Artificial Intelligence 327 (2024) 104050

Contents lists available at ScienceDirect

Artificial Intelligence

journal homepage: www.elsevier.com/locate/artint

Negotiation strategies for agents with ordinal preferences: Theoretical analysis and human study ☆

Noam Hazon [a,*], Sefi Erlich [a], Ariel Rosenfeld [b], Sarit Kraus [c]

[a] Department of Computer Science, Ariel University, Israel
[b] Department of Information Science, Bar-Ilan University, Israel
[c] Department of Computer Science, Bar-Ilan University, Israel

ARTICLE INFO

Keywords:
Automated negotiation
Negotiation with ordinal preferences
Negotiation strategies
Human study

ABSTRACT

Negotiation is a very common interaction between 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 focus on negotiation with ordinal preferences over a finite set of outcomes. We study an intuitive protocol for bilateral negotiations, where the two parties make offers alternately. We analyze the negotiation protocol under two settings: First, we consider the full information setting, where each party is fully aware of the other party's preference order. For this case, we provide elegant strategies that specify a sub-game perfect equilibrium. In addition, we show how the studied negotiation protocol almost completely implements a known bargaining rule. Second, we analyze the complementary no-information setting where neither party knows the preference order of the other party. For this case, we provide a Maxmin strategy and show that every pair of Maxmin strategies specifies a robust-optimization equilibrium. Finally, through a human study ($N = 150$), we empirically study the practical relevance of our full information analysis to people engaging in negotiations with each other and/or with an automated agent using the studied protocol. Surprisingly, our results indicate that people tend to arrive at the equilibrium outcomes despite frequently departing from the proposed strategies. In addition, in contrast to commonly held beliefs, we find that an equilibrium-following agent performs very well with people.

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. From an artificial intelligence perspective, the study of automated negotiation is primarily concerned with the creation of agents which will be able to proficiently negotiate on behalf of their human users or owners [16].

In order to adequately develop such agents, several interwoven aspects of the negotiation framework have to be addressed, including the representation of one's preferences and the negotiation protocol [24]. Focusing on the former, most negotiating agents

☆ This paper extends an earlier conference paper [15].

* Corresponding author.
E-mail addresses: noamh@ariel.ac.il (N. Hazon), erlichsefi@gmail.com (S. Erlich), ariel.rosenfeld@biu.ac.il (A. Rosenfeld), sarit@cs.biu.ac.il (S. Kraus).

https://doi.org/10.1016/j.artint.2023.104050
Received 1 January 2023; Received in revised form 27 November 2023; Accepted 29 November 2023

Available online 4 December 2023
0004-3702/

assume access to cardinal preferences represented as a utility function that adequately associates different outcomes with numerical values. Unfortunately, utility functions are not always readily available and the use of cardinal utilities for representing human preferences has been widely criticized on the grounds of cognitive complexity, the difficulty of elicitation, and various other concerns (e.g., [2]). On the other hand, ordinal preferences, commonly represented as ranking over the outcomes, are assumed to be easier to elicit from people. Unfortunately, agents which assume access to ordinal preferences are significantly less prevalent. When ordinal preferences are assumed, it is common to convert them to cardinal preferences by following some non-trivial assumptions [30,38,10,12,35].

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 negotiation protocol for bilateral negotiation introduced by Anbarci [3], 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, and thus if there are m possible outcomes, the negotiation will last at most m rounds.

Our analysis consists of several complementary components: First, we analyze the *full information* setting, where both parties are informed of the other party's preferences. We provide elegant strategies that specify a Sub-game Perfect Equilibrium (SPE) for the parties. Our strategies are easy to implement and improve the previous result of Anbarci [4] by finding an SPE strategy in linear time instead of quadratic time. The SPE result is then linked to the designed *Rational Compromise* (RC) bargaining rule [20], which is a centralized procedure useful in a cooperative environment. Specifically, we show that the SPE result of the negotiation protocol is always part of the set of results returned by the RC rule, even though the negotiation protocol does not force the parties to offer specific outcomes in a specific order as the RC rule does. This non-trivial and perhaps surprising connection also enables us to prove that the SPE result of the protocol satisfies several axiomatic properties such as monotonicity which are normatively desired. We further analyze the *no information* setting where neither party knows the preference order of the other party nor do they have any prior probability distribution over possible orders. We begin by showing 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 the negotiation protocol. We further show that any pair of Maxmin strategies also specifies a robust-optimization equilibrium. Last, we empirically investigate the practical relevance of our full information analysis through a human study with 150 participants. Using ten negotiation instances, which vary in the degree of disagreement between the parties' preferences, we study how ordinary people negotiate using the studied protocol, either with each other or with an automated agent that implements the SPE strategy. Our results show that ordinary people tend to arrive at the SPE outcomes when negotiating with each other or with the automated agent. At the same time, our results show that people do not closely follow the SPE strategy. In addition, in contrast with the commonly held belief that equilibrium-following agents are unsuccessful when negotiating with people, our results show that our equilibrium-following agent is, in fact, highly successful and significantly outperforms human negotiators.

Above all, the contribution of this work is four-fold: First, we introduce elegant strategies that specify an SPE and provide a substantial analysis for showing that they indeed form an SPE. We also provide an improved algorithm for computing an 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 [20], who studied the RC rule, the 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 work is the first to find such a connection. The third contribution of this paper is the analysis of the negotiation protocol under a no information setting, which have not been considered before. Finally, our user study demonstrates the practical relevance of the SPE strategy and outcome as well as the adequacy of the SPE strategy for human-agent negotiation.

2. Related work

The basic rules of engagement between the negotiating agents, be they automated or otherwise, are defined through a negotiation protocol. The protocol determines, among other things, how the negotiation will proceed, the space of possible proposals, the rules determining how such proposals may be made, and how agreement or failure is determined [5]. Focusing on bilateral automated negotiation, one of the first protocols was introduced by Rubinstein [37]. Since then, a lot of work has been done to develop and investigate a variety of negotiation protocols that extend this basic alternating-offers negotiation protocol (see the following comprehensive books [32,21,16]). The traditional assumption in the negotiation theory is that there is a continuum of feasible outcomes, but many real-life negotiation scenarios violate this assumption. Indeed, there are several works that consider problems with a finite number of outcomes. For example, see [42], [25], [29] and [31]. All of these works focus on negotiation when the preferences are represented by a cardinal utility, while we study negotiation with ordinal preferences.

There are several works that studied bargaining rules with ordinal preferences over a finite set of outcomes (for example, [41]). 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. Sequential procedures, in particular the *fallback bargaining* method, have attracted considerable interest [39,19,8,20,11], since they satisfy some nice theoretical properties. All of these works are inherently different from our setting: bargaining rules are useful in a cooperative environment or where there is a central authority that can force the parties to offer specific outcomes in a specific order, while we study a negotiation protocol that is useful in a non-cooperative environment. Indeed, we show that the studied negotiation protocol almost completely implements the individually rational variant of the *fallback bargaining* method, i.e.,

[1] Note that this is also the typical assumption in most voting literature [9].

the *RC* rule [20]. We note that the *RC* bargaining rule is equivalent to *Bucklin* voting [18] 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 a few works that study negotiation protocols with ordinal preferences over a finite set of outcomes. De Clippel et al. [13] study the problem of selection of arbitrators, and they concentrate on two-step protocols. Barberà and Coelho [6] study three protocols that are refinements of "rules of k names", whereby one of the parties proposes a shortlist and the other chooses from it. Recently, Bol et al. [7] have conducted a human study to examine the efficiency and fairness of three negotiation protocols with ordinal preferences over five outcomes. The most closely related works are Anbarci's papers. In [3] 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 [4] 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 an 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 [40]. Our work provides a more efficient algorithm that finds an elegant SPE strategy and establishes the relationship between the VAOV protocol and the *RC* rule. In addition, we analyze the no information setting, which has yet to be investigated thus far. Last, we provide the first human study for this protocol.

