Negotiation Strategies for Agents with Ordinal Preferences

Sefi Erlich, Noam Hazon, and Sarit Kraus

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 a linear time algorithm for computing elegant strategies in the SPE 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 both the case where neither party knows the preference order of the other party, and the case where only one party is uninformed. We present a linear time algorithm for computing the maxmin strategies for these cases.

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 [17]. Since then a lot of work has been done to develop other types of protocols, and to extend the basic bilateral negotiation protocol [8].

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 to 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., [1]). 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 [14, 18]. 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 utility [1]. We analyze an intuitive protocol for bilateral negotiation that was introduced by Anbarci [2], where the two parties

[1] This is also the assumption in most of the voting literature [5].

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 an efficient algorithm for computing elegant strategies that specify a Sub-game Perfect Equilibrium (SPE) for the parties. Specifically, we improve the previous result of [3] and find the strategies 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 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 algorithm finds 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 [10]. 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 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 analyze the properties of the maxmin solution concept that is the result of the protocol, and show a linear time algorithm for computing it. We also 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 provide an improved algorithm for computing a SPE for the studied negotiation protocol. Our algorithm finds elegant strategies, and we provide a substantial analysis for showing that it indeed finds a SPE. 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 [10], 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 [16], [11], and [8] for extensive coverage of the different approaches. The traditional assumption in the 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 [21], [12], [13] and recently [15]. 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. Sequential procedures, in particular the *fallback bargaining* method, have attracted considerable interest [19, 9, 4, 10, 6], 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 [10]. We note that the *RC* bargaining rule is equivalent to *Bucklin* voting with two voters, and thus our result can be also 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 Clippel et al. [7] study the problem of selection of arbitrators, and they concentrate on two-step protocols. The most closely related works are the papers of Anbarci. In [2] 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 [3] 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 specifies a SPE strategies. He also shows that under some assumption, if the number of outcomes tends to infinity the VAOV protocol converges to the equal area rule [20]. We provide a more efficient algorithm that finds elegant SPE strategies. 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 $O = \{o_1, ..., o_m\}$, 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 and transitive 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 [2]. 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 the 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 at 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 $|L_t^k|$ 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 $|L_t^k|$ lowest outcomes in her preference order. Therefore, $o \notin Low_t$. □

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 Sub-game Perfect Equilibrium (SPE). Since [2] showed that the SPE result is unique, it suffices to find any set of strategies that specify a SPE. We present an intuitive linear time algorithm, Algorithm 1, that computes elegant SPE strategies. The algorithm describes both the offering and the response strategies. We use the following definitions. Given a round t and a party p^j, let $o_{t,m}^j$ be the least preferred outcome in O_t according to $\succ_{p^j}$. In addition, let o_a be the result of the negotiation if both parties reject all of the offers that they get (except for the last offer) but still use the offering strategy that is specified by Algorithm 1, from the current round and on. To illustrate the strategies that are defined

Algorithm 1 Full information, SPE strategies

OFFERING STRATEGY

Input: the current round t, the set O_t.

```
if I_t = L_t^i ∩ L_t^j ≠ ∅ then
    Offer o ∈ I_t
else
    Offer o_{t,m}^j
```

RESPONSE STRATEGY

Input: the current round t, the set O_t, an offer o.

```
if o ⪰_{p^j} o_a then
    Accept o
else
    Reject o
```

by the algorithm, consider the following examples:

Example 1. *Suppose that*

$$\succ_{p^1} = o_6 \succ o_5 \succ o_4 \succ o_3 \succ o_2 \succ o_1$$

$$\succ_{p^2} = o_1 \succ o_3 \succ o_2 \succ o_6 \succ o_4 \succ o_5.$$

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 Algorithm 1 p^1 would offer p^2's least preferred outcome - o_5. Then, p^2 would reject, since $o_a = o_3 \succ_{p^2} o_5$, 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*

$$\succ_{p^1} = o_6 \succ o_5 \succ o_4 \succ o_3 \succ o_2 \succ o_1$$

$$\succ_{p^2} = o_1 \succ o_3 \succ o_6 \succ o_2 \succ o_4 \succ o_5.$$

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 Algorithm 1 p^1 would offer o_2. Then, p^2 will reject, since $o_a = o_6 \succ_{p^2} o_2$, 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.

