

# Reasoning about Goal Revelation in Human Negotiation

Sohan D'souza<sup>1</sup>, Ya'akov Gal<sup>2</sup>, Philippe Pasquier<sup>3</sup>, Sherief Abdallah<sup>4,5</sup>, and Iyad Rahwan<sup>1,6</sup>

<sup>1</sup>Masdar Institute, UAE

<sup>2</sup>Harvard University, USA

<sup>3</sup>Simon Fraser University, Canada

<sup>4</sup>British University in Dubai, UAE

<sup>5</sup>University of Edinburgh, UK

<sup>6</sup>Massachusetts Institute of Technology, USA

**Abstract**—This paper studies how people reveal private information in strategic settings in which participants need to negotiate over resources, but are uncertain about each other's objectives. The study compared two negotiation protocols which differed in whether they allowed participants to disclose their objectives in a repeated negotiation setting of incomplete information. Results show that most people agree to reveal their goals when asked, and this leads participants to more beneficial agreements. Machine learning was used to model the likelihood that people reveal their goals in negotiation, and this model was used to make goal request decisions in the game. In simulation, use of this model is shown to outperform people making the same type of decisions. These results demonstrate the benefit of this approach towards designing agents to negotiate with people under incomplete information.

**Index Terms**—computer-supported cooperative work, evaluation/methodology, decision-support

## I. INTRODUCTION

This paper studies the use of computer-based protocols for facilitating and modeling human negotiation in strategic settings where parties lack information about each other's goals and incentives. Such settings are endemic to many negotiation contexts, from electronic commerce to diplomatic relations [3, 5]. Often, the lack of available information about the underlying interests of participants prevents parties from reaching beneficial agreements, or from reaching agreements altogether. Evidence of the effect of information exchange on human negotiation is inconclusive. While information exchange can lead to more equitable outcomes among the negotiation parties [1], it may also result in the exploitation of a vulnerable party [6].

For example, consider a hypothetical scenario in which friends in an online social network seek to agree on a restaurant at which to meet in the evening. While all of them share the common goal of enjoying their activity, they also have individual preferences that may conflict with each other. One of the friends may insist on eating at a preferred restaurant that is far away, making it difficult for them to reach agreement. This person can choose to reveal that he or she is gluten intolerant, and that the preferred restaurant is gluten free.

However, revealing one's interests may also cost; the friends may suggest a gluten-free restaurant that is nearer, but that the person may not like.

This paper presents a negotiation protocol that allows people to reveal their goals (and to request others to reveal their goals) at fixed points during a repeated negotiation process. The protocol was inspired by interest-based negotiation protocols designed for computational agents that allow participants to exchange information about their underlying objectives [10]. The protocol was evaluated empirically using a testbed which provides an analogue to the ways in which goals, tasks and resources interact in the real world. The testbed consists of a computer board game in which participants take turns proposing take-it-or-leave-it offers to each other under time constraints, communication was associated with a cost, and there was a large number of possible agreements. Twenty-two subjects participated in the study and played over 60 different games that varied environmental conditions such as the dependency relationships that held between participants and the resources at their possession at the onset of the interaction.

Results from analyzing people's behavior in this game show that goal revelation occurred in a minority (43%) of the games played. However, most players reveal their goals once asked to do so by the other participant. This revelation is shown to facilitate agreement, and to increase people's performance in the game as compared to an alternative negotiation protocol that does not allow revelation.

The data collected from people was used to train models for predicting people's goal revelation behavior in the game. These models were integrated into a decision theoretic paradigm used to decide whether to ask people to reveal their goals. Through simulation, we show that using this paradigm to make goal revelation requests can outperform people, when measured by the benefit from subsequent offers in the game.

This study is significant for the study of computer-mediated negotiation in two ways. First, it shows that computer systems can facilitate people's goal revelation decisions in negotiation, allow them to reach more beneficial agreements, and improve their overall performance. Second, it demonstrates the efficacy

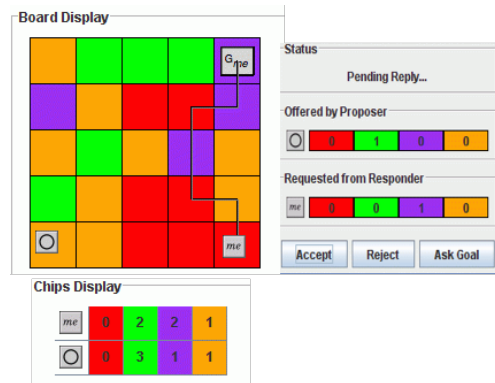
of using computational approaches to modeling people’s goal revelation behavior when negotiating under incomplete information.