Our human study follows a substantial line of research examining human negotiation decisions and outcomes. This research spans a wide range of disciplines, including psychology, business, economics, and others [34]. Common to most of these works is the observation that people's bargaining and negotiation behavior does not adhere to equilibrium [26,14]. From an automated negotiation perspective, this observation raises concerns over the adequacy of following equilibrium strategies when negotiating with people. Indeed, various empirical investigations have shown that equilibrium-following agents are often unsuccessful. For example, in Kraus et al. [22], the authors have found that despite the low complexity of finding an equilibrium strategy in their full information setting, an equilibrium-following agent was highly unsuccessful. In a no-information setting, Peled et al. [33] have provided similar results. Following these and similar works, it is often claimed that agents cannot rely on equilibrium strategies alone to negotiate well with people [23]. As a result, the design of such agents has relied primarily on machine learning techniques which often require extensive contextual data or hand-crafted heuristics which need not necessarily generalize well across domains and settings [36, Ch. 4.3].

Very few studies have demonstrated the usefulness of following an equilibrium strategy in human-agent interaction. Notably, Haim et al. [17] has demonstrated this in a negotiation setting over contracts in a three-player market. The authors attribute their equilibrium-following agent's success to the competitive nature of the market, however, they also claim that categorizing a priori the type of strategies that will succeed in a given negotiation environment is still an open challenge. In this respect, our human study is the first study to examine human-agent negotiation under ordinal preferences with an equilibrium-following agent. The importance of further promoting research into these and similar human-agent negotiation issues is also highlighted in the Automated Negotiating Agent Competition[2] which introduced the Human-Agent League since 2017 [27,28].

3. Formal model

We assume that there are two negotiation parties, p^1 and p^2, negotiating over a finite 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, 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 the outcome of the negotiation to be ranked as high as possible in her preference ordering. We assume that any agreement is preferred by both parties over a no-agreement outcome.

The negotiation protocol (VAOV, [3]) works as follows: The parties make offers alternately with no offer made twice. 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 (i.e., offers) 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 outcome 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 since both parties wish to avoid a no-agreement outcome.

4. Theoretical analysis

In the following, we provide a theoretical analysis of the formal model specified above under the full information (i.e., both parties know the preferences of the other) and no information (i.e., neither party knows the preference of the other) settings. For convenience, we denote p^i as the party whose turn it is to make an offer at a given round i, and p^j denotes the other party. That is, $p^i = p^1$ in odd round numbers and $p^i = p^2$ otherwise.

We start by providing a general result that is useful for both information settings. Consider the following definition:

[2] http://ii.tudelft.nl/negotiation/index.php/node/7.

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 acting rationally, regardless of the information setting. Intuitively, these are all the outcomes that are in the lower parts of the preference orders of either party, denoted by Low_t. We denote all of the other outcomes by JG_t (intuitively denoting the "Joint Goals").

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 acting rationally. Then, $o \notin Low_t$.*

Proof. Given a round t, let $m_t = |O_t|$. Starting from round 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 different information settings. In each case, we are interested in finding the best actions that a party should take, given the information that she has.

4.1. 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 deriving her best strategy. Therefore, in the full information setting, we are interested in finding an SPE. Since Anbarci [3] showed that the SPE result is unique, it suffices to find one set of strategies that specify an SPE.

4.1.1. SPE strategies

Recall that SPE is a function 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 $H_t^i = (o^1, o^2, ..., o^{t-1})$ be the history for p^i at round t, and recall 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 define the following offer strategy, which will later be shown to specify an SPE.

Strategy 1 *(Offer 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 $H_t^j = H_t^i + o = (o^1, o^2, ..., o^{t-1}, o)$ 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 $H_m^i = H_t^j + o_{t+1}^i + ... + o_{m-1}^i$. 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 offer strategy that is specified by Strategy 1 from round $t+1$ and on. Note that $o_a(H_t^j)$ remains the same regardless of the specific order in which outcomes from I_t are chosen. We define the following response strategy, which will later be shown to specify an 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, we consider two examples: an instance in which $I_1 = \emptyset$, and an instance in which $I_1 \neq \emptyset$.

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 the offer 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 $o_a((o_5)) = o_3$, as we will show, and $o_3 \succ_{p^2} o_5$ see Fig. 1(a)). In round 2, in which it is the turn of p^2 to make an offer, $L_2^1 = \{o_1, o_2\}$ and $L_2^2 = \{o_6, o_4\}$, and thus p^2 would offer o_1 which p^1 would reject for a similar reason (see Fig. 1(b)). In round 3 (Fig. 1(c)), it is p^1's turn to make an offer. In this round $L_3^1 = \{o_2\}$, $L_3^2 = \{o_6, o_4\}$, and p^1 would offer o_4. p^2 would then reject and, in turn, she would offer o_2. p^1 would

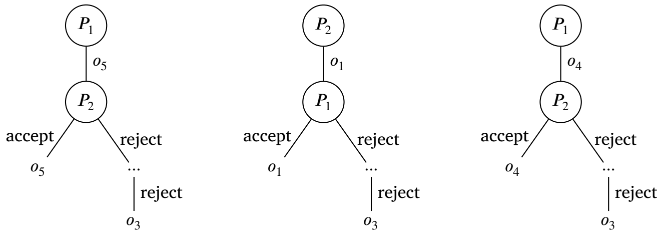

Fig. 1. An illustration of the first 3 rounds in Example 1.

reject and offer o_6, which 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 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 in which o_6 will be accepted as the result of the negotiation.

We now prove that our strategies specify an SPE. The key idea is to consider the outcomes each party could offer at each round and analyze the dynamics of the negotiation process. The proof proceeds by induction, starting from the base case of a negotiation with two outcomes. In this case, it is straightforward to show that the strategies specify an SPE. Then, the induction step considers negotiations with more than two outcomes. The proof first shows that there is no incentive to deviate from the response strategy, since the result of following the offer and response strategies is Pareto optimal. It then shows that there is no incentive to deviate from the offer strategy. To that end, we define functions that quantify the number of outcomes each party can offer until a specific outcome becomes part of their lower preference set.

We begin by noting that in the offer 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.

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,

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

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* $JG_t = \{o_{eq}\}$.

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^k_{t+1}, $k \in \{1,2\}$.

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* $|L^j_{t+1}| = |L^i_t|$.

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

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 Strategy 1. Specifically, Lemma 4 together with Lemma 5 cover the offer strategy where $I_t = \emptyset$, and Lemma 5 together with Lemma 6 cover the offer 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* $L^j_{t+1} \leftarrow L^i_t$.

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

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

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 $o \in L^j_t$ and p^j rejects it, o is the only outcome that becomes unavailable in round $t+1$, and we can thus be assured that $L^i_{t+1} = L^j_t \setminus \{o\}$. □

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

Proof. According to Lemma 3, the sets L^j_{t+1} and L^i_t have the same size. Therefore, if p^i offers $o \in L^i_t$ 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^j_{t+1}. □

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

Lemma 7. $o_a(H^j_t) = o_a(H^j_t + o^i_{t+1})$.

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

Lemma 8. *Given any history* H^i_t,

1. $o^i_t \prec_{p^j} o_{at}$.
2. $o_{at} \in JG_t$.
3. 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. □

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., $o_t^i, o_{t+1}^i, ..., o_{m-1}^i$), 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 $O_t, O_{t+1}, ..., O_m$.

Before we prove that Strategies 1 and 2 specify an 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 negotiation 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 an 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 offer 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 offer strategy.

Now, assume that if there are $m-1$ outcomes in round $t+1$, our strategies specify an SPE. We show that they specify an 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 an 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 $o_x \succ_{p^i} o_{at} = o_a(H_t^i + o_t^i)$. 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 $L^j_{t''}$, that is, $d_{j,x,t} \leq \ell_{j,t}$. In summary:

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

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

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

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, $o_x \succ_{p^i} o_{at}$, as demonstrated in Fig. 2. 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, $|L^i_{t+1}| + 1 = |L^j_t|$, 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:

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

If we combine (2) and (3), and apply Lemma 9, 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 once more 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, $|L^i_{t+1}| + 1 = |L^j_t|$, 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:

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

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 an SPE when there are $m-1$ outcomes in round $t+1$. Therefore, p^i does not have any incentive to deviate. □

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 an SPE strategy of [4] is not explicitly analyzed, but its running time is at least $O(m^2)$ since it requires finding all the Pareto optimal outcomes for a given state of the game tree (i.e., given any

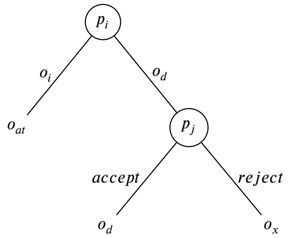

Fig. 2. p^i's decision node.

history H_t^i or H_t^j). We propose a completely different approach and provide 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 an 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.