We first note that trivial exploration of the whole game tree 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 algorithm of [3] for finding a SPE is not explicitly analyzed, but its running time is at least $O(m^2)$. Our algorithm only needs to simulate one branch of the tree (to find o_a) and then it only traces the intersection between L_t^i and L_t^j. Therefore, its running time is linear in m.

In order to prove that our algorithm computes a SPE, we need some lemmas. Due to space constraints, we defer the proofs of some of the lemmas to the appendix. We begin with a simple corollary of Lemma 1. Let o_{eq} be the SPE result. We get:

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

We note that in step 1 of the offering strategy of Algorithm 1, p^i offers an outcome from the set I_t. We now show the relation between the set I_t and the set JG_t, which will turn out to be very useful.

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

If we combine the findings from Corollary 1 and 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 $JG_t = \{o_{eq}\}$.*

Proof. From Lemma 2, $|JG_t| = 1$. Assume that $o_{eq} \notin JG_t$, then $o_{eq} \in Low_t$, in contradiction to Lemma 1. □

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 $t+1$, $|L^i_{t+1}| = |L^j_t| - 1$ and $|L^j_{t+1}| = |L^i_t|$*

We note that the number of outcomes in L^k_t 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^k_{t+1}. 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 Algorithm 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 $o \notin L^i_t$ and p^j rejects it, then $L^j_{t+1} \leftarrow L^i_t$*

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

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

We are now ready to prove that our algorithm finds a SPE, and we will split our proof into two cases. We first handle the case where $I_t = \emptyset$.

Theorem 7. *If $I_t = \emptyset$ then Algorithm 1 finds a SPE.*

Proof. We prove by induction on m. If $m = 2$, without loss of generality (WLOG) assume that $\succ_{p^j} = o_1 \succ o_2$. Thus, $JG_1 = \{o_1\}$, and according to Corollary 2, o_1 is the SPE result. On the other hand, according to our algorithm p^i will offer o_2 in the first round, p^j will reject it and offer o_1, and then the negotiation will end with o_1 as the negotiation result. Now, assume that if there are $m-1$ outcomes in round $t+1$, and $I_{t+1} = \emptyset$, our algorithm finds a SPE. We show that the algorithm finds a SPE when there are m outcomes in round t, and $I_t = \emptyset$. According to the algorithm, p^i will offer $o^j_{t,m}$ and p^j will reject it. According to Lemmas 4 and 5, $Low_{t+1} = Low_t \setminus \{o\}$, and therefore $I_{t+1} = \emptyset$ and $JG_{t+1} = JG_t$. In round $t+1$ there are $m-1$ outcomes, and according to the induction assumption our algorithm results with a SPE, o_{eq}. According to Corollary 2, $JG_{t+1} = o_{eq}$, and since $JG_t = JG_{t+1}$ after following our algorithm in round t, we get that the algorithm finds a SPE. □

Before we proceed to analyze the case where at some round $I_t \neq \emptyset$, we show that the result of our algorithm has a desirable property. Specifically, it is Pareto optimal, i.e., $\forall o \in O \setminus \{o_a\}$, $o_a \succ_{p^i} o$ or $o_a \succ_{p^j} o$, where o_a is the negotiation result by following our algorithm.

