Human-computer Coalition Formation in Weighted Voting Games

MOSHE MASH, Carnegie Mellon University, USA
ROY FAIRSTEIN, Ben-Gurion University of the Negev, Israel
YORAM BACHRACH, Google DeepMind, England
KOBI GAL, Ben-Gurion University of the Negev and the University of Edinburgh, U.K.
YAIR ZICK, National University of Singapore, Singapore

This article proposes a negotiation game, based on the weighted voting paradigm in cooperative game theory, where agents need to form coalitions and agree on how to share the gains. Despite the prevalence of weighted voting in the real world, there has been little work studying people's behavior in such settings. This work addresses this gap by combining game-theoretic solution concepts with machine learning models for predicting human behavior in such domains. We present a five-player online version of a weighted voting game in which people negotiate to create coalitions. We provide an equilibrium analysis of this game and collect hundreds of instances of people's play in the game. We show that a machine learning model with features based on solution concepts from cooperative game theory (in particular, an extension of the Deegan-Packel Index) provide a good prediction of people's decisions to join coalitions in the game. We designed an agent that uses the prediction model to make offers to people in this game and was able to outperform other people in an extensive empirical study. These results demonstrate the benefit of incorporating concepts from cooperative game theory in the design of agents that interact with people in group decision-making settings.

CCS Concepts: • **Computing methodologies → Multi-agent systems**; **Intelligent agents**; **Cooperation and coordination**; • **Theory of computation → Solution concepts in game theory**;

Additional Key Words and Phrases: Negotiation and contract-based systems, cooperative game theory

ACM Reference format:
Moshe Mash, Roy Fairstein, Yoram Bachrach, Kobi Gal, and Yair Zick. 2020. Human-computer Coalition Formation in Weighted Voting Games. *ACM Trans. Intell. Syst. Technol.* 11, 6, Article 73 (October 2020), 20 pages.
https://doi.org/10.1145/3408294

1 INTRODUCTION

Weighted voting games are cooperative games in which agents can form binding coalitions, but differ in the amount of resources that they contribute to the coalition. They are inspired by

Authors' addresses: M. Mash, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213; email: mmash@cmu.edu; R. Fairstein and K. Gal, Ben-Gurion University of the Negev, 1 Ben-Gurion Avenu, Be'er Sheva, Israel; emails: Roi rfire01@gmail.com, kobig@bgu.ac.il; Y. Bachrach, Google DeepMind, 6 Pancras Square, London, UK; email: yorambac@google.com; Y. Zick, National University of Singapore, 21 Lower Kent Ridge Rd, Singapore 119077; email: zick@comp.nus.edu.sg.

2157-6904/2020/10-ART73 $15.00
https://doi.org/10.1145/3408294

real-world settings such as parliamentary government systems where the number of votes of each member state is proportional to the size of that state's population [13].

While an agent's ability to influence the outcome of a voting game is related to the amount of resources in its possession, it is not necessarily directly proportional to it. For example, consider a parliament with three parties, A, B, and C: A and B both have 50 seats, while C has 20. Suppose that a government must control a majority of the house (i.e., at least 60 votes). If one equates voting power with weight, then A and B are significantly more powerful than C. However, as a government can be formed by any two coalitions, no single party can form a government on its own; one might reasonably argue that all parties are equally powerful. Thus, in many settings, it makes sense to talk about parties' electoral *power*, rather than *weight,* and past work has formally quantified voting power under various assumptions (see Felsenthal and Machover [9] for an overview).

In contrast to the theoretical work on weighted voting games, little is known about how people behave in such settings and the effect of electoral power on their strategies. This article addresses this gap by studying how human participants form coalitions in weighted voting games and proposes a new agent design for playing the game with people. It introduces a configurable software platform that allows people to play variants of weighted voting games with other people or with computer agents. Participants negotiate proposals for division of revenue under different weight configurations, forming a coalition if they reach an agreement. We collected hundreds of instances of people's negotiation dynamics, the coalitions they formed, and the way revenue was shared.

We designed software agents that use measures from cooperative game theory, namely, the Shapley value [27], Banzhaf [2], and Deegan-Packel [7] indices, to predict how people respond to offers to join coalitions in the game. We show that some people tend to exhibit biases when making offers in the game as proposers. For instance, some proposers asked for individual shares that were mis-aligned with the voting power of the responders; some proposers tried to form coalitions that were too large, forcing a thin payoff spread among the members. In contrast, the agents we designed built coalitions with the minimal number of participants necessary to succeed and made offers that were aligned with people's voting power in the game. The agents using the Deegan-Packel and Shapley value measures were able to significantly outperform their human counterparts.

These results demonstrate a novel use of cooperative game theoretic concepts for revenue division systems, comprising both people and computers. In the spirit of public repositories in computational social choice [16, 30], we have created a public library that includes all of the collected data and made freely available to the research community at https://tinyurl.com/mrna7w6.

This article extends an initial study reporting on the formation of winning coalitions in mixed human-computer settings [15] in several ways. First, it provides an equilibrium analysis of the weighted voting game. Second, it compares people's play to those of different types of agents using game theoretic indices (Banzhaf, Shapley, and Deegan-Packel) as well as agents using the equilibrium strategy. Last, it extends the behavioral analysis to provide a more detailed description of people's behavior in the game.

2 RELATED WORK

This work is related to two different strands of research: coalitional bargaining and multi-agent negotiation in systems comprising people and agents.

There is an extensive body of theoretical work studying cooperative games, and weighted voting games in particular; the interested reader may see Chalkiadakis and Wooldridge [6] or Chalkiadakis et al. [5, Chapter 4] for an overview. However, empirical work studying human coalition formation is more sparse [24, 26]. Nash et al. [20] empirically study the correlation between payments obtained via human negotiation in lab experiments and solution concepts from cooperative game theory. They reveal that in over 60% of cases the resulting payments were similar to (or positively correlated with) the Shapley value of the corresponding players. Bachrach et al. [1] propose an asynchronous cooperative negotiation game (any player can make an offer at any time). They show that the average payoffs offered in this scenario correlate to the Shapley value. Nash et al. [20] conducted a laboratory experiment on finitely repeated three-person coalition formation games. In these games players with different strength according to the coalition payoffs could transfer power to another player, who then distributes the coalition payoffs. They find that fair divisions that equally divide the payoffs describe over 80% of all instances, rather than power-based divisions (e.g., Shapley) and stability-based divisions (e.g., the Core).

We extend these works by using a machine-learning model to predict which offers are going to be accepted and use these predictions to build a negotiating agent that performs well against humans. In this respect, we are inspired by other works combining machine-learning models together with game-theoretical solution concept for predicting human behavior [10, 12].

d'Eon et al. [8] explore how game-theoretic solution concepts reflect actual human reward sharing in cooperative games. Cooperative solution concepts are often axiomatized: There is a list of desiderata (axioms) that they uniquely satisfy. However, there is some empirical evidence that humans do not actually follow the theoretical rewards and rather utilize alternative methods, with a bias towards equal split of rewards no matter what player contributions are. The authors examine how AMT workers choose to share rewards amongst a set of fictional characters in a stylized setting. The analysis of d'Eon et al. differs from ours in several ways. First, the participants in d'Eon et al. do not actively play the game, but rather are analyzing a given game. Our participants are incentivized towards *strategic* rather than *fair* play, i.e., towards rewards that will maximize their own benefit. Second, our rewards are also higher, which makes our participants likelier to be invested in the work. Last, their work does not include an agent learning to play with humans.

Prior works empirically analyze the behavior of voters picking between alternatives [18, 31]. Van der Straeten et al. [31] study a setting where voters are positioned on a line and are incentivized to pick candidates that are as close to them as possible. Voters may misreport their position to obtain a better outcome, a phenomenon studied in this work. Our work differs on several accounts: voters are represented by weights, rather than by positions on a line. Moreover, in our experimental setup voters have a clear incentive to collaborate with one another to arrive at a mutually agreeable revenue division.

Several lines of work study the voting power of agents in weighted voting systems. Penrose [23] and Banzhaf [2] independently propose similar measures of *a priori* power (see Felsenthal and Machover [9] for an overview of the history of power indices). Penrose [23] suggested a way of determining player weights in weighted voting settings that ensure proportional representation: Assuming that each player is a representative of a population (e.g., in the case of the EU Council), one needs to assign a number of seats proportional to the square root of the country's population; this ensures that the voting power of each country is proportional to its true weight. Banzhaf [2] uses the power index that is his namesake to argue that the voting system in Nassau County was unfair.

