

How to Form Winning Coalitions in Mixed Human-Computer Settings

Moshe Mash

Ben-Gurion University
mashm@post.bgu.ac.il

Ya'akov (Kobi) Gal

Ben-Gurion University
kobig@bgu.ac.il

Yoram Bachrach

DigitalGenius
yorambac@gmail.com

Yair Zick

National University of Singapore
zick@comp.nus.edu.sg

Abstract

This paper proposes a new 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. We show that 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 an online version of a weighted voting game. We design an agent that combines supervised learning with decision theory to make offers to people in this game. We show that the agent was able to obtain higher shares from coalitions than did people playing other people, without reducing the acceptance rate of its offers. We also find that people display certain biases in weighted voting settings, like creating unnecessarily large coalitions, and not rewarding strong players. These results demonstrate the benefit of incorporating concepts from cooperative game theory in the design of agents that interact with other people in weighted voting systems.

1 Introduction

Weighted voting games are types of cooperative games in which agents can form binding coalitions, but differ in the amount of resources that they contribute to the coalition. A simple example of such settings is the parliamentary government system used in many countries, and the EU council, where the number of votes of each member state is proportional to the size of that state's population [Leech, 2002].

While an agent's ability to influence the outcome of the game is related to its amount of resources, 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; thus, 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*. Many researchers tried to formally quantify voting power, under various assumptions (see Felsenthal and Machover [2005] for an overview). Despite their widespread study, little is known about how people actually make decisions in weighted voting settings. This paper addresses this gap by introducing a configurable software platform that allows users to play variants of weighted voting games with other people or with computer agents. In our setting, participants negotiate revenue division proposals under different weight configurations, eventually forming a coalition if they reach an agreement. Using this platform, we collected hundreds of instances of users' negotiation dynamics, the coalitions they formed, and the way revenue was shared.

We designed a negotiating software agent and tested its performance when interacting with other people playing this game. Our agent uses influence measures from the cooperative game theory literature [Banzhaf, 1965; Shapley and Shubik, 1954; Deegan and Packel, 1978] to predict how people respond to offers to join a coalition in the game. Our results show that the agent significantly outperforms its human counterparts. It can retain a relatively high revenue, without incurring a drop in acceptance rates. These results can be explained by the agent's ability to predict which coalitions would be accepted, but also by the fact that people tend to exhibit biases. For instance, some human proposers took low amount for themselves, just so that the coalition would form, and some tried to form coalitions that were too large, forcing a thin payoff spread among too many members.

Our agent uses machine learning and game theory to decide on proposals that maximize its expected revenue. These results demonstrate a novel use of cooperative game theoretic concepts for revenue division systems, comprising of both people and computers. In the spirit of public repositories in computational social choice [Mattei and Walsh, 2013; Tal *et al.*, 2015], we are making our platform open source, and have created a public library which will include all of the collected data, and made freely available to the research community at <https://tinyurl.com/mrna7w6>.

2 Related Work

There exists an extensive body of work on weighted voting games, and their applications, such as predicting nego-

tiation outcomes, pricing cloud services or crypto-currencies and evaluating contribution in crowdsourcing settings [Blocq *et al.*, 2014; Lewenberg *et al.*, 2015; Bachrach *et al.*, 2012]; Chalkiadakis and Wooldridge [2011] and Chalkiadakis *et al.* provide an overview of such applications [2016]. Most of these works handle computational and mathematical challenges raised by weighted voting, e.g. computing influence measures [Bachrach *et al.*, 2010; Elkind *et al.*, 2007]; our work, on the other hand, takes an empirical approach, analyzing human actors and their decision making in the weighted voting setting.

The round-based negotiation implemented by our work relates to work on coalitional bargaining. Some works in this realm focus on bargaining dynamics and the solutions they converge to [Ausubel *et al.*, 2002; Zick *et al.*, 2015; Maschler *et al.*, 2013; See *et al.*, 2014, chapter 15], while others focus on computational aspects of coalition formation (see overview by Rahwan *et al.* [2015]). However, as is the case for weighted voting games, empirical work studying human coalition formation is relatively sparse. One exception is the work by Bachrach *et al.* [2011] who propose an asynchronous cooperative negotiation game (any player can make an offer at any time); Bachrach *et al.* show that payoffs correlate to the Shapley value averaged across many games. In contrast, we predict which offers are going to be accepted, and use these predictions to build a negotiating agent that performs well against humans.

