

Rip-off: Playing the Cooperative Negotiation Game (Extended Abstract)

Yoram Bachrach, Pushmeet Kohli, Thore Graepel
Microsoft Research
{yobach,pkohli,thoreg}@microsoft.com

ABSTRACT

We propose “Rip-off”, a new multi-player bargaining game based on the well-studied weighted voting game (WVG) model from cooperative game theory. Many different solution concepts, such as the Core and the Shapley value have been proposed to analyze models such as WVGs. However, there is little work on analyzing how humans actually play in such settings. We conducted several experiments where we let humans play “Rip-off”. Our analysis reveals that although solutions of games played by humans do suffer from certain biases, a player’s average payoff over several games is roughly reflected by the Shapley value.

Categories and Subject Descriptors

I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligence—*Multiagent Systems*;

J.4 [Computer Applications]: Social and Behavioral Sciences—*Economics*

General Terms

Economics, Experimentation, Algorithms, Human Factors

Keywords

Cooperative Game Theory, Negotiation, Shapley Value

1. INTRODUCTION

Many domains involve both competition and cooperation. Researchers have coined the term “co-opetition” to describe such settings [6]. One example is the model of weighted voting games (WVG), where each player has a weight, and a coalition of players wins the game if the sum of the weights of its participants exceeds a certain quota. Agent behavior in such settings has been studied in cooperative game theory. Forming a stable coalition requires the agents to share the gains in an appropriate way. Cooperative game theory provides several *solution concepts* that define how these joint gains should be distributed, such as the core [4] and the Shapley value [8], which were also studied in the context of WVGs [3, 1, 9, 2]. These solutions model “co-opetition”, but it is unclear whether they *predict human behavior*. They

Cite as: Rip-off: Playing the Cooperative Negotiation Game (Extended Abstract), Yoram Bachrach, Pushmeet Kohli, Thore Graepel, *Proc. of 10th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2011)*, Tumer, Yolum, Sonenberg and Stone (eds.), May, 2–6, 2011, Taipei, Taiwan, pp. XXX-XXX.
Copyright © 2011, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

assume that agents are completely rational, however human rationality may be bounded, and humans may have social biases such as avoiding very unequal payoffs [5, 7].

We study how humans behave in “co-opetition” settings and compare the payoff distribution results with those predicted by existing solutions. We have developed a new on-line multi-player cooperative bargaining game called “Rip-off”, based on the WVG model. We conducted experiments where groups of people played this game to win money. Our analysis revealed that solutions agreed by humans contain some biases, but that player’s expectations of their payoff are roughly reflected by the Shapley value.

1.1 The Rip-off Game

A transferable utility (TU) coalitional game Γ is composed of a set of n agents, I , and a characteristic function $v_{\Gamma} : 2^I \rightarrow \mathbb{R}$, mapping any subset (coalition) of the agents to a real value, indicating the total utility they achieve together. A specific class of games are *weighted voting games* (WVGs). In WVGs each agent $i \in I$ has weight w_i , and the game has a threshold t . A coalition $C \subseteq I$ wins if its total weight exceeds t : $v(C) = 1$ if $\sum_{i \in C} w_i \geq t$ and otherwise $v(C) = 0$. We denote the WVG over the n agents with weights w_1, w_2, \dots, w_n and threshold t as $[w_1, w_2, \dots, w_n; t]$. Given a coalition $C \subseteq I$ we denote $w(C) = \sum_{i \in C} w_i$. Game theory provides solutions that define how the participants might distribute the gains. An *imputation* (p_1, \dots, p_n) is a division of the gains, where $p_i \in \mathbb{R}$ and $\sum_{i=1}^n p_i = v(I)$. The value p_i is the payoff of agent i , and a coalition’s payoff is $p(C) = \sum_{i \in C} p_i$. The Shapley value is an imputation fulfilling certain fairness axioms [8]. We denote by π a permutation of the agents, by Π the set of all such permutations and by $S_{\pi}(i)$ the predecessors of i in π . The Shapley value of a game Γ is $sh(v_{\Gamma}) = (sh_1(v_{\Gamma}), \dots, sh_n(v_{\Gamma}))$ where $sh_i(v_{\Gamma}) = \frac{1}{n!} \sum_{\pi \in \Pi} [v_{\Gamma}(S_{\pi}(i) \cup \{i\}) - v_{\Gamma}(S_{\pi}(i))]$.

“Rip-off” is an online instance of a WVG played by humans. Similarly to a WVG $[w_1, w_2, \dots, w_n; t]$ where $C \subseteq I$ wins if $w(C) = \sum_{i \in C} w_i \geq t$, in “Rip-off” each player $i \in I$ is endowed with a fixed random weight $0 \leq w_i \leq 1$ and a ‘desired-share’ $0 \leq s_i \leq 1$ which is specified by the player. The share represents the amount the player would win if she is part of the winning coalition when the game ends. Thus, players wish to have the highest possible share. However, the winning coalition is entitled to £1 *in total*, to be shared among all the members of the winning coalition. Each “Rip-off” player sees the entire board, which includes the weight, desired payoff and current team number of each player. There are as many teams as there are players. All players who choose the same team number are considered

as part of a single coalition. Given a team j , we denote the players whose current choice is team j as $C_j \subseteq I$. A coalition $C_j \subseteq I$ is *successful* if the sum of the weights of its players exceeds the threshold $t = 1$, *ie.* $w(C_j) = \sum_{i \in C_j} w_i \geq t$. A coalition $C_j \subseteq I$ of players is *in agreement* if the sum of the ‘desired-shares’ of its players is at most £1, *ie.* $\sum_{i \in C_j} s_i \leq 1$. A coalition *wins* if it is both *successful* and *in agreement*. We say that C_j is in the “negotiation phase” if it successful so $w(C_k) \geq 1$ but has not yet reached agreement so $s(C_k) > 1$. More formally, $w(C_j) = \sum_{i \in C_j} w_i \geq t$ and $\sum_{i \in C_j} s_i \leq 1$. Such a coalition C_j is a winning coalition in the underlying WVG. The player weights are chosen so no player can win the game on their own *ie.* for all i , $w_i < 1$. A successful team could potentially win £1, however its players must agree on how to split this reward. To negotiate how to share the reward, each “Rip-off” player $i \in I$ chooses a share $0 \leq s_i \leq 1$ by entering a number into a text field.

