

An Experimental Study of Classic Colonel Blotto Games

Kevin Modzelewski

Joel Stein

Julia Yu

MIT

6.207/14.15

December 2009

1. Introduction

1.1 Colonel Blotto Games

Introduced in 1921 by Borel, Colonel Blotto games are a class of two-person constant-sum games in which each player has limited resources to distribute onto n independent battlefields without knowledge of the opponent's actions. The player who allocates a higher number of resources on battlefield i wins that battlefield and its associated payoff. Both players seek to maximize the expected number of battlefields won. Games within this class often differ in parameters and rules; common variations are described in Table 1.1. The classic Colonel Blotto game, the chief version this project is concerned with, has symmetric player resources, discrete distribution, ordered battles, and battlefield wins determined by higher resource count.

TABLE 1.1: Example Variations within Class of Colonel Blotto Games.

<i>Parameter</i>	<i>Variation</i>	
Player Resources	Symmetric	Asymmetric
Distribution of Resources	Discrete	Continuous
Winning	(Auction) More resources in a battlefield win with certainty	(Lottery) Probability is equal to the ratio of a player's resources to sum of his allocated resources on that battlefield
Order of battle	Battlefields have a numbered order for battle	Each player has own fields that are randomly selected to be compared

Borel and Ville (1938) first solved the special case of $n = 3$ battlefields and symmetric resources, then Gross and Wagner (1950) extended it to any finite number of battlefields. These Nash equilibriums are pairs of n -variate distributions. Friedman (1958) published a partial characterization for n battlefields and asymmetric resources. Although many applications and extensions of Colonel Blotto were studied, there was a lull in significant developments until Roberson in 2006. Roberson established new and novel solutions that do not utilize the regular

n -gons of Borel, Gross, and Wagner's papers. Roberson applied the theory of copulas "to prove the uniqueness of Nash equilibrium payoffs for n battlefields and asymmetric resources and to prove that uniform univariate marginal distributions are necessary for equilibrium over a wide range of endowments of resources."

1.2 Applications

The Colonel Blotto game is a fundamental model for multidimensional strategic resource allocation, thereby widely applicable in fields from operations research (Bellman, 1969), to advertising (Friedman, 1958), to military and systems defense (Shubik and Weber, 1981). In conflict on multiple fronts, oftentimes a defending player needs to successfully win all battles, whereas the attacker only needs to win at one place. Colonel Blotto payoffs can be modified to accurately reflect these scenarios, e.g. war, supply chain, immune system (Clark & Konrad, 2007). One-dimension models, most frequently the all-pay auction, appears in economics to model political campaigns, lobbying, litigation, research and development races (Roberson, 2006). Also of interest is Laslier and Picard's redistributive politics model (2002), which employs the Colonel Blotto game with asymmetric resources to simulate political candidates who simultaneously announce their fixed budget allocations and citizens who vote according to their higher utility.

1.3 Experimental Study

Given the extensive mathematical analyses, practical applications, and simple setup of the game, we are surprised to find very few experimental studies on pure Colonel Blotto games. The main existing research is a working paper by Chowdhury et al. (2009).

Chowdhury investigated experimental versus theoretical equilibriums in Colonel Blotto games with discrete distribution, asymmetrical resources, and both auction and lottery style

winning. Trials were further separated into partnered and random matches. Experimental results supported all major theoretical predictions. Players also exhibited high serial correlation in resource distribution across time when playing in random matches, but significantly less when repeatedly playing with a partner.

Hence is our motivation for this experimental study. We parallel Chowdhury's research to investigate the evolution of player strategies – or lack thereof – in both repeated and random matching play of the classic version of Colonel Blotto. We utilize an original experimental game design, as detailed in Section 2.

2. Experimental Environment

2.1 Experimental Design

The experimental Colonel Blotto game consists of five battlefields and two players with symmetric resources of 50 units each. The players distribute their own units without knowledge of the other, trying to maximize the expected number of battlefields won. The player with a