Finally, our work is also related to previous works that analyze human behavior in negotiation systems [11, 22, 25] and strategic voting [3]. Oosterbeek et al. [21] study cultural differences in the ultimatum game and show that the behavior of the responder is influenced by the geographical regions of the participants. Oshrat et al. [22] propose an automated negotiator that uses opponent modeling to make offers; Rosenfeld and Kraus [25] show that human behaviors in negotiation domain can be predicted by using machine learning, and Haim et al. [11] study a market setting with three agents that use equilibrium strategies to negotiate with people. We extend these works to support human-agent interactions in weighted voting games. Bitan et al. [3] show that a computer agent using a best response strategy is able to outperform people in a repeated voting settings. Other works analyzing real-world voting systems include Leech [13], Słomczyński and Życzkowski [29], and Merrill [19].

3 WEIGHTED VOTING GAMES AND POWER INDICES

In *Weighted Voting Games (WVG)* a group of agents form coalitions to perform a task. Each agent has a certain amount of a resource; to achieve the task, a minimal amount of that resource is required. Any coalition whose members have a total weight exceeding the threshold is called *winning*, and is called *losing* otherwise.

A WVG is a three-tuple $(\vec{w}; t, r)$: We are given a set of agents $N = \{1, \ldots, n\}$, each agent $i \in N$ has a weight w_i. A coalition $S \subseteq N$ has a value of r if its total weight, $w(S) = \sum_{i \in S} w_i$, exceeds a given threshold t and has a value of 0 otherwise. Traditionally, WVGs are defined with the reward r set to 1, forming a subclass of *cooperative simple games.* We consider an arbitrary reward r and refer to the *value* of a coalition, $v(S)$, defined as

$$v(S) = \begin{cases} r & \text{if } w(S) \geq t \\ 0 & \text{otherwise.} \end{cases} \tag{1}$$

Power indices in WVGs capture the *influence* or *voting power* of agents. A power index is a function ϕ mapping weighted voting games to vectors in $\mathbb{R}^n$, where $\phi_i\langle \vec{w}; t, r\rangle$ should roughly correspond to i's ability to influence outcomes. To illustrate the application of these indices for weighted voting games, we use a simple WVG with three agents defined by the tuple $\langle (8, 2, 3); 10, 1\rangle$ Here there are three agents whose weights are $w_1 = 8, w_2 = 2, w_3 = 3$, a threshold of $t = 10$, and the reward is $r = 1$.

We introduce three well-known solution concepts (power indices)—Banzhaf, Shapley-Shubik, and Deegan-Packel indices—which are broadly related in studies that deal with domains of weighted voting game. The notion of marginal contribution is crucial for the definition of voting power. We say that i is *pivotal for a coalition* S if S is losing, but $S \cup \{i\}$ is winning. Let $m_i(S) = v(S \cup \{i\}) - v(S)$ be the *marginal contribution* of i to S. We say that i is pivotal for S if and only if the marginal contribution of i to S is r. The three methods determine the power of agent i by its average marginal contribution to randomly sampled coalitions; however, each approach considers different distributions over coalitions.

Under the *Banzhaf index* [2], the power of agent i is the expected marginal contribution of i for a coalition sampled uniformly at random from $N \setminus \{i\}$:

$$\beta_i(\vec{w}; t, r) = \mathbb{E}_{S \sim \mathcal{U}(N \setminus \{i\})}[m_i(S)] = \frac{r}{2^{n-1}} \sum_{S \subseteq N \setminus \{i\}} m_i(S). \tag{2}$$

In our example, the winning coalitions are $\{1, 2\}$, $\{1, 3\}$, and $\{1, 2, 3\}$. Agent 1 (whose weight is 8) is pivotal in all of these coalitions, agent 2 (with weight 2) is pivotal for $\{1, 2\}$, and agent 3 (with weight 3) is pivotal for $\{1, 3\}$. Thus, the Banzahaf index of the three agents is $(3/4, 1/4, 1/4)$.

The *Shapley-Shubik power index* [27] differs from the Banzhaf Index in that it measures the average marginal contribution of each agent to permutations (i.e., orderings of the agent set N) rather than coalitions. Given a permutation $\sigma : N \to N$, let $P_i(\sigma) = \{j \in N : \sigma(j) < \sigma(i)\}$ be the set of i's predecessors under σ; We define the marginal contribution of i to σ, denoted $m_i(\sigma)$, to be simply $m_i(P_i(\sigma))$: i's marginal contribution to its predecessors under σ. The Shapley value of agent i is the expected marginal contribution of i to a permutation chosen uniformly at random. Formally:

$$\varphi_i(\vec{w}; t, r) = \mathbb{E}_{\sigma \sim \mathcal{U}(\Pi(N))}[m_i(\sigma)] = \frac{r}{n!} \sum_{\sigma \in \Pi(N)} m_i(\sigma), \tag{3}$$

where $\Pi(N)$ denotes the set of all permutations of N. In our example, agent 1 is pivotal for the agent orderings $(2, 1, 3)$, $(2, 3, 1)$, $(3, 2, 1)$, and $(3, 1, 2)$; agent 2 is pivotal for the ordering $(1, 2, 3)$;

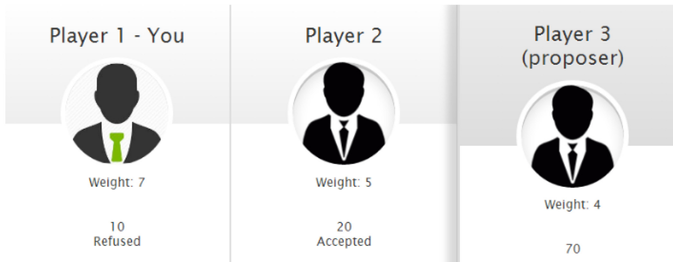

Fig. 1. Snapshot of the Cooperative Negotiation Game for three agents showing a proposal and associated responses.

and agent 3 is pivotal for the ordering $(1, 3, 2)$. Thus, the Shapley-Shubik power indices for our agents are $(2/3, 1/6, 1/6)$.

By assigning a positive probability to every coalition, both the Banzhaf and Shapley-Shubik power indices consider all coalitions that can form. The *Deegan-Packel index* [7], however, assumes that once a coalition is winning, it will not accept others. Deegan and Packel [7] measure power in the following manner: Whenever a minimal winning coalition forms, all of its members are equally powerful, and all minimal winning coalitions are equally likely to form. Formally, let $\mathcal{W}_{\min}(\vec{w}; t, r)$ be the set of all minimally winning coalitions in the WVG $\langle \vec{w}; t, r \rangle$ (we refer to $\mathcal{W}_{\min}(\vec{w}; t, r)$ as $\mathcal{W}_{\min}$ when $\langle \vec{w}; t, r \rangle$ is clear from context). Fixing an agent $i \in N$, we let $\mathcal{W}_{\min, i}(\vec{w}; t, r) = \{S \in \mathcal{W}_{\min}(\vec{w}; t, r) : i \in S\}$. The Deegan-Packel index is then

$$DP_i(\vec{w}; t, r) = \frac{r}{|\mathcal{W}_{\min}(\vec{w}; t, r)|} \sum_{S \in \mathcal{W}_{\min, i}(\vec{w}; t, r)} \frac{1}{|S|}. \tag{4}$$

In our example, the minimal winning coalitions are $\{1, 2\}$, and $\{1, 3\}$; thus, the Deegan-Packel indices are $(1/2, 1/4, 1/4)$.

4 METHODOLOGY

In this section, we describe our methodology for designing an agent for interacting with people in a WVG.

We designed an online variant of a WVG called the *Cooperative Negotiation Game*, in which players with different amounts of resources can cooperate and divide profits. The input to the game is a set of agents N, weights $\vec{w}$, threshold t, and reward r, which are common knowledge to all players. There are two phases in the game. In the *proposal phase*, a randomly chosen proposer p can suggest a coalition $S \subseteq N$. A coalition S is derived via a proposed payoff division $\vec{x} \in \mathbb{R}^n_+$ such that $supp(\vec{x}) = S$ and $\sum_{i \in S} x_i = r$ (i.e., S is the set of agents getting a positive payoff, and the total payoff is r). In addition, S must be winning and contain p.

In the *response phase*, every designated member of S can either accept or reject its offered share. If all agents in S accept their share, S forms, and its members receive their respective share. Otherwise, the coalition fails, and no agent receives any payoff. The proposal is visible only to the players who receive a positive share in the proposal, and these players can see each other's proposals and responses.

