or a couple choosing from among a few apartments. Even in negotiation over monetary payoffs, the number of outcomes is bounded by the indivisibility of the smallest monetary unit.

In this paper we study negotiation with ordinal preferences over a finite set of outcomes, without converting the ordinal preferences to a cardinal utility¹. We analyze an intuitive protocol for bilateral negotiation that was introduced by [Anbarci, 1993], where the two parties make alternating offers. Each offer is a possible outcome, and we allow the parties to make any offer they would like, in any order. The only restriction is that no offer can be made twice, thus if there are m possible outcomes the negotiation will last at most m rounds.

We analyze the negotiation protocol in different settings. First, we assume that each party has full information about the other party's preference order, and she will thus take this into account and act strategically when deciding what to offer. We provide elegant strategies that specify a Sub-game Perfect Equilibrium (SPE) for the parties. Specifically, our strategies are easy to implement, and we improve the previous result of [Anbarci, 2006] and find a SPE strategy in linear time instead of quadratic time.

We note that there are several works that studied bargaining rules with ordinal preferences over a finite set of outcomes, but they are inherently different from non-cooperative negotiation protocols. A bargaining rule is a function that assigns to each negotiation instance a subset of the outcomes, which are considered the result of the negotiation. These rules are

¹This is also the assumption in most of the voting literature [Brandt *et al.*, 2016], but see [Boutilier *et al.*, 2015] that analyzes social choice functions after extracting cardinal utilities from ordinal preferences.

useful only in a cooperative environment, or where there is a central authority that can force the parties to offer specific outcomes in a specific order.

The proof that our proposed strategies specify a SPE provides us with a deep understanding of the negotiation protocol, which enables us to establish a connection to the results of the designed *Rational Compromise (RC)* bargaining rule [Kıbrıs and Sertel, 2007]. Surprisingly, the SPE result of the negotiation protocol is always part of the set of results returned by the *RC* rule, even though the protocol does not force the parties to offer specific outcomes in a specific order as the *RC* rule does. This connection also enables us to prove that the SPE result of the protocol is monotonic.

We then move on to analyze the *no information* setting. We analyze the case where neither party knows the preference order of the other party nor do they know prior probability distribution over possible orders. In this setting we first show that an ex-post SPE does not exist. We then provide the maxmin strategy and the maxmin value of the game that is imposed by our protocol. We also show that in our setting, surprisingly, any pair of maxmin strategies also specifies a robust-optimization equilibrium. Finally, we consider the case where one party has full information while the other party has no information, and show how the informed party is able to use her information so that the negotiation result will be better for her.

The contribution of this work is threefold. First, we introduce elegant strategies that specify a SPE, and provide a substantial analysis for showing that they indeed form a SPE. We also provide an improved algorithm for computing a SPE strategy for the studied negotiation protocol. The second contribution of our work is that we show how the studied negotiation protocol almost completely implements the *RC* rule. As noted by [Kıbrıs and Sertel, 2007], who studied the *RC* rule, the descriptive relevance of the *RC* rule for real-life bargaining depends on the existence of non-cooperative games that implement it, and to the best of our knowledge our paper is the first to find such a connection. Finally, we provide an analysis of the negotiation protocol under a no information setting, which has not been considered before.

2 Related Work

Negotiation protocols and the strategic interaction they imply have been extensively studied. We refer to the books of [Osborne and Rubinstein, 1990], [Kraus, 2001], and [Fatima *et al.*, 2014] for extensive coverage of the different approaches. The traditional assumption in negotiation theory is that there is a continuum of feasible outcomes, but many real-life negotiation scenarios violate these assumptions. Indeed, there are several works that consider problems with a finite number of outcomes. For example, see [Zlotkin and Rosenschein, 1996], [Mariotti, 1998], [Nagahisa and Tanaka, 2002] and recently [Nunez and Laslier, 2015]. All of these works focus on negotiation when the preferences are represented by a cardinal utility, while we study negotiation with ordinal preferences.

Many other works study negotiation with ordinal preferences over a finite set of outcomes (for example, [Zhang

and Zhang, 2008]). Sequential procedures, in particular the *fallback bargaining* method, have attracted considerable interest [Sprumont, 1993; Hurwicz and Sertel, 1999; Brams and Kilgour, 2001; Kıbrıs and Sertel, 2007; Conley and Wilkie, 2012], since they satisfy some nice theoretical properties. All of these works study bargaining rules that are useful in a cooperative environment. We study a negotiation protocol that is useful in a non-cooperative environment, and show that its SPE almost completely implements the individually rational variant of the *fallback bargaining* method, i.e., the *RC* rule [Kıbrıs and Sertel, 2007]. We note that the *RC* bargaining rule is equivalent to *Bucklin* voting with two voters, and thus our result can also be interpreted as a (weak) SPE implementation of the *Bucklin* rule where there are two voters.

There are few works that study negotiation protocols with ordinal preferences over a finite set of outcomes. [De Clippeel *et al.*, 2014] study the problem of selection of arbitrators, and they concentrate on two-step protocols. The most closely related works are Anbarci’s papers. In [Anbarci, 1993] he introduces the *Voting by Alternating Offers and Vetoes (VAOV)* negotiation protocol, which we study here, and shows the possible SPE results in different scenarios. Implicitly, this work shows that the SPE result is unique and Pareto optimal. In [Anbarci, 2006] he introduces three additional negotiation protocols. Moreover, he sharpens his previous result by exactly identifying the SPE result of the VAOV protocol, and by providing an algorithm that computes a SPE strategy. He also shows that if the outcomes are distributed uniformly over the comprehensive utility possibility set and as the number of outcomes tends to infinity the VAOV protocol converges to the equal area rule [Thomson, 1994]. We provide a more efficient algorithm that finds an elegant SPE strategy. In addition, we were able to establish the relationship between the VAOV protocol and the *RC* rule, which works with a *finite* number of outcomes, and we also analyze the no information setting.

3 The Negotiation Protocol

We assume that there are two negotiation parties, p^1 and p^2 , negotiating over a set of potential outcomes $O = \{o_1, \dots, o_m\}$, where p^1 is the party that makes the first offer. Each party has a preference order over the potential outcomes that does not permit any ties. Formally, the preferences of a party p are a strict order, \succ_p , which is a complete, transitive and irreflexive binary relation on O . We write $o' \succ_p o$ to denote that party p strictly prefers o' to o , and $o' \succeq_p o$ to denote that $o' \succ_p o$ or $o' = o$ (i.e o' is the exact same outcome as o). Clearly, each party would like to maximize her utility, i.e., that the result of the negotiation will be the outcome that is placed as high as possible according to her preferences.

We analyze the following negotiation protocol, which is the VAOV protocol of [Anbarci, 1993]. The parties make offers alternately. No offer can be made twice, but an agreement must be reached since we assume that any agreement is preferred by both parties over a no-agreement result. We also assume that lotteries are not valid offers, as in most real-life negotiations. Formally, denote by O_t the set of available outcomes at round t , and let $O_1 = O$. At round 1, party p^1

offers an outcome $o \in O_1$ to p^2 . If p^2 accepts, the negotiation terminates successfully with o as the result of the negotiation. Otherwise, party p^2 offers an outcome $o' \in O_2 = O_1 \setminus \{o\}$. If p^1 accepts, the negotiation terminates successfully with o' as the result of the negotiation. Otherwise, p^1 offers an outcome $o'' \in O_3 = O_2 \setminus \{o'\}$ to p^2 , and so on. If no offer was accepted until round m then the last available outcome is accepted in the last round as the result of the negotiation. We denote by p^i the party whose turn it is to make an offer at a given round, and by p^j the other party. That is, $p^i = p^1$ in odd round numbers and $p^i = p^2$ otherwise.

