Behavioral Game Theory on Online Social Networks: Colonel Blotto is on Facebook

Pushmeet Kohli, Yoram Bachrach, Thore Graepel, Gavin Smyth, Michael Armstrong
Microsoft Research Cambridge, UK

David Stillwell
University of Nottingham, UK

Ralf Herbrich
Facebook Inc., USA

Michael Kearns
University of Pennsylvania, USA

Abstract

We show how online social networks such as Facebook can be used in Behavioral Game Theory research. We report the deployment of a Facebook application 'Project Waterloo' that allows users to play the Colonel Blotto game against their friends and strangers. Unlike conventional studies performed in the laboratory environment, which rely on monetary incentives to attract human subjects to play games, our framework does not use money and instead relies on reputation and entertainment incentives. We describe the Facebook application we created for conducting this experiment, and perform a preliminary analysis of the data collected in the game. We conclude by discussing the advantages of our approach and list some ideas for future work.

1 Introduction

For a long time, economists and game theorists have been interested in understanding how individuals (people) or institutions (businesses, corporation, and countries) behave in different economic situations. This knowledge is extremely valuable and can be used, for instance, to build more accurate and robust economic models. Classical game theory predicts how rational agents behave in strategic settings, such as in the domain of advertising, during business interactions, and on the job market. Although it does allow for making predictions regarding human behaviour, it makes very strong assumptions. For example, the agents in classical game theory are assumed to be fully rational: they base their decisions solely on maximising utility, are capable of performing very complex reasoning and assume that their adversaries are equally rational.

Humans in the real-world, on the other hand, are quite different. Their behaviour is sometimes emotional, they sometimes base decisions on concepts such as fairness and reciprocity (rather than only on the monetary amount they get), and are bounded in their reasoning capabilities and thus often use heuristic reasoning. One prominent example is the Ultimatum Game, where two players interact to determine how to divide a sum of money given to them. In this game, the first player makes a "take it or leave it" offer to the other player, suggesting how the sum should be divided. The second player may either accept the offer, in which case the money is divided according to the proposal, or she can reject the offer, in which case both players get nothing. When humans, from many cultures, play this game they often offer equal shares (50:50 offer), and offers below 20% are rejected quite often [9, 15]. This behavior is quite different from game theoretic solutions according to which the first player should offer the minimal non-zero amount and the second player should accept all solutions with non zero payoffs.

In order to study how agents (individuals or groups) behave in social and economic situations, researchers have conducted a number of empirical studies. Research in this space falls in the very active field of Behavioral Game Theory [6], which examines how humans behave in various game theoretic settings [4, 13, 12, 7, 16]. However, due to logistical constraints, most such studies have been limited to the laboratory environment and to a small number of people. This introduces a number of biases in the data collected in these studies. For instance, Arnett [2] in a survey of empirical studies in psychology found that *'96% of subjects were from Western industrialized countries which house just 12% of the world's population'*. Henrich *et al.*[10] argue that as human subjects used in most such studies are Western, educated, industrialized, rich, democratic (WEIRD), it would be inappropriate to generalize the findings to people from other societies.

This paper reports our first attempt at gathering human behavior data from online social networks. We have created an application that allows users of a popular online social network to play a two player turn-based zero-sum game called Colonel Blotto (or simply, Blotto). The Colonel Blotto game is well known in Game Theory (See [1] for further discussion of the game and its origins), and has been used to model many political and economic situations. Calculating the equilibrium of this game is complicated and so is the choice of the optimal strategy. We analyze how users play the game and compare their behaviour with that reported in previous studies.

The use of online social networks for Behavioral Game Theory experiments overcomes some of the problems associated with laboratory experiments, and also has some unique benefits. First, the popularity of social networks such as Facebook has meant that that they have users from all over the world[1]. This allows researchers to study the effect of regional and cultural differences in the behaviour of subjects. Secondly, social networks capture how users are related to each other. This data can help behavioral game theorists to find, for instance, if people play differently with people they know as opposed to strangers. Furthermore, the subjects participate in their normal "habitat" rather than in an artificial lab setting. And lastly, but most importantly, the experiments can be conducted at much larger scales with the ability to deal with thousands of subjects to obtain results at a finer granularity of subject attributes while maintaining statistical significance.

2 The Colonel Blotto Game

The Colonel Blotto game was initially proposed by Borel [5] and attracted a lot of research after it was used to model a game between two presidential candidates who have to allocate their limited budgets to campaigns in the battlefield states. A related vote-buying interpretation of the Blotto game has also been proposed (see Myerson [14]).