4.1.2. 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 an SPE, and the SPE result is unique, we can infer that the SPE result is Pareto optimal.

Next, we analyze the relationship between the SPE result and the results of the designed *Rational Compromise* (*RC*) bargaining rule [20]. By establishing such a connection, one can transfer additional axiomatic properties that were previously proven for the RC outcome(s) to the SPE result of the VAOV protocol. Note that 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 $RC = \{o_x | \max_{o_x \in O} \min_{k\in\{1,2\}} (d_{k,x,1} + |L_1^k| - 1)\}$. It can also be computed by the following steps:

1. Let $v = 1$.
2. For each $k \in \{1,2\}$, let $B_v^k = \{$the v 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 strategies always result in a single outcome. Surprisingly, the SPE result is always part of the set returned by the RC rule. The intuition is that our strategies specify an 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 $JG_t = B^i_{|O_t|-|L_t^i|} \cap B^j_{|O_t|-|L_t^j|}$. If $|L_t^j| = |L_t^i|$, then for any v where $v \le |O_t| - |L_t^j|$, $B_v^i \cap B_v^j = \{o_{eq}\}$ or $B_v^i \cap B_v^j = \emptyset$. If $|L_t^j| = |L_t^i| + 1$, then for any v where the $v \le |O_t| - |L_t^j|$, $B_v^i \cap B_v^j = \{o_{eq}\}$ 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, $o_{eq} \in RC$. $\square$

Fig. 3 illustrates the relation between the step where RC stops, v, and the SPE result o_{eq} in three different scenarios. The orange color represents the outcomes of Low_t, where $I_t = \emptyset$, and the other outcomes are shown in blue color. In the first scenario (Fig. 3(a)) $|L_t^j| = |L_t^i|$ and $v < |O_t| - |L_t^j|$. That is, RC returns a single outcome, o_{eq}. In the second scenario (Fig. 3(b)) $|L_t^j| = |L_t^i| + 1$ and $v = |O_t| - |L_t^j|$. Therefore, RC returns a single outcome, o_{eq}. In the last scenario (Fig. 3(c)) $|L_t^j| = |L_t^i| + 1$, $v = |O_t| - |L_t^i|$, and RC returns two outcomes, $\{o_{eq}, o_x\}$.

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

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 $o_y \succ_{p^2} o_x$.*

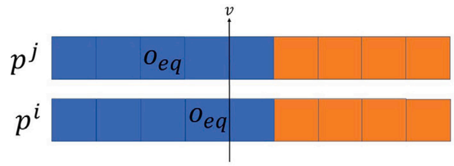

(a) First scenario: $|L_t^j| = |L_t^i|$ and $v < |O_t| - |L_t^j|$.

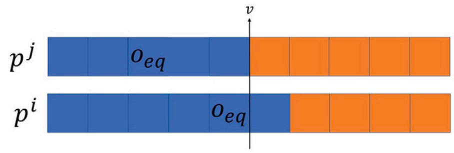

(b) Second scenario: $|L_t^j| = |L_t^i| + 1$ and $v = |O_t| - |L_t^j|$.

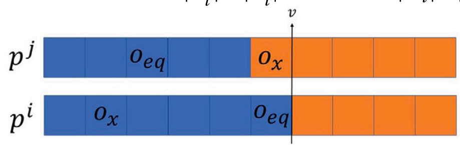

(c) Third scenario: $|L_t^j| = |L_t^i| + 1$ and $v = |O_t| - |L_t^i|$.

Fig. 3. An illustration of the relations between the step where RC stops, v, and the SPE result o_{eq}. (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

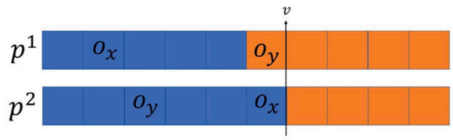

(a) $\ell_{1,1} + \ell_{2,1}$ is odd.

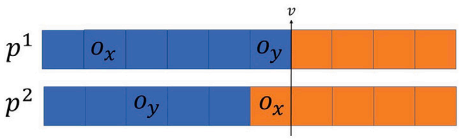

(b) $\ell_{1,1} + \ell_{2,1}$ is even.

Fig. 4. An illustration of the two scenarios where $|RC| = 2$.

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, $B^1_{|O_t| - |L_t^1|} \cap B^2_{|O_t| - |L_t^2|} = \{o_{eq}\}$. Therefore, $|L_t^1| \neq |L_t^2|$, and thus m_t is even. 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 $o_x \succ_{p^1} o_y$, $o_y \in L_t^1$, as illustrated in Fig. 4(a). 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 $o_y \succ_{p^2} o_x$, $o_x \in L_t^2$, as illustrated in Fig. 4(b). 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 VAOV 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. □

That is, adding outcomes that are preferred over o_{eq} by both parties causes RC to stop before reaching o_{eq}. To illustrate this, consider the following example. Assume that $\succ_i = o_5 \succ o_4 \succ o_3 \succ o_2 \succ o_1$ and $\succ_j = o_1 \succ o_3 \succ o_5 \succ o_2 \succ o_4$. Then, RC returns $\{o_3, o_5\}$ and the SPE result is o_3. Now assume that we add an outcome o_6, where $\succ'_i = o_5 \succ o_4 \succ o_6 \succ o_3 \succ o_2 \succ o_1$ and $\succ'_j = o_1 \succ o_6 \succ o_3 \succ o_5 \succ o_2 \succ o_4$. Clearly, RC returns o_6 that is also the SPE result, and $o_6 \succ o_3$ by both parties.

4.2. 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 an 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 model.*

Proof. Clearly, every ex-post SPE is also an SPE (i.e., in the full information setting), we can apply Strategies 1 and 2 to identify the SPE result. 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 an 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. □

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^k_t, i.e., the sets of the 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, which is composed of offer and response strategies, is defined as follows:

Strategy 3 *(Maxmin Strategy). Given a history H^i_t, offer any $o \in U^i_t$. Given a history H^j_t,* ***if*** *$o \in U^j_t$* ***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^k_1. 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 $o \notin L^j_t$ and p^j rejects it, then $L^i_{t+1} \leftarrow L^j_t \setminus \{o'\}$, 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. $\square$

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 $U^{max}_1 \succ_{p^{max}} o$.

Theorem 16. *Strategy 3 specifies a Maxmin strategy, and the Maxmin value of the game is the set U^{max}_1.*

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 $U^{max}_1 = \{o_1, o_2\}$ and clearly one of them will be the negotiation result. If $p^{max} = p^2$ then $U^{max}_1 = \{o_1\}$. 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 $U^{max}_1 \succ_{p^{max}} o_2$. 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^{max}_{t+1}. We show that our strategy specifies a Maxmin strategy, and the Maxmin value of the game is the set U^{max}_t 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 $U^{max}_t \succ_{p^{max}} o$ then p^{min} can accept it, and the negotiation results in o. On the other hand, if p^{max} offers any $o \in U^{max}_t$ 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^{max}_{t+1} by following our strategy. However, according to Lemma 4, $L^{max}_{t+1} = L^{max}_t$ and thus $U^{max}_{t+1} \cup \{o\} = U^{max}_t$. Overall, the Maxmin value of the game is the set U^{max}_t.

Now assume that it is p^{min}'s turn to offer, and p^{min} offers $o \in U^{max}_t$. Clearly, if p^{max} accepts then the negotiation result is from U^{max}_t. If p^{max} deviates and rejects, then according to induction assumption p^{max} can guarantee the Maxmin value of U^{max}_{t+1}. However, according to Lemma 15, $L^{max}_{t+1} = L^{max}_t \setminus \{o'\}$, and thus $U^{max}_{t+1} = U^{max}_t \setminus \{o\} \cup \{o'\}$. That is, o' is a possible result of the negotiation even though $U^{max}_t \succ_{p^{max}} o'$. Finally, assume that p^{min} offers $o \notin U^{max}_t$. 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^{max}_{t+1}. However, according to Lemma 5, $L^{max}_{t+1} = L^{max}_t \setminus \{o\}$, and thus $U^{max}_{t+1} = U^{max}_t$. $\square$

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^k_1) 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 acting rationally and will thus follow the Maxmin strategy? Clearly, the negotiation result will be an outcome o such that $o \in U^1_1 \cap U^2_1$. 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 $U \subset U^i_1$, since we proved that this is the Maxmin value. However, she can heuristically offer in each round t the best outcome in U^i_t, instead of an arbitrarily chosen $o \in U^i_t$. Since $|U^j_t| \geq |L^j_t|$, if the other party p^j is also acting rationally and plays the Maxmin strategy, there are more cases in which 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^i_t.