Lemma 8. *Suppose the parties follow the strategies specified in Algorithm 1 and the negotiation ends with o_a, then o_a is Pareto optimal.*

Proof. Assume by contradiction that o_a is not Pareto optimal. Then, there exists $o \in O$ such that $o \succ_{p^i} o_a$ and $o \succ_{p^j} o_a$. Let t be the round where one party accepted o_a, and first assume that $o \in O_t$. Then, in this round $o_a \notin Low_t$ and thus $o \notin Low_t$. Therefore, $|JG_t| \geq 2$ and by lemma 2 $|I_t| \geq 1$, in contradiction to the algorithm acceptance criterion. Now assume that $o \notin O_t$. Then, there exists a round $t' < t$ where party p^i offered o. If o was offered because $o = o^j_{t',m}$ then is must be that $o_a \notin O_{t'}$ since $o \succ_{p^j} o_a$, and thus $o_a \notin O_t$, in contradiction to the definition of t. If o was offered because $o \in I_{t'}$ then also $o_a \in I_{t'}$, and thus $o_a \in I_t$ due to Lemmas 5 and 6. In addition, from these lemmas we get that $JG_{t+1} \subseteq JG_t$. Therefore, if both parties reject all of the offers that they get (except for the last offer) but still use the offering strategy that is specified by Algorithm 1, the resulting outcome o' is guaranteed to be in JG_t. However, since $o' \in JG_t$ and $o_a \in I_t$, it must be that $o' \succ_{p^j} o_a$, and according to Algorithm 1 p^j should have reject the offer of o_a, in contradiction to the definition of o_a. □

The proof for the case where $I_t \neq \emptyset$ is more involved, and 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. $d_{x,k,t} = |\{o \in O_t : o_x \succeq_{p^k} o \land o \notin L_t^k\}|$ *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 Algorithm 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 $d_{3,1,1} = 1$ for party p^1 and $d_{3,2,1} = 2$ 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 $d_{6,1,1} = 4$ and $d_{6,2,1} = 1$, and the distance of o_3 at round 1 for the parties is $d_{3,1,1} = 1$ and $d_{3,2,1} = 2$.

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. *If $I_1 \neq \emptyset$ then Algorithm 1 finds a SPE.*

Proof. Clearly, p^j has no incentive to deviate. Indeed, if according to Algorithm 1 p^j should reject an offer o, it is because $o_a \succ_{p^j} o$. Therefore, it is not worthwhile for p^j to deviate and accept o instead of the result of following the algorithm, o_a. Similarly, if according to Algorithm 1 p^j should accept an offer o, it is because $o \succeq_{p^j} o_a$. Therefore, it is not worthwhile for p^j to deviate and reject o in order to get as a result the outcome o_a. We thus concentrate on the strategy of p^i, but we first derive some general inequalities. Let o_a be the result of the negotiation if both parties follow our algorithm from round t. Let o_x be an outcome such that $o_x \succ_{p^i} o_a$. According to Lemma 8, since $o_x \succ_{p^i} o_a$, $o_a \succ_{p^j} o_x$. Let t' be the round in which $JG_{t'} = \{o_a\}$. Then, in round t, $\ell_{i,t} < d_{i,a,t}$ and $\ell_{j,t} < d_{j,a,t}$ (otherwise, $o_a \notin JG_{t'}$). By definition, $t' = \ell_{i,t} + \ell_{j,t}$. In addition, since $JG_{t'} = \{o_a\}$, o_x must be part of $Low_{t''}$ for some $t'' < t'$ (otherwise, $o_x \in JG_{t'}$). Since $o_x \succ_{p^i} o_a$ and $\ell_{i,t} < d_{i,a,t}$, it must be that o_x is part of $L_{t''}^j$, that is, $d_{j,x,t} \leq \ell_{j,t}$. In summary:

$$\begin{aligned} \ell_{i,t} < d_{i,a,t} < d_{i,x,t} \\ d_{j,x,t} \leq \ell_{j,t} < d_{j,a,t} \end{aligned} \tag{1}$$