## II. THE COLORED TRAILS GAME

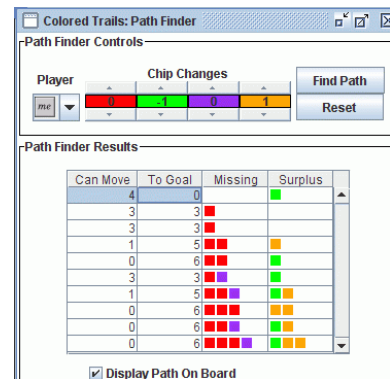
Colored Trails (CT) is a testbed developed for investigating the decision-making that arises in task settings, where the key interactions are among goals, tasks required to accomplish those goals, and resources needed to perform the tasks<sup>1</sup>. The empirical investigations described in this paper used a particular configuration of CT that is played by 2 players on a  $5 \times 5$  board of colored squares. Each player had a designated goal square and a piece on the board, initially located in one of the non-goal squares. At the onset of a CT game, players are issued a set of 7 colored chips chosen from the same palette as the squares. To move a piece into an adjacent square a player must turn in a chip of the same color as the square. Players had full view of the board and each others’ chips and positions, but they could only see their own goal location. CT provides a realistic analog to task settings, highlighting the interaction among goals, tasks required to achieve these goals, and resources needed for completing tasks. For example, chips correspond to agent capabilities and skills required to fulfill tasks, and different squares on the board represent different types of tasks.

A CT game comprises three phases. In the *communication phase*, players alternated between one of two roles: *proposer* players could offer some subset of their chips to be exchanged with some subset of the chips of responder players; *responder* players could in turn accept or reject proposers’ offers. If no offer was made, or if each offer was declined, then both players are left with their initial allocation of chips. The game controller prevents players from offering chips they do not have, or asking for chips the other does not have. In the *exchange phase*, chip exchanges were enforced by the game controller (if an agreement was reached). Finally, in the *movement phase*, the game controller automatically moved both players as close as possible to the goal square. The scoring function for each player depended solely on its individual performance: 100 points for reaching the goal; 10 points for each chip left in a player’s possession; 15 points deducted for any square in the shortest path between a player’s final position and goal-square (in case the goal was not reached). These parameters were chosen so that getting to the goal was by far the most important component, but if a player could not get to the goal it was preferable to get as close to the goal as possible. The score that each player received if no offer was made was identical to the score each player received if the offer was rejected by the responder.

Snapshots of the CT GUI for the interest-based protocol of one of the games used in the experiment are shown in Figure 1. The Main “Window” panel, shown in Figure 1a, includes the board game, the goal square, represented by an icon displaying the symbol  $G_{me}$ , and two icons, “me” and “O”, representing the location of the two players on the board at the onset of the



(a) Main and pending offer panels



(b) Decision-support tool

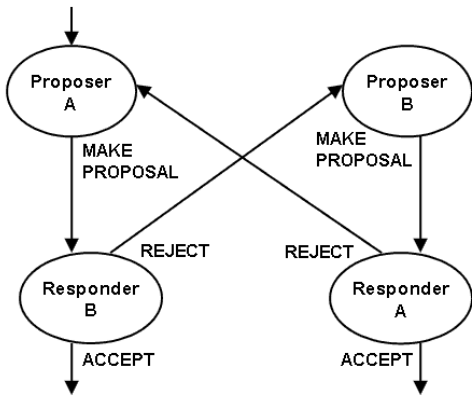
Fig. 1: Snapshot of a CT Game using an Interest-Based Negotiation protocol

game. The bottom part of the “Main Window” panel, titled “chips”, shows the chip distributions for the players. In the game shown here, the “me” player can get to the goal square, using the path that is outlined on the board, but the “O” player is lacking the chips to get to the goal (note that O’s goal is not shown here). The “me” player has received an offer asking it for 1 purple chip in return for 1 green chip. A proposer uses the “Propose Exchange” panel, to make an offer to a responder, or to ask for the other’s goal. The “Path Finder” panel, shown in Figure 1b, provides decision support tools to be used during the game. It displays a list of path suggestions to the goal, the missing chips required to fulfil a potential path, the surplus chips left over once a potential path has been fulfilled, and the best position the agent can reach relative to its scoring function. These paths are optimized for a given chip distribution and player, as queried by the player, such that they represent the best route given a player’s objectives. Players can view this information for the chip set that is currently in their possession, or for any hypothetical chip set.