Lastly, previous works have analyzed human behavior in ultimatum games [Güth, 1995; Oosterbeek *et al.*, 2004; Azoulay *et al.*, 2014], bilateral negotiation [Oshrat *et al.*, 2009; Rosenfeld and Kraus, 2016; Haim *et al.*, 2017] and strategic voting [Bitan *et al.*, 2013]. The weighted voting setting (and the cooperative negotiation game it induces) offers a more complex interaction space.

3 Weighted Voting Games

Our work studies *Weighted Voting Games (WVG)* which reflect situations in which each agent has a certain amount of a resource; in order to achieve a task (e.g. pass a bill, generate revenue), 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. More formally, a WVG is a tuple $\langle \vec{w}; t, r \rangle$: we are given a set of agents $N = \{1, \dots, n\}$, 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*. Our formalism allows an arbitrary reward r . We often refer to the *value* of a coalition, $v(S)$, defined as

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

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 $\phi_i(\vec{w}, t, r)$ 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 3 agents defined by the tuple $\langle 8, 2, 3; 10, 1 \rangle$ (i.e. the threshold is $t = 10$, and the reward is $r = 1$).

3.1 The Shapley-Shubik Power Index and the Banzhaf Index

Given a coalition S , we say that i is *pivotal* for S if S is losing, but $S \cup \{i\}$ is winning. Formally, i is pivotal for S iff the *marginal contribution* of i to S , defined as $m_i(S) = v(S \cup \{i\}) - v(S)$, is r . The *Banzhaf index* [Banzhaf, 1965] of agent i is the expected marginal contribution of i for a coalition sampled uniformly at random from $N \setminus \{i\}$. Formally:

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

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 Banzhaf index of the three agents is $(3/4, 1/4, 1/4)$.

The Shapley-Shubik power index [Shapley and Shubik, 1954] differs from the Banzhaf Index in that it measures the average marginal contribution of each agent to permutations on the set of coalitions (i.e. orderings of the agent set N). Given a permutation $\sigma : N \rightarrow N$, let $P_i(\sigma) = \{p \in N : \sigma(j) < \sigma(i)\}$ 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:

$$\varphi_i(\vec{w}; t) = \mathbb{E}_{\sigma \sim \mathcal{U}(\Pi(N))} [m_i(\sigma)] = \frac{1}{n!} \sum_{\sigma \in \Pi(N)} m_i(\sigma) \quad (1)$$

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

3.2 The Deegan-Packel Index

By assigning a positive probability to every coalition, both the Banzhaf and Shapley-Shubik power indices implicitly assume that all coalitions might form. The *Deegan-Packel index* [Deegan and Packel, 1978], on the other hand, assumes that once a coalition has sufficiently many members as to ensure that it has a value of 1, it will not accept others. Deegan and Packel [1978] 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 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 a 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

$$DP_i(\vec{w}; t, r) = \frac{r}{|\mathcal{W}_{\min}(\vec{w}; t, r)|} \sum_{S \in \mathcal{W}_{\min, i}(\vec{w}; t, r)} \frac{1}{|S|} \quad (2)$$

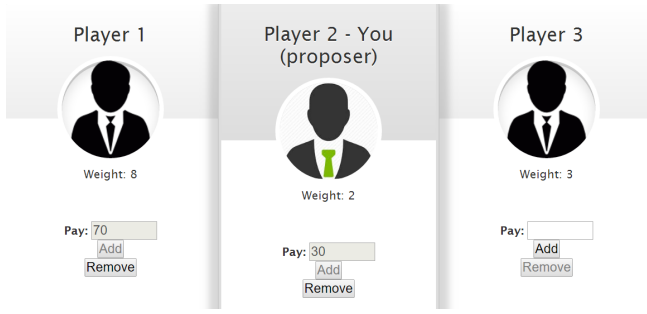


Figure 1: Snapshot of the Cooperative Negotiation Game for three agents showing Proposal Phase

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 The Cooperative Negotiation Game

In the real world, coalition formation is a process of negotiation between multiple parties who combine their resources [Dupont, 1996; Bachrach *et al.*, 2013]. To reflect this aspect, we designed an online version of a WVG called the *Cooperative Negotiation Game*. The game consists of two phases; in the *proposal phase*, a randomly chosen proposer p can suggest a coalition $S \subseteq N$. A coalition S is derived via a payoff division $\vec{x} \in \mathbb{R}_+^n$ such that $\text{supp}(\vec{x}) = S$ 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.