We first provide a general result that is useful with any model of information. Consider the following definition:

Definition 1. *In each round t , let L_t^j be the $\lfloor |O_t|/2 \rfloor$ lowest ranked outcomes in \succ_{p^j} . If $|O_t|$ is odd, then let L_t^i be the $\lfloor |O_t|/2 \rfloor$ lowest ranked outcomes in \succ_{p^i} . If $|O_t|$ is even, then L_t^i is the $|O_t|/2 - 1$ lowest ranked outcomes in \succ_{p^i} .*

We show that in each round t we can identify a set of outcomes that cannot be the negotiation result if the parties are rational, regardless of the information they have. Intuitively, these are all the outcomes that are in the lower parts of the preference orders of both parties, denoted by Low_t . We denote all of the other outcomes by JG_t .

Definition 2. *Given a round t , let $Low_t = \{o : o \in L_t^i \cup L_t^j\}$, and $JG_t = O_t \setminus Low_t$.*

Lemma 1. *Let o be the result of the negotiation if both parties are rational. Then, $o \notin Low_t$.*

Proof. Starting from round t where $|O_t| = m_t$, each party will be able to reject all of the offers that she would receive from the other party, except for the offer she would receive in the last round. Specifically, if m_t is odd, p^i and p^j can reject at most $\lfloor m_t/2 \rfloor$ offers. If m_t is even, p^i can reject at most $\lfloor m_t/2 \rfloor - 1$ offers (since it is p^i 's turn to offer) and p^j can reject at most $\lfloor m_t/2 \rfloor$ offers. That is, each party p^k , $k \in \{1, 2\}$, can reject at most $\lfloor L_t^k \rfloor$ offers. Therefore, each party will always be able to guarantee that the result of the negotiation will be an outcome that is placed higher than the $\lfloor L_t^k \rfloor$ lowest outcomes in her preference order. Therefore, $o \notin Low_t$. \square

We now analyze the negotiation protocol under two different models of information: full information and no information. In each case we are interested in finding the best actions that a party should take, given the information that she has.

4 Full Information

In this setting we assume that each party has full information about the other party's preference order, and she will thus take this information into account when calculating her best strategy. Therefore, in the full information setting we are interested in finding a SPE. Since [Anbarci, 1993] showed that the SPE result is unique, it suffices to find one set of strategies that specify a SPE.

Recall that SPE is a mapping that maps the histories of players' choices. Note that in our case if an offer was accepted the game is over. Therefore, a history for p^i , the party

whose turn it is to make an offer at a given round, consists of a sequence of outcomes that were proposed and rejected in the previous rounds. Let $H_t^i = (o^1, o^2, \dots, o^{t-1})$ be the history for p^i at round t , and note that $O_t = O \setminus H_t^i$. Let o_t^- be the least preferred outcome in O_t according to \succ_{p^j} . We show that the following strategy specifies the offering strategy in our SPE.

Strategy 1 (OFFERING STRATEGY). *Given a history H_t^i , if $I_t = L_t^i \cap L_t^j \neq \emptyset$ then offer $o \in I_t$, else offer o_t^- .*

A history for p^j , the party whose turn it is to decide whether to accept or reject an offer at a given round, consists of a sequence of outcomes that were proposed and rejected in the previous rounds and an additional outcome o that was offered by p^i in the current round. Let $H_t^j = H_t^i + o = (o^1, o^2, \dots, o^{t-1}, o)$ be the history for p^j at round t . In addition, given a round t and history for p^i , H_t^i , let o_t^i be an outcome $o \in I_t = L_t^i \cap L_t^j$ if $I_t \neq \emptyset$, and o_t^- otherwise. Given a round t and history for p^j , H_t^j , let $o_a(H_t^j)$ be the single outcome in $O_m = O \setminus H_m^i$, where $H_m^i = H_t^j + o_{t+1}^i + \dots + o_{m-1}^i$. That is, $o_a(H_t^j)$ is the result of the negotiation if both parties reject all of the offers that they get (except for the last offer) from round t and on, but use the offering strategy that is specified by Strategy 1 from round $t+1$ and on. We show that the following strategy specifies the response strategy in our SPE.

Strategy 2 (RESPONSE STRATEGY). *Given a history H_t^j , if $o \succeq_{p^j} o_a(H_t^j)$ then accept o , else reject o .*

To illustrate the strategies of our SPE, consider the following examples:

Example 1. *Suppose that*

$$\begin{aligned} \succ_{p^1} &= o_6 \succ o_5 \succ o_4 \succ o_3 \succ o_2 \succ o_1 \\ \succ_{p^2} &= o_1 \succ o_3 \succ o_2 \succ o_6 \succ o_4 \succ o_5. \end{aligned}$$

Following Definition 1, $L_1^1 = \{o_2, o_1\}$ and $L_1^2 = \{o_6, o_4, o_5\}$. Therefore, $I_1 = \emptyset$ and according to the offering strategy (Strategy 1) p^1 would offer p^2 's least preferred outcome - o_5 . Then, according to the response strategy (Strategy 2) p^2 would reject, since $o_a((o_5)) = o_3 \succ_{p^2} o_5$, as we will show. In round 2, $L_2^1 = \{o_1, o_2\}$ and $L_2^2 = \{o_6, o_4\}$, and thus p^2 would offer o_1 . Now p^1 would reject, and offer o_4 , p^2 would reject and offer o_2 , p^1 would reject and offer o_6 , and in the final round p^2 would reject and offer o_3 which is accepted as the result of the negotiation since no other outcome is available.

Example 2. *Now suppose that*

$$\begin{aligned} \succ_{p^1} &= o_6 \succ o_5 \succ o_4 \succ o_3 \succ o_2 \succ o_1 \\ \succ_{p^2} &= o_1 \succ o_3 \succ o_6 \succ o_2 \succ o_4 \succ o_5. \end{aligned}$$

Following Definition 1, $L_1^1 = \{o_2, o_1\}$ and $L_1^2 = \{o_2, o_4, o_5\}$. Therefore, $I_1 = \{o_2\}$ and according to Strategy 1 p^1 would offer o_2 . Then, p^2 will reject, since $o_a((o_2)) = o_6 \succ_{p^2} o_2$, as we will show. In round 2, $L_2^1 = \{o_3, o_1\}$ and $L_2^2 = \{o_4, o_5\}$, $I_2 = \emptyset$, and thus p^2 would offer o_1 . In each subsequent round the parties would offer each other the least preferred outcomes, until the final round where o_6 will be accepted as the result of the negotiation.

In order to prove that our strategies specify a SPE, we need a deeper understating of the offering and response strategies. We note that in the offering strategy (Strategy 1), p^i offers an outcome from the set I_t if it is not empty. We now show the relation between the set I_t and the set JG_t . Due to space constraints, we defer the proofs of some of the lemmas to the full version of the paper [Erlich *et al.*, 2018].

Lemma 2. $|JG_t| = |I_t| + 1$

Considering Lemma 1, we show a simple corollary. Let o_{eq} be the SPE result. We get:

Corollary 1. $o_{eq} \notin Low_t$.

