

Colonel Blotto On Facebook: The Effect of Social Relations On Strategic Interaction

Pushmeet Kohli
Microsoft Research
Cambridge, UK
pkohli@microsoft.com

Yoram Bachrach
Microsoft Research
Cambridge, UK
yobach@microsoft.com

David Stillwell
University of Cambridge
Cambridge, UK
david@mypersonality.org

Michael Kearns
University Of Pennsylvania
Philadelphia, USA
mkearns@cis.upenn.edu

Ralf Herbrich
Facebook
Mountain View, USA
ralf@fb.com

Thore Graepel
Microsoft Research
Cambridge, UK
thoreg@microsoft.com

ABSTRACT

We study how social relations between people affect the way they play the famous resource allocation game called Colonel Blotto. We report the deployment of a Facebook application called “Project Waterloo” which allows users to invite both friends and strangers to play Colonel Blotto against them. Most previous empirical studies of Blotto have been performed in a laboratory environment and have typically employed monetary incentives to attract human subjects to play games. In contrast, our framework relies on reputation and entertainment incentives to attract players. Deploying the game on a social network allows us to capture the social relations between players and analyze their impact on the used strategies.

Following [1] we examine player strategies and contrast them with game theoretic predictions. We then investigate how strategies are affected by social relations. Our analysis reveals that knowledge of the opponent affects the strategies chosen by players and how well they perform in the game. We show that players with few Facebook friends tend to play more games and have a higher probability of winning, that players responding to a challenge in the game have a higher probability of winning than those initiating the game, and that the initiators of a game have a higher probability of defeating their friends than strangers.

INTRODUCTION

Economists and game theorists are interested in understanding how agents (individuals or institutions) behave in strategic situations. This knowledge is extremely valuable and can

be used, for instance, to build more accurate and robust economic models. Game theory predicts how agents behave in strategic settings such as auctions and business interactions. 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 behavior 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 frequently rejected [11, 17]. This behaviour is very 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 pay-offs.

In order to study how people behave in social and economic situations, researchers have conducted 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, 15, 14, 7, 18]. Due to logistical constraints, most such studies were limited to the laboratory environment and to a small number of subjects. This introduces a number of biases in the data collected in these studies. For instance, Arnett [2] in a survey of empir-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to publish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

WebSci 2012, June 22–24, 2012, Evanston, Illinois, USA.
Copyright 2012 ACM 978-1-4503-1228-8...\$10.00.

Game	Player	F1	F2	F3	F4	F5	BW	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 different battlefields (represented as F1 to F5) in an example Blotto game. (BW = Battlefields Won)

ical 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.* [12] argue that since human subjects used in most such studies are from Western, educated, industrialized, rich, democratic (WEIRD) countries it would be inappropriate to generalize the findings to people from other societies.

Another drawback of much of the empirical research in behavioral game theory is that it ignores social relations between players [6]. In typical laboratory experiments the subjects are strangers interacting in game theoretic settings. However, in many real-world economic settings, the interaction is between people who know one another, and social relations may affect their behavior.

This paper reports our 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). Blotto is well known in Game Theory (See [1] for further discussion of the game and its origins), and has been used to model political and economic situations. Calculating the equilibrium of this game is a hard problem and so is the choice of the optimal strategy. We analyze how social relations affect the players’ behavior in this game.

Using online social networks for Behavioral Game Theory experiments overcomes many problems associated with laboratory experiments. Popular social networks such as Facebook have users from all over the world¹. This allows researchers to study the effect of regional and cultural differences in the the players’ behaviour. The players participate in their normal ‘habitat’ rather than in an artificial lab setting. Additionally, experiments can be conducted at much larger scales, supporting finer grade results while maintaining statistical significance. And lastly, but most importantly, social networks capture how users are related to each other. This data enables the study of the affect of social relations on the way people play with each other.

THE COLONEL BLOTTO GAME

The Colonel Blotto game was proposed by Borel [5]. It 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 [16]). The mechanics of

¹According to Facebook, 70% of their users are from outside USA. See: <http://www.facebook.com/press/info.php?statistics>.



Figure 1. Choosing a troop allocation and viewing the game result. The figure shows the different 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. (2) The user distributes the troops and makes the move. (3) The opponent does the same and the result is revealed to both players.

the game are as follows. Each of the two players is endowed with a certain number of resources, which we call “troops”. They must simultaneously distribute these troops across a set of “battlefields” (without knowing how their opponent player allocates her troops). A player wins a battlefield if they have more troops there than their opponent. The player who wins in more battlefields wins the game. Examples of some game instances are shown in Table 1.

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 m troops and n battlefields has been investigated analytically in several papers. For a complete discussion of solutions to Blotto games see [19] and [10]. Under one definition, the optimal strategy in the Blotto game 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 $[0, 2m/n]$.

Previous Experimental Analysis

A number of studies have been conducted on how people play different variants of Blotto [3, 1, 8, 13]. However, these are based on a setting where each subject plays the game once and hence has only one distribution of troops. This selected distribution is then compared to the troop distribution of all other players to find rankings for different troop distributions. In contrast, our research focuses on the impact of social relations between participants on the way they play the game. In our application, participants may challenge either people they know or strangers to play the game. Also, in our setting the same pair of players may play the

game many times, allowing them to learn about how each of them play the game and adapt their used strategies over time. This allows us to analyze questions such as whether people play Blotto differently with people they know as opposed to strangers, or whether an initiator of a challenge is more likely to win than the responder.