The mechanics of the game are as follows. The two players of the game are tasked to simultaneously distribute limited resources over several, in our case 5, objects. These objects could represent cities, products, or battlefields. To make the game appealing to users, we decided to used battlefields/territories as objects, and troops as resources. In our implementation, each player has 100 troops and distributes them over five battlefields. Once a player has distributed the troops, they send this distribution as a challenge to another player who cannot see the allocation. The responding player distributes their 100 troops among the battlefields and finalizes the game. If a player allocates more troops in a territory than their opponent does, they win that territory. The player who wins more battlefields wins the game.

Theoretical Analysis and Optimal Strategies The analysis of the optimal strategy for Blotto games is quite difficult. The general formulation of discrete Blotto games with N players and K battlefields has been analyzed analytically by Hart [8]. The optimal strategy is a randomized strategy which allocates troops to battlefields symmetrically while maintaining that the marginal distribution of the troops in each battlefield is uniform in the interval $[0, 2N/K]$.

Previous Experimental Analysis A number of studies have been conducted on how people play Blotto [3, 1, 11]. Most such studies, however, deal with the case where the player chooses the distribution of troops once. This selected distribution is then compared to the troop distribution of all other players to find rankings for different troop distributions. Such studies, however, do not allow researchers to analyze questions such as whether people play Blotto differently with people they know as opposed to strangers.

The analysis of Arad and Rubinstein [1] is particularly noteworthy. They analyze how people play a Blotto game which has 120 troops that have to be distributed among six battlefields. Their aim was to explain the behavior of players using a decision procedure based on multi-dimensional iterative reasoning. They used two datasets of played Blotto games. The first dataset, called *Classes*, was collected from game theory students who were asked to play this game by their teachers. The second

[1] 70% of Facebook users are from outside USA. See `http://www.facebook.com/press/info.php?statistics`.

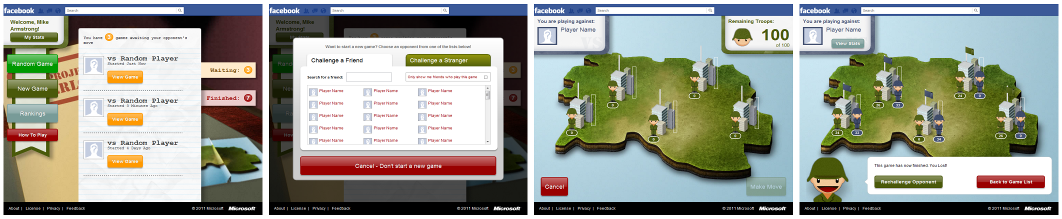

Figure 1: Steps of the Project Waterloo Facebook application. From left to right: (1) The user chooses a game with either a random hidden opponent or a visible one. (2) For a visible opponent game, the user chooses an opponent from his Facebook friend list or the list of current users of our application. (3) The user distributes the troops and makes the move. (4) The opponent does the same and the result is revealed to both players.

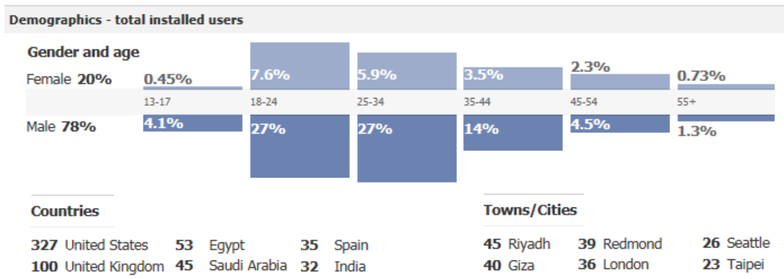

Figure 2: Demographic data on Project Waterloo users.

Game	Player	Field 1	Field 2	Field 3	Field 4	Field 5	Cities Won	Result
G1	P1	**22**	13	**22**	15	28	2	Loss
	P2	0	**30**	0	**35**	**35**	3	Win
G2	P1	**21**	**15**	**34**	26	4	3	Win
	P2	10	5	25	**40**	**20**	2	Loss

Table 1: Troop distributions in some sample games.

dataset was collected from readers of the Hebrew business daily, Calcalist. In section 4, we compare the results of our analysis with those obtained in [1].

3 The Project Waterloo Application on Facebook

We have developed a game, called Project Waterloo[2], which allows users of the online social network Facebook, to play Blotto with friends and strangers. Our implementation allows users to play either random players whose identity is hidden from the player, or known players from their friend list or from a list of players who have played the game before. Figure 1 shows screenshots of the various steps of the Project Waterloo application.