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

$$\begin{aligned} \ell_{i,t+1} < d_{i,x,t+1} < d_{i,a,t+1} \\ d_{j,a,t+1} \leq \ell_{j,t+1} < d_{j,x,t+1} \end{aligned} \tag{2}$$

We now prove by induction on m. If $m = 2$ then $I_1 = \emptyset$. If $m = 3$ and $I_1 \neq \emptyset$ then $|I_1| = 1$ and WLOG assume that $I_1 = \{o_3\}$. There are two cases: either p^i and p^j agree that $o_1 \succ o_2$ or not. If they agree that $o_1 \succ o_2$, it is clear that the SPE result will be o_1. Since o_1 is the only Pareto optimal outcome, it will also be the result of our algorithm. If

p^i and p^j disagree on o_1 and o_2, i.e., $\succ_{p^i} = o_1 \succ o_2 \succ o_3$ and $\succ_{p^j} = o_2 \succ o_1 \succ o_3$, then in the SPE p^i can offer o_3, since she knows that p^j will reject it, and in the resulting game with 2 outcomes, p^i can guarantee that her most preferred outcome, o_1, will be the SPE result. Following our algorithm we get the same behavior.

Now, assume that if there are $m-1$ outcomes in round $t+1$, our algorithm finds a SPE. We show that the algorithm finds a SPE when there are m outcomes in round t, and $I_t \neq \emptyset$. 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, $o_x \succ_{p^i} o_a$. 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_a and o_x, from round t to round $t+1$. According to Lemma 3, $|L^i_{t+1}| + 1 = |L^j_t|$, and since $o_x \succ_{p^j} o_d$ and $o_a \succ_{p^j} o_d$, $d_{j,x,t}$ and $d_{j,a,t}$ do not change when moving to round $t+1$. Let c be an integer. Then, $d_{i,a,t} + c = d_{j,a,t+1}$, $d_{j,a,t} = d_{i,a,t+1}$, and $d_{j,x,t} = d_{i,x,t+1}$. If we combine these 3 equations and equations 2 we get that $\ell_{j,t+1} < d_{j,x,t} < d_{j,a,t}$ and $d_{i,a,t} + c \leq \ell_{i,t+1}$. Adding equation 1 we get that $\ell_{i,t} < \ell_{i,t+1} - c$ and $\ell_{i,t+1} < \ell_{j,t}$. 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, $\ell_{i,t+1} \leq \ell_{i,t+1} - c - 2$ and thus $c \leq -2$. However, the distance function cannot decrease by more than 1 when moving from round t to $t+1$, and 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_a and o_x, from round t to round $t+1$. Note that since p^i deviates, $o_d \succ_{p^i} o_a$. According to Lemma 8, $o_a \succ_{p^j} o_d$. According to Lemma 3, $|L^i_{t+1}| + 1 = |L^j_t|$, and since $o_a \succ_{p^j} o_d$, $d_{j,a,t}$ 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, $d_{i,a,t} + c = d_{j,a,t+1}$, $d_{j,a,t} = d_{i,a,t+1}$, and $d_{j,x,t} + 1 = d_{i,x,t+1}$. If we combine these 3 equations and equations 2 we get that $\ell_{i,t+1} < d_{j,x,t} + 1 < d_{j,a,t}$ and $d_{i,a,t} + c \leq \ell_{j,t+1}$. Adding equation 1 we get that $\ell_{i,t} < \ell_{j,t+1} - c$ and $\ell_{i,t+1} - 1 < \ell_{j,t}$. 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, $\ell_{j,t+1} - 1 \leq \ell_{j,t+1} - c - 2$ and thus $c \leq -1$. However, in order for $d_{i,a,t}$ to decrease by at least one, $o_a \succ_{p^i} o_d$, but in our case $o_d \succ_{p^i} o_a$.