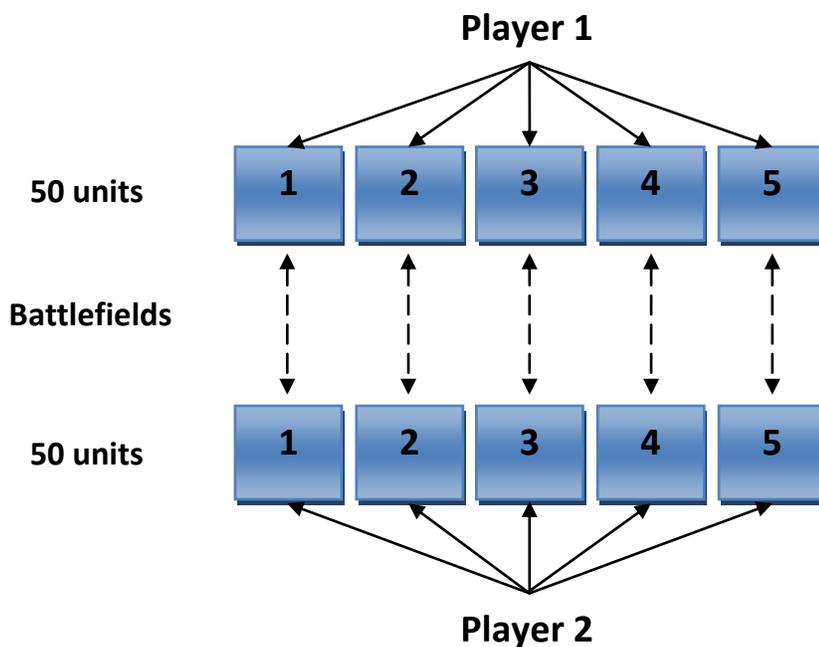


FIGURE 2.1: Colonel Blotto Game Structure.

higher unit allocation to a battlefield wins that battlefield with certainty; ties are possible. In this type of game play, there is no pure strategy Nash equilibrium (Roberson 2006).

2.2 Experimental Procedures

The experiment game was created for simple online play. Contestants played 1v1 games against each other using an online interface allowing players to place armies in battlefields and see their history of previous games. Contestants had a timer with their remaining time for the game, but could choose to submit their actions early. The website was created with HTML and Javascript using the Django framework. The experiment consisted of three independent sessions and the 24 participants were randomly assigned.

In each session, eight participants play in a round-robin tournament of five games per competitor. Specifically, each subject engages in best of five games, collectively termed as one round, with Opponent 1 before being matched with Opponent 2 through Opponent 7, playing 35 games total. Each game has a 45 second time limit and a leader board displays each player's number of round wins, ties, and losses. Through this design, we are consequently able to study formation of individual strategies when in repeated play with one partner as well as when matched against other new players.

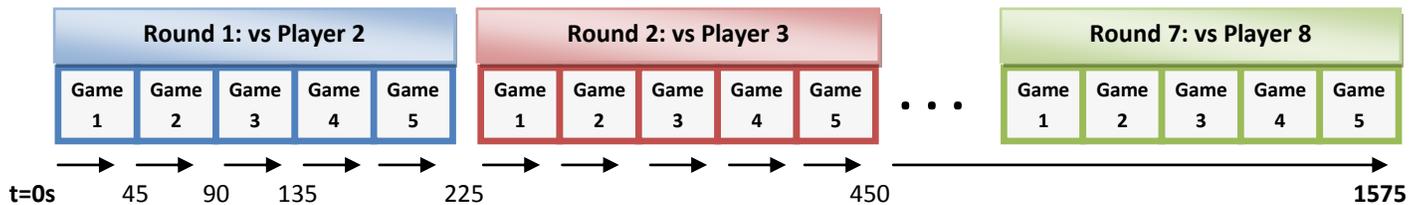


FIGURE 2.2: Order of Game Play for Player 1.

2.3 Theoretical Predictions

For the classic Colonel Blotto game, the unique symmetric Nash equilibrium is for each player to allocate his units by the univariate distribution function (Roberson 2006):

$$\forall j \in \{1, \dots, n\} \quad F^j(x) = \frac{x}{\frac{2}{n}X} \quad x \in \left[\frac{2}{n}X \right]$$

n = number of battlefields (greater than 2), *X* = total number of units

In this case, *n* = 5 and *X* = 50 so the theoretical equilibrium is characterized by

$$F(x) = \frac{x}{20} \quad x \in [0, 20]$$

We hypothesize that the evolution of strategies when an individual is playing against a familiar opponent (i.e. within a round of five games) is different from when playing against a new opponent (i.e. game 1 of new rounds) as well as divergent from the theoretical predictions of uniformity. Specifically, we hypothesize that allocation strategies of players will be correlated between games, implying that players strategies remain somewhat consistent, even in repeated play against the same opponent. We expect this correlation to be different depending on whether the player won or lost the previous game.

3. Results

3.1 General Results

Table 3.1 demonstrates that a player's average allocation of resources to Battlefields 1-4 does not vary much from the expected mean of 10 ($p > 0.05$) although middle battlefields are slightly favored, but there is a distinct heavy allocation bias against Battlefield 5 ($p < 0.01$).