Our first part of the analysis follows [1], allowing us to compare our finding to previous work. The analysis in [1] examines how people play a Blotto game which has 120 troops and 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 dataset was collected from readers of the Hebrew business daily, Calcalist. In the analysis section we compare our results to those of [1]. We then turn to focus on the impact of social relations on the strategies used.

“PROJECT WATERLOO” ON FACEBOOK

We have developed a game, called Project Waterloo² which allows users of the online social network Facebook to play the Blotto game with friends and strangers. In the general Blotto game, each of the players may be endowed with a different number of troops. However, in our Facebook (Project Waterloo) implementation of the game, each player had the same number of $m = 100$ troops, which are to be distributed across $n = 5$ battlefields. The implementation allows users to play with three types of partners: a) random players whose identity is hidden from the player, b) known players from their friend list or c) players from a general list of players who have played the game before (but are not necessarily their Facebook friends). Figure 1 shows screenshots of the various steps of the Project Waterloo application.

Every instance of the Blotto game in our implementation starts with a player initiating a challenge against another player by distributing her troops across the battlefields. The responding player distributes her troops among the battlefields and finalizes the game. If a player allocates more troops in a battlefield than their opponent does, they win that battlefield. The player who wins more battlefields wins the game.

Most work in behavioral game theory is based on monetary incentives to recruit human subjects. In contrast, and similarly to [1], our methodology is based on an internet game, so we rely on entertainment and reputation incentives to attract users to play the game. The reputation incentive is realized by showing users their rankings (based on performance) among a) all players and b) players in their friend network. Users are ranked according to their rating (R) which is computed as: $R = \frac{\#Games\ Won + 0.5 * \#Games\ Drawn}{\#Games\ Played + 10} \times 100$.

As the $\#Games\ Played$ approaches infinity, the measure converges towards the player’s average score. However, it also encourages subjects to play more games because for two

²Available at <http://apps.facebook.com/msrwaterloo/>.

players with the same average score it prefers the player who has ‘proven’ their skill in more games.

STRATEGIC ABSTRACTION

The number of different (pure) strategies in the Blotto game is enormous. A game with n battlefields and m troops has $\binom{m+n-1}{n-1}$ strategies³, so for our implementation with $m = 100$ and $n = 5$ we obtain $\binom{104}{4} \approx 4.6M$ different strategies. Due to this large strategy space it is not feasible for a computer or human to evaluate all possible strategies, so deciding on a strategy requires abstracting away some details regarding the strategies. Consider how a subject might play Blotto.

One obvious strategy that can serve as a “focal point” is spreading troops evenly across the battlefields. In our games with 100 troops and 5 battlefields, this corresponds to the strategy [20, 20, 20, 20, 20], which we call the “uniform strategy”. If both sides play this strategy, there will be a draw on each battlefield, so the entire game would be a draw. Suppose you believe your opponent is likely to play this strategy. By focusing your troops on four battlefields at the expense of the last battlefield, you would win in each of these battlefields and lose the last battlefield. For example, by playing [25, 25, 25, 25, 0] you would be guaranteed to beat the uniform strategy. However, if your opponent also focuses their troops, you may again lose. For example, if your opponent plays an even more focused strategy such as [34, 33, 33, 0, 0] (i.e. putting troops roughly equally on only three battlefields), they would beat the strategy of [25, 25, 25, 25, 0].

The two previous example strategies focus troops in the left battlefields: the strategy [25, 25, 25, 25, 0] leaves the last battlefield completely unguarded, and [34, 33, 33, 0, 0] leaves the last two battlefields completely unguarded. These unguarded battlefields make easy and lucrative targets. For example, the strategy [35, 35, 28, 1, 1] beats the previous two strategies, by putting some troops on those unguarded battlefields. Seeing this example you might think it unwise to leave a battlefield completely undefended, and propose to slightly change the strategy [34, 33, 33, 0, 0] into [33, 33, 30, 2, 2] (which would defeat [35, 35, 28, 1, 1] by slightly strengthening your allocation to the last two battlefields). But what if your opponent decided to attack these last two battlefields with slightly more troops, for example using [35, 33, 26, 3, 3]? How well-guarded would you like these least-guarded battlefields to be? Further, the above example strategies place more troops on the battlefields to the left. But perhaps it might be a good idea to place more emphasis on the battlefields to the right? Or perhaps the center?

One might hope that the Blotto game has a strategy that defeats all the other strategies. However, this is easy to disprove. If your opponent knows your troop allocation, they can always beat you, by placing a single troop more than

³Note that a strategy can be represented as a binary string with $n-1$ “ones” and m “zeros”, where the number of “zeros” before the first “one” represents the number of troops on the first battlefield, the number of “zeros” between the first “one” and the second “one” represents the number of troops in the second battlefield and so on.

you on your three least-defended battlefields, at the expense of the two remaining battlefields. Thus, if you always use the *same* strategy against your opponent, you might defeat them at first, but once they learn your never-changing strategy, they can easily defeat you. As a consequence, to play the game well over time, you must use some form of randomization.