Figure 1 shows a snapshot of the proposal phase of the cooperative negotiation game with three agents, with weight configuration $\langle 8, 2, 3 \rangle$. The snapshot is shown from the proposer’s perspective (here, the proposer p is agent 2, and the reward r is set to 100). The proposer is attempting to form a coalition $\{1, 2\}$, where her share is 30 and the share of agent 1 is 70.

4.1 Data Collection

We recruited 111 subjects (2nd year software engineering undergraduates) with no prior background in game theory. 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 the institution running the study. All subjects played a 5-agent configuration of the cooperative negotiation game, in which agent weights varied between 1 and 9, the threshold t was set to 10, and the coalition value r was set to 100.

All subjects received the equivalent of an \$8 show-up fee, as well as a bonus that depended on their performance in the game (see Section 5), 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.

Each subject played a series of 5-agent cooperative negotiation games. The weight for each agent varied from 1 (weakest) to 9 (strongest), and was sampled from a normal distribution. The weights of all agents 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 members of a coalition could observe the proposals (including the proposer’s own share in the coalition), as well as the others’ responses. 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 ensues for the same participants, and a new proposer is chosen at random. The maximal number of rounds was set to 3 for all games (this information was not conveyed to any of the agents to avoid backward induction type reasoning, and was not used by the agent to make proposals in the game). In all, we collected 180 games and 343 coalition proposals.

4.2 The Extended Deegan-Packel Index

In this section we describe an extension of the original Deegan-Packel index, adapted to the cooperative negotiation game. The new measure differs in two ways: first, 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 she proposes. Second, it shares the revenue of winning coalitions in proportion to agent weights (rather than assume that all members are equally powerful). 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. Let us define $\mathcal{W}'_{\min,i}(\vec{w}, t, r)$ as

$$\mathcal{W}'_{\min,i}(\vec{w}, t, r) = \left\{ S \subseteq N : \forall S' \subset S : i \in S', w(S') < t \right\}$$

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 $\mathcal{W}_{\min,1} = \emptyset$: $\{2, 3\}$ is the unique minimally winning coalition.

We define $EDP_{i,p}(\vec{w}; t, r)$ to be the extended Deegan Packel index of an 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 still refer to it as a power index because it provides a measure of influence for a coalition members in the cooperative negotiation game.

$$EDP_{i,p}(\vec{w}; t, r) = \frac{1}{|\mathcal{W}'_{\min,p}|} \sum_{C \in \mathcal{W}'_{\min,p}, i \in S} \frac{r \cdot w_i}{w(S)} \quad (3)$$

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 which 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. Note that $EDP_i = \frac{1}{n} \sum_{p \in N} EDP_{i,p}$ has some interesting properties (for example, $EDP_i > 0$ for all $i \in N$, unlike most other power indices), which we leave for future work.

4.3 Predictive Model

In this section, we describe a supervised learning model that was used to predict responder acceptance in the cooperative negotiation game. We begin with 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 always assume that $x_p > 0$. The proportional power index of a responder i given proposer p , denoted $PP_{i,p}$, is the ratio between the power index for i and for p ($PP_{i,p} = \phi_i(\vec{w}, t, r) / \phi_p(\vec{w}, t, r)$). This measures the extent to which the proposer is more powerful in the game.

In our example $\langle 8, 2, 3; 10, 1 \rangle$, suppose that agent 2 is elected to be the proposer. For the extended Deegan-Packel index, we have that $PP_{1,2} = 4$ and that $PP_{3,2} = 0$. Note that by definition, it is always the case that $EDP_{p,p} > 0$, so $PP_{i,p}$ is well-defined when using EDP ; for the other power indices, if $\phi_i = 0$, we add a small $\varepsilon = 10^{-6}$ to ϕ_p to ensure that PP is well defined. The proportional share of a responder i , denoted $PS_{i,p}$ is the ratio between the share for i and the proposer p given \vec{x} : ($PS_{i,p} = x_i / x_p$). In our example, if $\vec{x} = (70, 30, 0)$, then $PS_{1,2} = 7/3$.

We use the following set of features to predict the probability of acceptance by responder i given the offer \vec{x} :

1. The power index of the responder i . ($\phi_i(\vec{w}, t, r)$).
2. The power index of the proposer p . ($\phi_p(\vec{w}, t, r)$).
3. The share of the proposer p in the coalition \vec{x} : (x_p).
4. The share of the responder the coalition \vec{x} . (x_i).
5. The ratio between the proportional share and the proportional power of the responder ($PS_{i,p} / PP_{i,p}$).