Figure 1 shows a snapshot of the proposal phase in one of the game configurations in the study. There are three players, with corresponding weights $\langle 7, 5, 4 \rangle$. The proposer p is Player 3; the threshold t is 10; and the reward r is set to 100. The snapshot is shown from the perspective of Player 1. The proposed is attempting to form the grand coalition with all three players, with corresponding shares $\langle 10, 20, 70\rangle$. We note that player 1 is a veto player, as it is a member of all winning coalitions in this game. In this example, Player 1 rejected the proposal, and the proposed coalition failed.

4.1 Extending the Deegan-Packel Index

We begin by providing an extension to the classic Deegan-Packel index (Equation (4)) and apply it to measuring power in the Cooperative Negotiation Game. The extended Deegan-Packel measure differs from the original index in two ways:

- It is defined with respect to a specific agent acting as a proposer and assumes that the proposer is always a member of the coalition it proposes.
- It shares the revenue of winning coalitions in proportion to agent weights (rather than assume that all members are equally powerful as in the original Deegan-Packel index).

Let $\mathcal{W}'_{\min,i}(\vec{w};t,r)$ be the set of all minimally winning coalitions, under the condition that agent i may not be excluded. Consider the following scenario: Player i is selected to be a proposer in the cooperative negotiation game; intuitively, it should form a coalition that is winning, but does not include extraneous players (the collation will still win without them). This leads us to the following definition:

$$\mathcal{W}'_{\min,i}(\vec{w};t,r) = \left\{ S \subseteq N : \begin{array}{c} w(S) \geq t, i \in S, \\ \forall S' \subset S : i \in S', w(S') < t \end{array} \right\}. \quad (5)$$

Note that $\mathcal{W}'_{\min,i}$ may not necessarily equal $\mathcal{W}_{\min,i}$ (the set of all minimally winning coalitions that contain i), nor does it necessarily contain any minimally winning coalition. To illustrate, consider the WVG $\langle(1,4,6);10,1\rangle$. In this case, $\mathcal{W}'_{\min,1}$ contains only $\{1,2,3\}$, but $\mathcal{W}_{\min,1} = \emptyset$, because $\{2,3\}$ is the unique minimally winning coalition.

We define $EDP_{i,p}(\vec{w};t,r)$ to be the extended Deegan Packel index of agent i given that p is the proposer. This is the expected revenue of i from a coalition S in $\mathcal{W}'_{\min,p}$ chosen uniformly at random, in which p allocates each member of S a share proportional to her weight.

Note that, strictly speaking, *EDP* is not a function from a WVG to a vector, thus it is not a power index: It uses additional information, namely, the identity of the proposer, which is common knowledge among the coalition members. We abuse notation and refer to it as a power index, because it provides a measure of influence of each coalition's member in the cooperative negotiation game:

$$EDP_{i,p}(\vec{w};t,r) = \frac{1}{|\mathcal{W}'_{\min,p}|} \sum_{C \in W'_{\min,p},\, i \in S} \frac{r \cdot w_i}{w(S)}. \quad (6)$$

In our example $\langle(8,2,3);10,1\rangle$, suppose that agent 2 is chosen to be the proposer. The only minimal winning coalition that contains agent 2 is $\{1,2\}$; thus, the Extended-Deegan-Packel power index for the agents is $(4/5, 1/5, 0)$. One can make $EDP_{i,p}$ into a power index by selecting the identity of p uniformly at random.

A player i is considered a *dummy* in a WVG $\langle\vec{w};t,r\rangle$, if $v(i \cup S) = v(S)$ for all coalitions, S. That is, the player does not add value to any coalition. Player i is considered a *veto* player in a WVG $\langle\vec{w};t,r\rangle$, if $v(N \setminus \{i\}) = 0$. That is, the player is a necessary participant of all possible coalitions. Note that $EDP_i = \frac{1}{n}\sum_{p \in N} EDP_{i,p}$ has some interesting properties. For example, unlike most power indices, EDP_i does not satisfy the dummy property: even dummy players have a strictly positive influence on the game.

4.2 Least-core Solution

We also considered a baseline model that used a power index that is based on the standard epsilon-core solution from the cooperative game theoretic literature [14]. This is the set of all coalitions

Table 1. Distribution of Player Weights in Game Configurations

W	1	2	3	4	5	6	7	8	9
P	0.31	0.27	0.13	0.11	0.08	0.03	0.03	0.02	0.01

that are immune to defection assuming players pay a penalty ε for leaving the coalition (e.g., none of the participants in a coalition in the epsilon-core can improve their performance by joining an alternative winning coalition and paying penalty ε):

$$C_\epsilon(\vec{w}; t, r) = \left\{ x \in R^N : \forall S \subseteq N, \left\{ \begin{array}{ll} \sum_{i \in N} x_i = r; \sum_{i \in S} x_i \geq r - \varepsilon, & \text{if } w(S) \geq t \\ \forall_{i \in S} x_i \geq 0, & \text{if } w(S) < t \end{array} \right\} \right. . \tag{7}$$

The least-core solution is the epsilon-core for the smallest ε value such that the epsilon-core of Equation (7) is non-empty. For example, consider the following voting game: $\langle(6, 4, 4, 3); 10, 1\rangle$. The least-core solution for this game is the payoff vector $\langle 1/3, 1/3, 1/3, 0\rangle$ with a minimal epsilon $\varepsilon \sim= 1/3$. According to this baseline approach, we assign the power index of the responder to its weight in an arbitrary coalition in the least-core solution.

When a solution exists with $\varepsilon = 0$, there exist coalitions in the strong-core, which is immune to defection of participants with no penalties.[1] For example, in the voting game $\langle(8, 2, 3); 10, 1\rangle$, the first player is a veto player who must exist in all possible coalitions. In this game, the strong core solution is the payoff vector $\langle 1, 0, 0\rangle$ that assigns all of the reward to the veto player. Note that all conditions of the least-core will hold for $\varepsilon = 0$, i.e., any coalition that does not pass the threshold necessarily do not include the veto player, and all successful coalitions necessarily contain the veto player.

4.3 Data Collection

We recruited 111 subjects (second-year software engineering undergraduates from Ben-Gurion University who have not taken a game theory course). Subjects were given a detailed tutorial of the game; participation in the study was contingent on passing a comprehension quiz. IRB approval was obtained from Ben-Gurion University for conducting this study. All subjects played a five-agent cooperative negotiation game with other people, in which the weight configuration satisfied the following constraints: Player weights varied between 1 and 9, and the sum of the weights for all agents was between 13 and 17, so as to create games with several possible successful coalitions. For all games, the threshold t was set to 10, and the coalition value r was set to 100. These constraints were imposed so there would be several possible coalitions in each game, the weight distribution over players would be sufficiently wide, and games would vary in the number of veto and dummy players. Table 1 shows the empirical weight distribution over players in the games, which were randomly generated. The weights of all players were common knowledge between participants.

At each round of the game, one of the participants was randomly chosen to be a proposer, while the other participants were responders. All the participants could observe the proposals (including the proposer's own share in the coalition), as well as the responses of all responders; however, the decisions of the players were displayed only after all the responders have replied to their proposals (to prevent influence of the decisions of the other responders in the current round, on the decisions of a particular responder).

When a coalition succeeded, the game ended, and a new game started with different participants and weight configurations; otherwise, a new round of the game commenced for the same

[1] A necessary condition for a strong-core solution is a pivot player that exists in all winning coalitions. This was the case in some of our game settings.

Table 2. Performance of the Logistic Regression Model When Using Different Power Indices to Predict Proposal Acceptance

Method	AUC	STD
EDP	**0.721**	0.088
Shapley-value	**0.719**	0.088
Banzhaf-index	0.717	0.090
Deegan-Packel	0.717	0.089
Least-Core	0.715	0.091
Always accept	0.5	0

participants, and a new proposer was chosen at random. The maximal number of rounds was set to three for all games (this was not conveyed to any of the participants). In all, we collected 180 games and 343 coalition proposals.

All subjects received the equivalent of an $8 show-up fee, as well as a bonus that depended on their performance in the game, computed as follows: For each successful coalition, participants received a payoff that was equal to their share in the coalition. At the end of the experiment, the total payoff for each participant was converted to a bonus payment. For example, a participant who received a total payoff of 322 points would receive a cash bonus of 3 dollars and 22 cents.

4.4 Predicting Human Responses

In this section, we describe a model for using power indices to predict how people respond to offers to join coalitions in the game. Before we describe the features used in the model, we make the following definitions: Let $\vec{x} \in \mathbb{R}_+^n$ be a vector of shares for all agents; that is, $\sum_{i \in N} x_i = r$. We use $(supp(\vec{x}), p)$ to refer to a proposed coalition between a proposer p and a set of responders $\{i \in supp(\vec{x}), i \neq p\}$. We assume that $x_p > 0$.