The Blotto game is complex, and the above discussion shows that there is no universal strategy for winning the game. Still one may expect to gain an insight regarding playing Blotto using Game Theory. Such an analysis would perhaps allow deriving two probability distributions over the strategy space which are the best responses to one another. In other words, if your opponent were to play by choosing her strategy at random from one of these given distributions, your best response would be to choose your strategy at random from the other distribution. However, how could you find these distributions? And even if you could find them, as researchers have done in the past [10], do you have reason to believe that other players would play the game this way? Prior research such as [1] shows that a natural way for humans to tackle the difficulty posed by the large size of the strategy space is to only consider high level features of the various strategies. The brief discussion of strategies above reveals several possible features.

The naïve focal point strategy of $[20, 20, 20, 20, 20]$ assigns each battlefield an equal number of $t = m/n = 100/5 = 20$ troops. To increase the number of troops on one battlefield, a player must decrease the number of troops in another. Following Arad and Rubinstein [1] we call a battlefield to which a player assigns more than m/n troops a *reinforced battlefield*. An obvious high-level feature of a strategy is the number of such reinforced battlefields. For example, the strategy $[25, 25, 25, 25, 0]$ only reinforces the first 4 battlefields, whereas $[5, 5, 30, 30, 30]$ reinforces only the last 3 battlefields. A strategy may never have more than 4 reinforced battlefields (as that requires putting at least 21 troops in all 5 battlefields, which requires at least 105 troops rather than the $m = 100$ troops a player has), and only a single strategy has 0 reinforced battlefields — $[20, 20, 20, 20, 20]$.

As discussed above, since any strategy can be defeated by some other strategy, successfully playing the game over time requires choosing strategies from a certain distribution of strategies. Indeed, the game theoretic analysis of the Blotto game examines mixed-strategy equilibria (i.e. equilibria where each player chooses a strategy from a certain distribution of strategies), as no pure Nash equilibria exist for the game [10, 19]. Such troop distributions may put more emphasis on some battlefields rather than on others. If no special emphasis is placed on any field, the expected number of troops under the distribution must be equal in all the fields. We refer to distributions in which some battlefields have a higher expected number of troops than others as having *position bias*. A simple feature that characterizes the position bias is the mean number of troops in each battlefield under the distribution. Yet another distribution feature of a similar nature is the variance of the number of troops in each battlefield.

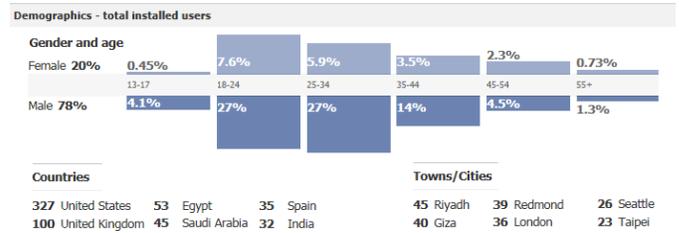


Figure 2. Demographic data on Project Waterloo users.

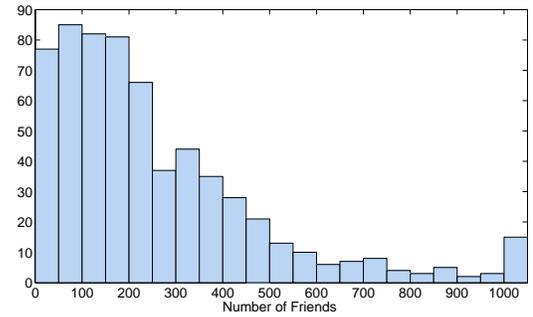


Figure 3. Distribution of the number of friends. The last histogram bucket contains all users that have 1000 or more friends.

The game theoretic analysis of the Blotto game predicts that equilibrium solutions will have no position bias [19, 10] and focuses on mixed strategy equilibria where each player uses a distribution with no position bias. The theoretical analysis in [10] begins by abstracting away specific battlefield allocations and focuses on a game where players’ strategies are *unordered* sets of the number of troops per battlefield, which are then allocated to the specific battlefields by a random permutation. Beginning with the section on common strategies we adopt a similar convention.

In our analysis in the following sections we empirically examine the above-mentioned high-level features in games played by users on our implementation of the Blotto game on Facebook. Specifically we examine the number of reinforced battlefields, the unit digit allocation in the battlefields, the position bias and the variance in the number of troops. We show which features are typically selected by players and correlate these choices to the characteristics of the players, such as their properties in the social network structure.

ANALYSIS OF THE DATA

We now present the results of our analysis of the game data collected from the users of the ‘Project Waterloo’ Facebook application. We first explain which choices users can make regarding their challenged opponents.

The Project Waterloo implementation allowed challenging three types of opponents. One way to initiate a game is when a player challenges one of her Facebook friends. We call such opponents *friend opponents*. Another way is selecting a “Random Game” option from the menu, which allows the player to initiate a game against an unknown opponent. In