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 [1], 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 $\prec = o_1 \prec o_2 \prec \ldots \prec o_m$ then $\prec_{op} = o_1 \succ o_2 \succ \ldots \succ o_m$. According to Theorem 16, if s_k is a Maxmin strategy then the negotiation result is $o \in U^k_1$. That is, the worst negotiation result for p^k is the least preferred outcome in U^k_1, denoted by o_{wo}. Since the other party p^{3-k} is also using a Maxmin strategy, $o \in U^k_1 \cap U^{3-k}_1$. For every preference order, $\prec$, if the preference order of $p^{3-k} = \prec_{op}$, $U^k_1 \cap U^{3-k}_1 = \{o_{wo}\}$. 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. □

5. Human study

In the following, we present a study on how ordinary people negotiate, either with each other or with an automated agent, using the studied protocol. We focus on the full information setting, where an SPE strategy and outcome were identified and analyzed earlier in this article; the empirical investigation of the no-information setting is left for future work. The reason is that an ex-post SPE does not exist and the Maxmin strategy may entail various actions and outcomes which, from a theoretical standpoint, are considered equally appropriate. As such, an additional strategy has to be defined in order to select one offer from the various possible ones for any agent implementation. The evaluation of such possible strategies is outside the scope of this work.

In our context, we are mainly interested in three key issues: First, whether the negotiation outcomes coincide with the SPE; Second, whether the individual decisions made by the participants are aligned with those prescribed by the SPE strategy (Strategies 1 and 2). It is important to note in this context, that reaching an equilibrium outcome need not necessarily mean that the negotiators followed an equilibrium strategy, hence one needs to examine both issues. Last, whether an equilibrium-following agent performs well with people. To that end, we devised the following human study, which consists of two experimental setups. In both setups, participants were presented with the following negotiation task:

Motivating Scenario: *A colleague and you were asked to select the next venue for the company's retreat. There are seven possible venues: 'Los Angeles', 'Buenos Aires', 'London', 'New York', 'Rome', 'Tokyo' and 'Paris'. In the following, you will be presented with 10 instances, in each you will be given a different preferences profile (i.e., ordering over the seven possible venues). For each instance, you will be asked to negotiate with a colleague. Your goal is that the negotiation outcome will be ranked as high as possible in your preference orderings. Clearly, your colleague tries to achieve the same while considering her own preference ordering. Both of you are informed of your own preference order as well as your counterpart's preference order.*

Ten negotiation instances were devised such that each encompasses a different level of disagreement between the negotiators' preferences, ranging from having the exact same preference order to having reversed ones. The ten instances, along with their characteristics (e.g., which agent starts the negotiation) and the computed SPE and RC outcomes (for completeness), are provided in Appendix B.

In Experimental Setup 1, participants negotiated with each other (i.e., Human vs. Human). Experimental Setup 2 replicates Experiment 1 while replacing one of the negotiating parties with an automated agent that implements strategies 1 and 2 (i.e., Human vs. Agent). In both experiments, participants did not know who they negotiate with.[3] In both setups, participants were not provided with any advice or recommendation on how they should negotiate in order to pursue their objective.

5.1. Participants and procedure

All participants ($N = 150$, 91 males, average age 24) were recruited via ads posted on the academic platform used by the Computer Science undergraduate programs of the authors' universities. All participants took part in this study of their own free will. The study was approved by the corresponding IRB.

Participants were asked to log in to a designated system, at a specified time, using their personal computer. Once participants logged into the system, they first had to fill out a standard informed consent form, followed by a few basic demographic questions. Then, the motivating scenario described above was presented along with an explanation of the VAOV negotiation protocol. Then, the participants were presented with detailed instructions on how they respond to an offer and make an offer. Participants then had to pass a short quiz, ensuring they understood the negotiation protocol and the task at hand. Both the instructions and the quiz are available in Appendix C (see Figs. C.9, C.10, C.11, and C.12). Then, each participant was randomly assigned to one of the two experimental setups discussed above. In each setup, the ten instances were presented, one after the other, in a random order. Overall, 100 participants were assigned to Experimental Setup 1 and 50 participants were assigned to Experimental Setup 2, resulting in 1,500 negotiations in each setup.[4] Recall that participants were unaware of who or what they were negotiating with.

In order to avoid possibly under-quality data, negotiations in which at least one response move was extremely fast (less than 1 second) or unreasonably slow (more than 100 seconds) were omitted. Less than 10% of the negotiations were omitted from further analysis.

For each negotiation, we first identify the negotiation outcome and contrast it with the SPE and RC outcomes. Second, we examine the individual decisions made by the participants and compare them with those prescribed by our proposed strategies (Strategies 1 and 2).

[3] Short informal interviews with some of the participants indicated that, indeed, the participants were unaware of the fact they negotiated with a person or an agent.

[4] 50 human pairs x 10 instances in setup 1 plus 50 human participants who negotiated with an agent x 10 instances in setup 2.

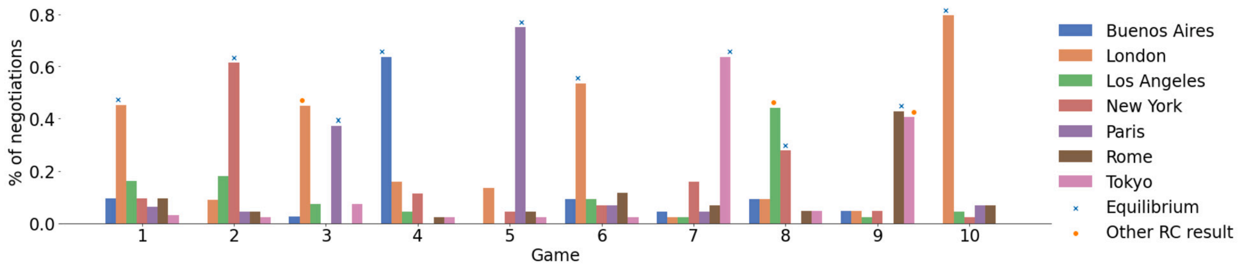

Fig. 5. Negotiation outcomes for Experimental Setup 1 - Human vs. Human. The X-Axis denotes the negotiation instance and the Y-Axis represents the percentage of negotiations that resulted in the corresponding outcome. The SPE and any additional RC outcomes are marked by x and o, respectively.

Table 1
Negotiation outcomes for Experimental Setup 1 - Human vs. Human. The first column denotes the negotiation instance. The SPE outcome is marked in bold and any additional RC outcome is marked in italic.

#	Buenos Aires	London	Los Angeles	New York	Paris	Rome	Tokyo	Total
1	5	**23**	8	5	3	5	1	50
2	0	4	9	**30**	2	2	1	48
3	1	**18**	3	0	*14*	0	3	39
4	**30**	7	2	5	0	1	1	46
5	0	7	0	2	**37**	2	1	49
6	4	**25**	4	3	3	5	1	45
7	2	1	1	8	2	3	**30**	47
8	4	4	*19*	**12**	0	2	2	43
9	2	2	1	2	0	**20**	*19*	46
10	0	**40**	2	1	3	3	0	49

5.2. Experimental setup 1 - human vs. human

Starting with the negotiation outcome, as can be seen in Fig. 5 and Table 1, in eight out of ten negotiation instances, the majority of negotiations have resulted in the predicted SPE outcome. In the remaining two instances, the SPE outcome is second only to the additional RC outcome, which exists only in three instances. Recall that a non-SPE RC outcome is strongly related to the SPE outcome as shown in Theorem 12.