We wish to predict the probability of acceptance by responder i given the offer $\vec{x}$. We considered two types of features. The first type of features depended on the power index of the players:

- The power index of the responder i. $(\phi_i(\vec{w}; t, r))$.
- The power index of the proposer p. $(\phi_p(\vec{w}; t, r))$.

We varied the type of power index used (Banzhaf, Deegan-Packel, Shapley-Shubik, Extended Deegan-Packel), such that each power index generated a different set of features.

The second type of features depended on the weight and proposed share of players:

- The share of the proposer p in the coalition $\vec{x}$: (x_p).
- The share of the responder in the coalition $\vec{x}$: (x_i).
- The weight of the proposer (w_p).
- The weight of the responder (w_i).

We compared different machine learning models, including logistic regression, a multilayer neural network, a Decision Tree model (J48), and a Naive Bayes model, training each of the models separately. The best result was obtained with logistic regression for all possible sets of features that we considered.

Table 2 describes the AUC score (measured by 10-fold cross-validation) achieved by logistic regression for the different indices, as well as standard deviation measures over the 10 folds. We also include an "always accept" predictor. As shown in the table, all power indices were beneficial for predicting the acceptance of responders in the game, with the Extended Deegan-Packel and the Shapely-value indices achieving the best performance. The Least-Core solution achieved the lowest performance.

Other features that relate to the proposed shares of the other responders, their power indices, as well as proportional shares and proportional power indices between responders and between responders and the proposer, were examined as well but did not provide a significant contribution to the model's performance.

We also computed the correlations between the different power indices. The EDP and Shapley-value power indices exhibited the highest correlation (0.97), while the Least-Core and EDP exhibited the lowest correlation (0.63). These results are compatible with the performance of the logistic regression model. Interestingly, the correlation between the EDP and Deegan-Packel indices (0.83) were lower than the correlation between the EDP and any other index. A possible reason for this is that DP is the only index that considers minimal coalitions, while the EDP index requires players to form non-minimal coalitions.

4.5 Agent Strategies

In this section, we describe three types of agents that used the indices described in Sections 3 and 4.1 to play the Cooperative Negotiation Game. The agents combined decision-theoretic reasoning with the supervised learning model that was described in the previous section.

All agents used the predictive model to compute the acceptance probability $\Pr(Acc_i \mid \vec{x}, p)$ of a proposed share for a responder i. They differed in how they computed the features for the power index ($\phi_p(\vec{w}; t, r)$) and the ratio between the proportional share and the proportional power indices of the responder ($PS_{i,p}/PP_{i,p}$). Specifically, the *EDP agent* used the Extended Deegan-Packel Index ($EDP_{i,p}(\vec{w}; t, r)$) of Equation (6); the *Shapley agent* used the Shapley Index ($\varphi_i(\vec{w}; t, r)$) of Equation (3); and the *Banzhaf agent* used the Banzhaf Index ($\beta_i(\vec{w}; t, r)$) of Equation (2).

Given a payoff division $\vec{x}$ proposed by p, we write

$$\mathbb{E}[rev_p \mid \vec{x}, p] = x_p \cdot \prod_{i \in supp(\vec{x}),\, i \neq p} \Pr(Acc_i \mid \vec{x}, p) \tag{8}$$

to be the expected revenue received by p if they propose $\vec{x}$. Naturally, agents choose a payoff division $\vec{x}^*$ that maximizes the expected revenue:

$$\vec{x}^* \in \arg\max_{\vec{x}} \mathbb{E}[rev_p \mid \vec{x}, p]. \tag{9}$$

The following theorem states that a payoff division $\vec{x}^*$ that maximizes expected revenue is necessarily a minimally winning coalition that contains the proposer p.

Theorem 4.1. *Given a negotiation game $\langle \vec{w}; t, r\rangle$ let $\vec{x^*}$ be a payoff division, if $\vec{x}^* \in \arg\max_{\vec{x}} x_p \cdot \prod_{i \in supp(\vec{x}), i \neq p} \Pr(Acc_i \mid \vec{x}, p)$, then $supp(\vec{x^*}) \in \mathcal{W}'_{\min,p}(\vec{w}; t, r)$.*

Proof. We first make two assumptions on $\Pr(Acc_i \mid \vec{x}, p)$: First, if $x_i = 0$, then $\Pr(Acc_i \mid \vec{x}, p) = 0$ (agents will never accept a proposal that offers them no reward); second, $\Pr(Acc_i \mid \vec{x}, p)$ is strictly monotone increasing in x_i (agents are likelier to accept proposals that offer them a higher reward). Also note that $\Pr(Acc_i \mid \vec{x}, p)$ depends only on the payoffs to i and the proposer and the underlying WVG (vis-à-vis the power indices).

Given a negotiation game $\langle \vec{w}; t, r\rangle$ let $\vec{x}$ be a payoff division; we show that if $supp(\vec{x}) \notin \mathcal{W}'_{\min,p}(\vec{w}; t, r)$, then there exists another payoff division $\vec{x'}$ such that $\mathbb{E}[rev_p \mid \vec{x}, p] < \mathbb{E}[rev_p \mid \vec{x}', p]$.

If $supp(\vec{x})$ is a losing coalition, then we are done; thus, we assume that $supp(\vec{x})$ is winning. Since $supp(\vec{x})$ is not in $\mathcal{W}'_{\min,p}$, there is a player $j \neq p$, such that $supp(\vec{x}) \setminus \{j\}$ is a winning coalition as well. If $x_j = 0$, then j will not accept $\vec{x}$ and $\mathbb{E}[rev_p \mid \vec{x}, p] = 0$. Thus, we assume that $x_j > 0$. If $supp(\vec{x}) = \{p, j\}$, then p has veto power and forming a coalition alone is optimal; otherwise, we choose an arbitrary player k in $supp(\vec{x}) \setminus \{j, p\}$ and offer her $x_k + x_j$. Since the rewards to other players and the proposer are unchanged, their acceptance probabilities remain the same. Moreover, the acceptance probability for k is strictly increasing (as k's payoff increased). This implies that $\mathbb{E}[rev_p \mid \vec{x}, p] < \mathbb{E}[rev_p \mid \vec{x}', p]$ for every $\vec{x}$ such that $supp(\vec{x}) \notin \mathcal{W}'_{\min,p}$; in particular, if $\vec{x}^*$ maximizes expected revenue for p, then $supp(\vec{x}^*) \in \mathcal{W}'_{\min,p}$. □

4.6 Evaluation

We evaluated the agents by comparing their performance to that of people playing against other people. To this end, we recruited an additional 107 human subjects to play the cooperative negotiation game. In all, we collected 594 games including 681 proposals. All games included four humans and computer agents. As a proposer, the software agents implemented the proposal strategies of the EDP agent, Shapley agent, or Banzhaf agent according to Equation (9). People were told that they may be playing other people or computer agents, and the strategies used by the computer agents were not disclosed. In all other respects, the setting was the same as the one used for data collection in which people played other people (Section 4.3).

To compute this proposal, we note that the search space defined by Equation (9) includes all possible payoff divisions. To find an optimal payoff division, we iterate over all possible divisions in 5 unit increments (i.e., approximately 45K divisions in each configuration). The reason for the brute-force search approach is twofold: First, due to the relatively small number of payoffs, the brute-force search can be concluded in a short amount of time; second, over 95% of shares made by human proposers were percentage values that divided by five (e.g., 20%, 35%); a software agent making other types of proposals would easily stand out from its human counterparts.

As a responder, all of the software agents used a simple strategy: Accept all proposals offering it at least 5% of total revenue; i.e., those that it perceived as offering it a strictly positive utility. The EDP, Shapley, and Banzhaf agents played 281, 152, 161 games and were elected to be the proposer 84, 44, 57 times, respectively.

5 RESULTS AND DISCUSSION

We measure the performance of a player by the total revenue gained, averaged over all games in which the player was chosen to be a proposer. For a given game, the revenue was defined as follows: If no successful coalition was formed, then the revenue for all agents equaled zero. If the proposed coalition was successful, the revenue for all agents equaled their respective share in the proposal. All results reported in Section 5 are statistically significant in the $p < 0.05$ range using Mann-Whitney tests [17].[2]

Figure 2 shows the performance (measured by average revenue over all games) and the average requested shares of the agents and human proposers. As shown by the figure, both the EDP and Shapley agents outperformed people (Mann-Whitney test: $\alpha \approx 0$ and $\alpha = 0.033$), while there was no significant difference between the performance of the Banzhaf agent and people's performance.