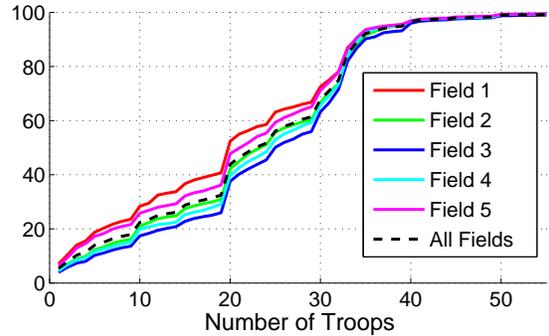
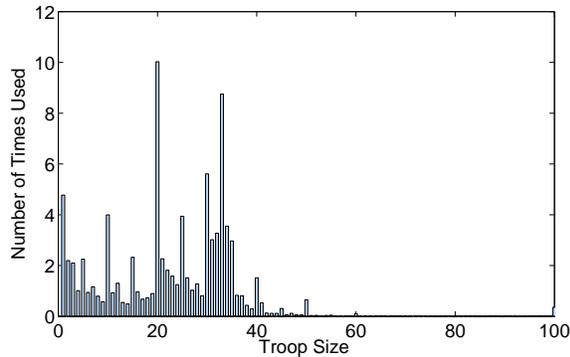


Figure 4. (LEFT) Frequency of troop group size (in percentage of all troop group sizes observed). (RIGHT) Cumulative distributions of field assignments.

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	9	13	6	12	4	4	3	3

Table 2. Distribution of unit digits (in percentages) in all the single-field assignments in the data collected by us and the datasets used by Arad and Rubinstein. Note that in their study the game involved $m = 120$ troops to be distributed among $n = 6$ battlefields, so the numbers are not exactly comparable.

this case, no information is revealed regarding the opponent player (the label only contains the words “Random Opponent”). We call such opponents *hidden opponents* (as all the information regarding the opponent is hidden). The final way to initiate a game is the “Stranger Mode”. When a user selects this option, our platform creates a list of 50 “Project Waterloo” players who are not the Facebook friends of the initiator. It displays their names, images and ranking (based on previous games), and allows the initiator to select one of them as the opponent. The initiator is unlikely to know the chosen opponent personally (as they are not Facebook friends), but does get some minimal information about the opponent — their name, profile picture and how well they have played in the past. We call such opponents *stranger opponents*.

Demographics

The dataset used for our analysis contains 1,883 games played by 632 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 players who were not Facebook friends (i.e. stranger opponents or hidden opponents).

The users of our Facebook application come from a wide range of cities, countries and age groups, and most of them are male.⁴ The gender statistic is interesting because most games on social networks tend to have more female users,

⁴Self-reported information from their Facebook profile.

the popular exception being Texas Hold’em which has more male users [20]. Detailed demographic data for users is shown in Figure 2. In the analysis below, we study if and how the number of friends of a user affects the number of games they play and the number of games they win. The average number of friends of subject in our study was 269 which is higher than 130, which is the average number of friends of all Facebook users. Figure 3 shows the distribution of the number of friends for Project Waterloo users.

Troop Allocations and Strategies

We now analyze the different troop group sizes and strategies used by the players in the Project Waterloo application. We call each troop distribution a strategy. We examined the 3,776 strategies that were submitted by players. Some of the strategies were used more than once (either several times by the same player, or the same strategy submitted by several different players), and in total we had 1,402 unique strategies. While computing this number we counted all the permutations of a troop distribution as a single strategy. Most unique strategies were rarely used (only once or twice) and very few strategies were frequently used (for example submitted over than 20 times). Figure 5 shows how many strategies we had of every given frequency. On the X-axis we have the frequency of a strategy (as measured by the number of times it was submitted), and on the Y-axis we have the number of strategies we had of this frequency. Both the X-axis and Y-axis are logarithmic. Table 4 shows the most popular strategies and the number of times they were used. It can be seen that the frequency of strategies roughly follows a power law, i.e. few strategies are very common, with a long tail of strategies (large number of unique strategies) that are infrequently used.

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 battlefields. 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

%Reinforced	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 3. Proportion of strategies with given number of reinforced battlefields (i.e. battlefields with more than 20 troops).

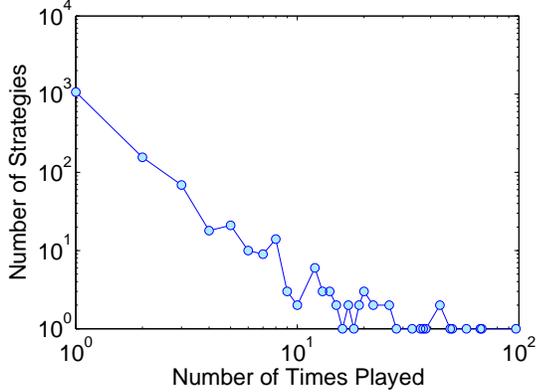


Figure 5. Frequency of strategies: the number of unique strategies of any given frequency.

Strategy					Number of Times Used
34	33	33	0	0	271
20	20	20	20	20	235
33	33	33	1	0	127
33	33	32	1	1	97
35	35	30	0	0	68
100	0	0	0	0	67
35	35	10	10	10	58
25	25	25	25	0	50

Table 4. Most frequently used strategies.

hints about the number of levels of iterated reasoning they are employing while playing the game. Figure 4 and Table 2 provide statistics about the number of troops used by players of the game. 97% of the troop groups used by players had sizes in the interval $[0, 40]$. Recall that the theoretical optimal way to play the game is to choose troop group sizes in the range $[0, 2\frac{m}{n}] = [0, 40]$.