The last feature measures the extent to which the relative difference in shares between the proposer and the responder agrees with their relative difference in power. Suppose the responder is more powerful than the proposer (as in our example), i.e. $PP_{i,p} > 1$. A proposal that respects this difference would offer a more equal share to the responder. In our example, $PP_{1,2} = 4$ and $PS_{1,2} = 7/3$, so $PS_{1,2} / PP_{1,2} = (7/3) / 4 = 7/12$. Intuitively this ratio captures a notion of payment fairness (with respect to a given power index): no responder should reasonably agree to an offer that offers it a small share relative to the proposer, when its power is much greater.

We compare several predictive models using the above features, varying the type of power index used (Banzhaf, Shapley-Shubik, Banzhaf, Deegan-Packel, Extended Deegan-Packel). For each power index configuration, we

Method	AUC
EDP	0.71
Deegan-Packel	0.669
Shapley-value	0.68
Banzhaf-index	0.65
Always accept	0.5

Table 1: Performance of Logistic regression model when using different power indices to predict acceptance of proposal

implement several supervised machine learning models: logistic regression, a multilayer neural network (3 hidden layers, 3 decision nodes in each layer), and a Naive Bayes model. We report the receiver-operator characteristic curve (AUC), which measures the sensitivity of performance to the choice of the threshold for determining acceptance. AUC is a useful performance measure when evaluating unbalanced datasets (Although 85% of proposals were accepted, just 70% of coalition formation attempts were successful, see section 5) [Hanley and McNeil, 1982; Bradley, 1997; Provost *et al.*, 1998].

Table 1 describes the AUC score the logistic regression for the different indices using ten-fold cross validation. We also include an “always accept” predictor as a baseline. As shown in the figure, all power indices were beneficial for predicting the acceptance of responders in the game. However, the Extended Deegan-Packel index achieved the best performance by a small margin. The most important features, determined by their weights in the regression model, were as follows, ordered in decreasing order: the extended Deegan-Packel index of the proposer, the extended Deegan-Packel index of the responder, the proposed share of the proposer, the proposed share of the responder, the ratio between the proportional share and the proportional power of the proposer and the responder.

5 The EDP Agent

In this section we describe an agent termed EDP, which combines a decision-theoretic approach with the predictive model (using Deegan-Packel) that was described in the last section. Assuming each responder makes an independent decision whether to accept or reject the offer, the agent chooses a payoff division \vec{x}^* that maximizes its expected revenue:

$$\vec{x}^* \in \arg \max_{\vec{x}} x_p \cdot \prod_{i \in supp(\vec{x}), i \neq p} \Pr(Acc_i | \vec{x}, p) \quad (4)$$

Finding an approximately optimal \vec{x}^* is done by iterating over all possible payoff divisions in 5 unit intervals. The reason for this was twofold: There are approximately 45,000 possible payoff divisions to consider in this configuration, so brute-force search can be achieved in a short amount of time. Second, over 95% of shares made by human proposers were multiples of five; a software agent making arbitrary proposals would easily stand out from its human counterparts.

We evaluate the EDP agent by comparing its performance to that of people playing against other people. To this end we

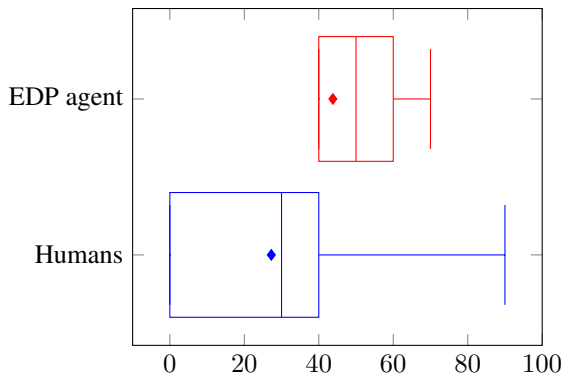


Figure 2: Comparison of total revenue gained on average in proposals.




Figure 3: Average share requested by proposer

recruited an additional 32 human subjects to play the cooperative negotiation game. All games included either five humans, or four humans and the EDP agent. In all, we collected 120 games including 163 proposals. The EDP agent was chosen to be the proposer 32 times. All results reported in this section are statistically significant in the $p < 0.01$ range using Mann-Whitney tests.