Also shown in Figure 2 is that all the agents (EDP, Shapley, and Banzhaf) requested significantly greater shares for themselves than did people (Mann-Whitney test: $\alpha \approx 0$). The Banzhaf agent requests more for itself than the EDP agent (Mann-Whitney test: $\alpha = 0.026$).

Figure 3 presents the percentage of successful coalitions (when all the members accept their proposals) and the average acceptance rate of proposals. As seen in the figure, the number of

[2] We used the Mann-Whitney test, because it does not require the data to be normally distributed.

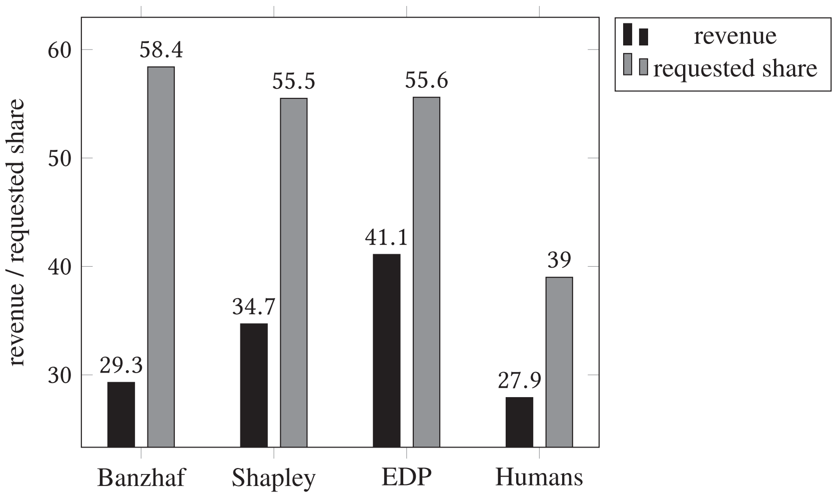

Fig. 2. Average revenue and shares requested by proposer.

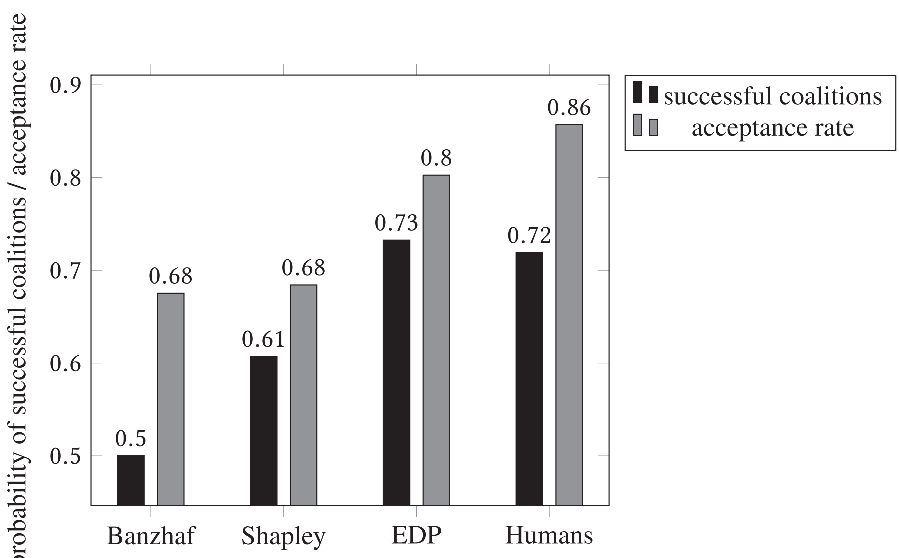

Fig. 3. Ratio of successful coalitions and the acceptance rates of proposals.

successful coalitions formed by people was higher than those formed by Shapley and Banzhaf agents (Mann-Whitney test: $\alpha \approx 0$). There was no statistically significant difference between the number of successful coalitions formed by human proposers and the EDP agent.

The acceptance rate of people playing with the EDP agent was higher than that of people playing with the Shapley and Banzhaf agents (Mann-Whitney test: $\alpha < 0.04$). There was no statistically significant difference between the acceptance rates of people playing the EDP agent and people playing other people.

To summarize, the EDP (and Shapley agents) achieved the highest performance, with the EDP agent having the best coalition success rate out of all the other agents. The Banzhaf agent exhibited

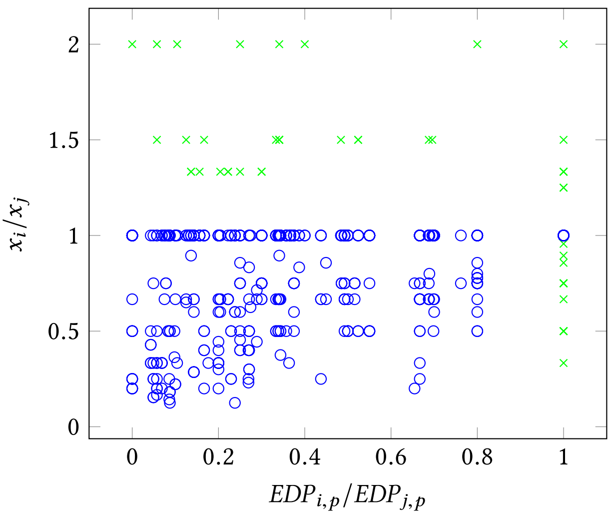

Fig. 4. Ratio between the index of any human responder pair (i, j) in a human proposed coalition (x axis) and the ratio between the shares proposed to (i, j) (y axis).

the lowest coalition success rate out of all of the other agents. We present an example from the data that illustrates the difference in behavior between human proposers and the EDP agent.

Example 5.1. Consider the weight configuration $\langle 4, 4, 3, 3, 3 \rangle$ when $p = 1$. The Extended Deegan-Packel power index for the agents is $(0.381, 0.181, 0.145, 0.145, 0.145)$. When the EDP-agent was elected to be the proposer it formed the coalition $supp(\vec{x}) = \{1, 3, 4\}$ with the shares $(50, 25, 25)$, respectively. This coalition proposal was always accepted by human responders. When humans formed the same coalition $\{1, 3, 4\}$, they always included agent 2 in the coalition. Since agent 2 is more powerful than agent 3 ($EDP_{2,1} > EDP_{3,1}$), it was offered a higher proposed share, at the expense of the proposer and agent 3.

We provide two possible explanations for the lower human performance that was exhibited in our experiments. First, human proposers make offers that do not align with the power index of responders in the game. To see this, consider the scatter-plot in Figure 4, which shows the ratio between the *EDP* power indices of responder pair (i, j) (x axis) and the ratio between the actual shares proposed to the responder pair (y axis) by proposer p.

For each coalition, any responder pair (i, j) contributed a single point to the scatter-plot, with the constraint that $EDP_{j,p} \geq EDP_{i,p}$. Thus, all points on the $x = 1$ line represent equal *EDP* power between responders i and j. For all points to the left of this line, $EDP_{j,p} > EDP_{i,p}$. Similarly, points on the $y = 1$ line represent equal shares proposed to i and j. Points above this line represent offers that propose more to responder i than to j. The offers marked in green are not "power preserving," in that $EDP_{i,p} < EDP_{j,p}$ but $x_i > x_j$, or $EDP_{i,p} = EDP_{j,p}$ but $x_i \neq x_j$.

Many of the human-proposed offers were non power preserving (41% of all offers), and most of them were declined. To illustrate, consider the following game collected from the data:

Example 5.2. Consider the game $\langle 6, 2, 2, 2, 1 \rangle$ in which $p = 5$. The extended Deegan-Packel power indices of the participants are $EDP_{i,5} = (0.58, 0.12, 0.12, 0.12, 0.09)$. The proposed coalition by one of the human participants was $\vec{x} = (15, 15, 20, 20, 30)$, which is not power preserving: We can see that $EDP_{1,5} \geq EDP_{i,5}$ for $2 \leq i \leq 4$ but the share for agent 1 is smaller or equal to the shares for agents 2, 3, and 4.