Previous experimental studies of Blotto have used monetary incentives to recruit human subjects. In contrast, we rely on entertainment and reputation incentives to attract users to play the game. The reputation incentive is realized by showing users their ranking (based on performance) among all players and players in their friend network. Users are ranked according to their rating (R) which is computed as: $R = \frac{\text{\#Games Won}}{\text{\#Games Played}+10} \times 100$. This measure encourages users to play more games as it makes sure that people who have played a very small number of games are not given a high rating.

4 Analysis of the Data

We now present an analysis of the game data we collected from the users of the 'Project Waterloo' Facebook application. The dataset used for our analysis contains 1,883 games played by 685 players. In 1,027 of the 1,883 games, the players initiating the game did not know the identity of their opponent. For the remaining 856 games, both players knew each other's identities. 524 of these 856 games were played between people who were friends on Facebook, and the remaining 332 games were played between strangers. Examples of some game instances are shown in Table 1. Detailed demographic data for users is shown in Figure 2.

[2]Available at `http://apps.facebook.com/msrwaterloo/`.

	Field **1**	Field **2**	Field **3**	Field **4**	Field **5**
Mean, Median (Hidden)	19.32, 20	20.66, 20	22.00, 20	19.91, 20	18.07, 20
Mean, Median (Friends)	17.44, 16	21.50, 22	21.88, 23	20.46, 22	18.70, 20
Mean, Median (All Games)	17.81, 20	20.67, 21	22.13, 24	20.83, 22	18.56, 20

Table 2: Position bias in the allocation of troops to different territories. The table shows the mean and median of the troop allocations to different battlefields.

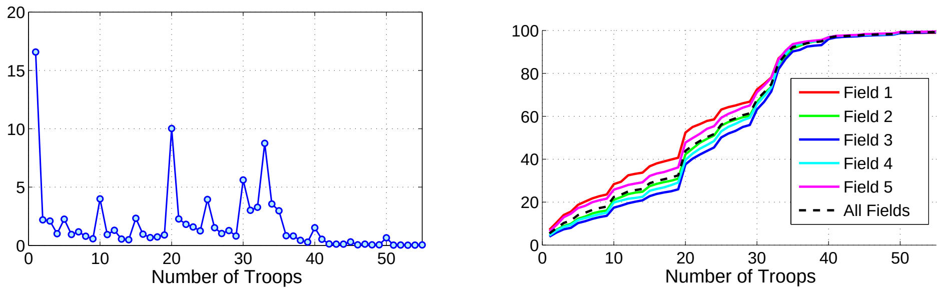

Figure 3: (**LEFT**) Frequency of troop group size. (**RIGHT**) Cumulative distributions of field assignments. Observe that there are three spikes at the numbers 1, 20 and 33. These correspond to the intentions: users do not want to leave a field unguarded and allocate atleast 1 troop, users play the uniform strategy and allocate 20 troops, and finally users go for the strategy of going for 3 highly reinforced territories with atleast 33 troops in each city.

Dataset	0	1	2	3	4	5	6	7	8	9
Classes	62%	10%	3%	2%	4%	12%	1%	1%	2%	4%
Calcalist	56%	13%	5%	2%	5%	11%	1%	1%	2%	4%
Our	34.12%	11.56%	8.76%	13.13%	6.44%	11.85%	4.33%	3.79%	3.32%	2.65

Table 3: The percentage of troop allocations to fields which have a particular unit digit. We compare the distribution for the data collected by us with the datasets used by Arad and Rubinstein [1].

#Reinforced Cities	0	1	2	3	4
(Hidden)	11.98%	14.51%	26.29%	43.23%	3.99%
(Friends)	5.15%	8.21%	27.48%	49.8092 %	9.35%
(All Games)	6.24 %	7.67%	23.76%	56.85%	5.47%

Table 4: The figure shows the what percentage of played troop splits had a specific number of reinforced battlefields *ie.* battlefields with more than 20 troops.

Positional Bias in Troop Allocation As mentioned before, the optimal strategy in a Blotto game is to allocate troops to battlefields symmetrically at random. However, earlier studies have observed that the order in which troops are allocated to territories introduces a bias in the troop allocations [1]. Arad and Rubinstein [1] have conjectured that this may be due to a player's instincts to over-assign troops leading to 'residual' allocations at the fringe territories. We observe this bias in the data collected by us (see Table 2 and Figure 3). Interestingly, we also observe that the mean and median allocation of troops to the fringe territories is different in games played against 'anonymous' users as compared to games played against friends.