Overall, we showed that p^i does not have an incentive to deviate in round t. According to the induction assumption, our algorithm finds a SPE when there are $m-1$ outcomes in round $t+1$. Therefore, p^i does not have any incentive to deviate. □

4.1 Properties

We first note that since we showed that the result of Algorithm 1 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 [10]. 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 $RC = \{o_x | \max_{o_x \in O} \min_{k \in \{1,2\}} (d_{k,x,1} + |L^k_1| - 1)\}$. It can also be computed by the following steps:

1. Let $v = 1$
2. For each $k \in \{1, 2\}$, let $B^k_v = \{$the v most preferred outcomes in $\succ_{p^k}\}$.
3. If $B^1_v \cap B^2_v > 0$ then return $B^1_v \cap B^2_v$ 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 algorithm finds 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 algorithm. By Corollary 2, $JG_t = \{o_{eq}\}$. Rephrasing the definition of JG_t we get that $JG_t = B^i_{|O_t|-|L_t^i|} \cap B^j_{|O_t|-|L_t^j|}$. Now, 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, $o_{eq} \in RC$. □

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 $o_x \succ_{p^i} o_y$. If $\ell_{1,1} + \ell_{2,1}$ is even then $o_{eq} = o_y$ and $o_y \succ_{p^j} o_x$.*

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. $|RC| = 2$ and 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, $B^1_{|O_t|-|L_t^1|} \cap B^2_{|O_t|-|L_t^2|} = \{o_{eq}\}$. Therefore, $|L_t^1| \neq |L_t^2|$, and thus m_t is even. Now, 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, $o_{eq} = o_y$.

□

Finally, we adapt the monotonicity criterion that the RC rule satisfies to 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. □

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 the possible preference orders of each other. Following a conservative approach, a party p^k, $k \in \{1,2\}$, who wants to maximize her utility will have 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 is in our case 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^k_t, i.e., the sets of highest ranked outcomes.

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

The maxmin strategy is presented by Algorithm 2, which describes both the offering and the response strategies. A party p^k that follows our algorithm can guarantee the maxmin value of the game, which is the set U^k_1. We now prove that Algorithm 2 specifies a maxmin

Algorithm 2 No information, maxmin strategies

<u>OFFERING STRATEGY</u>

Input: the current round t, the set O_t.

1: Offer any $o \in U^i_t$

<u>RESPONSE STRATEGY</u>

Input: the current round t, the set O_t, an offer o.

1: **if** $o \in U^j_t$ **then**
2: Accept o
3: **else**
4: Reject o

strategy, and that the maxmin value of the game is the set U^k_1. We denote the party that uses the algorithm 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 14. *In round t, if p^i offers $o \notin L^j_t$ and p^j rejects it then $L^i_{t+1} \leftarrow L^j_t \setminus \{o'\}$, where $o \neq o'$.*

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 Algorithm 2, p^{min} is able to make the negotiation result in an outcome o, such that $U_1^{max} \succ_{p^{max}} o$.

Theorem 15. *Algorithm 2 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 algorithm p^{max} should accept it. If p^{min} offers o_2 in the first round, according to our algorithm 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 algorithm specifies a maxmin strategy, and the maxmin value of the game is the set U_{t+1}^{max}. We show that the algorithm 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 algorithm. 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 14, $L_{t+1}^{max} = L_t^{max} \setminus \{o'\}$, 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 algorithm 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 $U_{t+1}^{max} = U_t^{max}$. □

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 her upper part of the preference order (i.e., U_1^k) by following our algorithm. 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, and thus the negotiation result is the same both for 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 arbitrary 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.

Finally, suppose that 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 (Algorithm 2). p^{info} would like to take advantage of her knowledge, so the negotiation result will be better for her. However, according to Theorem 15, 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 improved upon previous results by providing a linear time algorithm that specifies SPE strategies. We provided substantial analysis of our algorithm, 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. 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.

References

[1] S. Ali and S. 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.