If we combine the findings from Corollary 1 and Lemma 2, we get that if the set I_t is empty, i.e., the intersection between the lower parts of the preference orders of the parties is empty, then the set JG_t contains only one outcome, o_{eq} .

Corollary 2. If $I_t = \emptyset$ then $JG_t = \{o_{eq}\}$.

Next, we show how the transition from round t to round $t + 1$ affects the number of outcomes in L_{t+1}^k , $k \in \{1, 2\}$.

Lemma 3. Suppose that in round t , p^i offered an outcome o and p^j rejected it, then in round $t + 1$, $|L_{t+1}^i| = |L_t^j| - 1$ and $|L_{t+1}^j| = |L_t^i|$

Proof. Assume $|O_t| = m_t$ is even, then by definition $|L_t^i| = \frac{m_t}{2} - 1$ and $|L_t^j| = \frac{m_t}{2}$. After that p^i offered the outcome o and p^j rejected it, m_{t+1} is odd, and the roles are switched between p^i and p^j . Therefore, $|L_{t+1}^i| = |L_{t+1}^j| = \lfloor \frac{m_{t+1}}{2} \rfloor = \lfloor \frac{m_t - 1}{2} \rfloor = \lfloor \frac{m_t}{2} - \frac{1}{2} \rfloor = \frac{m_t}{2} - 1$. Now assume that m_t is odd, then $|L_t^i| = |L_t^j| = \lfloor \frac{m_t}{2} \rfloor = \frac{m_t - 1}{2}$. After that p^i offered the outcome o and p^j rejected it, m_{t+1} is even, and the roles are switched between p^i and p^j . Therefore, $|L_{t+1}^j| = \frac{m_{t+1}}{2}$ and $|L_{t+1}^i| = \frac{m_{t+1}}{2} - 1$. \square

We note that the number of outcomes in L_t^k is important, since we already showed in Corollary 1 that these are the outcomes that cannot be an equilibrium result. Indeed, it is more important to understand how the transition from round t to round $t + 1$ affects which outcomes become part of L_{t+1}^k . Obviously, it depends on the offer that was made in round t . The following three lemmas analyze this transition, based on the offers that are made according to Strategy 1. Specifically, Lemma 4 together with Lemma 5 cover the offering strategy where $I_t = \emptyset$, and Lemma 5 together with Lemma 6 cover the offering strategy where $I_t \neq \emptyset$.

Lemma 4. In round t , if p^i offers $o \notin L_t^i$ and p^j rejects it, then $L_{t+1}^j \leftarrow L_t^i$

Lemma 5. In round t , if p^i offers $o \in L_t^j$ and p^j rejects it, then $L_{t+1}^i \leftarrow L_t^j \setminus \{o\}$

Lemma 6. In round t , if p^i offers $o \in L_t^i$ and p^j rejects it, then $L_{t+1}^j \leftarrow L_t^i \setminus \{o\} \cup \{o'\}$

Note that when p^j follows Strategy 2 she computes the outcome $o_a(H_t^j)$ to decide whether to accept or reject the offer that she gets from p^i . By definition,

Lemma 7. $o_a(H_t^j) = o_a(H_t^j + o_{t+1}^i)$.

We now show that if, in a given round t , p^i follows Strategy 1, i.e., $o_a(H_t^j) = o_a(H_t^i + o_t^i)$, then $o_a(H_t^j)$ has some desirable properties. For example, it is Pareto optimal in O_t , i.e., $\forall o \in O_t \setminus \{o_a(H_t^i + o_t^i)\}$, $o_a(H_t^i + o_t^i) \succ_{p^i} o$ or $o_a(H_t^i + o_t^i) \succ_{p^j} o$. For ease of notation, let $o_{at} = o_a(H_t^i + o_t^i)$.

Lemma 8. Given any history H_t^i , $o_t^i \prec_{p^j} o_{at}$, $o_{at} \in JG_t$, and o_{at} is Pareto optimal in O_t .

Proof. We prove by induction on m . If $m = 2$ and $t = 1$, without loss of generality (WLOG) assume that $\succ_{p^2} = o_1 \succ o_2$. Since $I_t = \emptyset$, $o_t^1 = o_t^- = o_2$ and $H_t^1 = ()$ by definition. Thus, $o_{at} = o_a(H_t^1 + o_t^1) = o_1$. In addition, $JG_t = \{o_1\}$, and therefore $o_t^1 \prec_{p^2} o_{at}$ and $o_{at} \in JG_1$ as required. Since $O_t \setminus \{o_{at}\} = \{o_2\}$ and $o_2 \prec_{p^2} o_{at}$ then o_{at} is also Pareto optimal in O_t . Now, assume that if there are $m - 1$ outcomes in round $t + 1$, $o_{t+1}^i \prec_{p^j} o_a(H_{t+1}^i + o_{t+1}^i)$, $o_a(H_{t+1}^i + o_{t+1}^i) \in JG_{t+1}$ and $o_a(H_{t+1}^i + o_{t+1}^i)$ is Pareto optimal in O_{t+1} . We show that when there are m outcomes in round t , $o_t^i \prec_{p^j} o_a(H_t^i + o_t^i) = o_{at}$, $o_{at} \in JG_t$ and o_{at} is Pareto optimal in O_t . According to Lemma 7, if $H_t^j = H_t^i + o_t^i$ we get that $o_{at} = o_a(H_t^i + o_t^i) = o_a(H_t^i + o_t^i + o_{t+1}^i) = o_a(H_{t+1}^i + o_{t+1}^i)$. Now, if $I_t = \emptyset$ and thus $o_t^i = o_t^- \in L_t^j$ then according to Lemmas 4 and 5 $JG_{t+1} = JG_t$. If $I_t \neq \emptyset$ and thus $o_t^i \in I_t$ then according to Lemmas 5 and 6 $JG_{t+1} \subseteq JG_t$. According to the induction assumption, $o_a(H_{t+1}^i + o_{t+1}^i) \in JG_{t+1}$ and thus $o_{at} = o_a(H_{t+1}^i + o_{t+1}^i) \in JG_t$. In addition, $o_t^i \in L_t^j$ by definition and since we showed that $o_{at} \in JG_t$, we get that $o_t^i \prec_{p^j} o_{at}$. Finally, since o_{at} is Pareto optimal in O_{t+1} , $O_t = O_{t+1} \cup \{o_t^i\}$, and $o_t^i \prec_{p^j} o_{at}$, we conclude that o_{at} is Pareto optimal in O_t . \square

Rephrasing Lemma 8, we showed that given any history H_t^i , if both parties follow Strategies 1 and 2 from round t and on, p^j would always reject the offers that she gets from p^i (i.e., $o_t^i, o_{t+1}^i, \dots, o_{m-1}^i$), and the negotiation result would be o_{at} , which would be accepted in the last round. Moreover, the negotiation result o_{at} is Pareto optimal in O_t, O_{t+1}, \dots, O_m .

Before we prove that Strategies 1 and 2 specify a SPE we need to add some definitions. We first define a distance function for each party p^k , that given an outcome $o_x \notin L_t^k$ counts the number of outcomes $o \notin L_t^k$ such that $o_x \succeq_{p^k} o$. Intuitively, this is the number of outcomes a party can offer until a round t' where o_x becomes part of $L_{t'}^k$. Formally:

Definition 3. $d_{k,x,t} = |\{o \in O_t : o_x \succeq_{p^k} o \wedge o \notin L_t^k\}|$ where $k \in \{1, 2\}$

We also define the number of offers that are made before reaching a round t' where $I_{t'} = \emptyset$.