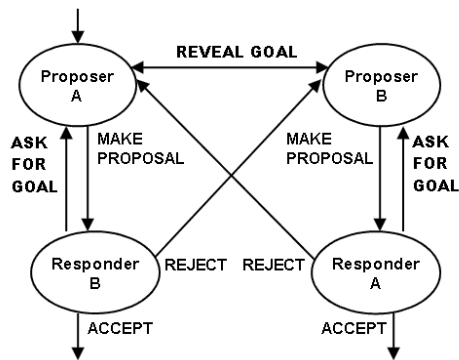
### A. Interest- and Position-Based Protocols in CT

We now describe the implementation of interest- and position-based protocols in CT. In both protocols, neither player can see the goal of the other at the onset of the game, and players are randomly allocated as proposers or responders. In the communication phase, a proposer can make an offer to the responder, as shown in Figure 1a.

<sup>1</sup>CT is Free Software and can be downloaded at <http://www.eecs.harvard.edu/ai/ct>



(a) Position-Based Negotiation protocol (Goal revelation disallowed)



(b) Interest-Based Negotiation protocol (Goal revelation allowed)

Fig. 2: Two alternative protocols for repeated negotiation

In the position-based protocol, once a responder receives an offer, it can accept it, in which case the offer is realized, both players automatically move towards the goal, and the game ends. If the responder rejects the offer, the game controller reverses the players' roles, and the new proposer player (formerly the responder) can make an offer to the new responder player (formerly the proposer). A state-based representation of this protocol is shown in Figure 2a.

The interest-based protocol is an extension of the position-based protocol that allows players, in a controlled fashion, to ask about, and reveal, their goals. Once a responder receives an offer from the proposer, it has the option to ask the proposer for its goal, in addition to rejecting or accepting the offer. If the responder chooses not to ask for the goal, the game proceeds as in the position-based negotiation case. If the responder chooses to ask the proposer for its goal, the proposer now has the option to agree to reveal its goal, or to make another offer to the responder, which is effectively a rejection of the revelation request. Responders may ask proposers for their goals numerous times, but once a goal is revealed, it cannot be asked about, or revealed, again. Goal revelations are always truthful. It is not possible to misreport one's goal in the game. A state-based representation of this protocol is presented in Figure 2b.




Fig. 3: Average benefit in IBN/PBN condition

### III. EMPIRICAL METHODOLOGY

We refer to the session involving the interest-based negotiation protocol as the IBN condition, and the session involving position-based negotiation protocol as the PBN condition. Twenty-two subjects participated in the experiment, drawn from a pool of students and adults residing in the Boston area. Twelve people participated in the IBN condition while ten people participated in the PBN condition. Each person was given an identical 30 minute tutorial on CT and played a series of games in succession. Subjects' scores were not revealed at any point during the experiment. Each subject was identified by a serial number, was seated in front of a terminal for the entire length of the experiment, and could not see or speak to any of the other participants.<sup>2</sup> No subject was paired up with any other subject more than once, and subjects were not told about the identity of their counterparts. Participants were paid in a manner consistent with their aggregate scores in all of the games they played. Between games, players engaged in a neutral activity which did not affect their payment (answering questions about their mood), designed to minimize the effects of past games on their future performance.

The games played were generated from a distribution to meet the following criteria: (1) At least one player could reach its goal, possibly independently, or by some exchange with the other player. (2) It was not possible for *both* players to reach their respective goals independently. This ensures that it is worthwhile for players to negotiate. For each game, we recorded the board and chip settings, as well as the actions of both players and their associated scores in the game.

### IV. RESULTS

We present an analysis of the *same* 65 games that were played in both conditions. In 14 of these games, players were co-dependent (both players needed each other to get to the goal), and in the other 51 games, one of the players needed the other player. Most IBN games did not feature a single goal revelation. However, when they were solicited (76% of the offers), players revealed their goals more often than not. There were 39 games in which one goal was revealed, and 10 games in which two goals were revealed, making for a total

<sup>2</sup>Approval was obtained from the Institutional Review Board (IRB) on the use of human subjects at Harvard University.