Number of Reinforced Battlefields

Another indicator of the level of reasoning employed by players in the Blotto game is the number of *reinforced* battlefields. A battlefield is called *reinforced* if it has more troops than the uniform allocation of troops, i.e., 20 troops in our case. Our analysis in Table 3 shows that the number of reinforced battlefields in the distributions made by players of Project Waterloo was large. Following the discussion on strategic abstraction, this may be interpreted in favor of the hypothesis that players of the game are employing multiple levels of iterative reasoning.

Relation between strategies

To analyze the winning and losing relations between different strategies we computed the scores of some frequent strategies when they are played against all the possible permutations of other strategies. The result of the analysis are

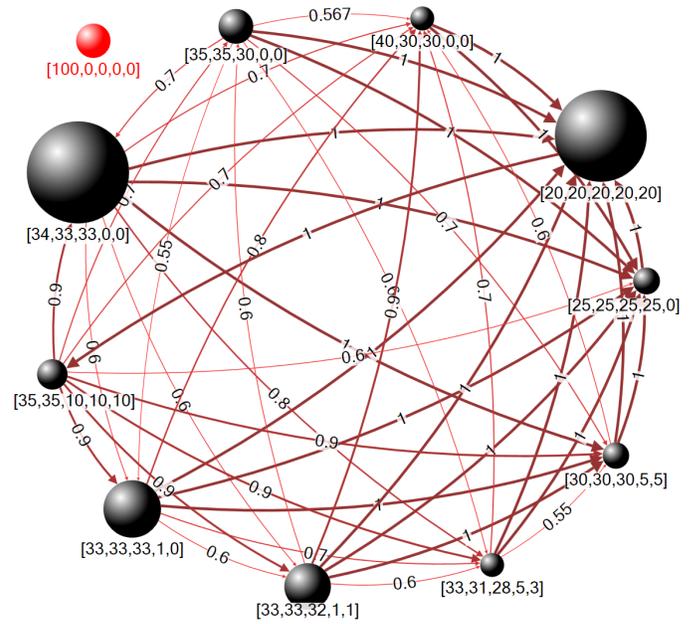


Figure 6. The strategy map. The figure shows the most frequently used strategies and their winning and losing relations. Different strategies are represented by circles. The area of each circle is proportional to the number of times the corresponding strategy was used. An edge labelled with the value x from a strategy to another strategy indicates that the ‘from’ strategy beats the ‘to’ strategy in a fraction x of the possible permutations of the two strategies. Note that $\{100,0,0,0\}$ is always beaten by all the other strategies and the arrows are left out for clarity.

shown in Figure 6. A first interesting observation is that — somewhat surprisingly — the $\{100,0,0,0,0\}$ strategy, which is beaten by almost every other strategy, is played by a significant number of players. Note that the strategy map nicely illustrates the intransitivity of the relation A beats B , which holds when the number of permutations of battlefields in which A wins over B is greater than the number of permutations in which B wins over A . Observe that there are almost deterministic cycles such as: $\{34, 33, 33, 0, 0\}$ beats $\{20, 20, 20, 20, 20\}$ beats $\{35, 35, 10, 10, 10\}$ beats $\{34, 33, 33, 0, 0\}$.

These cycles are also present in non-transitive dice games, see, e.g., [9]. Assuming stationarity of the strategy distribution, knowledge of this plot for a given player population would enable a player entering that population to choose a strategy that would be most successful given the particular strategies and their frequencies played in that population.

Best Strategies

For each strategy we computed the proportion of other strategies it defeats. The top strategies are listed in Table 5 where strategies which were submitted more than once are counted multiple times (multiple counting).

Table 5 shows that the best performing strategies had three reinforced battlefields, typically where two were very strongly reinforced and one was reinforced but not quite as strongly. In these strategies the unit digit of the non-reinforced battle-

Strategy	Win Proportion	Frequency
[36, 35, 24, 3, 2]	0.74	1
[37, 37, 21, 3, 2]	0.73	17
[35, 35, 26, 2, 2]	0.73	1
[37, 36, 24, 2, 1]	0.73	1
[37, 36, 21, 5, 1]	0.72	1
[36, 34, 28, 1, 1]	0.71	2
[35, 35, 28, 1, 1]	0.71	2
[37, 33, 21, 5, 4]	0.71	1
[34, 34, 25, 4, 3]	0.71	1
[37, 37, 22, 2, 2]	0.71	3
[35, 34, 25, 3, 3]	0.70	3
[34, 34, 28, 2, 2]	0.70	3
[37, 37, 23, 2, 1]	0.70	2
[35, 35, 24, 3, 3]	0.70	13

Table 5. Top performing strategies (multiple counting)

Strategy	Win Proportion	Frequency
[36, 35, 24, 3, 2]	0.74	1
[36, 34, 28, 1, 1]	0.73	2
[37, 36, 24, 2, 1]	0.73	1
[34, 34, 30, 2, 0]	0.73	2
[34, 33, 32, 1, 0]	0.72	26
[35, 33, 29, 2, 1]	0.71	1
[34, 31, 31, 2, 2]	0.71	7
[37, 33, 30, 0, 0]	0.71	8
[36, 27, 37, 0, 0]	0.71	4
[33, 31, 31, 3, 2]	0.71	2
[35, 35, 28, 1, 1]	0.70	2
[37, 37, 21, 3, 2]	0.70	17

Table 6. Top performing strategies (unique counting)

fields was between one and five. Table 6 lists the top performing strategies under unique counting (i.e. several wins against the same strategy count as a single win).