We measure the performance of the EDP agent by the total revenue gained, averaged over all games played. For each game, the EDP agent share was equal to zero (if no successful coalition was formed) or the proposed share of the EDP agent (if the proposed coalition was successful). We compare the total number of shares obtained by the EDP agent to that obtained by human proposers, averaged over all games.

5.1 Results and Discussion

We first describe the EDP agent’s performance as a proposer. Figure 2 compares between the performance of agent and human proposers, as well as summary statistics of the distribution (quartiles). As shown by the figure, the total average share obtained by the EDP agent (43.78) is significantly higher than that obtained by people (27.26).

Figure 3 shows the average shares requested by human and computer proposers for themselves. As seen in the figure, the EDP agent requested a much higher share for itself on average than did people; moreover, people’s proposals were more di-




Figure 4: Offers made by the people; note that the green dots represent non-power-preserving offers.

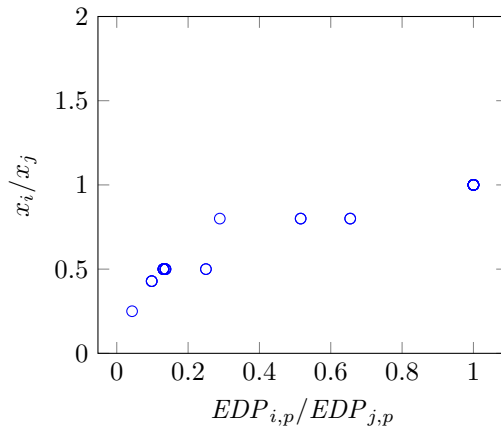


Figure 5: Offers made by the EDP agent; note that the agent always makes power-preserving offers

verse than the agent’s, with some requesting very low shares for themselves (low quartile for people’s requested share is 20, vs. 40 for the EDP agent).

The EDP agent also outperforms humans in forming coalitions (i.e. when all responders accept their individual share in the proposal): 79% of coalitions proposed by the EDP agent were successful, compared to 70% of human-proposed coalitions. The acceptance rate for individual responders to proposals made by the EDP agent (86%) which was not significantly different from that of people playing other people (85%).

These statistics show that on the one hand, the EDP agent made offers that were less advantageous to human responders than did humans; however, people were as likely to accept these offers as those made by other people. We offer several possible explanations for this discrepancy, by analyzing the behavior of human proposers and responders in the game.

The first explanation for the lower performance is that people make offers that do not align with responders’ power. The scatter-plot in Figure 4 shows the ratio between the EDP index of any responder pair (i, j) in a proposed coalition (x axis) and the ratio between the shares proposed to (i, j)

(y axis) by proposer p . For each coalition, any given 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 to the 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 $EDP_{i,p} \geq EDP_{j,p}$ but $p_i < p_j$, or $EDP_{i,p} > EDP_{j,p}$ but $p_i = p_j$. Many of the human-proposed offers were non power preserving (41% of all offers), and most of them were declined. As an example from the collected data, consider the game $\langle 6, 2, 2, 2, 1 \rangle$ in which $p = 5$. The extended Deegan-Packel power indices of the participants are $EDP_{i,5} = (0.6, 0.15, 0.1, 0.05, 0.1)$. The proposed coalition, $\vec{x} = (15, 15, 20, 20, 30)$ was 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.

In contrast, Figure 5 shows a scatter-plot of offers made by EDP agent according to the same criteria. As shown by the figure, there were 8 classes of offers made by the agent, all of them were “power preserving.” In particular, when the power of responder i and j were equal, the agent gave them equal shares. As the power of j grows, it receives a higher share, with a “jump” from 0.3 to 0.8 in the relative difference between the shares j and i when j ’s power increases to three times higher than that of i .

When acting as a responder, we measured performance by totaling the average share over all successful coalitions in which the responder was a member. The agent’s performance (34.7 average total share) was significantly larger than that of human responders (27.9 average total share). Here, the EDP agent 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. Figure 6 shows the cumulative distribution over human acceptance rates with games played with people. The figure shows that 35% of people reject offers with shares of 20% or lower. This bias, also documented in the ultimatum game [Camerer, 2003], explains the success of the EDP agent’s strategy as a responder.

A final explanation that can explain 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. In addition, spreading the reward among more responders results smaller shares on average, further decreasing the likelihood of acceptance.