Fig. 6 presents the individual actions taken by the participants compared to those prescribed by Strategies 1 and 2. Starting with the response actions (i.e., accepting or rejecting an offer), we see that participants' actions generally follow Strategy 2 with more than 60% of all response decisions in all instances adhering to the proposed strategy. In other words, participants generally tend to accept the SPE outcome and reject other proposals. However, considering the participants' offers (i.e., possible outcomes which were yet to be rejected), we see that these are only weakly aligned with Strategy 1, with all instances presenting less than 20% adherence. The combination of these two results is somewhat surprising since it suggests that participants did arrive at the predicted outcomes, yet they only seem to follow the response strategy to a reasonable extent. One possible explanation for this result is that the participants were able to identify the SPE as desired outcomes, and hence were able to accept them while rejecting most other offers. On the other hand, given the complexity of the offer strategy and the long negotiation process it entails (i.e., if both parties follow the proposed strategies then the negotiation process would consist of many offers being exchanged and rejected until the SPE is proposed and accepted), "shortcuts" were practiced and the SPE outcome was proposed and accepted much sooner than it should under the equilibrium strategy. Indeed, the data supports this explanation – the average number of offers made in this setup was statistically significantly lower than that of Experimental Setup 2 (i.e., Human vs. Agent) for all 10 instances at $p < 0.05$. Specifically, in this Human vs. Human setup, people exchanged 2.2 offers on average before the negotiation terminated whereas in the Human vs. Agent setup, more than twice the number of offers were exchanged (4.7 offers on average). Note if both parties follow the SPE strategies then 7 offers are to be exchanged.

5.3. Experimental setup 2 - human vs. agent

Similarly to Experimental Setup 1, as can be seen in Fig. 7 and Table 2, in six negotiation instances, the vast majority of negotiations have resulted in the predicted SPE outcome. In three of the remaining four instances, the SPE outcome is second to the additional RC outcome. Only in instance 4, the SPE outcome is second to a non-RC outcome. A closer examination of this instance reveals a possible explanation for this abnormality – the agent, who starts the negotiation, first offers Tokyo according to Strategy 1. The vast majority of human negotiators rejected the offer (since it is ranked very low in their preference order), as they should according to Strategy 2. Then, about 25% of the participants have offered their highest ranking choice - Buenos

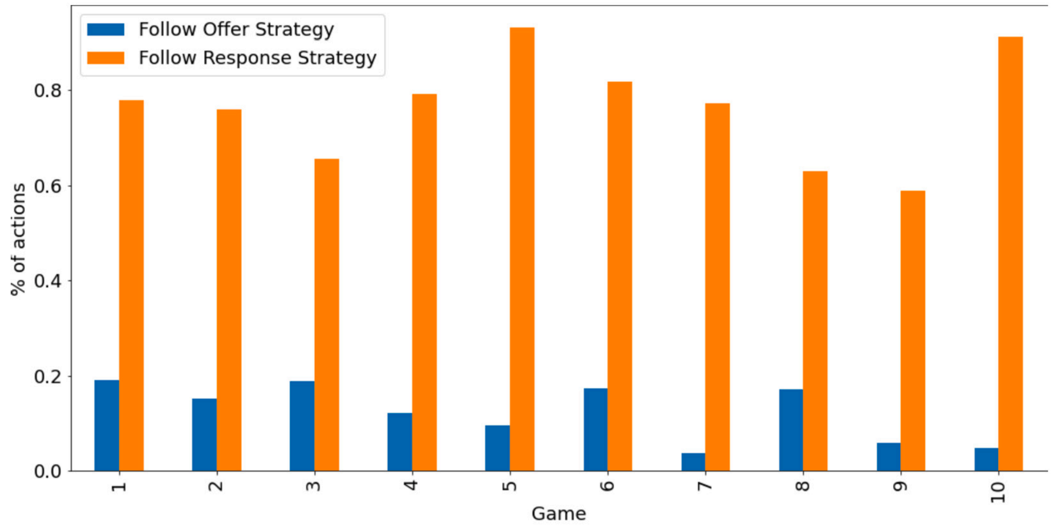

Fig. 6. Negotiation actions in Experimental Setup 1 compared to the actions prescribed by the SPE strategy.

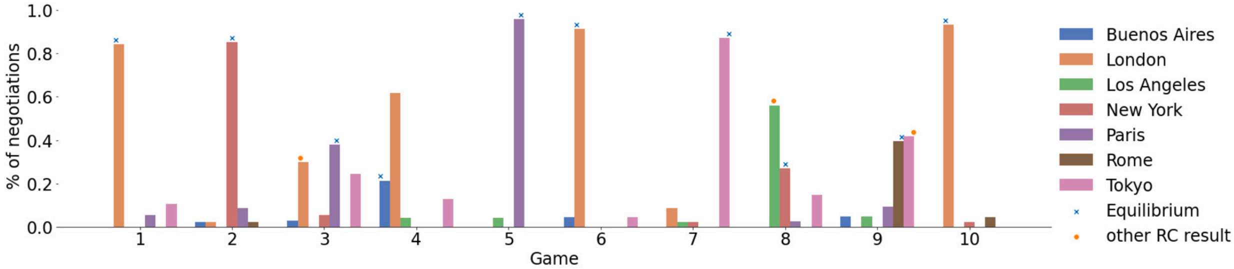

Fig. 7. Negotiation outcomes for Experimental Setup 2 - Human vs. Agent. The X-Axis denotes the negotiation instance and the Y-Axis represents the percentage of negotiations that resulted in the corresponding outcome. The SPE and any additional RC outcomes are marked by x and o, respectively.

Aires, deviating from Strategy 1. By rejecting this offer, the agent is able to "force" the negotiation to result in its most preferred option – London – which is not the SPE outcome. In other words, the second offer made by the human participants often allowed the agent to capitalize on their deviation from Strategy 1 (i.e., "mistake") and improve its own outcome. Further support for this explanation can be found when examining the differences in outcome distribution between Experimental Setups 1 and 2 (Figs. 5 and 7). In most instances, the differences are associated with better outcomes for the agent (and in some cases, also for the human), i.e., outcomes that are preferred by the agent are more prevalent in setup 2 (Human vs. Agent) than in setup 1 (Human vs. Human). For example, considering instance 3, we see that the Tokyo outcome, which is not an SPE or RC outcome, was much more prevalent in setup 2 while the London outcome, which is in RC, is much less prevalent. Tokyo is preferred by the agent to both the SPE and RC outcomes, while both are preferred by the human participant. As such, in the context of outcome distribution, the differences between setups 1 and 2 seem to be indicative of the strategic capabilities of the agent rather than any perplexing deviation from the SPE outcomes.

Considering the individual actions taken by the human participants, we see that participants' actions are, generally, better aligned with those of Strategies 1 and 2 compared to Experimental Setup 1. Starting with the participants' responses, in all instances, more than 83% of the responses coincide with those prescribed by Strategy 2. This result is not surprising as the agent, which implements Strategy 1, starts by offering low-ranking outcomes which are naturally rejected by most participants (as per Strategy 2). Considering the participants' offers, we see that a significantly larger portion of offers is aligned with Strategy 1 compared to Experimental Setup 1. Specifically, in 8 out of 10 instances, more than 35% of the offers coincide with those prescribed by Strategy 1. Since the participants of this study are assigned to the different setups completely at random, it is unreasonable to suspect that the participants in this setup are "more strategic" than those of setup 1. One plausible explanation is that participants were affected by the agent's offer strategy. Specifically, given that the agent's offers are lower-ranked in the human participant's preference order, the participant might have "retaliated" by doing the same. It is important to note that we did not find any significant temporal effect for the repeated interaction, namely, the percentage of participant's offers that coincide with Strategy 1 is not correlated with the number of offers or rounds played by that participant.

In order to evaluate the agent's performance, we further compare the negotiated outcomes in setups 1 and 2 with respect to the agent's and human's preferences. Specifically, for each negotiation instance, we compare the agent's average ranking of the negotiated

Table 2
Negotiation outcomes for Experimental Setup 2 - Human vs. Agent. The first column denotes the negotiation instance. The SPE outcome is marked in bold and any additional RC outcome is marked in italic.