Table 6 shows that strategies that do well in terms of defeating unique strategies rather than winning the most games still typically have 3 reinforced battlefields, but in many cases all of these are strongly reinforced (rather than only 2). Also, the unit digit in some of these is 0 (i.e. some battlefields are completely unguarded). Interestingly, quite a few of these strategies are very rare, with very few users playing them (and only a few times). However, there are examples of top performing strategies that are quite common, indicating that some “focal point” strategies can be very strong strategies.

EFFECTS OF BIAS, SOCIAL HABITAT, AND KNOWLEDGE OF OPPONENT

Having explained the theoretical underpinnings of Blotto games and shown the results of our empirical analysis on how users on Facebook play these games, we now investigate the presence of biases and effects of social habitat in game play.

Broken Symmetry: Positional Bias in Troop Allocations

As mentioned earlier, 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 battlefields introduces a bias in the troop allocations [1]. Arad and Rubinstein [1] have con-

cluded that this may be due to a player’s instincts to over-assign troops leading to ‘residual’ allocations at the fringe battlefields. We observe this bias in the data collected by us (see Table 7 and Figure 4).

On average, players of our Facebook application place more focus on some battlefields. Table 7 shows that more emphasis is placed at the center. It shows that there the mean and median number of troops are quite different across fields. We have used a Mann-Whitney-Wilcoxon test ⁵ (MWW test for short) to examine whether the differences in means across battlefields are statistically significant. Each row and column in Table 8 represents a battlefield, and the table shows “True” if the difference between the mean number of troops in the row battlefield and the mean number of troops in the column battlefield is statistically significant, at a significance level of $p < 1\%$ (and “False” if the difference in means is not significant at the level of $p < 1\%$).

Table 7 and Table 8 provide strong evidence for a significant bias towards certain battlefields when allocating troops. The focus appears to be on the center rather than the sides. Interestingly, it appears the difference between fields 2 and 4 (the center left and center right) is small and insignificant, but the difference between fields 1 and 5 (extreme left and extreme right) is significant. Apparently, more troops are allocated to the extreme right ⁶.

Another interesting question is whether the people vary the number of troops they put in some battlefields more than in others. More precisely, we wish to know whether the variances in the numbers of troops are significantly different across battlefields. Table 7 shows the variances in troop allocations in each battlefield.

Table 7 indicates slight differences in variance. For example, variance in field 1 (extreme left) appears to be quite higher than in the other fields. To test whether these differences are statistically significant we used Levene’s test. Each row and column in Table 9 represents a battlefield, and the table shows “True” if the difference between variances in troop allocation in the row battlefield and column battlefield is statistically significant, at a significance level of $p < 5\%$ (and “False” if the difference in variances is not significant at the level of $p < 5\%$).

Table 9 shows that in many cases the differences in variances are significant. For example, Battlefield 1 has a significantly higher variance in troops allocated to it than other fields. However, these differences seem quite small.

Effect of Knowledge About the Opponent

⁵This test is also known as the Mann-Whitney test or Wilcoxon’s rank-sum test.

⁶It would be interesting to see whether this left versus right bias is flipped in some locations. For examples, participants coming from countries where the writing direction is right-to-left rather than left-to-right, such as the middle-east, might have a reverse bias. However, in the interest of privacy, we did not store such information regarding each participant, and only examined the aggregate demographics information available from Facebook.

Opponent	Field 1	Field 2	Field 3	Field 4	Field 5
	Mean, Median				
H	19.3, 20	20.7, 20	22.0, 20	19.9, 20	18.1, 20
S	17.9, 20	20.0, 21	22.9, 27	20.7, 23	18.4, 20
F	17.8, 18	21.3, 22	21.5, 22	20.8, 23	18.7, 20
NF=H+S	18.8, 20	20.4, 20	22.4, 22	20.2, 20	18.2, 20
All	17.8, 20	20.7, 21	22.1, 24	20.8, 22	18.6, 20
	Variance				
All	225.4	199.5	201.9	189.4	192.2

Table 7. Position bias in the allocation of troops to different battlefields. Mean and variances in troop allocation across battlefields. (H = Opponent is hidden from player, S = Opponent is a stranger, F = Opponent is a Facebook friend of the player)

	Field 1	Field 2	Field 3	Field 4	Field 5
Field 1	-	True	True	True	True
Field 2	True	-	True	False	True
Field 3	True	True	-	True	True
Field 4	True	False	True	-	True
Field 5	True	True	True	True	-

Table 8. Significance of position bias (difference in means) across battlefields.

Playing the Blotto game involves reasoning about the strategy of your opponent. It is thus reasonable to assume that knowledge of the opponent in the game would affect the way in which 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 their game opponents or are their friends. Deploying the game on an online social network provides access to the friend relationships of the players.

