The Social Ultimatum Game and Adaptive Agents (Demonstration)

Yu-Han Chang
University of Southern California
4676 Admiralty Way #1001
Marina Del Rey, CA 90292, USA.
ychang@isi.edu

Rajiv Maheswaran
University of Southern California
4676 Admiralty Way #1001
Marina Del Rey, CA 90292, USA.
maheswar@usc.edu

ABSTRACT

The Ultimatum Game is a key exemplar that shows how human play often deviates from "rational" strategies suggested by game-theoretic analysis. One explanation is that humans cannot put aside the assumption of being in a multi-player multi-round environment that they are accustomed to in the real world. We introduce the Social Ultimatum Game (SUG), where players can choose their partner among a society of agents, and engage in repeated interactions of the Ultimatum Game. We develop mathematical models of human play that include "irrational" concepts such as fairness and adaptation to the expectations of the society. We will display a system where people can play SUG against a mixed system of other humans and autonomous agents based on our mathematical models.

Categories and Subject Descriptors

I.1.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence—*Intelligent agents, Multiagent systems*

General Terms

Algorithms,Economics,Experimentation

Keywords

Multi-Agent Systems, Game Theory, Ultimatum Game, Mathematical Models of Human Behavior, Learning, Adaptation

1. INTRODUCTION

The Ultimatum Game has been studied extensively and is a prominent example of how human behavior deviates from game-theoretic predictions that use the "rational actor" model. The classical game involves two players who are given the opportunity to split \$10. One player proposes a potential split, and the other can accept, in which case the players receive the amounts in the proposal, or reject, in which case, both players receive nothing. The subgame perfect Nash equilibrium (or Stackelberg equilibrium) for this game, has the first player offering \$1 to the other player and keeping \$9, and the second player accepting, because \$1 is better than nothing. However, when experiments are conducted with human players, this behavior is rarely observed.

One seemingly intuitive and straightforward explanation that has not received much treatment in the literature is that humans engage in similar endeavors in many real-life situations, and may not view the experimenter's game independently of these other, more familiar situations. When faced with an isolated Ultimatum Game in the lab, humans bring in these experiences and act in the way that is familiar and habitual to them. To understand this behavior, then, we need to examine the settings of these real-life interactions. One key feature of these interactions is that there are multiple potential game partners and many games to be played over time, that is, life is a multi-player and repeated game. This makes the strategy space much more complex, and introduces many new possible equilibrium strategies. To design multi-agent systems that interact with humans or model human behavior, we must understand the nature of strategic interactions in such games.

2. RELATED WORK

Economists and sociologists have proposed many variants and contexts of the Ultimatum Game that seek to address the divergence between the "rational" Nash equilibrium strategy and observed human behavior [3, 6, 5]. These papers show that various cultural factors along with other human properties bias human players away from classically "rational" play. In the machine learning and theoretical computer science communities, over the past decade, there has been interest in (1) design of algorithms that compute or converge to Nash equilibrium, and (2) design of agent strategies that achieve good results when interacting with other independently designed agents [8] Other researchers have formulated efficient solution methods for games with special structures, such as limited degree of interactions between players linked in a network, or limited influence of their action choices on overall payoffs for all players [4, 7]. When profit maximization is the key metric, adaptation policies have been proposed that can be shown to be optimal against certain opponents, or that minimize a regret metric when playing against arbitrary opponents [2, 1].

3. SOCIAL ULTIMATUM GAME

The Ultimatum Game, is a two-player game where a player, P_1 proposes a split of an endowment $e \in \mathbb{N}$ to another player P_2 where P_2 would receive $q \in \{0, \delta, 2\delta, \ldots, e-\delta, e\}$ for some value $\delta \in \mathbb{N}$. If P_2 accepts the offer, they receive q and P_1 receives $e-q$. If P_2 rejects, neither player receives anything.

Cite as: The Social Ultimatum Game and Adaptive Agents (Demonstration), Yu-Han Chang and Rajiv Maheswaran, *Proc. of 10th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2011)*, Yolum, Tumer, Stone and Sonenberg (eds.), May, 2–6, 2011, Taipei, Taiwan, pp. 1313-1314.

The subgame-perfect Nash or Stackelberg equilibrium states that P_1 offer $q = \delta$, and P_2 accept. This is because a "rational" P_2 should accept any offer of $q > 0$, and P_1 knows this. Yet, humans make offers that exceed δ, even making "fair" offers of $e/2$, and reject offers greater than the minimum.

To represent the characteristics that people operate in societies of multiple agents and repeated interactions, we introduce the Social Ultimatum Game. There are N players, denoted $\{P_1, P_2, \ldots, P_N\}$, playing K rounds, where $N \geq 3$. The requirement of having at least three players in necessary to give each player a choice of whom to interact with.

In each round k, every player P_m chooses a single potential partner P_n and makes an offer $q_{m,n}^k$. Each player P_n then considers the offers they have received and makes a decision $d_{m,n}^k \in \{0, 1\}$ with respect to each offer $q_{m,n}^k$ to either accept (1) or reject (0) it. If the offer is accepted by P_m, P_m receives $e - q_{m,n}^k$ and P_n receives $q_{m,nj}^k$, where e is the endowment to be shared. If an offer is rejected by P_n, then both players receive 0 for that particular offer in round k. Thus, P_m's reward in round k is the sum of the offers they accept from other players (if any are made to them) and their portion of the proposal they make to another player, if accepted, $r_m^k = (e - q_{m,n}^k)d_{m,n}^k + \sum_{j=1\ldots N, j \neq m} q_{j,m}^k d_{j,m}^k$. The total rewards for P_m over the game is the sum of per-round winnings, $r_m \sum_{k=1}^{K} r_m^k$.