	Goal Revelations		All Games
	0	1, 2	
indep. player	<b>19</b> / 0	-9 / -15	<b>15</b> / -2
dep. player	49 / 59	<b>56</b> / <b>35</b>	40 / 50

TABLE I: Average benefit in IBN/PBN conditions for different player dependencies (significant difference in bold)

of 59 revelations. In all, at least one goal was revealed in 43% of the games.

Figure 3 shows the average benefit to participants in the IBN condition (left entry) and the PBN condition (right entry) when playing the same 65 games. The benefit to a player in a game is defined as the difference between the final score in the game and the no negotiation alternative score, computed using the scoring function described in Section II. If no agreement is reached, a player’s benefit is zero. The results are measured with respect to the games in which one or two goals were revealed in the IBN condition (marked “Goal Revelations”) and the total set of games played in both conditions (marked “All Games”).

As shown in the Figure, for those games in which there was at least one goal revelation, the combined average benefit for players in the IBN condition (49 points) was significantly larger than the average benefit for players in the PBN condition (32 points, paired t-test  $t(29) = 1.7, p = 0.04$ ). Not shown in the Figure is that there was no significant difference in players’ benefit in the two conditions for those games for which there was no revelation. However, when considering all games (including those for which there was no revelation), the combined average benefit for players in the IBN condition (72 points) was significantly larger than the benefit for players in the PBN condition (62 points, paired test  $t(64) = 1.60, p = 0.04$ ).

Next, we present a pairwise comparison between the number of games that reached agreement in both of the conditions. Our analysis reveals that 16 of the games that resulted in agreement in the IBN condition had failed to reach agreement in the PBN condition. In contrast, only 7 of the games that succeeded in the PBN condition failed to reach agreement in the IBN condition, and this difference was statistically significant ( $\chi^2(1, N = 65) = 3.92, p = 0.04$ ). We can thus conclude that goal revelation had a positive effect on the performance of players in the IBN condition, in that it led to higher scores and agreement ratios as compared to the corresponding games in the PBN condition.

#### A. The Effect of Players’ Dependency Relationships

Table I shows the benefit to players as a function of their dependency relationships and the number of goal revelations. As shown by the table, across both conditions the benefit for dependent players was consistently higher than the benefit for independent players. When one or two revelations occurred, the dependent player in the IBN condition gained significantly more benefit than in the PBN condition (56 points versus 35 points,  $SE = 2.3, t(26) = 2, p = 0.02$ ).<sup>3</sup> The overall benefit to independent players in the IBN condition was significantly

<sup>3</sup>The difference in benefit between the independent player in the IBN and PBN condition (-9 versus -15) was not statistically significant.

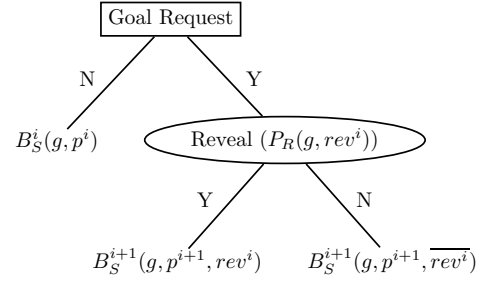


Fig. 4: Decision tree for goal revelation requests

higher than in the PBN condition. (15 versus -2 points,  $SE = 2.3, t(48) = 2.3, p = 0.01$ ). This is primarily due to their significant gain of 19 points over their PBN scores in IBN games where no goals were revealed.

#### B. A Decision Theoretic Paradigm for Goal Revelation Requests

To demonstrate the significance of our study for agent-designers, we used a decision theoretic approach to make goal request decisions in the game. The model integrates standard machine learning techniques for predicting and emulating people’s goal revelation behavior. We refer to participants who queried their partners about their goal as “goal solicitors,” and to participants that subsequently revealed their goal as “goal revealers”. Let  $g$  denote a CT game, and let  $NNAS(g)$  denote the no-negotiation alternative score for the solicitor  $S$ . Let  $PO_S^i(g, p^i)$  denote the score for the solicitor that is associated with a proposal  $p$  at round  $i$  in  $g$ . The benefit to the solicitor for this proposal is defined as  $B_S^i(g, p^i) := PO_S^i(g, p^i) - NNAS(g)$ .