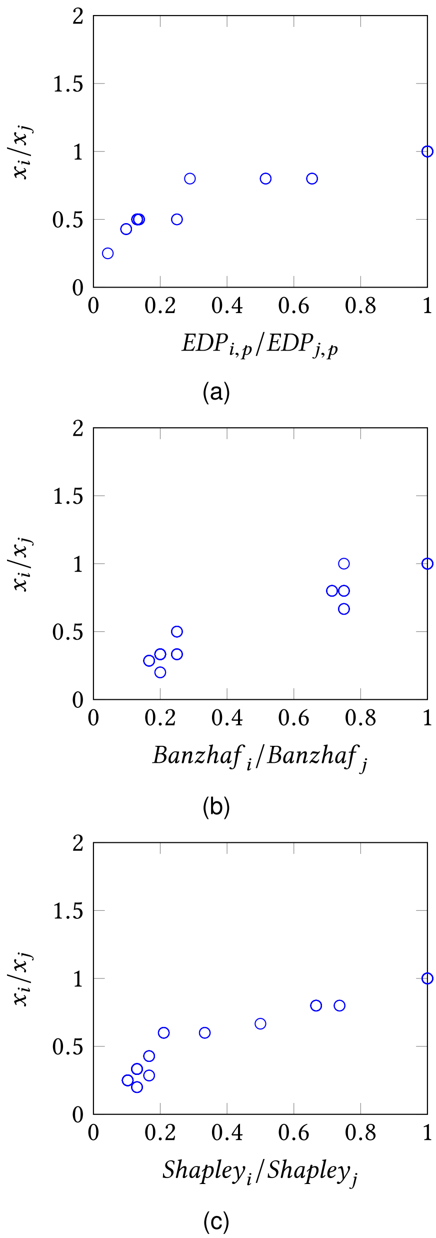

Fig. 5. Ratio between the different power indices of any responder pair $(i; j)$ in a proposed coalition of the different types of agents (x axis) and the ratio between the shares proposed to $(i; j)$ (y axis).

Figure 5 shows a scatter-plot of offers made by the different agents according to the same criteria. The power preserving ratio is computed for proposers made by the EDP agent (top), the Banzhaf agent (middle), and the Shapley agent (bottom). The terms $Banzhaf_i$, $Shapley_i$ in the figure refer, respectively, to the Banzhaf and Shapley-Shubik power index of the responder i. As shown by the figure, all of the offers made by the agents were power preserving. When the power of responders i and j were equal, the agents always gave them equal shares. As the power of j grows, it receives

Table 3. Size and Number of Formed Coalitions

coalition size	2	3	4	5
Humans	15.54%	49.56%	19.06%	15.83%
EDP	36.05%	51.16%	12.79%	0
Banzhaf	40.91%	43.18%	15.91%	0
Shapley	44.64%	41.07%	14.29%	0

Table 4. Performance of Human Proposers (Playing Other Humans)

Round	Requested Share	Accept Rate	Succ. Coalitions	Avg. Revenue
1	39.6	86%	73%	29.1
2	38.5	83%	69%	26.3
3	37.38	79%	57%	19.7
Avg	39.07	85%	71%	27.9

a higher share than does i. Interestingly, the Banzhaf power index ratios are clustered around low values (0.25) and high values (0.75). Thus, the Banzhaf agent perceived players as either "weak" or "strong," which limited the range of offers it made to responders, while the Shapley and EDP agents exhibited a wider range of values. This can be a possible reason for the poorer performance of the Banzhaf agent compared to the Shapley and EDP agents.

Another explanation for lower human performance is that 23% of the coalitions formed by people were non-minimal, i.e., the coalitions were not in $\mathcal{W}'_{\min,p}$. Larger coalitions are less likely to succeed than smaller coalitions, as coalitions require all members to agree to the proposals. Spreading the reward among more responders results in smaller shares on average, further decreasing the likelihood of acceptance.

Table 3 shows the number of coalitions the agents and humans formed by coalition size. As shown by the table, 19.06% of the coalitions formed by people included four members and 15.83% included all all five members (the grand coalition). In contrast, only 12.79%, 15.91%, and 14.29% of the coalitions formed by the EDP, Banzhaf, and Shapley agents included four members (respectively) and they never formed the grand coalition.

Table 4 analyzes the behavior of proposers for the cohort of human participants. It shows the average share requested by proposers, the acceptance rate of proposals, and the success rate of the resulting coalitions. As shown by the table, the acceptance rate of all proposals is generally high (on average, 85% of proposals were accepted). Also, as rounds progress, there is a slight decline in the average share requested by the proposer, and a significant decline in the acceptance rate and successful coalition rate decreases with the number of rounds. A possible reason for this behavior is that some responders expect offers to be more generous as rounds progress, and although there is a slight increase in offers made by proposers, it was not high enough to increase the acceptance rate.

Finally, we compare the performance of the different agents as responders. Here, we measured performance by the average share over all successful coalitions in which the responder was a member. The performance of the EDP, Banzhaf, and Shapley agents (30.2, 36.23, 30.02 average total shares, respectively) were significantly greater than that of human responders (27.9 average total share). We also found that 35% of people reject proposals offering them $\leq 20\%$ of the revenue, which is a similar result to that reported for canonical games in behavioral economics such as the ultimatum game [4].

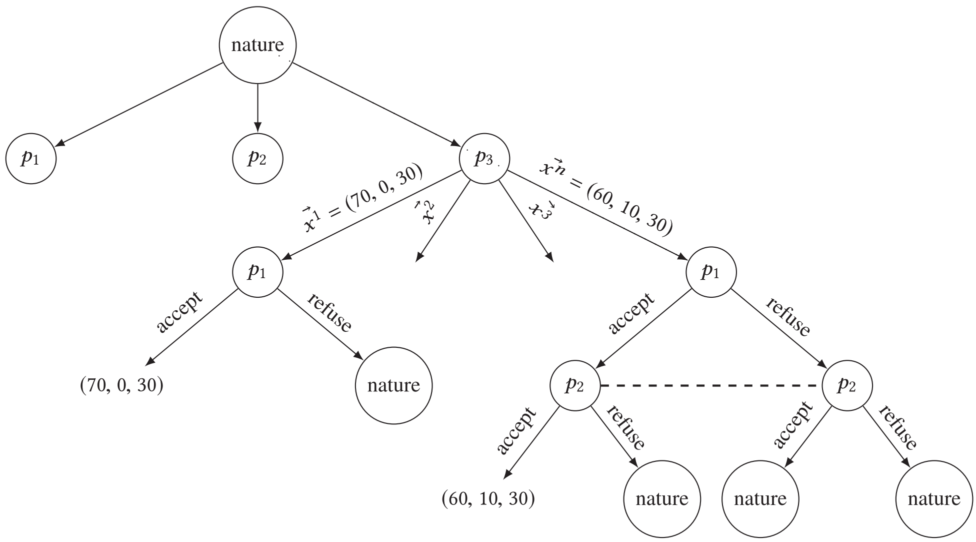

Fig. 6. Extensive form game of the cooperative negotiation game.

6 EQUILIBRIUM ANALYSIS

In this section, we provide an equilibrium analysis of the Cooperative Negotiation Game. This equilibrium was subsequently used by a computer agent to play people in the game.

The Cooperative Negotiation Game can be modeled as an extensive-form game. At each time-step t, nature makes a move (choosing a proposer), following which the proposer chooses a set of players to which it proposes a payoff division. Each of these players then decides whether to accept or reject the proposal. If any one of them rejects the proposal, the game proceeds to the next time-step. Figure 6 shows a partial tree in which the chosen proposer is p_3 offers the following shares $(60, 10, 30)$ in round 1. Both responders p_1 and p_2 have to decide whether to accept or refuse the coalition; if both of them accept it the payments are $(60, 10, 30)$; otherwise, another proposer will be elected at random.

Let n be the number of players and T denote the maximal number of rounds in the game. Suppose player i is the proposer at round t. Let $\mathcal{N}_i = \{S \subseteq N : i \in S\}$ be the set of all coalitions that contain player i. Let C_i^t be the set of winning coalitions (containing i) that minimize the total shares that i offers to the other players in round t. The acceptance reward threshold of responder k in round t is r_k^t. Player i will join any coalition that rewards it with at least r_i^t reward:

$$C_i^t = \arg\min_{S \in \mathcal{N}_i} \sum_{k \in S/i} r_k^t. \tag{10}$$

Proposer player i will propose any coalition in $C \in C_i^t$.

Now suppose player i is the responder in round t. The set of coalitions that can be formed by proposer player j in round t is C_j^t. The number of winning coalitions that include both players i and j is $|C_j^t \cap \mathcal{N}_i|$. The probability that i will be chosen by j to join a coalition is defined as $\frac{|C_j^t \cap \mathcal{N}_i|}{|C_j^t|}$.