Definition 4. Let $\ell_{k,t}$ be the number of offers a party p^k offers according to Strategy 1 from round t until round t' where $I_{t'} = \emptyset$.

Recall our previous examples. In Example 1 at round 1, $I_1 = \emptyset$ and thus $\ell_{1,1} = \ell_{2,1} = 0$. The distance of o_3 at round 1 is $d_{1,3,1} = 1$ for party p^1 and $d_{2,3,1} = 2$ for party p^2 . In

Example 2, $I_1 \neq \emptyset$ but $I_2 = \emptyset$ and thus $\ell_{1,1} = 1$ and $\ell_{2,1} = 0$. The distance of o_6 at round 1 for the parties is $d_{1,6,1} = 4$ and $d_{2,6,1} = 1$, and the distance of o_3 at round 1 for the parties is $d_{1,3,1} = 1$ and $d_{2,3,1} = 2$.

We also make the following simple observation, which is true since we use an alternating offers protocol:

Lemma 9. *At any round t , $\ell_{j,t} \leq \ell_{i,t}$.*

Our main theorem is as follows:

Theorem 10. *Strategies 1 and 2 specify a SPE.*

Proof. We prove by induction on m . If $m = 2$ and $t = 1$, WLOG assume that $\succ_{p^2} = o_1 \succ o_2$. Thus, $JG_t = \{o_1\}$, and according to Corollary 2, o_1 is the SPE result. Indeed, according to Strategy 1 p^1 will offer o_2 in the first round and p^2 will reject it according to Strategy 2 since $o_a(o_2) = o_1 \succ_2 o_2$. In the next round $p^i = p^2$ will offer o_1 , p^1 will accept it and then the negotiation will end with o_1 as the negotiation result. Clearly there are only two states where p^2 has the option to deviate: on the equilibrium path, i.e., where $H_t^2 = (o_2)$, and off the equilibrium path, i.e., where $H_t^2 = (o_1)$. Where $H_t^2 = (o_2)$, p^2 has no incentive to deviate from the response strategy and accept the offer of o_2 from p^1 , since $o_1 \succ_2 o_2$. Where $H_t^2 = (o_1)$, p^2 has no incentive to deviate and reject the offer of o_1 , since then o_2 would become the last available outcome and thus the negotiation result, but $o_1 \succ_2 o_2$. Similarly, there is only one state where p^1 has the option to deviate, i.e., $H_t^1 = ()$. In this state p^1 has no incentive to deviate from the offering strategy and offer o_1 , since p^2 will accept it (because $o_1 \succ_2 o_2$) and o_1 is already the SPE result if p^1 follows the offering strategy.

Now, assume that if there are $m - 1$ outcomes in round $t + 1$, our strategies specify a SPE. We show that they specify a SPE when there are m outcomes in round t . We first consider the strategy of p^j at round t . Note that $o_a(H_t^j) = o_a(H_t^i + o) = o_a(H_{t+1}^i)$ by definition, and $o_a(H_{t+1}^i)$ is the SPE result of following our strategies from state H_{t+1}^i according to Lemma 8 combined with the induction assumption. Clearly, if according to the response strategy (Strategy 2) p^j should reject the offer o , it is because $o_a(H_t^j) \succ_{p^j} o$. Therefore, it is not worthwhile for p^j to deviate and accept o instead of $o_a(H_{t+1}^i) = o_a(H_t^j)$. Similarly, if according to the response strategy p^j should accept an offer o , it is because $o \succeq_{p^j} o_a(H_t^j)$. Therefore, it is not worthwhile for p^j to deviate and reject o in order to get as the negotiation result the outcome $o_a(H_{t+1}^i) = o_a(H_t^j)$. Overall, p^j does not have an incentive to deviate in round t . According to the induction assumption, Strategies 1 and 2 specify a SPE when there are $m - 1$ outcomes in round $t + 1$. Therefore, p^j does not have any incentive to deviate.

We now concentrate on the strategy of p^i at round t , but we first derive some general inequalities. Given a history H_t^i , suppose that there is an outcome $o_x \in O_t$ such that $o_x \succ_{p^i} o_{at} = o_a(H_t^i + o_x^i)$. According to Lemma 8, since $o_x \succ_{p^i} o_{at}$, $o_{at} \succ_{p^j} o_x$. Suppose that both parties follow strategies 1 and 2, and let t' be the round in which $JG_{t'} = \{o_{at}\}$. Then, in round t , $\ell_{i,t} < d_{i,at,t}$ and $\ell_{j,t} < d_{j,at,t}$ (otherwise, $o_{at} \notin JG_{t'}$). By definition, $t' = \ell_{i,t} + \ell_{j,t}$. In

addition, since $JG_{t'} = \{o_{at}\}$, o_x must be part of $Low_{t'}$ for some $t'' < t'$ (otherwise, $o_x \in JG_{t'}$). Since $o_x \succ_{p^i} o_{at}$ and $\ell_{i,t} < d_{i,at,t}$, it must be that o_x is part of $L_{t''}^j$, that is, $d_{j,x,t} \leq \ell_{j,t}$. In summary:

$$\begin{aligned} \ell_{i,t} &< d_{i,at,t} < d_{i,x,t} \\ d_{j,x,t} &\leq \ell_{j,t} < d_{j,at,t} \end{aligned} \quad (1)$$

Now assume that in round t p^i deviates, and the result of the negotiation, if both parties follow our strategies from round $t + 1$, is o_x . Note that p^i in round $t + 1$ is p^j in round t , and thus $o_{at} \succ_{p^i} o_x$. Therefore, we use the same arguments as above to get

$$\begin{aligned} \ell_{i,t+1} &< d_{i,x,t+1} < d_{i,at,t+1} \\ d_{j,at,t+1} &\leq \ell_{j,t+1} < d_{j,x,t+1} \end{aligned} \quad (2)$$

Now, assume by contradiction that there is an outcome $o_d \notin I_t$ such that if p^i offers o_d the negotiation result will be o_x , $o_x \succ_{p^i} o_{at}$. We first analyze the case where p^j rejects the offer of o_d , since $o_x \succ_{p^j} o_d$ (otherwise, p^j would have accepted). We examine the change in the distance function for p^i and p^j , for outcomes o_{at} and o_x , from round t to round $t + 1$. According to Lemma 3, $|L_{t+1}^i| + 1 = |L_t^i|$, and since $o_x \succ_{p^j} o_d$ and $o_{at} \succ_{p^j} o_d$, $d_{j,x,t}$ and $d_{j,at,t}$ do not change when moving to round $t + 1$. Let c be an integer. Then:

$$\begin{aligned} d_{i,at,t} + c &= d_{j,at,t+1} \\ d_{j,at,t} &= d_{i,at,t+1} \\ d_{j,x,t} &= d_{i,x,t+1} \end{aligned} \quad (3)$$

If we combine (2) and (3) we get that:

$$\begin{aligned} \ell_{j,t+1} &< d_{j,x,t} < d_{j,at,t} \\ d_{i,at,t} + c &\leq \ell_{i,t+1} \end{aligned}$$

Adding (1) we get that:

$$\begin{aligned} \ell_{i,t} &< \ell_{i,t+1} - c \\ \ell_{i,t+1} &< \ell_{j,t} \end{aligned}$$

Adding Lemma 9 we can conclude that: $\ell_{i,t+1} < \ell_{j,t} \leq \ell_{i,t} < \ell_{i,t+1} - c$. That is, $\ell_{i,t+1} \leq \ell_{i,t+1} - c - 2$, thus $c \leq -2$. However, the distance function cannot decrease by more than 1 when moving from round t to $t + 1$, thus $c \geq -1$.