Consider a solicitor that is reasoning about whether to ask the other participant to reveal its goal after receiving some offer  $p^i$  at round  $i$ . The outcome of this decision depends on whether the other participant agrees to reveal its goal. Figure 4 shows this reasoning process as a decision tree from the point of view of the solicitor agent. The leaves of the tree represent the expected benefit to the solicitor from offers in round  $i + 1$  in case the other participant revealed its goal in round  $i$  (denoted  $B_S^{i+1}(g, p^{i+1}, rev^i)$ ) or did not reveal its goal in round  $i$ , (denoted  $B_S^{i+1}(g, p^{i+1}, \overline{rev^i})$ ). If the solicitor decides not to ask the other participant to reveal its goal, it receives the benefit  $B_S^i(g, p^i)$  that is associated with the offer  $p^i$ .

The expected benefit to the solicitor from asking the other participant to reveal its goal is defined as

$$EU_S(g, ask^i) := (P_R(g, rev^i) \cdot B_S^{i+1}(g, p^{i+1}, rev^i) + P_R(g, \overline{rev^i}) \cdot B_S^{i+1}(g, p^{i+1}, \overline{rev^i})) \quad (1)$$

The expected benefit to the solicitor for not asking to receive the goal is just the benefit of the offer it is given in round  $i$ .

$$EU_S(g, \overline{ask^i}) := B_S^i(g, p^i, \overline{rev^i})$$

The optimal strategy for the solicitor is to make a goal revelation request at round  $i$  if  $EU_S(g, ask^i) > EU_S(g, \overline{ask^i})$

There are two challenges to using this decision tree to make a goal revelation request. First, the other player’s revelation decision in round  $i + 1$  is not known to the solicitor at round  $i$ . We addressed this by employing a Naive Bayes classifier to estimate the likelihood that the other participant will reveal its goal in game  $g$  at round  $i$ . This is represented by the probability  $P_R(g, rev^i)$ . For each game  $g$ , the features for this classifier represented the information that was available to the solicitor at round  $i$ . These features included  $NNA_S(g)$  (the no-negotiation alternative score for the solicitor),  $B_S(g, p^i)$  (the solicitor’s benefit from the proposal at round  $i$ ), and the round number  $i$ . The second challenge is that the proposals in round  $i + 1$  are not known to the solicitor at round  $i$ . We addressed this by estimating the benefit of these offers, analyzing data collected on the offers made during the games.

We computed the expected performance of the solicitor agent from using the tree to decide whether to ask the other participant for its goal. We limited this evaluation to making goal requests after receiving the first offer in the game. The performance was measured as the expected benefit from offers made or received by the solicitor, given the solicitor’s decision whether to request the other’s goal. The Naive Bayes classifier was used to estimate the probability  $P_R(g, rev^i)$ .

We compared the performance of the tree-using solicitor to that of people, by simulating human behavior in the game. To this end, we constructed Naive Bayes classifiers for emulating people’s goal request and goal revelation behavior. The features for the classifier for emulating people’s goal request behavior included  $NNA_S$  (the no-negotiating alternative score for the solicitor), and  $B_S^i(g, p^i)$  (the benefit to the solicitor that is associated with the offer at round  $i$ ). The features for the classifier for emulating people’s goal revelation behavior represented information that was available to the other participant (the “potential revealer”) at round  $i$ . These features included  $NNA_R$  (the no-negotiating alternative score for the potential revealer), and  $B_R^i(p^i, rev^i)$  (the benefit to the potential revealer associated with the offer at round  $i$ ). We measured people’s performance by computing the expected benefit from offers in the second round using an Expectimax tree [13]. This tree had identical structure to the decision tree of Figure 4, but the probability of asking and revealing goals were computed using the emulation models described above.

We evaluated the decision tree by using it to make goal revelation requests. For each of the games, we computed the expected utility to solicitors using the decision theoretic paradigm described above. We compared this expected utility with that incurred by people, using the emulation model to compute the likelihood that people actually reveal their goals. A computer agent using the decision theoretic paradigm would choose not to ask for goal revelation in a game if the likelihood of revelation was lower than 44%. We used ten-fold cross validation to learn the parameters for the classifiers; all of the classifiers achieved precision and recall measures above 70%, significantly better than random guessing. The average benefit to goal solicitors using the decision-theoretic model to make decisions was -5.04, which was significantly higher than the average benefit to people (-6.4, t-test,  $p = 0.03$ ). This shows that combining decision theoretic modeling with

standard machine learning techniques can form the basis of an agent-design for making decisions with people in our setting of choice.