Lastly, we present an example from the data that illustrates the difference in behavior between human proposers and the EDP agent. 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 DP-agent was elected to be the proposer it formed the coalition $\text{supp}(\vec{x}) = \{1, 3, 4\}$ with the shares $(50, 25, 25)$ respectively. The agent received a 100% success rate for this coalition proposal. When human proposers formed the same coal-

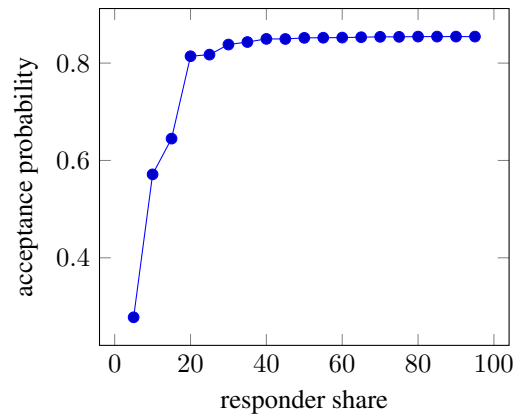


Figure 6: Cumulative distribution over the humans’ acceptance rate. The x axis indicates the responder’s share.

tion $\{1, 3, 4\}$, they awarded themselves a lower average share (35). For the same weight configuration, people also formed the coalition $\text{supp}(\vec{x}) = \{1, 2, 3\}$ that included agent 2 instead of agent 3. Since agent 2 is more powerful than agent 3 ($EDP_{2,1} > EDP_{3,1}$), it generally received a higher proposed share, at the expense of the proposer and agent 3. These coalitions were significantly less likely to succeed (75%) than the coalitions proposed by the agent (100%).

6 Discussion and Conclusions

The performance of the EDP agent makes a compelling argument for the combination of game-theoretic and ML based agents in coalitional bargaining domains. 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 our model to include repeated settings in which participants interact over time and participants need to consider the effects of reciprocity on their voting strategies. We are also studying how to create environments where humans can easily negotiate with one another (and with software agents) would be a challenge as the game grows more complex.

Acknowledgements

Gal and Mash are supported by ISF grant no. 773/16; Zick is supported by NRF Fellowship no. R-252-000-750-733.

References

[Ausubel *et al.*, 2002] Lawrence M Ausubel, Peter Cramton, and Raymond J Deneckere. Bargaining with incomplete information. *Handbook of game theory with economic applications*, 3:1897–1945, 2002.