We now analyze the case where p^j accepts the offer of o_d , since $o_d \succeq_{p^j} o_x$. We examine the change in the distance function for p^i and p^j , for outcomes o_{at} and o_x , from round t to round $t + 1$. Note that since p^i deviates, $o_d \succ_{p^i} o_{at}$. According to Lemma 8, $o_{at} \succ_{p^j} o_d$. According to Lemma 3, $|L_{t+1}^i| + 1 = |L_t^i|$, and since $o_{at} \succ_{p^j} o_d$, $d_{j,at,t}$ does not change when moving to round $t + 1$. However, since $o_d \succ_{p^j} o_x$, $d_{j,x,t}$ increases by one when moving to round $t + 1$. Let c be an integer. Then:

$$\begin{aligned} d_{i,at,t} + c &= d_{j,at,t+1} \\ d_{j,at,t} &= d_{i,at,t+1} \\ d_{j,x,t} + 1 &= d_{i,x,t+1} \end{aligned} \quad (4)$$

If we combine (2) and (4) we get that:

$$\begin{aligned} \ell_{i,t+1} &< d_{j,x,t} + 1 < d_{j,at,t} \\ d_{i,at,t} + c &\leq \ell_{j,t+1} \end{aligned}$$

Adding (1) we get that:

$$\begin{aligned} \ell_{i,t} &< \ell_{j,t+1} - c \\ \ell_{i,t+1} - 1 &< \ell_{j,t} \end{aligned}$$

Adding Lemma 9 we can conclude that: $\ell_{i,t+1} - 1 < \ell_{j,t} \leq \ell_{i,t} < \ell_{j,t+1} - c$. That is, $\ell_{j,t+1} - 1 \leq \ell_{j,t+1} - c - 2$, thus $c \leq -1$. However, in order for $d_{i,at,t}$ to decrease by at least one, $o_{at} \succ_{p^i} o_d$, but in our case $o_d \succ_{p^i} o_{at}$.

Overall, we showed that p^i does not have an incentive to deviate in round t . According to the induction assumption, Strategies 1 and 2 specify a SPE when there are $m - 1$ outcomes in round $t + 1$. Therefore, p^i does not have any incentive to deviate. \square

Finally, note that trivial exploration of the whole game tree in order to derive the SPE would take at least $O(2^m)$ operations, since there can be $m - 1$ rounds in which a party p^i can offer any outcome from the available outcomes and the other party p^j can decide either to accept the offer or reject it. The complexity of finding a SPE strategy of [Anbarci, 2006] is not explicitly analyzed, but its running time is at least $O(m^2)$. Our approach provides elegant strategies that are easy to implement and are (computationally) more efficient: given a state in the game tree (i.e. given any history H_t^i or H_t^j), we compute a SPE strategy from the current state in time that is linear in m . Indeed, in our approach we only need to simulate one branch of the tree (to find $o_a(H_t^i)$ or $o_a(H_t^j)$) and then trace the intersection between L_t^i and L_t^j .

4.1 Properties

We first note that since we showed that the result of following Strategies 1 and 2 is Pareto optimal, we proved that they specify a SPE, and the SPE result is unique, we can infer that the SPE result of the protocol is Pareto optimal. We now move to analyze the relationship between the SPE result of the protocol and the results of the designed *Rational Compromise (RC)* bargaining rule [Kıbrıs and Sertel, 2007]. The *RC* rule is a private case of the *Unanimity Compromise* rule, where any agreement is preferred by both parties over a no-agreement result, as we assume. With our notations, the *RC* rule can be rephrased as the set $RC = \{o_x \mid \max_{o_x \in O} \min_{k \in \{1,2\}} (d_{k,x,1} + |L_1^k| - 1)\}$. It can also be computed by the following steps:

1. Let $v = 1$
2. For each $k \in \{1, 2\}$, let $B_v^k = \{\text{the } v \text{ most preferred outcomes in } \succ_{p^k}\}$.
3. If $|B_v^1 \cap B_v^2| > 0$ then return $B_v^1 \cap B_v^2$ as the result.
4. Else, $v \leftarrow v + 1$ and go to line 2.

We note that the *RC* rule may return either one or two outcomes, while our protocol always results with a single outcome. Surprisingly, the SPE result of the negotiation protocol is always part of the set returned by the *RC* rule. The intuition is that our strategies specify a SPE by making offers and

rejecting them until $I_t = \emptyset$. At this stage $JG_t = \{o_{eq}\}$, and by definition the set JG_t is the intersection of the upper parts of the preferences of both parties, which corresponds to the $B_v^1 \cap B_v^2$ returned by *RC*.

Theorem 11. $o_{eq} \in RC$

Proof. Let t be the round where $I_t = \emptyset$ after both parties follow our strategies. By Corollary 2, $JG_t = \{o_{eq}\}$. Rephrasing the definition of JG_t we get that $JG_t = B_{|O_t|-|L_t^i|}^i \cap B_{|O_t|-|L_t^j|}^j$. If $|L_t^j| = |L_t^i|$, then for any v where $v \leq |O_t| - |L_t^j|$, $B_v^i \cap B_v^j = \{o_{eq}\}$ or $B_v^i \cap B_v^j = \emptyset$. If $|L_t^j| = |L_t^i| + 1$, then for any v where the $v \leq |O_t| - |L_t^j|$, $B_v^i \cap B_v^j = \{o_{eq}\}$ or $B_v^i \cap B_v^j = \emptyset$, and for $v = |O_t| - |L_t^i|$ it is possible that $B_v^i \cap B_v^j = \{o_{eq}, o_x\}$, for some outcome o_x . Overall, $o_{eq} \in RC$. \square

Based on Theorem 11, we can derive interesting results regarding the relationship between the *RC* rule and the SPE result of the negotiation protocol:

Theorem 12. 1. If $RC = \{o\}$ then $o_{eq} = o$.

2. If o_{eq} is the SPE result let $o_{eq'}$ be the SPE result if p^1 and p^2 switch their rules (i.e., p^2 starts the negotiation). If $o_{eq} \neq o_{eq'}$, then $RC = \{o_{eq}, o_{eq'}\}$
3. If m is odd and $\ell_{1,1} + \ell_{2,1}$ is even or if m is even and $\ell_{1,1} + \ell_{2,1}$ is odd, then $|RC| = 1$.
4. If $|RC| = \{o_x, o_y\}$ and $\ell_{1,1} + \ell_{2,1}$ is odd then $o_{eq} = o_x$ and $o_x \succ_{p^1} o_y$. If $\ell_{1,1} + \ell_{2,1}$ is even then $o_{eq} = o_y$ and $o_y \succ_{p^2} o_x$.

Proof. 1. An easy corollary of Theorem 11.