#	Buenos Aires	London	Los Angeles	New York	Paris	Rome	Tokyo	Total
1	0	**41**	0	0	3	0	5	49
2	1	1	0	**41**	4	1	0	48
3	1	*14*	0	2	**18**	0	11	46
4	**10**	29	2	0	0	0	6	47
5	0	0	0	2	**46**	0	0	48
6	2	**45**	0	0	0	0	2	49
7	3	1	1	0	0	0	**39**	44
8	0	0	*27*	**13**	1	0	7	48
9	2	0	2	0	4	**20**	*22*	50
10	0	**46**	0	1	0	2	0	49

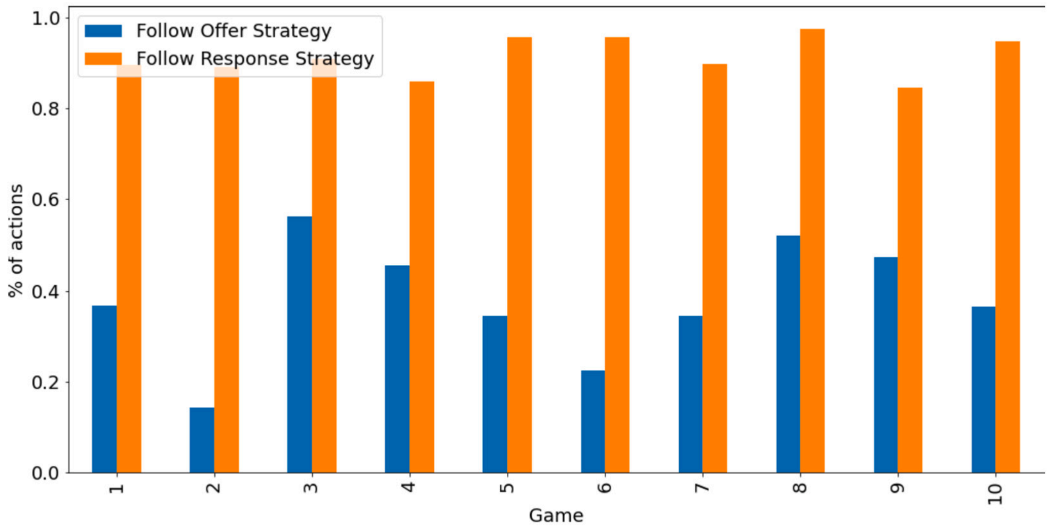

Fig. 8. Negotiation actions in Experimental Setup 2 compared to the actions prescribed by Strategies 1 and 2.

outcomes in setup 2 to the corresponding average ranking of the negotiated outcomes for the same negotiation role in setup 1. That is, we compare the agent's performance only to the people who assumed the same role as the agent in setup 1. Surprisingly, despite the scarce evidence to support the adequacy of equilibrium-following strategies for the agent in human-agent negotiation, in 7 out of the 10 instances, the agent's average ranking was statistically significantly higher than that of the corresponding human negotiators in setup 1, with the difference ranging across negotiation instances from 0.3 to 1.1 ranking positions, $p < 0.05$. In the remaining 3 instances, no statistically significant differences were encountered.

Overall, we believe that the success of our agent should be primarily attributed to the nature of the SPE outcome and strategy. First and foremost, we see that people tend to reach the SPE outcome when negotiating with other people (Fig. 5) and that this outcome has favorable axiomatic properties (Theorem 11). As such, it is reasonable to expect people to perceive this outcome as "fair" or "desirable" and thus people are unlikely to be deterred from reaching this outcome when negotiating with an agent. From the SPE strategy perspective, intuitively, the agent's offers "move up" the human negotiator's preference order over the negotiation process. Since the human negotiator has no incentive to accept any low-ranking offers, these are mostly rejected by people as can be seen in Fig. 8, thus avoiding clearly unsuccessful outcomes. In addition, any deviation by a human negotiator from the offer strategy, for example by trying to use "shortcuts" in reaching the SPE outcome (or any other outcome) as observed in setup 1 (Fig. 5), is likely to be costly as the agent can capitalize on their behavior and force the negotiation to end in a better outcome for the agent as demonstrated earlier using instance 4. Taken jointly, people have no clear incentive to deviate from the SPE strategy, which clearly leads to the SPE outcome, nor do they have any clear incentive to refrain from reaching that outcome. It is important to note, however, that these properties are associated with the VAOV negotiation protocol. As such, equilibrium-following agents applied to other negotiation protocols and strategic interactions, which are not associated with the same properties, may not be as successful as documented in previous literature (e.g., [22,33]).

6. Conclusion and future work

We investigated the negotiation over a finite set of outcomes, assuming ordinal preferences, using the VAOV negotiation protocol. We introduced strategies that specify an SPE and improved upon previous results by providing a linear time algorithm that computes an SPE strategy. We provided a substantial analysis of our strategies, which establishes a link between the SPE result of the protocol in a non-cooperative setting, to the result of the RC rule in a cooperative setting. We further analyzed the no information setting and show that, in our setting, every pair of Maxmin strategies specifies a robust-optimization equilibrium. Finally, through a human study of the full information settings, we have revealed an intriguing phenomenon where people tend to reach the SPE outcomes despite very frequently deviating from the proposed equilibrium strategies. Moreover, contrary to commonly held belief, an agent which follows the identified SPE strategy is shown to be highly successful when negotiating with people. Taken jointly, our results indicate that the studied VAOV negotiation protocol with ordinal preferences is especially suitable for non-cooperative, multi-agent systems with or without human negotiators.

In future work, we plan to extend this work in five directions: First, we plan to perform an extensive empirical evaluation of the no information setting as discussed earlier in Section 5. Second, we plan to replicate our human study with participants of different backgrounds and cultures, which may differ from our participant pool. Third, we plan to extend the protocol to a multi-party setting (i.e., more than two negotiators) and analyze the resulting SPE and its properties. Note that our fundamental Lemma 1 does not apply in the multi-party case as acceptance criteria have to be defined to determine when the negotiation ends. Given different acceptance criteria, our lemma and the subsequent analysis need not necessarily hold. Fourth, we plan to examine a variant of the VAOV protocol where the negotiation must terminate after m' rounds, where $m' < m$. As in the case above, our fundamental Lemma 1 does not hold in this case. Last, we intend to investigate the additional implementation of other bargaining rules by negotiation protocols, similar to the implementation of the RC bargaining rule by the VAOV protocol.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Sarit Kraus reports financial support was provided by Israel Science Foundation. Sarit Kraus reports financial support was provided by European Union. Noam Hazon reports financial support was provided by Israel Ministry of Innovation Science & Technology.

Data availability

No data was used for the research described in the article.

Acknowledgements

Sarit Kraus has been supported in part by the Israel Science Foundation under grant 1958/20 and the EU Project TAILOR under grant 952215. Noam Hazon has been supported in part by the Ministry of Science, Technology and Space (MOST), Israel.

Appendix A. Uniqueness of the SPE result

Even though the uniqueness of the SPE result was proven elsewhere [3], we provide direct and simpler proof.

Theorem 18. *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, p^j either 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 an SPE she will choose either an outcome that p^j will accept or a sub-tree of the game, that results 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. □

Appendix B. Negotiation instances

In Table B.3 we present the ten negotiation instances used in our human study. The instances were generated automatically such that each represents a different degree of disagreement between the preferences of the two parties. Specifically, instance 1 presents a completely reversed order, whereas instance 2 presents an identical order. Instances 3 and 4 represent the setting in which both parties agree on the top 4 (and bottom 3) venues, but rank them differently. Instances 5 and 6 represent a setting in which the two parties do not agree on any venue in the top 4 (and bottom 3). Instances 7 and 8 follow the same rationale, with the two parties agreeing on 2 out of the 4 top (and 1 out of 3) venues and rank them differently. Instances 9 and 10 complete the picture, with the two parties agreeing on 3 out of the 4 top (and 2 out of 3) venues and rank them differently.

Table B.3
The ten negotiation instances used in the human study. Player p^i begins the negotiation process in instances 1,3,6,8 and 9. p_j is played by the automated agent in Experimental Setup 2.