[2] N. Anbarci. Noncooperative foundations of the area monotonic solution. *The Quarterly Journal of Economics*, 108(1):245–258, 1993.

[3] N. Anbarci. Finite alternating-move arbitration schemes and the equal area solution. *Theory and decision*, 61(1):21–50, 2006.

[4] S. J. Brams and D. M. Kilgour. Fallback bargaining. *Group Decision and Negotiation*, 10(4):287–316, 2001.

[5] F. Brandt, V. Conitzer, U. Endriss, A. D. Procaccia, and J. Lang. Handbook of computational social choice. Cambridge University Press, 2016.

[6] J. P. Conley and S. Wilkie. The ordinal egalitarian bargaining solution for finite choice sets. *Social Choice and Welfare*, 38(1):23–42, 2012.

[7] G. De Clippel, K. Eliaz, and B. Knight. On the selection of arbitrators. *American Economic Review*, 104(11):3434–58, 2014.

[8] S. Fatima, S. Kraus, and M. Wooldridge. *Principles of automated negotiation*. Cambridge University Press, 2014.

[9] L. Hurwicz and M. R. Sertel. Designing mechanisms, in particular for electoral systems: the majoritarian compromise. In *Contemporary Economic Issues*, pages 69–88. Springer, 1999.

[10] Ö. Kıbrıs and M. R. Sertel. Bargaining over a finite set of alternatives. *Social Choice and Welfare*, 28:421–437, 2007.

[11] S. Kraus. *Strategic negotiation in multiagent environments*. MIT press, 2001.

[12] M. Mariotti. Nash bargaining theory when the number of alternatives can be finite. *Social choice and welfare*, 15(3):413–421, 1998.

[13] R.-i. Nagahisa and M. Tanaka. An axiomatization of the kalai-smorodinsky solution when the feasible sets can be finite. *Social Choice and Welfare*, 19(4):751–761, 2002.

[14] J. F. Nash Jr. The bargaining problem. *Econometrica*, 18(2):155–162, 1950.

[15] M. Nunez and J.-F. Laslier. Bargaining through approval. *Journal of Mathematical Economics*, 60:63–73, 2015.

[16] M. J. Osborne and A. Rubinstein. *Bargaining and Markets*. Academic Press, 1990.

[17] A. Rubinstein. Perfect equilibrium in a bargaining model. *Econometrica: Journal of the Econometric Society*, pages 97–109, 1982.

[18] L. S. Shapley. *Utility comparison and the theory of games*. Cambridge: Cambridge University Press. Originally published in La Decision, 1969.

[19] Y. Sprumont. Intermediate preferences and rawlsian arbitration rules. *Social Choice and Welfare*, 10(1):1–15, 1993.

[20] W. Thomson. Cooperative models of bargaining. *Handbook of game theory with economic applications*, 2:1237–1284, 1994.

[21] G. Zlotkin and J. S. Rosenschein. Mechanism design for automated negotiation, and its application to task oriented domains. *Artificial Intelligence*, 86(2):195–244, 1996.

Sefi Erlich
Ariel University
Ariel, Israel
Email: `erlichsefi@gmail.com`

Noam Hazon
Ariel University
Ariel, Israel
Email: `noamh@ariel.ac.il`

Sarit Kraus
Bar-Ilan University
Ramat Gan, Israel
Email: `sarit@cs.biu.ac.il`

7 Appendix

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

Proof. Suppose that in round t, $|O_t| = m_t$ is odd. Then

$$|L_t^i \setminus I_t| = \left\lfloor \frac{m_t}{2} \right\rfloor - |I_t| = \frac{m_t - 1}{2} - |I_t|$$

$$|L_t^j \setminus I_t| = \left\lfloor \frac{m_t}{2} \right\rfloor - |I_t| = \frac{m_t - 1}{2} - |I_t|$$