The acceptance threshold r_i^t of player i in round t can now be defined as follows: At the last round T, the acceptance threshold is ϵ. For any other round $t < T$, this threshold depends on the expected value for the player in round $t + 1$. This value depends in turn on whether the player will

Table 5. SPNE Equilibrium of the *Cooperative Negotiation Game* for the Setting: $\langle 8, 2, 3; 10, 1\rangle$, $T = 4$

Player / Round	1	2	3	4
1	$\{0.727, 0.233, 0\}$, $\{0.727, 0, 0.233\}$	$\{0.723, 0.277, 0\}$, $\{0.723, 0, 0.277\}$	$\{0.666, 0.333, 0\}$, $\{0.666, 0, 0.333\}$	$\{1-\epsilon, \epsilon, 0\}$, $\{1-\epsilon, 0, \epsilon\}$
2	$\{0.537, 0.463, 0\}$	$\{0.444, 0.666, 0\}$	$\{0.333, 0.666, 0\}$	$\{\epsilon, 1-\epsilon, 0\}$
3	$\{0.537, 0, 0.463\}$	$\{0.444, 0, 0.666\}$	$\{0.333, 0, 0.666\}$	$\{\epsilon, 0, 1-\epsilon\}$

Player / Round	1	2	3	4
1	0.537	0.444	0.333	ϵ
2	0.231	0.277	0.333	ϵ
3	0.231	0.277	0.333	ϵ

Top table refers to proposer strategies, bottom table refers to responder strategies.

be chosen as a proposer (with probability $1/n$) or a responder in round $t + 1$:

$$r_i^t = \begin{cases} \epsilon & \text{if } t = T \\ \frac{1}{n}\left((r - p_i^{t+1}) + \sum_{j \neq i} \frac{|C_j^{t+1} \cap N_i| \cdot r_i^{t+1}}{|C_j^{t+1}|}\right) & \text{if } t < T \end{cases}. \tag{11}$$

Here, $p_i^{t+1} = \sum_{j \in S, j \neq i} r_j^{t+1}$ is the sum of shares offered by a proposer i to the responders in coalition S in round t, where $S \in C_i^{t+1}$ is any coalition in the set of possible coalitions made by the proposer. In particular, if $t = T - 1$ (second to last round), the acceptance threshold is r/n.

We now provide the following equilibrium definition for the cooperative negotiation game:

Proposition 6.1. *The subgame perfect Nash equilibrium of the WVG is as follows: If player i is a responder, then it will accept any proposal awarding it with positive utility in the final round T and accept any proposal with reward of at least r_i^t in round t when $1 \leq t < T$. If player i is a proposer, it will form a coalition C, where $C \in C_i^t$ and offer ϵ to all responders in the final round T, and r_i^t to each responder in round t when $1 \leq t < T$.*

To illustrate the equilibrium, consider the cooperative negotiation game example of Section 3, $\langle 8, 2, 3; 10, 1\rangle$. The strategies for proposer players are described in Table 5 (top), by listing the shares made to responders in each round. A share of 0 means that the corresponding agent is not included in the coalition. The strategies for responder players are described in Table 5 (bottom) by listing the expectance threshold for each player.

In round 4 (final round) the responders accept any positive proposal and the proposer takes $(1 - \epsilon)$ for itself while it offers ϵ to the coalition's members (e.g., player 1 offers ϵ to player 2 or player 3, and player 2 and 3 offer ϵ to player 1). The expected revenue of player 1 is $(r - \epsilon)/n + \epsilon \cdot ((n - 1)/n)$ and of players 2 and 3 is $(r - \epsilon)/n + \epsilon \cdot (1/2) \cdot (1/n)$ (player 1 has two options to form a coalition, so the probability that it will choose player 1 or 2 to the coalition is 1/2). We assume $\epsilon \approx 0$, so the expected revenues of all the players are r/n, which is $1/3 \approx 0.333$. In round 3, the players accept any proposal above (or equal) their expected revenues in round 4 (0.333). Therefore, the proposer offers 0.333 to each member of the coalition and it forms a minimal coalition (therefore, it pays a minimal value). For example, player 1 forms a coalition with player 1 or 2, pays 0.333 for each of them, and takes 0.666 for itself. Players 2 and 3 offer 0.333 to player 1 and take 0.666 for themselves. Therefore, the expected revenue of player 1 is $1/3 \cdot 0.666 + 2/3 \cdot 0.333 \approx 0.444$ in this round. When player 1 is elected to be the proposer it has two options to form a coalition while in each of them it needs to pay 0.333 to player 2 or 3, therefore the probability for each coalition is 1/2 and the expected revenues of players 2 and 3 are $1/3 \cdot 0.666 + 1/3 \cdot 1/2 \cdot 0.333 \approx 0.277$ in this round. Rounds 2 and 1 are calculated in the same manner. Furthermore, we expect that the proposer in the first round will form the coalition of the equilibrium and the responders will accept it, so the payments will be determined according to this round. For example, if player 2 is elected to be the proposer in the first round it offers 0.531 for player 1 (and takes 0.469 for herself) and player 1 accepts it. We now present a proof for the SPNE specified in Proposition 6.1.

Proof. We now prove Proposition 6.1. The subgame perfect Nash equilibrium of the game is obtained using backward induction. The proof relies on the finite horizon nature of the game.[3] In the last round, whether the coalition is not successful the payment of all the players is zero, so the responders accept any proposal that they received (they are indifferent between acceptance and refusal of a null proposal); thus, the coalition succeeds in any case, so the proposer offers zero to all members (it takes r for itself). In this round, the proposer forms any winning coalition, because its share is still r (it offers null proposals to all members). Hence, given a uniform probability distribution of the proposer role in round T, the expected revenue of a player i in this round is r/n.

We assume that at round $t+1$, all players use the strategy specified in Proposition 6.1. We prove that the proposition holds at round t.

Suppose player i is a responder at round t. Any offer providing less than r_i^t should be rejected, because the player can guarantee the expected value in the next round as specified in Equation (11). Specifically, if the player is a proposer at round $t+1$ (with probability $1/n$), it will receive a reward of $r - p_i^{t+1}$) where $p_i^{t+1} = \sum_{j \in S, j \neq i} r_j^{t+1}$. By the induction assumption, any responder j will accept the share of r_j^{t+1} in round $t+1$. If it is a responder at round $t+1$ with probability $n - 1/n$, it will be chosen for a coalition with probability $\frac{|C_j^{t+1} \cap \mathcal{N}_i|}{|C_j^{t+1}|}$ and will receive a reward of r_i^{t+1}. We can thus conclude that r_i^t is the acceptance threshold for the responder at round t.

Suppose player i is the proposer at round t. Any proposal less than r_i^t will be refused by the responders, therefore player i should offer to any coalition member j at least r_j^t and a rational player will offer them excatly this value. Specifically, it will propose a coalition such that p_i^{t+1} is minimal (a coalition in the set C_i^t). □

6.1 Using Equilibrium Strategy When Playing with Humans

In this section, we directly evaluate the equilibrium strategy with that of people playing the game.

To this end, we recruited 20 subjects (fourth-year software engineering undergraduates from Ben-Gurion University) to play the Cooperative Negotiation Game. The configurations of the game are the same as of those used in Section 4.3 with one exception: People were told that there were three rounds in the game, to provide them with the same information as for the equilibrium agent.

Participants were divided into two pools. The first pool included 11 people playing 5-player cooperative negotiation games with other people. In total, we collected 74 games and 105 coalition proposals from this pool. The second pool included 8 people playing 5-player games consisting of 4 people and an SPNE agent. In total, we collected 75 games and 106 coalition proposals; of these, 28 proposals were made by the agent in this pool.

The equilibrium agent played according to the subgame perfect Nash equilibrium of Proposition 6.1. In the last round, the agent will accept any proposal above (or equal) a predetermined epsilon value of 5 cents (to align with the type of minimal offers made by people). All the results were evaluated using Mann-Whitney test with p-value < 0.05.

Tables 6 and 7 compare the performance of the SPNE agent and human proposers in each round as well as on all the rounds together. As can be seen in the tables, the average share that the agent

[3]For additional information about backward induction and Sub-game Perfect Nash Equilibrium, see Shoham et al. [28].

Table 6. Performance of Humans as Proposers

Round	Requested Share	Acceptance Rate	Successful Coalitions	Average Revenue
1	32	81%	71%	22.48
2	32.66	69%	52%	19.57
3	45.5	81%	70%	25.5
Avg	32.46	78%	67%	22.19

Table 7. Performance of the SPNE Agent as Proposer

Round	Average Request	Acceptance Rate	Successful Coalitions	Average Revenue
1	48.26	57%	31%	15.78
2	58	10%	0%	0
3	90	12%	0%	0
Avg	55.96	42%	21%	10.71

requested for itself as the proposer (55.96) was significantly higher than the requests of human proposers (32.46). Furthermore, The average agent request is higher than the human average request in each round. The failure of the agent is particularly apparent in rounds 2 and 3 where all the agent-formed coalitions failed. This result, in turn, can explain why humans were more willing to accept other humans' proposals (78%) than agent proposals (42%). The average ratio of successful coalitions formed by humans (67%) was also significantly higher than the success rate obtained from SPNE agent proposals (21%). We present an example from the data that illustrate the proposals offered by the SPNE agent and people's responses.