#	Pref. p^i	Pref. p^j	SPE	RC
1	$NewYork > LosAngeles > Rome > London > Paris > BuenosAires > Tokyo$	$Tokyo > BuenosAires > Paris > London > Rome > LosAngeles > NewYork$	London	London
2	$NewYork > Los_Angeles > Rome > London > Paris > BuenosAires > Tokyo$	$NewYork > LosAngeles > Rome > London > Paris > BuenosAires > Tokyo$	New York	New York
3	$Paris > LosAngeles > London > Tokyo > NewYork > BuenosAires > Rome$	$Tokyo > London > Paris > LosAngeles > BuenosAires > Rome > NewYork$	Paris	Paris London
4	$BuenosAires > NewYork > London > LosAngeles > Tokyo > Rome > Paris$	$London > BuenosAires > LosAngeles > NewYork > Paris > Tokyo > Rome$	Buenos Aires	Buenos Aires
5	$NewYork > Paris > Tokyo > BuenosAires > Rome > London > LosAngeles$	$Paris > Rome > LosAngeles > London > Tokyo > NewYork > BuenosAires$	Paris	Paris
6	$NewYork > LosAngeles > Rome > London > Paris > BuenosAires > Tokyo$	$Tokyo > BuenosAires > Paris > London > Rome > LosAngeles > NewYork$	London	London
7	$Tokyo > BuenosAires > Rome > NewYork > London > Paris > LosAngeles$	$Paris > Tokyo > New_Y ork > LosAngeles > BuenosAires > London > Rome$	Tokyo	Tokyo
8	$NewYork > London > LosAngeles > Paris > Tokyo > Rome > BuenosAires$	$LosAngeles > Tokyo > NewYork > Rome > London > Paris > BuenosAires$	New York	New York Los Angeles
9	$Rome > Tokyo > Paris > London > LosAngeles > BuenosAires > NewYork$	$Tokyo > Rome > London > NewYork > LosAngeles > BuenosAires > Paris$	Rome	Rome Tokyo
10	$Paris > London > BuenosAires > LosAngeles > NewYork > Rome > Tokyo$	$Tokyo > London > NewYork > BuenosAires > Paris > Rome > LosAngeles$	London	London

Appendix C. The instructions and the quiz from the experiment

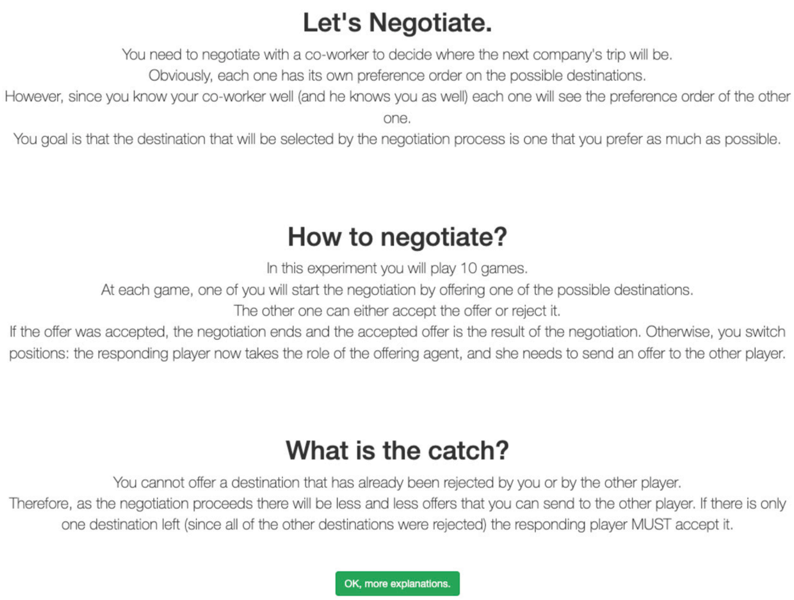

Fig. C.9. The motivating scenario, and the explanation of the VAOV negotiation protocol.

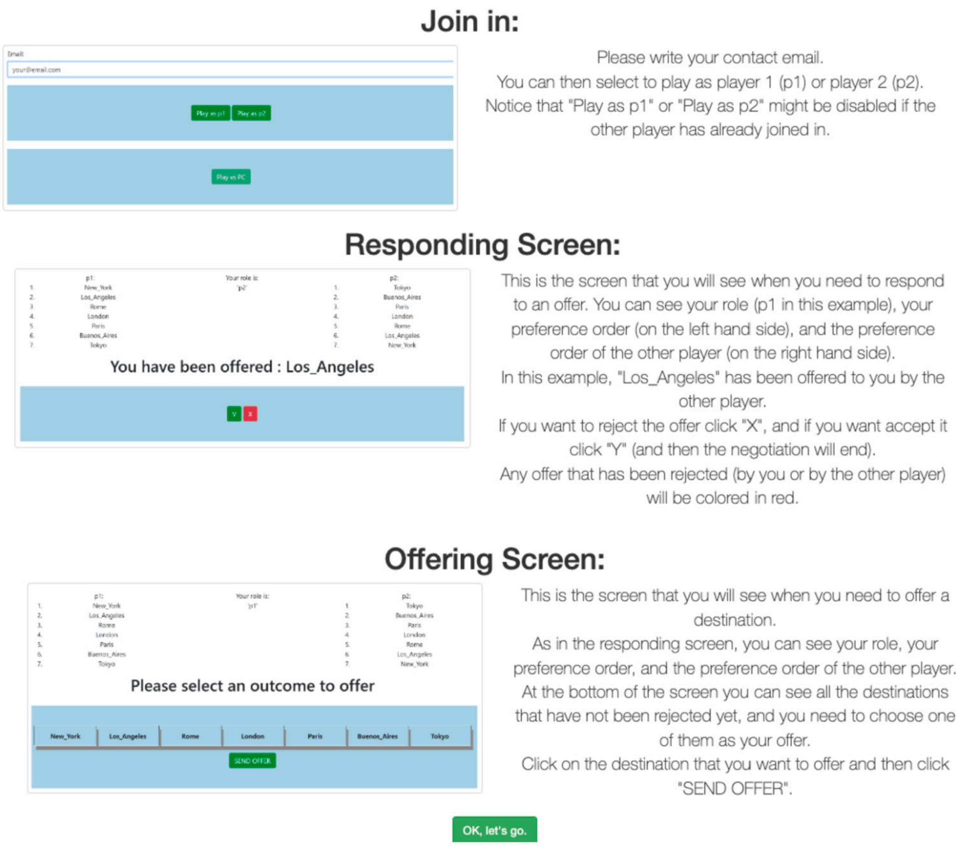

Fig. C.10. The instruction on how to respond to an offer and make an offer.

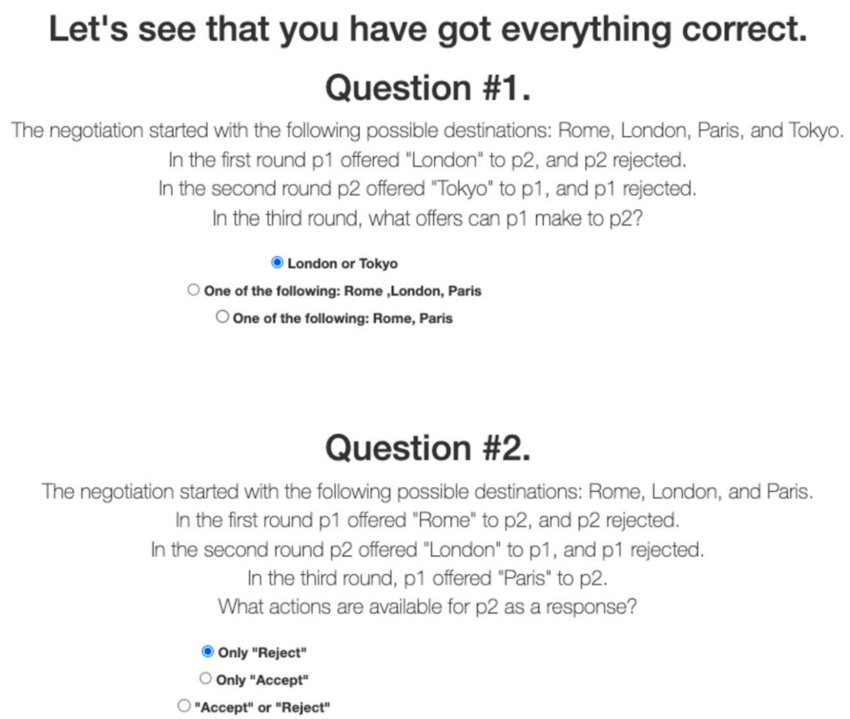

Fig. C.11. Quiz, part 1.