Troop Distributions against Friends and Strangers Table 7 and Table 3 show that the distribution of troops selected by the subjects is somewhat different in situations when the opponent is “hidden” (i.e. either a random opponent whose identity is completely unknown to the initiator, or an opponent who is not a Facebook friend of the initiator but whose name and profile picture are revealed) as opposed to situation where the opponent is a Facebook friend. An MWW test shows that the differences in the mean number of troops is statistically significant at the $p < 5\%$ level only for the second and fourth battlefield. This provides some evidence that players use different strategies when playing against friends than when playing against strangers. One possible explanation for this is that when playing against a friend players are more likely to play several games against the same opponent. In this case, the players may examine the previous games and select an appropriate strategy based on what they have learned from previous games. On the other hand, when requesting a “hidden opponent” game (one of the game modes in our “Project Waterloo” platform), an opponent is selected at random from the set of all players. In this case the initiator has no knowledge of the opponent, and is likely to choose a “generic” strategy that she believe is suitable to a random opponent. We have also tried to see whether the variances of troop allocation in each battlefield is different when playing against friends and non-friends, however the test showed no significant difference in the variance between these two conditions.

	Field 1	Field 2	Field 3	Field 4	Field 5
Field 1	-	True	True	True	True
Field 2	True	-	False	False	True
Field 3	True	False	-	False	True
Field 4	True	False	False	-	True
Field 5	True	True	True	True	-

Table 9. Significance of difference in variances across battlefields .

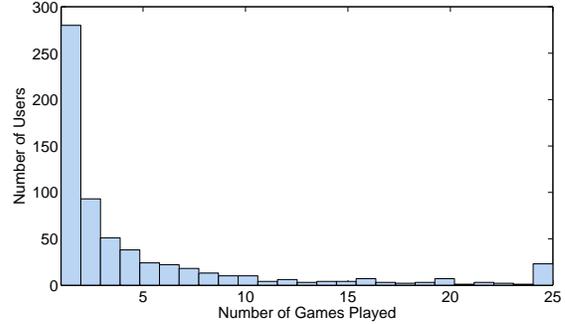


Figure 7. Distribution of number of games. The last histogram bucket represents all users that have played more than 25 or more games.

Winning Against Friends and Strangers One key question is whether the knowledge of a player about their opponent affects their chances of winning the game. As discussed in the analysis section, we have considered three categories — friend opponents, hidden opponents and stranger opponents. We have first examined players who have initiated games in *all* these categories. For these players, the mean score when initiating a game against friend opponents was 53.02%, whereas their mean score when playing against hidden opponents was only 40.76%. This difference in probabilities is quite large, and an MWW test shows that this difference is significant at the level of $p < 10\%$, providing some evidence that indeed players are more likely to win against a hidden opponent.

One possible explanation for this effect is that the initiators decide who to invite, and use the knowledge of their opponents to challenge more predictable players. They can then use a strategy tailored to defeating these opponents.

We also examined *aggregate* statistics regarding games initiated against friends and strangers. In total our dataset contained 524 games initiated against friends, 332 initiated against strangers and 1,027 against hidden opponents. The probability of the initiator winning against friends was 51.34%, only 47.89% against strangers and 38% against hidden opponents. This also supports the conjecture that initiators manage to find opponents that they find easy to defeat (note that when initiating a game against a hidden opponent, the hidden opponent *does* get to examine information regarding the initiator).

Effect of social network on games played and won

Figure 7 shows that the number of users which play a specific number of games obeys a rough power law distribution.

Threshold T	Users < T friends (Average #Games)	Users ≥ T friends (Average #Games)
100	162, (6.97)	470, (5.61)
300	428, (5.96)	204, (5.95)
500	556, (6.28)	76, (3.64)

Table 10. The table shows the number of users which have a specific number of friends (controlled by threshold T) and the average number of games played by these users.

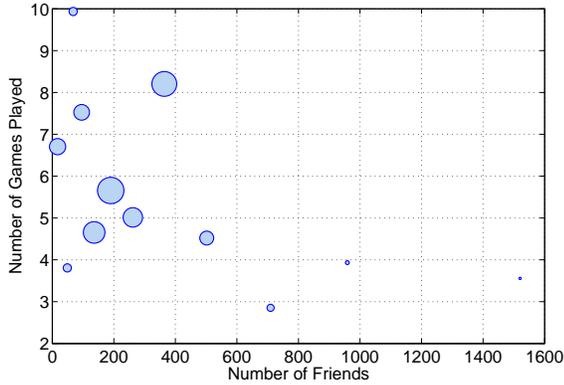


Figure 8. The average number of friends and games played by some user groups. The groups were constructed by clustering (grouping) users with similar numbers of friends. The size of the group can be seen by the radius of the circle.

Threshold T	Users < T friends (Average Score)	Users ≥ T friends (Average Score)
100	162, (3.52)	470, (2.79)
300	428, (3.00)	204, (2.93)
500	556, (3.17)	76, (1.59)

Table 11. The table shows the number of users which have a specific number of friends and the mean score of these users. The mean score is shown in brackets.

We first investigated whether the number of friends of a user affects the number of games played by them. We observed that users with more friends on Facebook in general played less games on our application. The results are summarized in Table 10 and Figure 8. It can be observed that people with many friends tend to play relatively fewer games on average. We performed a similar analysis with the number of games won by users. The results are shown in Table 11.

Effect of experience on winning

In most games of skill one would expect the players to become better as they play more games. Figure 9 shows the winning rates and games played for all users. It can be seen that the average winning rate of player groups increases as they play more games.