## V. DISCUSSION AND RELATED WORK

The results shown in the last section establish the role of interest-based negotiation protocols as mechanisms of cooperation in settings of incomplete information. It also demonstrates the efficacy of using decision theoretic and standard machine learning techniques to computationally model people’s behavior in such settings. Dependent players are likely to agree to reveal their goals once asked, and this information is not abused by solicitors. Indeed, they choose to use this information as a tool for assisting the revealers, while succeeding not to incur a loss themselves, as compared to the no-negotiation alternative score. This results in a net gain to goal revealers, but also increases the social benefit of both participants. This mechanism was not overused by participants. Solicitors generally dislike to ask for others’ goals, and choose to do so mainly in cases where there are few avenues open for beneficial exchanges.

There are few empirical studies of human negotiation strategies in repeated interactions. Rubinstein [12] has provided a theoretical model for prescribing negotiating strategies that are optimal under certain conditions (e.g., participants are rational and consistent in their beliefs about each other’s objectives). Work in the psychological literature about strategic interaction has focused on specific domains (e.g., seller-buyer disputes [7] or completely abstract settings such as the Prisoners Dilemma. Loewenstein and Brett [8] conducted a study which studied how goal framing prior to the negotiation procedure affects strategy revision. None of these works have compared the effects of goal revelation directly within repeated negotiation. We show that when using a protocol for goal solicitation and revelation, negotiators are willing to disclose private information to others and that this allows them to reach mutually-beneficial agreements.

Work in automated negotiation in Artificial Intelligence (AI) has proposed algorithms for argumentative strategies which support or attack the different positions of parties in a negotiation [9] [4]. These algorithms have been used by computational agents and several works to study conditions under which such strategies outperform position-based protocols [2, 11]. This work directly extends these studies by showing that argumentative-type protocols are advantageous to people.

## REFERENCES

- [1] P. Carnevale, D. G. Pruitt, and S. D. Seilheimer. Looking and competing: Accountability and visual access in integrative bargaining. *Journal of Personality and Social Psychology*, 40(1):111 – 120, 1981. ISSN 0022-3514. doi: 10.1037/0022-3514.40.1.111.
- [2] A. Kakas and P. Moraitis. Adaptive agent negotiation via argumentation. In *Proc. 5th International Joint Conference on Autonomous Agents and Multi-agent Systems (AAMAS)*, 2006.

- [3] P. Kollock. The emergence of exchange structures: An experimental study of uncertainty, commitment, and trust. *American Journal of sociology*, pages 313–345, 1994. ISSN 0002-9602.
- [4] S. Kraus, K. Sycara, and A. Evenchik. Reaching agreements through argumentation: A logical model and implementation. *Artificial Intelligence*, 104:1–69, 1998.
- [5] S. Kraus, P. Hoz-Weiss, J. Wilkenfeld, D.R. Andersen, and A. Pate. Resolving crises through automated bilateral negotiations. *Artificial Intelligence*, 172(1):1–18, 2008.
- [6] G. Loewenstein and D. A. Moore. When ignorance is bliss: Information exchange and inefficiency in bargaining. *The Journal of Legal Studies*, 33(1):37–58, 2004.
- [7] G.F. Loewenstein, M.H. Bazerman, and L. Thompson. Social utility and decision making in interpersonal contexts. *Journal of Personality and Social Psychology*, (57)(3):426–441, 1989.
- [8] J. Loewenstein and J. Brett. Goal framing predicts strategy revision: When and why negotiators reach integrative agreements. In *Proceedings of the Cognitive Science Society*, 2007.
- [9] P. Pasquier, R. Hollands, I. Rahwan, F. Dignum, and L. Sonenberg. An empirical study of interest-based negotiation. *Autonomous Agents and Multi-Agent Systems*, 22(2):249–288, 2011.
- [10] I. Rahwan, S. D. Ramchurn, N. R. Jennings, P. McBurney, S. Parsons, and L. Sonenberg. Argumentation-based negotiation. *The Knowledge Engineering Review*, 18(4): 343–375, 2003.
- [11] I. Rahwan, P. Pasquier, S. Sonenberg, and F. Dignum. On the benefits of exploiting underlying goals in argumentation-based negotiation. In *Proceedings of the 22nd Conference on Artificial Intelligence*, 2007.
- [12] A. Rubinstein. A bargaining model with incomplete information about preferences. *Econometrica*, 53(5): 1151–1172, 1985.
- [13] S.J. Russell, P. Norvig, J.F. Canny, J.M. Malik, and D.D. Edwards. *Artificial intelligence: a modern approach*. Prentice hall Englewood Cliffs, NJ, 1995.