4. ADAPTIVE AGENTS MODEL

To create mathematical models of human player for the Social Ultimatum Game that can yield results that match observed phenomena, we need to incorporate some axioms of human behavior that may be considered "irrational". The desiderata that we address include assumptions that people will (1) start with some notion of a fair offer, (2) adapt these notions over time at various rates based upon their interactions, (3) have models of other agents, (4) choose the best option while occasionally exploring for better deals. Each player P_m is characterized by three parameters: (1) α_m^0 : Player m's initial acceptance threshold, (2) β_m : Player m's reactivity and (3) γ_m : Player m's exploration likelihood

The value of $\alpha_m^0 \in [0, e]$ is P_m's initial notion of what constitutes a "fair" offer and is used to determine whether an offer to P_m, i.e., $q_{n,m}^k$, is accepted or rejected. The value of $\beta_m \in [0, 1]$ determines how quickly the player will adapt to information during the game, where zero indicates a player who will not change anything from their initial beliefs and one indicates a player who will solely use the last data point. The value of $\gamma_m \in [0, 1]$ indicates how much a player will deviate from their "best" play in order to discover new opportunities where zero indicates a player who never deviates and one indicates a player who always does.

Each player P_m keeps a model of other players in order to determine which player to make an offer to, and how much that offer should be. The model is composed as follows:

- $a_{m,n}^k$: P_m's estimate of P_n's acceptance threshold
- $\bar{a}_{m,n}^k$: Upper bound on $a_{m,n}^k$
- $\underline{a}_{m,n}^k$: Lower bound on $a_{m,n}^k$

Thus, P_m has a collection of models for all other players $\{[\underline{a}_{m,n}^k a_{m,n}^k \bar{a}_{m,n}^k]\}_n$ for each round k. The value $a_{m,n}$ is the P_m's estimate about the value of P_n's acceptance threshold, while $\underline{a}_{m,n}^k$ and $\bar{a}_{m,n}^k$ represent the interval of uncertainty over which the estimate could exist.

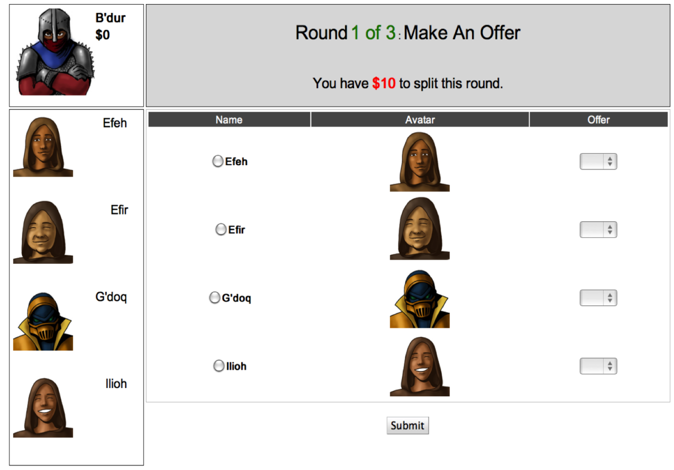

Figure 1: The Social Ultimatum Game Interface

5. DEMONSTRATION

People will able to be play the Social Ultimatum Game in hybrid environments against other people along with the adaptive agents described above along with classical rational agents. The interface is shown in Figure 1. All participants and agents will have avatars so that one cannot tell if a player is a human, adaptive or rational agent. Human players will be rewarded based on their performance in the game. In addition, we will keep a running tally board of how humans have performed with respect to adaptive and rational agents as well as the top-performing human players.

6. REFERENCES

[1] Y.-H. Chang. No regrets about no-regret. *Artificial Intelligence*, 171(7), 2007.

[2] Y.-H. Chang and L. P. Kaelbling. Hedged learning: Regret minimization with learning experts. In *International Conference on Machine Learning (ICML)*, 2005.

[3] J. Henrich. Does culture matter in economic behavior? ultimatum game bargaining among the machiguenga. *American Economic Review*, 90(4):973–979, 2000.

[4] M. Kearns, M. Littman, and S. Singh. Graphical models for game theory. In *Conference on Uncertainty in Artificial Intelligence*, pages 253–260, 2001.

[5] A. G. S. Mascha van't Wout, René S. Kahn and A. Aleman. Affective state and decision-making in the ultimatum game. *Experimental Brain Research*, 169(4):564–568, 2006.

[6] H. Oosterbeek, R. Sloof, and G. van de Kuilen. Differences in ultimatum game experiments: Evidence from a meta-analysis. *Experimental Economics*, 7:171–188, 2004.

[7] L. Ortiz and M. Kearns. Nash propagation for loopy graphical games. In *Neural Information Processing Systems*, 2003.

[8] Y. Shoham, R. Powers, and T. Grenager. If multi-agent learning is the answer, what is the question? *Artificial Intelligence*, 171(7):365–377, 2007.