Example 6.2. Consider the weight configuration $\langle 5, 5, 3, 2, 2 \rangle$ when $p = 3$ and $t = 2$ (second round). When the SPNE-agent was elected to be the proposer it formed the coalition $supp(\vec{x}) = \{2, 3, 5\}$ with the shares $(22, 59, 19)$, respectively (the agent request 59 for itself). Both the responders (agent 2 and agent 4) refused the proposal.

Last, we compare between the performance of human proposers who played the equilibrium agent (Table 6) with human proposers who played other people (Table 4). We see that human proposers playing other humans generated more revenue than those playing the equilibrium agent. A possible reason for this is that the SPNE agent proposed significantly lower shares to responders than did human proposers. These offers were more likely to be rejected by human responders than those made by other people. Interestingly, for the final round, human proposers offered significantly less shares to responders than in previous rounds. This trend was not observed for the case in which human proposers played other people, where the number of rounds was not known to particpants. This shows that human proposers may have been affected by backward induction type reasoning in their play.

7 CONCLUSIONS

In this article, we have studied the problem of coalition formation in settings that include both human and computer players. The performance of the EDP and Shapley agents makes a compelling argument for the combination of game-theoretic and machine learning based agents in coalitional bargaining domains. We showed that including concepts from cooperative game theory was able to significantly increase the predictive power of our models and to inform the design of successful agents for playing the game. We also showed that machine learning is a necessary condition for successful agent performance in weighted voting games, as people play quite differently than what is predicted by Nash equilibrium. The best-performing EDP proposer agent was able to increase its own performance when compared to other agents and human proposers, without hurting people's performance.

Our results can inform the design of future voting systems in which people and computers interact, by (1) creating agents that serve as proxies for people in future voting systems, or as training tools for people to improve their bargaining skills in voting settings; (2) modeling how people vote in computerized environments; (3) using these models to inform the design of improved voting systems that lead voters to better outcomes (whether for individuals or society). We are currently extending the model to include repeated settings in which participants interact over time and need to consider the effects of reciprocity on their voting strategies.

REFERENCES

[1] Y. Bachrach, P. Kohli, and T. Graepel. 2011. Ripoff: Playing the cooperative negotiation game. In *Proceedings of the 10th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS'11)*. 1179–1180.
[2] J. F. Banzhaf. 1964. Weighted voting doesn't work: A mathematical analysis. *Rutgers Law Rev.* 19 (1964), 317–343.
[3] M. Bitan, Y. Gal, S. Kraus, E. Dokow, and A. Azaria. 2013. Social rankings in human-computer committees. In *Proceedings of the 27th AAAI Conference on Artificial Intelligence (AAAI'13)*. 116–122.
[4] C. F. Camerer. 2003. *Behavioral Game Theory: Experiments in Strategic Interaction.* Princeton University Press.
[5] G. Chalkiadakis, E. Elkind, and M. Wooldridge. 2011. *Computational Aspects of Cooperative Game Theory*. Morgan-Claypool.
[6] G. Chalkiadakis and M. Wooldridge. 2016. Weighted voting games. In *Handbook of Computational Social Choice*, F. Brandt, V. Conitzer, U. Endriss, A. D. Procaccia, and J. Lang (Eds.). Cambridge University Press, UK, Chapter 16.
[7] J. Deegan and E. W. Packel. 1978. A new index of power for simple n-person games. *Int. J. Game Theor.* 7, 2 (1978), 113–123.
[8] Greg d'Eon, Kate Larson, and Edith Law. 2019. The effects of single-player coalitions on reward divisions in cooperative games. In *Proceedings of the 1st Workshop on Games, Agents and Incentives (GAIW) held at International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS'19)*.
[9] Dan S. Felsenthal and M. Machover. 2005. Voting power measurement: A story of misreinvention. *Soc. Choice Welf.* 25, 2 (2005), 485–506.
[10] S. Ganzfried and T. Sandholm. 2011. Game theory-based opponent modeling in large imperfect-information games. In *Proceedings of the 10th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS'11)*. 533–540.
[11] G. Haim, Y. Gal, B. Ann, and S. Kraus. 2017. Human-computer negotiation in a three player market setting. *Artif. Intell.* 246 (2017), 34–52.
[12] J. S. Hartford, J. R. Wright, and K. Leyton-Brown. 2016. Deep learning for predicting human strategic behavior. In *Proceedings of the 30th Annual Conference on Neural Information Processing Systems (NIPS'16)*. 2424–2432.
[13] D. Leech. 2002. Designing the voting system for the Council of the European Union. *Pub. Choice* 113, 3–4 (2002), 437–464.
[14] Michael Maschler, Bezalel Peleg, and Lloyd S. Shapley. 1979. Geometric properties of the kernel, nucleolus, and related solution concepts. *Math. Op. Res.* 4, 4 (1979), 303–338.
[15] M. Mash, Y. Bachrach, Y. Gal, and Y. Zick. 2017. How to form winning coalitions in mixed human-computer settings. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI'17)*. 465–471.
[16] N. Mattei and T. Walsh. 2013. PrefLib: A library of preference data. In *Proceedings of the 3rd International Conference on Algorithmic Decision Theory (ADT'13)*. 259–270.
[17] Patrick E. McKnight and Julius Najab. 2010. Mann-Whitney U test. In *The Corsini Encyclopedia of Psychology.* Wiley.
[18] Timo Mennle, Michael Weiss, Basil Philipp, and Sven Seuken. 2015. The power of local manipulation strategies in assignment mechanisms. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence.*
[19] S. Merrill. 1978. Citizen voting power under the electoral college: A stochastic model based on state voting patterns. *SIAM J. Appl. Math.* 34, 2 (1978), 376–390.
[20] J. F. Nash, R. Nagel, A. Ockenfels, and R. Selten. 2012. The agencies method for coalition formation in experimental games. *Proc. Nat. Acad. Sci. United States Amer.* 109, 50 (2012), 20358–20363.
[21] H. Oosterbeek, R. Sloof, and K. G. Van De Kuilen. 2004. Cultural differences in ultimatum game experiments: Evidence from a meta-analysis. *Exper. Econ.* 7, 2 (2004), 171–188.

[22] Y. Oshrat, R. Lin, and S. Kraus. 2009. Facing the challenge of human-agent negotiations via effective general opponent modeling. In *Proceedings of the 8th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS'09)*. 377–384.
[23] L. S. Penrose. 1946. The elementary statistics of majority voting. *J. Roy. Statist. Soc.* 109, 1 (1946), 53–57.
[24] A. Rapoport, J. P. Kahan, S. G. Funk, and A. D. Horowitz. 2012. *Coalition Formation by Sophisticated Players.* Springer-Verlag.
[25] A. Rosenfeld and S. Kraus. 2016. Providing arguments in discussions on the basis of the prediction of human argumentative behavior. *Trans. Interact. Intell. Syst.* 6, 4 (2016), 30:1–30:33.
[26] H. Sauermann. 1978. *Coalition Forming Behavior.* Vol. 8. Mohr Siebrek Ek.
[27] L. S. Shapley and M. Shubik. 1954. A method for evaluating the distribution of power in a committee system. *Amer. Polit. Sci. Rev.* 48, 3 (1954), 787–792.
[28] Yoav Shoham, Kevin Leyton-Brown, et al. 2009. Multiagent systems. *Algorithmic, Game-Theoretic, and Logical Foundations* (2009).
[29] W. Słomczyński and K. Życzkowski. 2006. Penrose voting system and optimal quota. *Acta Phys. Polon. B* 37, 11 (2006), 3133–3143.
[30] M. Tal, R. Meir, and Y. Gal. 2015. A study of human behavior in online voting. In *Proceedings of the 14th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS'15)*. 665–673.
[31] Karine Van der Straeten, Jean-François Laslier, Nicolas Sauger, and André Blais. 2010. Strategic, sincere, and heuristic voting under four election rules: An experimental study. *Soc. Choice Welf.* 35, 3 (2010), 435–472.

Received March 2019; revised May 2020; accepted June 2020