Therefore,

$$|JG_t| = m_t - 2 \cdot (\frac{m_t - 1}{2} - |I_t|) - |I_t| = |I_t| + 1$$

Now suppose that m_t is even. Then

$$|L_t^i \setminus I_t| = \left\lfloor \frac{m_t}{2} \right\rfloor - 1 - |I_t| = \frac{m_t}{2} - 1 - |I_t|$$

$$|L_t^j \setminus I_t| = \left\lfloor \frac{m_t}{2} \right\rfloor - |I_t| = \frac{m_t}{2} - |I_t|$$

Therefore,

$$|JG_t| = m_t - (\frac{m_t}{2} - 1 - |I_t|) - (\frac{m_t}{2} - |I_t|) - |I_t| = |I_t| + 1$$

□

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 rolls are switched between p^i and p^j. Therefore, $|L_{t+1}^i| = |L_{t+1}^j| = \left\lfloor \frac{m_{t+1}}{2} \right\rfloor = \left\lfloor \frac{m_t - 1}{2} \right\rfloor = \left\lfloor \frac{m_t}{2} - \frac{1}{2} \right\rfloor = \frac{m_t}{2} - 1$. Now assume that m_t is odd, then $|L_t^i| = |L_t^j| = \left\lfloor \frac{m_t}{2} \right\rfloor = \frac{m_t - 1}{2}$. After that p^i offered the outcome o and p^j rejected it, m_{t+1} is even, and the rolls are switched between p^i and p^j. Therefore, $|L_{t+1}^j| = \frac{m_t - 1}{2}$ and $|L_{t+1}^i| = \frac{m_t - 1}{2} - 1$. □

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

Proof. According to Lemma 3, the sets L_{t+1}^j and L_t^i have the same size. Therefore, if p^i offers $o \notin L_t^i$ and p^j rejects it, we can be assured that $L_{t+1}^j = L_t^i$. □

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

Proof. According to Lemma 3, the set L_{t+1}^i contains one outcome less than the set L_t^j. Therefore, if p^i offers $o \in L_t^j$ and p^j rejects it, o is the only outcome that becomes unavailable in round $t+1$, and we can thus be assured that $L_{t+1}^i = L_t^j \setminus \{o\}$. □

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

Proof. According to lemma 3, the sets L_{t+1}^j and L_t^i have the same size. Therefore, if p^i offers $o \in L_t^i$ and p^j rejects it, o is the only outcome that becomes unavailable in round $t+1$, and thus there must be another outcome $o' \in O_{t+1}$ that becomes part of L_{t+1}^j. □

Lemma 14. *In round t, if p^i offers $o \notin L_t^j$ and p^j rejects it then $L_{t+1}^i \leftarrow L_t^j \setminus \{o'\}$, where $o \neq o'$.*

Proof. According to Lemma 3, the set L^i_{t+1} contains one outcome less than the set L^j_t. Therefore, if p^i offers $o \notin L^j_t$ and p^j rejects it, there must be another outcome $o' \in O_t$ that left the set L^j_t. □

Even though the uniqueness of the SPE result was proven by Anbarci [2], we provide a direct and simpler proof.

Theorem 16. *The SPE result is unique.*

Proof. We prove by induction on m. If $m = 2$, then no matter what p^1 offers, the negotiation results with the most preferred outcome of p^2, and thus the SPE is unique. Now, assume that if there are $m - 1$ outcomes in round $t + 1$, the SPE is unique. We show that the SPE is unique when there are m outcomes in round t. p^i is able to offer an outcome $o \in O_t$. For any such o, either p^j accepts o or rejects it and the game moves to round $t + 1$ with $m - 1$ outcomes. According to the induction assumption, the SPE is unique in each sub-tree of the game where there are $m - 1$ outcomes. Since p^i has strict preferences, in a SPE she will choose either an outcome that p^j will accept or a sub-tree of the game, that result with the best outcome according to p^i's preferences. That is, in all of the offers of p^i in round t that are in SPE, the SPE result is the same. □