DISCUSSION AND FUTURE WORK

Our data gathered from behavioral game theory experiments conducted using online social networks is somewhat consistent with previous studies, indicating that this research tool may allow overcoming the scalability limitation of laboratory experiments. More importantly, we have shown that this new source of data can help uncover interesting rela-

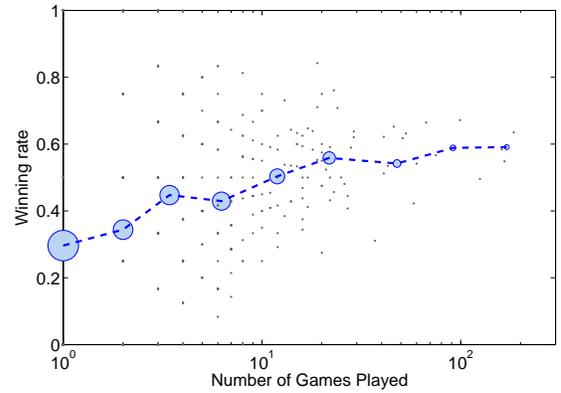


Figure 9. The winning rates and games played for all users. The winning rate was computed as: $\frac{\#wins+0.5*\#draws}{\#games-played}$. The figure shows also how the average winning rate of groups of users (grouped by number of games played) changes with the average number of games played in the group.

tions between player strategies and the game context. We found that the strategies adopted by players change if their opponents are their friends. Our results suggest that players with few Facebook friends tend to play more games and have higher probabilities of winning, that players responding to a challenge in the game have higher probabilities of winning than those initiating the game, and that the initiators of a game have higher probabilities of defeating their friends than strangers.

There are a number of interesting issues that remain unexplored. Questions on how users learn and change their strategies with more games, how their play is affected by their age, gender, location, etc., are promising directions for future work. To conclude, we hope that our successful demonstration of how a large-scale behavioural game theory experiment can be carried on an online social network will motivate others to conduct similar empirical studies using this medium. We believe that online social networks have the potential to become a very useful resource for empirical research in fields such as behavioral game theory and experimental economics.

ACKNOWLEDGEMENTS

We would like to thank Gavin Smyth and Michael Armstrong for their contributions in designing and implementing the Project Waterloo game Facebook application, and Ayala Arad for valuable discussions on game design and research questions.

REFERENCES

1. Arad, A., and Rubinstein, A. Colonel Blotto’s top secret files: Multi-dimensional iterative reasoning in action. *Working paper* (2010).
2. Arnett, J. J. The neglected 95%: Why American psychology needs to become less American. *American Psychologist* 63 (2008), 602–614.

3. Avrahami, J., and Kareev, Y. Do the weak stand a chance? Distribution of resources in a competitive environment. *Cognitive Science* 33, 5 (2009), 940–950.
4. Bachrach, Y., Kohli, P., and Graepel, T. Rip-off: Playing the cooperative negotiation gam. In *AAMAS* (2010).
5. Borel, É. The theory of play and integral equations with skew symmetric kernels (La théorie du jeu et les équations intégrales à noyau symétrique). *Econometrica* 21 ([1921] 1953), 97–100.
6. Camerer, C. *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton University Press Princeton, NJ, 2003.
7. Chakraborty, T., Judd, J. S., Kearns, M., and Tan, J. A behavioral study of bargaining in social networks. In *ACM Conference on Electronic Commerce* (2010), 243–252.
8. Chowdhury, S. M., J., D. K., and Sheremeta, R. M. An experimental investigation of colonel blotto games. *Economic Theory* (2011), 1–28.
9. Finkelstein, M., and Thorp, E. O. Nontransitive dice with equal means. In *Optimal Play: Mathematical Studies in Games and Gambling*, S. Ethier and W. Eadington, Eds., Institute for the Study of Gambling and Commercial Gaming (2000), 293–310.
10. Hart, S. Discrete Colonel Blotto and General Lotto games. Discussion paper series, Center for Rationality and Interactive Decision Theory, Hebrew University, Jerusalem, October 2006.
11. Henrich, J., Boyd, R., Bowles, S., Camerer, C., Fehr, E., Gintis, H., and McElreath, R. In search of homo economicus: behavioral experiments in 15 small-scale societies. *American Economic Review* 91, 2 (2001), 73–78.
12. Henrich, J., Heine, S. J., and Norenzayan, A. Most people are not WEIRD. *Nature* 466 (2010), 29.
13. Hortala-Vallve, R., and Llorente-Saguer, A. A simple mechanism for resolving conflict. *Games and Economic Behavior* 70, 2 (2010), 375–391.
14. Judd, J. S., and Kearns, M. Behavioral experiments in networked trade. In *ACM Conference on Electronic Commerce* (2008), 150–159.
15. Kearns, M., Suri, S., and Motfort, N. An experimental study of the coloring problem on human subject networks. In *Science* (2006).
16. Myerson, R. B. Incentives to cultivate favored minorities under alternative electoral systems. *The American Political Science Review* 87, 4 (1993), 856–869.
17. Oosterbeek, H., Sloof, R., and Van De Kuilen, G. Cultural differences in ultimatum game experiments: Evidence from a meta-analysis. *Experimental Economics* 7, 2 (2004), 171–188.
18. Ray, D., King-Casas, B., Montague, P., and Dayan, P. Bayesian model of behaviour in economic games. In *NIPS* (2008).
19. Roberson, B. The colonel blotto game. *Economic Theory* 29, 1 (2006), 1–24.
20. Su, S. Who’s using facebook’s top apps? demographic data indicate diverse audiences. *InsideFacebook, June 1st* (2010).