2. An easy corollary of Theorem 11.

3. If m is odd and $\ell_{1,1} + \ell_{2,1}$ is even or if m is even and $\ell_{1,1} + \ell_{2,1}$ is odd, then m_t is odd. Therefore, $|L_t^1| = |L_t^2|$ by definition. Then, by Theorem 11, for any v where $v \leq |O_t| - |L_t^1|$, $B_v^1 \cap B_v^2 = \{o_{eq}\}$ or $B_v^1 \cap B_v^2 = \emptyset$. That is, $RC = \{o_{eq}\}$.

4. $|RC| = 2$, thus there exists v such that $B_v^1 \cap B_v^2 = \{o_x, o_y\}$, and for every $v' < v$, $B_{v'}^1 \cap B_{v'}^2 = \emptyset$. From Theorem 11, $o_{eq} = o_x$ or $o_{eq} = o_y$. Let t be the round such that $I_t = \emptyset$ and $JG_t = \{o_{eq}\}$. That is, $B_{|O_t|-|L_t^1|}^1 \cap B_{|O_t|-|L_t^2|}^2 = \{o_{eq}\}$. Therefore, $|L_t^1| \neq |L_t^2|$, and thus m_t is even. If $\ell_{1,1} + \ell_{2,1}$ is odd then it is p^2 's turn to offer. That is, $|L_t^2| + 1 = |L_t^1|$, and since $o_x \succ_{p^1} o_y$, $o_y \in L_t^1$. Therefore, $o_{eq} = o_x$. Similarly, if $\ell_{1,1} + \ell_{2,1}$ is even then it is p^1 's turn to offer. That is, $|L_t^1| + 1 = |L_t^2|$, and since $o_y \succ_{p^2} o_x$, $o_x \in L_t^2$. Therefore, $o_{eq} = o_y$. \square

Finally, we adapt the monotonicity criterion that the *RC* rule satisfies for our domain, and show that the negotiation protocol is monotonic.

Definition 5. A negotiation protocol is monotonic if given an instance $(O, \succ_{p^1}, \succ_{p^2})$ where the SPE result is o_{eq} , then for any instance $(O', \succ'_{p^1}, \succ'_{p^2})$ such that:

1. $O \subset O'$.
2. For any $o_1, o_2 \in O, o_1 \neq o_2$, and for $k \in \{1, 2\}$, if $o_1 \succ_{p^k} o_2$ then $o_1 \succ'_{p^k} o_2$.
3. For any $o \in O' \setminus O$, and for $k \in \{1, 2\}$, $o \succ'_{p^k} o_{eq}$.

we have that $o'_{eq} \succ'_{p^k} o_{eq}$.

Theorem 13. *The negotiation protocol is monotonic.*

Proof. Given an instance $(O, \succ_{p^1}, \succ_{p^2})$, we know from Theorem 11 that $o_{eq} \in RC$. If we add a set of outcomes $O' \setminus O$ such that for every outcome $o \in O' \setminus O$, $o \succ o_{eq}$ for both parties, then for every outcome o' in the set returned by the RC rule on the modified instance $(O', \succ'_{p^1}, \succ'_{p^2})$, $o' \succ o_{eq}$ by both parties. Since $o'_{eq} \in RC$ on $(O', \succ'_{p^1}, \succ'_{p^2})$, we get that $o'_{eq} \succ o_{eq}$ for both parties, as required. \square

5 No Information

We now consider the case of no information, where we assume that neither party knows the preference order of the other party. Moreover, the parties do not even hold any prior probability distribution over each other's possible preference orders. A common solution concept for this case is an ex-post equilibrium, or in our case, an ex-post SPE. Intuitively, this is a strategy profile in which the strategy of each party depends only on her own type, i.e., its preference order, and it is a SPE for every realization of the other party's type (i.e. her private preference order). Formally, let $s_k(\prec)$ be a strategy for player $k \in \{1, 2\}$ given a preference order \prec , and let $\mathcal{F}([s_1(\prec), s_2(\prec')])$ be the negotiation result if both parties follow their strategies. In the ex-post setting, a strategy for party $k \in \{1, 2\}$, s_k , is a best response to s_{3-k} if for every strategy s'_k and for every preference orders \prec, \prec' , $\mathcal{F}([s_k(\prec), s_{3-k}(\prec')]) \succeq_k \mathcal{F}([s'_k(\prec), s_{3-k}(\prec')])$. A strategy profile $[s_1, s_2]$ is an ex-post equilibrium if s_1 is a best response to s_2 and s_2 is a best response to s_1 , and it is an ex-post SPE if it is an ex-post equilibrium in every subgame of the game. We show that ex-post SPE is too strong to exist in our setting.

Theorem 14. *There are no two strategies that specify an ex-post SPE for our protocol.*

Proof. Clearly, every ex-post SPE is also a SPE (i.e., in the full information setting), and we can thus use our previous results that characterize the SPE. Assume by contradiction that there are two strategies s_1, s_2 for parties p^1, p^2 , respectively, such that $[s_1, s_2]$ is an ex-post SPE. Let $\prec_1 = o_1 \prec o_2 \prec o_5 \prec o_4 \prec o_3 \prec o_6$, and let $\prec_2 = o_1 \prec o_2 \prec o_6 \prec o_3 \prec o_5 \prec o_4$. Following our strategies we get that the SPE result is o_4 . Since the SPE result is unique and every ex-post SPE is also a SPE, $\mathcal{F}([s_1(\prec_1), s_2(\prec_2)]) = o_4$. Now consider $\prec'_1 = o_4 \prec o_5 \prec o_1 \prec o_2 \prec o_3 \prec o_6$. According to Corollary 1, $\mathcal{F}([s_1(\prec'_1), s_2(\prec_2)]) = o_3$. Note that $o_4 \prec_{p^1} o_3$. Consider the following strategy: $s'_1(\prec) = s_1(\prec'_1)$ if $\prec = \prec_1$, and $s'_1(\prec) = s_1(\prec)$ otherwise. That is, $\mathcal{F}([s_1(\prec_1), s_2(\prec_2)]) \prec_{p^1} \mathcal{F}([s'_1(\prec_1), s_2(\prec_2)])$, and thus s_1 is not a best response to s_2 . \square

We note that Theorem 14 also implies that there is no solution in dominant strategies. Another approach to uncertainty, which follows a conservative attitude, is that a party $p^k, k \in \{1, 2\}$, who wants to maximize her utility may want to play a maxmin strategy. That is, since the preference order and the strategy of the other party p^{3-k} are not known, it is sensible to assume that p^{3-k} happens to play a strategy that causes the greatest harm to p^k , and to act accordingly. p^k then guarantees the maxmin value of the game for her, which in our case is a set of outcomes such that no other outcome that is ranked lower than all of the outcomes in this set will be accepted as the result of the negotiation, regardless of the preferences of p^{3-k} . Before we show the maxmin strategy we define the complement sets for the sets L_t^k , i.e., the sets of highest ranked outcomes.

Definition 6. *In each round t , for each party $p^k, k \in \{1, 2\}$, $U_t^k = O_t \setminus L_t^k$.*

The maxmin strategy, which is composed of offering and response strategies, is defined as follows:

Strategy 3 (MAXMIN STRATEGY). *Given a history H_t^i , offer any $o \in U_t^i$. Given a history H_t^j , if $o \in U_t^j$ then accept o , else reject o .*

We now prove that our strategy specifies a maxmin strategy, and that a party p^k that follows it can guarantee the maxmin value of the game, which is the set U_1^k . We denote the party that uses Strategy 3 by p^{max} and the other party, which might try to minimize the utility of p^{max} , by p^{min} . Note that we need to handle both the case where p^{max} starts the negotiation (i.e, $p^{max} = p^1$) and the case where p^{min} starts it (i.e., $p^{min} = p^1$). We re-use Lemmas 4, 5 and 6, since they do not depend on the full-information assumption. Furthermore, we add a fourth lemma, which complements these three lemmas by considering the fourth possible offer type.