Frequency of Troop Allocations A natural allocation of troops to battlefields is a uniform allocation of 20 troops to all battlefields. Some users do follow this strategy, but given this information, it would be better to allocate slightly more than 20 troops to at least some of the cities. In fact, the most efficient allocation of troops to a particular city is one more than the expected allocation of troops made by the opponent. Players of Blotto typically use iterative reasoning to come up with a troop allocation. The size of troop groups used by players in their allocations give hints about the number of levels of iterated reasoning they are employing while playing the game. Figure 3 and Table 3 provide statistics about the number of troops used by players of the game.

Number of Reinforced Cities Another indicator of the level of reasoning employed by players in the Blotto game is the number of *reinforced* territories. A territory is called *reinforced* if it has more troops than the uniform allocation of troops *ie.* 20 troops in our case. Our analysis shows that the number of reinforced territories in the distributions made by players of our game was large. This may be interpreted in favour of the hypothesis that players of the 'Project Waterloo' game are employing multiple levels of iterative reasoning.

Effect of Knowledge about the Opponent Playing the Blotto game involves reasoning about the strategy of your opponent. Thus, it would be reasonable to assume that knowledge of the opponent in the game would affect the way a player plays the game. Previous experimental studies have not been able to capture such effects. This has primarily been due to the difficulty in obtaining game data where all players know, or are friends of, their opponents. Deploying the game on an online social network provides access to the friend relationships of the players. We observe that the distribution of troops selected by the users is different in situations when the opponent is 'hidden' as compared to being a friend (see Table 2 and Table 4).

5 Discussion and Future Work

We have shown that data gathered from behavioral game theory experiments conducted using online social networks is consistent with previous studies. We have also shown that this new source of data can help uncover interesting relationships between player strategies and the game context. We found that the strategy adopted by players changes if their opponent is their friend.

There are a number of interesting issues that remain unexplored. Questions on how users learn, how their play is affected by their age, gender, location, etc., are promising directions for future work. To conclude, we believe that online social networks will become a very useful resource for empirical research in fields such as behavioral game theory and experimental economics.

References

[1] Ayala Arad and Ariel Rubinstein. Colonel Blotto's top secret files: Multi-dimensional iterative reasoning in action. *Working paper*, 2010.

[2] Jeffrey J. Arnett. The neglected 95%: Why American psychology needs to become less American. *American Psychologist*, 63:602–614, 2008.

[3] Judith Avrahami and Yaakov Kareev. Do the weak stand a chance? Distribution of resources in a competitive environment. *Cognitive Science*, 33(5):940–950, 2009.

[4] Y. Bachrach, P. Kohli, and T. Graepel. Rip-off: Playing the cooperative negotiation gam. In *AAMAS*, 2010.

[5] Émile Borel. The theory of play and integral equations with skew symmetric kernels (La théorie du jeu et les équations intégrales à noyau symétrique). *Econometrica*, 21:97–100, [1921] 1953.

[6] C. Camerer. *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton University Press Princeton, NJ, 2003.

[7] Tanmoy Chakraborty, J. Stephen Judd, Michael Kearns, and Jinsong Tan. A behavioral study of bargaining in social networks. In *ACM Conference on Electronic Commerce*, pages 243–252, 2010.

[8] Sergiu Hart. Discrete Colonel Blotto and General Lotto games. Discussion paper series, Center for Rationality and Interactive Decision Theory, Hebrew University, Jerusalem, October 2006.

[9] 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.

[10] Joseph Henrich, Steven J. Heine, and Ara Norenzayan. Most people are not WEIRD. *Nature*, 466:29, 2010.

[11] Rafael Hortala-Vallve and Aniol Llorente-Saguer. A simple mechanism for resolving conflict. *Games and Economic Behavior*, 70(2):375–391, 2010.

[12] J. Stephen Judd and Michael Kearns. Behavioral experiments in networked trade. In *ACM Conference on Electronic Commerce*, pages 150–159, 2008.

[13] M. Kearns, S. Suri, and N. Motfort. An experimental study of the coloring problem on human subject networks. In *Science*, 2006.

[14] Roger B. Myerson. Incentives to cultivate favored minorities under alternative electoral systems. *The American Political Science Review*, 87(4):856–869, 1993.

[15] 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.

[16] D. Ray, B. King-Casas, P. R. Montague, and P. Dayan. Bayesian model of behaviour in economic games. In *NIPS*, 2008.