Initial State: The game starts with player i starts in team i , so all players are assigned to different teams. The shares of all players are initialized to 1. Figure 1 shows the initial state for the WVG [0.25, 0.25, 0.4, 0.4, 0.25; 1], from the perspective of Player 1. Each player can identify who she is, as the active player is marked with a box. A player can only change her team and share and not those of the other players, but the selections of all players are displayed.

Progress: At any time a player may change her selection of a team, thereby choosing to join a different coalition. A player may also change her share at any time. However, a player is not allowed to join already successful teams.



Figure 1: Example of a “Rip-off” game board.

Termination: The game ends when a *winning* coalition C_j is formed, *ie.* a team which is both *successful* and *in agreement*. Upon termination, each player $i \in C_j$ in the winning coalition obtains a reward of s_i . Any agent $i \in I \setminus C_j$ obtains a reward of zero, regardless of her share s_i .

2. RESULTS AND ANALYSIS

We invited 20 volunteers to play “Rip-off”. The participants were divided into 4 groups of 5 participants each. Each group played for 90 minutes and players were awarded the sum of their payoffs through all the games played. We picked games with 9 different weight settings, to cover a broad range of Shapley values. The configurations were chosen uniformly at random for the various games, and the weights were randomly assigned to the players. The “Rip-off” game are directly based on WVGs so one can view their game theoretic solutions as predictions regarding the results of such games.

Not all “Rip-off” players are equally powerful: depending on the weights, some coalitions are winning while others are losing. Players are aware of all the weights and the team selections of other players, although they cannot change them. The Shapley value is considered a powerful tool to analyze

a player’s power in such settings, which may not be proportional to her weight. It reflects the fair share each player (weight) should get in a WVG.

One can interpret the Shapley value as the *average* amount a weight is likely to get over *many* “Rip-off” games. We denote the set of all weights as W . For each playing group and each board (weight configuration) we logged the total rewards each weight has won, and denote this as $tot(w)$. Given a weight w , its proportional gain is $p(w) = \frac{tot(w)}{\sum_{w \in W} tot(w)}$.

Figure 2 is a scatter plot, showing the relation between a weight’s Shapley value and its proportional gain. An experiment is the session of all the games a single group of 5 human participants played. Each data point represents a single WVG board configuration in an experiment. The x-axis is the Shapley value of the weight and the y-axis is the proportional gains of that weight.

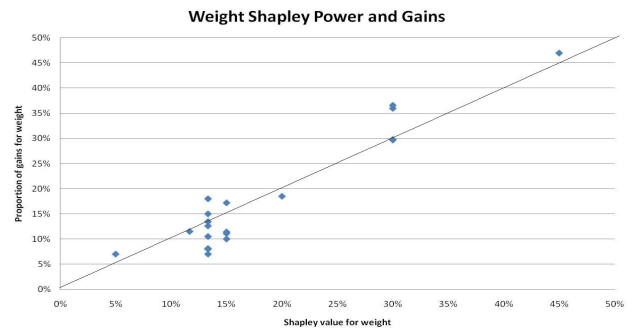


Figure 2: Scatter plot showing correlation between a weights Shapley value and its gains proportion.

Figure 2 shows that the Shapley value is quite accurate as a prediction of a weight’s proportional gains. If it fully predicted the gains, all points should be on the line $y = x$. The points are indeed close, and the correlation coefficient is 95%. Although the Shapley value was designed as a theoretical tool for fair allocation, it can be a useful tool for predicting human negotiation in “co-opetition” settings.

3. REFERENCES

- [1] Y. Bachrach and E. Elkind. Divide and conquer: False-name manipulations in weighted voting games. In *AAMAS*, 2008.
- [2] Y. Bachrach, R. Meir, M. Zuckerman, J. Rothe, and J. Rosenschein. The cost of stability in weighted voting games. In *AAMAS*, 2009.
- [3] E. Elkind and D. Pasechnik. Computing the nucleolus of weighted voting games. In *SODA*, 2009.
- [4] D. B. Gillies. *Some theorems on n-person games*. PhD thesis, Princeton University, 1953.
- [5] J. Henrich, R. Boyd, S. Bowles, C. Camerer, E. Fehr, H. Gintis, and R. McElreath. In search of homo economicus: behavioral experiments in 15 small-scale societies. *American Economic Review*, 91(2):73–78, 2001.
- [6] B. Nalebuff and A. Brandenburger. *Co-opetition*. HarperCollinsBusiness, 1996.
- [7] H. Oosterbeek, R. Sloof, and G. Van De Kuilen. Cultural differences in ultimatum game experiments: Evidence from a meta-analysis. *Experimental Economics*, 7(2):171–188, 2004.
- [8] L. S. Shapley. A value for n-person games. *Contrib. to the Theory of Games*, pages 31–40, 1953.
- [9] M. Zuckerman, P. Faliszewski, Y. Bachrach, and E. Elkind. Manipulating the quota in weighted voting games. In *AAAI 2008*.