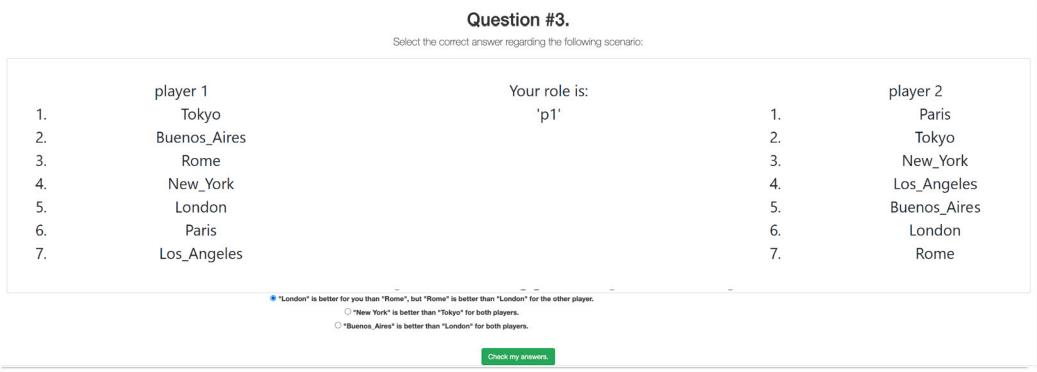

Fig. C.12. Quiz, part 2.

References

[1] Michele Aghassi, Dimitris Bertsimas, Robust game theory, Math. Program. 107 (1–2) (2006) 231–273.

[2] Shehzad Ali, Sarah Ronaldson, Ordinal preference elicitation methods in health economics and health services research: using discrete choice experiments and ranking methods, Br. Med. Bull. 103 (1) (2012) 21–44.

[3] Nejat Anbarci, Noncooperative foundations of the area monotonic solution, Q. J. Econ. 108 (1) (1993) 245–258.

[4] Nejat Anbarci, Finite alternating-move arbitration schemes and the equal area solution, Theory Decis. 61 (1) (2006) 21–50.

[5] Reyhan Aydoğan, David Festen, Koen V. Hindriks, Catholijn M. Jonker, Alternating offers protocols for multilateral negotiation, in: Modern Approaches to Agent-Based Complex Automated Negotiation, Springer, 2017, pp. 153–167.

[6] Salvador Barberà, Danilo Coelho, Compromising on compromise rules, Rand J. Econ. 53 (1) (2022) 95–112.

[7] Damien Bol, Jean-François Laslier, Matías Núñez, Two person bargaining mechanisms: a laboratory experiment, Group Decis. Negot. 31 (6) (2022) 1145–1177.

[8] Steven J. Brams, D. Marc Kilgour, Fallback bargaining, Group Decis. Negot. 10 (4) (2001) 287–316.

[9] Felix Brandt, Vincent Conitzer, Ulle Endriss, Ariel D. Procaccia, Jérôme Lang, Handbook of Computational Social Choice, Cambridge University Press, 2016.

[10] Michal Chalamish, Sarit Kraus, Automed: an automated mediator for multi-issue bilateral negotiations, Auton. Agents Multi-Agent Syst. 24 (2012) 536–564.

[11] John P. Conley, Simon Wilkie, The ordinal egalitarian bargaining solution for finite choice sets, Soc. Choice Welf. 38 (1) (2012) 23–42.

[12] Andreas Darmann, Christian Klamler, Proportional Borda allocations, Soc. Choice Welf. 47 (3) (2016) 543–558.

[13] Geoffroy De Clippel, Kfir Eliaz, Brian Knight, On the selection of arbitrators, Am. Econ. Rev. 104 (11) (2014) 3434–3458.

[14] Ido Erev, Alvin E. Roth, Predicting how people play games: reinforcement learning in experimental games with unique, mixed strategy equilibria, Am. Econ. Rev. (1998) 848–881.

[15] Sefi Erlich, Noam Hazon, Sarit Kraus, Negotiation strategies for agents with ordinal preferences, in: The 27th International Joint Conference on Artificial Intelligence, 2018, pp. 210–218.

[16] Shaheen Fatima, Sarit Kraus, Michael Wooldridge, Principles of Automated Negotiation, Cambridge University Press, 2014.

[17] Galit Haim, Bo An, Sarit Kraus, et al., Human–computer negotiation in a three player market setting, Artif. Intell. 246 (2017) 34–52.

[18] Clarence Gilbert Hoag, George Hervey Hallett, Proportional Representation, The Macmillan Company, New York, 1926.

[19] Leonid Hurwicz, Murat R. Sertel, Designing mechanisms, in particular for electoral systems: the majoritarian compromise, in: Contemporary Economic Issues, Springer, 1999, pp. 69–88.

[20] Özgür Kıbrıs, Murat R. Sertel, Bargaining over a finite set of alternatives, Soc. Choice Welf. 28 (2007) 421–437.

[21] Sarit Kraus, Strategic Negotiation in Multiagent Environments, MIT Press, 2001.

[22] Sarit Kraus, Penina Hoz-Weiss, Jonathan Wilkenfeld, David R. Andersen, Amy Pate, Resolving crises through automated bilateral negotiations, Artif. Intell. 172 (1) (2008) 1–18.

[23] Raz Lin, Sarit Kraus, Can automated agents proficiently negotiate with humans?, Commun. ACM 53 (1) (2010) 78–88.

[24] Fernando Lopes, Michael Wooldridge, Augusto Q. Novais, Negotiation among autonomous computational agents: principles, analysis and challenges, Artif. Intell. Rev. 29 (1) (2008) 1–44.

[25] Marco Mariotti, Nash bargaining theory when the number of alternatives can be finite, Soc. Choice Welf. 15 (3) (1998) 413–421.

[26] Richard D. McKelvey, Thomas R. Palfrey, An experimental study of the centipede game, Econometrica (1992) 803–836.

[27] Johnathan Mell, Jonathan Gratch, Tim Baarslag, Reyhan Aydoğan, Catholijn M. Jonker, Results of the first annual human-agent league of the automated negotiating agents competition, in: Proceedings of the 18th International Conference on Intelligent Virtual Agents, 2018, pp. 23–28.

[28] Johnathan Mell, Jonathan Gratch, Reyhan Aydoğan, Tim Baarslag, Catholijn M. Jonker, The likeability-success tradeoff: results of the 2nd annual human-agent automated negotiating agents competition, in: 2019 8th International Conference on Affective Computing and Intelligent Interaction (ACII), IEEE, 2019, pp. 1–7.

[29] Ryo-ichi Nagahisa, Makoto Tanaka, An axiomatization of the Kalai-Smorodinsky solution when the feasible sets can be finite, Soc. Choice Welf. 19 (4) (2002) 751–761.

[30] John F. Nash Jr., The bargaining problem, Econometrica 18 (2) (1950) 155–162.

[31] Matias Nunez, Jean-François Laslier, Bargaining through approval, J. Math. Econ. 60 (2015) 63–73.

[32] Martin J. Osborne, Ariel Rubinstein, Bargaining and Markets, Academic Press, 1990.

[33] Noam Peled, Ya'akov Kobi Gal, Sarit Kraus, A study of computational and human strategies in revelation games, Auton. Agents Multi-Agent Syst. 29 (1) (2015) 73–97.

[34] Howard Raiffa, The Art and Science of Negotiation, Harvard University Press, 1982.

[35] Ariel Rosenfeld, Avinatan Hassidim, Too smart for their own good: trading truthfulness for efficiency in the Israeli medical internship market, Judgm. Decis. Mak. 15 (5) (2020) 727–740.

[36] Ariel Rosenfeld, Sarit Kraus, Predicting Human Decision-Making: From Prediction to Action, in: Synthesis Lectures on Artificial Intelligence and Machine Learning, 2018, pp. 1–150.

[37] Ariel Rubinstein, Perfect equilibrium in a bargaining model, Econometrica (1982) 97–109.

[38] Lloyd S. Shapley, Utility Comparison and the Theory of Games, Cambridge University Press, Cambridge, 1969, Originally published in La Decision.

[39] Yves Sprumont, Intermediate preferences and Rawlsian arbitration rules, Soc. Choice Welf. 10 (1) (1993) 1–15.

[40] William Thomson, Cooperative models of bargaining, in: Handbook of Game Theory with Economic Applications, vol. 2, 1994, pp. 1237–1284.

[41] Dongmo Zhang, Yan Zhang, An ordinal bargaining solution with fixed-point property, J. Artif. Intell. Res. 33 (2008) 433–464.

[42] Gilad Zlotkin, Jeffrey S. Rosenschein, Mechanism design for automated negotiation, and its application to task oriented domains, Artif. Intell. 86 (2) (1996) 195–244.