Lemma 15. *In round t , if p^i offers $o \notin L_t^i$ and p^j rejects it, then $L_{t+1}^i \leftarrow L_t^i \setminus \{o\}$, where $o \neq o^1$.*

For ease of notation, we write $U \succ_p o$ for $U \subset O$ to denote that party p strictly prefers all of the outcomes in the set U over o . The intuition of our proof is as follows. We show that if p^{max} deviates from the strategy specified by Strategy 3, p^{min} is able to make the negotiation result in an outcome o , such that $U_1^{max} \succ_{p^{max}} o$.

Theorem 16. *Strategy 3 specifies a maxmin strategy, and the maxmin value of the game is the set U_1^{max} .*

Proof. We will prove by induction on m . If $m = 2$ WLOG assume that $\succ_{p^{max}} = o_1 \succ o_2$. If $p^{max} = p^1$ then $U_1^{max} = \{o_1, o_2\}$ and clearly one of them will be the negotiation result. If $p^{max} = p^2$ then $U_1^{max} = \{o_1\}$. If p^{min} offers o_1 in the first round, according to our strategy p^{max} should accept it. If p^{min} offers o_2 in the first round, according to our strategy p^{max} should reject it, and offer o_1 in the next round. Since this is the last round, o_1 will be accepted. In any case, the negotiation result is o_1 . On the other hand, if p^{max} deviates and rejects the offer of o_1 or accepts the offer of o_2 , then o_2 will be the result of the negotiation, but $U_1^{max} \succ_{p^{max}} o_2$. Now, assume that if there are $m - 1$ outcomes in round $t + 1$ our

strategy specifies a maxmin strategy, and the maxmin value of the game is the set U_{t+1}^{max} . We show that our strategy specifies a maxmin strategy, and the maxmin value of the game is the set U_t^{max} when there are m outcomes in round t .

Assume that it is p^{max} 's turn to offer. Clearly, if p^{max} deviates and offers an outcome o such that $U_t^{max} \succ_{p^{max}} o$ then p^{min} can accept it, and the negotiation results in o . On the other hand, if p^{max} offers any $o \in U_t^{max}$ then p^{min} can either accept or reject it. If p^{min} rejects it then there are $m-1$ outcomes in the next round, and according to the induction assumption p^{max} can guarantee the maxmin value of U_{t+1}^{max} by following our strategy. However, according to Lemma 4, $L_{t+1}^{max} = L_t^{max}$ and thus $U_{t+1}^{max} \cup \{o\} = U_t^{max}$. Overall, the maxmin value of the game is the set U_t^{max} .

Now assume that it is p^{min} 's turn to offer, and p^{min} offers $o \in U_t^{max}$. Clearly, if p^{max} accepts then the negotiation result is from U_t^{max} . If p^{max} deviates and rejects, then according to induction assumption p^{max} can guarantee the maxmin value of U_{t+1}^{max} . However, according to Lemma 15, $L_{t+1}^{max} = L_t^{max} \setminus \{o'\}$, and thus $U_{t+1}^{max} = U_t^{max} \setminus \{o\} \cup \{o'\}$. That is, o' is a possible result of the negotiation even though $U_t^{max} \succ_{p^{max}} o'$. Finally, assume that p^{min} offers $o \notin U_t^{max}$. Clearly, if p^{max} deviates and accepts, then the negotiation results in o . On the other hand, if p^{max} follows our strategy and rejects, then according to the induction assumption p^{max} can guarantee the maxmin value of U_{t+1}^{max} . However, according to Lemma 5, $L_{t+1}^{max} = L_t^{max} \setminus \{o\}$, and thus $U_{t+1}^{max} = U_t^{max}$. \square

We note that even though a party does not hold any information regarding the preference order of the other party, she can still guarantee that the negotiation result will be from the upper part of her preference order (i.e., U_1^k) by following Strategy 3. This is possible since both parties have some important common knowledge, which is the number of outcomes m , as formally captured in Lemma 1.

Now, what will be the negotiation result if neither party knows the preference order of the other party, but both are rational and will thus follow the maxmin strategy? Clearly, the negotiation result will be an outcome o such that $o \in U_1^1 \cap U_1^2$. That is, an outcome from the set JG_1 as defined in Definition 1. We then get an interesting observation: if $I_1 = \emptyset$, $JG_1 = \{0_{eq}\}$ according to Corollary 2, thus the negotiation result is the same for both the case of full information and the case of no information.

In addition, we note that a party p^i cannot guarantee that the negotiation result will be from a subset $U \subset U_1^i$, since we proved that this is the maxmin value. However, she can heuristically offer in each round t the best outcome in U_t^i , instead of an arbitrarily chosen $o \in U_t^i$. Since $|U_t^j| \geq |L_t^j|$, if the other party p^j is also rational and plays the maxmin strategy, there are more cases where p^j will accept this offer, and it is thus beneficial for p^i to heuristically offer in each round t the best outcome in U_t^i .

The idea of the maxmin strategy is that a party, not knowing the preferences of the other party, makes a worst case assumption about the behavior of that party (i.e., that she does not need to be rational). This assumption may seem too restrictive, and we therefore also consider the robust-optimization equilibrium solution concept from [Aghassi and

Bertsimas, 2006], which we adapt to our setting. Intuitively, in this solution concept each party makes a worst case assumption about the preference order of the other party, but each party still assumes that the other party will play rationally and thus her aim is to maximize her utility. Formally, given a strategy profile $[s_1, s_2]$ and a preference order \prec , let $w_{s_1, \prec, s_2} = \mathcal{F}([s_1(\prec), s_2(\prec)])$, where \prec' is a preference order such that for all \prec'' , $\mathcal{F}([s_1(\prec'), s_2(\prec')]) \preceq_{p^1} \mathcal{F}([s_1(\prec), s_2(\prec)])$. In the robust-optimization setting, a strategy for party $k \in \{1, 2\}$, s_k , is a best response to s_{3-k} if for all s'_k and for all \prec , $w_{s_k, \prec, s_{3-k}} \succeq_{p^k} w_{s'_k, \prec, s_{3-k}}$. A strategy profile $[s_1, s_2]$ is a robust-optimization equilibrium if s_1 is a best response to s_2 and s_2 is a best response to s_1 . We show that in our setting, surprisingly, every pair of maxmin strategies specifies a robust-optimization equilibrium.