- [Azoulay *et al.*, 2014] R. Azoulay, R. Katz, and S. Kraus. Efficient bidding strategies for cliff-edge problems. *JAA-MAS*, 28:290–336, 2014.
- [Bachrach *et al.*, 2010] Yoram Bachrach, Evangelos Markakis, Ezra Resnick, Ariel D Procaccia, Jeffrey S Rosenschein, and Amin Saberi. Approximating power indices: theoretical and empirical analysis. *JAAMAS*, 20(2):105–122, 2010.
- [Bachrach *et al.*, 2011] Yoram Bachrach, Pushmeet Kohli, and Thore Graepel. Ripoff: playing the cooperative negotiation game. In *AAMAS*, 2011.
- [Bachrach *et al.*, 2012] Yoram Bachrach, Thore Graepel, Gjergji Kasneci, Michal Kosinski, and Jurgen Van Gael. Crowd iq: aggregating opinions to boost performance. In *AAMAS*, 2012.
- [Bachrach *et al.*, 2013] Yoram Bachrach, David C Parkes, and Jeffrey S Rosenschein. Computing cooperative solution concepts in coalitional skill games. *Artificial Intelligence*, 2013.
- [Banzhaf, 1965] John F Banzhaf. Weighted voting doesn't work: Mathematical analysis. *Rutgers Law Review*, 1965.
- [Bitan *et al.*, 2013] Moshe Bitan, Ya'akov Gal, Sarit Kraus, Elad Dokow, and Amos Azaria. Social rankings in human-computer committees. In *AAAI*, 2013.
- [Blocq *et al.*, 2014] Gideon Blocq, Yoram Bachrach, and Peter Key. The shared assignment game and applications to pricing in cloud computing. In *AAMAS*, 2014.
- [Bradley, 1997] Andrew P Bradley. The use of the area under the ROC curve in the evaluation of machine learning algorithms. *Pattern recognition*, 30(7):1145–1159, 1997.
- [Camerer, 2003] Colin F Camerer. *Behavioral game theory: Experiments in strategic interaction*. Princeton University Press, 2003.
- [Chalkiadakis and Wooldridge, 2016] Georgios Chalkiadakis and Michael Wooldridge. Weighted voting games. In *Handbook of computational social choice*, chapter 16. Cambridge University Press, 2016.
- [Chalkiadakis *et al.*, 2011] Georgios Chalkiadakis, Edith Elkind, and Michael Wooldridge. *Computational Aspects of Cooperative Game Theory*. Morgan-Claypool, 2011.
- [Deegan and Packel, 1978] John Deegan and Edward W Packel. A new index of power for simple n -person games. *IJGT*, 1978.
- [Dupont, 1996] Christophe Dupont. Negotiation as coalition building. *International Negotiation*, 1(1):47–64, 1996.
- [Elkind *et al.*, 2007] Edith Elkind, Leslie Ann Goldberg, Paul Goldberg, and Michael Wooldridge. Computational complexity of weighted threshold games. In *AAAI*, 2007.
- [Felsenthal and Machover, 2005] Dan S Felsenthal and Moshé Machover. Voting power measurement: A story of reinvention. *Social Choice and Welfare*, 2005.
- [Güth, 1995] Werner Güth. On ultimatum bargaining experiments - a personal review. *Journal of Economic Behavior & Organization*, 27(3):329–344, 1995.
- [Haim *et al.*, 2017] Galit Haim, Ya'akov Gal, Bon An, and Sarit Kraus. Human-computer negotiation in a three player market setting. *Artificial Intelligence*, 246:34–52, 2017.
- [Hanley and McNeil, 1982] James A Hanley and Barbara JI McNeil. The meaning and use of the area under a receiver operating characteristic (roc) curve. *Radiology*, 143(1):29–36, 1982.
- [Leech, 2002] Dennis Leech. Designing the voting system for the council of the European Union. *Public Choice*, 113(3-4):437–464, 2002.
- [Lewenberg *et al.*, 2015] Yoav Lewenberg, Yoram Bachrach, Yonatan Sompolinsky, Aviv Zohar, and Jeffrey S Rosenschein. Bitcoin mining pools: A cooperative game theoretic analysis. In *AAMAS*, 2015.
- [Maschler *et al.*, 2013] Michael Maschler, Eilon Solan, and Shmuel Zamir. *Game Theory*. Cambridge University Press, 2013.
- [Mattei and Walsh, 2013] Nicholas Mattei and Toby Walsh. Preflib: A library of preference data. In *ADT*, 2013.
- [Oosterbeek *et al.*, 2004] Hessel Oosterbeek, Randolph Sloof, and Gijs Van De Kuilen. Cultural differences in ultimatum game experiments: Evidence from a meta-analysis. *Experimental Economics*, 7(2):171–188, 2004.
- [Oshrat *et al.*, 2009] Yinon Oshrat, Raz Lin, and Sarit Kraus. Facing the challenge of human-agent negotiations via effective general opponent modeling. In *AAMAS*, 2009.
- [Provost *et al.*, 1998] Foster J Provost, Tom Fawcett, and Roni Kohavi. The case against accuracy estimation for comparing induction algorithms. In *ICML*, 1998.
- [Rahwan *et al.*, 2015] Talal Rahwan, Tomasz P Michalak, Michael Wooldridge, and Nicholas R Jennings. Coalition structure generation: A survey. *Artificial Intelligence*, 229:139–174, 2015.
- [Rosenfeld and Kraus, 2016] Ariel Rosenfeld and Sarit Kraus. Providing arguments in discussions on the basis of the prediction of human argumentative behavior. *TiS*, 6:30, 2016.
- [See *et al.*, 2014] Abigail See, Yoram Bachrach, and Pushmeet Kohli. The cost of principles: analyzing power in compatibility weighted voting games. In *AAMAS*, 2014.
- [Shapley and Shubik, 1954] Lloyd S Shapley and Martin Shubik. A method for evaluating the distribution of power in a committee system. *APSR*, 1954.
- [Tal *et al.*, 2015] Maor Tal, Reshef Meir, and Ya'akov Gal. A study of human behavior in online voting. In *AAMAS*, 2015.
- [Zick *et al.*, 2015] Yair Zick, Yoram Bachrach, Ian A Kash, and Peter Key. Non-myopic negotiators see what's best. In *IJCAI*, 2015.