Theorem 17. *If s_1 and s_2 are maxmin strategies, then $[s_1, s_2]$ is a robust-optimization equilibrium.*

Proof. Given a preference order, \prec , let \prec_{op} be the opposite preference order, i.e., if $\prec = o_1 \prec o_2 \prec \dots \prec o_m$ then $\prec_{op} = o_1 \succ o_2 \succ \dots \succ o_m$. According to Theorem 16, if s_k is a maxmin strategy then the negotiation result is $o \in U_1^k$. That is, the worst negotiation result for p^k is the least preferred outcome in U_1^k , denoted by o_{wo} . Since the other party p^{3-k} is also using a maxmin strategy, $o \in U_1^k \cap U_1^{3-k}$. For every preference order, \prec , if the preference order of $p^{3-k} = \prec_{op}$, $U_1^k \cap U_1^{3-k} = \{o_{wo}\}$. That is, $w_{s_1, \prec, s_{3-k}} = o_{wo}$. Assume by contradiction that there is another strategy s'_k and a preference order \prec , such that $w_{s'_k, \prec, s_{3-k}} \succ_{p^k} o_{wo}$. However, if the preference order of $p^{3-k} = \prec_{op}$, $w_{s'_k, \prec, s_{3-k}} \notin U_1^{3-k}$, in contradiction to Theorem 16. \square

Finally, consider an asymmetric information setting, where there exists one party that has full information about the other party's preference order, while the other party does not have this information. Let p^{info} be the party that has the full information, and p^{null} be the other party. p^{null} has no information and she will thus act according to the maxmin strategy (Strategy 3). p^{info} would like to take advantage of her knowledge, so the negotiation result will be better for her. However, according to Theorem 16, the maxmin value of the game is U_1^{null} . Therefore, the best strategy for p^{info} is as follows. If p^{info} starts the negotiation, she should offer the best outcome from U_1^{null} according to her preferences, and p^{null} will accept it. If p^{info} starts the negotiation, she will offer an outcome from U_1^{null} . If this is the best outcome according to p^{info} 's preferences, she should accept it. Otherwise, in the second round p^{info} should offer the best outcome from U_1^{null} according to her preferences, and p^{null} will accept it.

6 Conclusion

We investigated the VAOV negotiation protocol, which is suitable for ordinal preferences over a finite set of outcomes. We introduced strategies that specify a SPE, and improved upon previous results by providing a linear time algorithm that computes a SPE strategy. We provided substantial analysis of our strategies, which showed the equivalence of the

SPE result of the protocol in a non-cooperative setting, to the result of the *RC* rule in a cooperative setting. Finally, we analyzed the no information setting. We believe that our approach is especially suitable for non-cooperative, multi-agent systems, since we provide easy to implement strategies that can be computed only once if both agents follow the SPE strategy on the equilibrium path. Moreover, there is no need for a central authority to guarantee that the negotiation result will be Pareto optimal, if both agents are rational and follow the SPE strategy. For future work, we would like to extend the protocol to a multi-party setting and analyze the resulting SPE. In addition, it is important to find additional implementation of other bargaining rules by negotiation protocols, similar to the implementation that we showed for the *RC* rule by the SPE of the VAOV protocol.

Acknowledgments

This work was supported by the Israel Science Foundation, Grant No. 1488/14.

References

- [Aghassi and Bertsimas, 2006] Michele Aghassi and Dimitris Bertsimas. Robust game theory. *Mathematical Programming*, 107(1-2):231–273, 2006.
- [Ali and Ronaldson, 2012] Shehzad Ali and Sarah Ronaldson. Ordinal preference elicitation methods in health economics and health services research: using discrete choice experiments and ranking methods. *British medical bulletin*, 103(1):21–44, 2012.
- [Anbarci, 1993] Nejat Anbarci. Noncooperative foundations of the area monotonic solution. *The Quarterly Journal of Economics*, 108(1):245–258, 1993.
- [Anbarci, 2006] Nejat Anbarci. Finite alternating-move arbitration schemes and the equal area solution. *Theory and decision*, 61(1):21–50, 2006.
- [Boutilier *et al.*, 2015] Craig Boutilier, Ioannis Caragiannis, Simi Haber, Tyler Lu, Ariel D Procaccia, and Or Sheffet. Optimal social choice functions: A utilitarian view. *Artificial Intelligence*, 227:190–213, 2015.
- [Brams and Kilgour, 2001] Steven J Brams and D Marc Kilgour. Fallback bargaining. *Group Decision and Negotiation*, 10(4):287–316, 2001.
- [Brandt *et al.*, 2016] Felix Brandt, Vincent Conitzer, Ulle Endriss, Ariel D Procaccia, and Jérôme Lang. *Handbook of computational social choice*. Cambridge University Press, 2016.
- [Conley and Wilkie, 2012] John P Conley and Simon Wilkie. The ordinal egalitarian bargaining solution for finite choice sets. *Social Choice and Welfare*, 38(1):23–42, 2012.
- [De Clippel *et al.*, 2014] Geoffroy De Clippel, Kfir Eliaz, and Brian Knight. On the selection of arbitrators. *American Economic Review*, 104(11):3434–58, 2014.
- [Erlich *et al.*, 2018] Sefi Erlich, Noam Hazon, and Sarit Kraus. Negotiation Strategies for Agents with Ordinal Preferences. *ArXiv e-prints arXiv:1805.00913*, 2018.
- [Fatima *et al.*, 2014] Shaheen Fatima, Sarit Kraus, and Michael Wooldridge. *Principles of automated negotiation*. Cambridge University Press, 2014.
- [Hurwicz and Sertel, 1999] Leonid Hurwicz and Murat R Sertel. Designing mechanisms, in particular for electoral systems: the majoritarian compromise. In *Contemporary Economic Issues*, pages 69–88. Springer, 1999.
- [Kıbrıs and Sertel, 2007] Özgür Kıbrıs and Murat R Sertel. Bargaining over a finite set of alternatives. *Social Choice and Welfare*, 28:421–437, 2007.
- [Kraus, 2001] Sarit Kraus. *Strategic negotiation in multi-agent environments*. MIT press, 2001.
- [Mariotti, 1998] Marco Mariotti. Nash bargaining theory when the number of alternatives can be finite. *Social choice and welfare*, 15(3):413–421, 1998.
- [Nagahisa and Tanaka, 2002] Ryo-ichi Nagahisa and Makoto Tanaka. An axiomatization of the kalai-smorodinsky solution when the feasible sets can be finite. *Social Choice and Welfare*, 19(4):751–761, 2002.
- [Nash Jr, 1950] John F Nash Jr. The bargaining problem. *Econometrica*, 18(2):155–162, 1950.
- [Nunez and Laslier, 2015] Matias Nunez and Jean-François Laslier. Bargaining through approval. *Journal of Mathematical Economics*, 60:63–73, 2015.
- [Osborne and Rubinstein, 1990] Martin J. Osborne and Ariel Rubinstein. *Bargaining and Markets*. Academic Press, 1990.
- [Rubinstein, 1982] Ariel Rubinstein. Perfect equilibrium in a bargaining model. *Econometrica: Journal of the Econometric Society*, pages 97–109, 1982.
- [Shapley, 1969] Lloyd S Shapley. *Utility comparison and the theory of games*. Cambridge: Cambridge University Press. Originally published in *La Decision*, 1969.
- [Sprumont, 1993] Yves Sprumont. Intermediate preferences and rawlsian arbitration rules. *Social Choice and Welfare*, 10(1):1–15, 1993.
- [Thomson, 1994] William Thomson. Cooperative models of bargaining. *Handbook of game theory with economic applications*, 2:1237–1284, 1994.
- [Zhang and Zhang, 2008] Dongmo Zhang and Yan Zhang. An ordinal bargaining solution with fixed-point property. *Journal of Artificial Intelligence Research*, 33:433–464, 2008.
- [Zlotkin and Rosenschein, 1996] Gilad Zlotkin and Jeffrey S Rosenschein. Mechanism design for automated negotiation, and its application to task oriented domains. *Artificial Intelligence*, 86(2):195–244, 1996.