TABLE 3.1: Average Player Resource Allocation per Battlefield

Battlefield	Resource Allocation
1	9.57
2	10.42
3	10.44
4	10.13
5	7.61

Chowdhury’s results also noticed an allocation bias to its leftmost boxes, albeit very minor in their case, and attributed it to Schelling’s “focal points” theory (1960). Similarly, in our case, all participants were native left-to-right language speakers, and therefore, when battlefields are aligned horizontally in a row, the leftmost battlefield was the natural starting point. Allocation typically begins with large values and decreases until all resources are employed, so distribution in a left to right direction would frequently leave Battlefield 5 with a low number of units.

Figure 3.2 provides more detail in the form of allocation histograms for each battlefield. Battlefield 5 is readily seen to have left skewness. Overall, experimental behavior corresponds very well to the Nash equilibrium predictions, given the relatively small number of games run in this study. In theory, players will employ a joint distribution which generates a uniform marginal distribution over the interval $[0, 20]$ for each battlefield. From the experiment, only 1.9% of allocations were above 20 units. Of that minority, 51.3% were allocations of size 21 units.

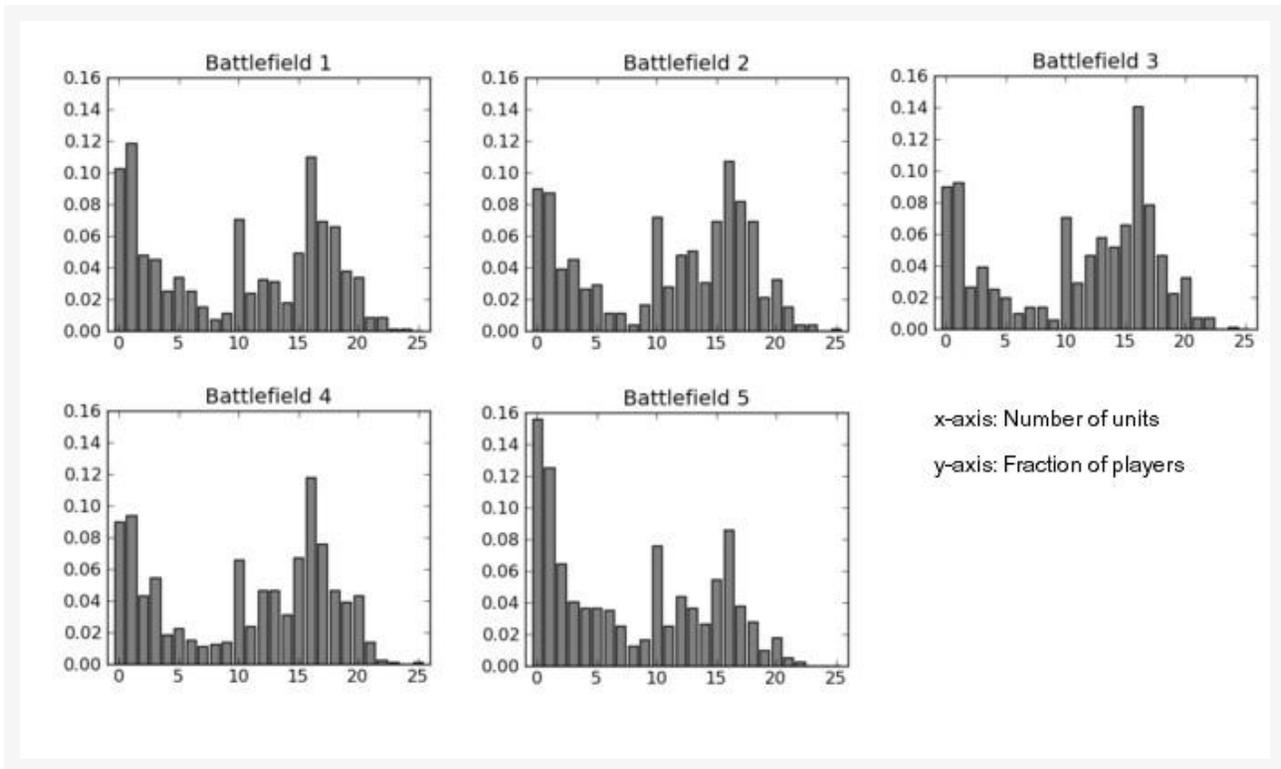


FIGURE 3.1: Distribution of Units across Battlefields. This figure illustrates how players distribute their resources across the five battlefields.

From Figure 3.2 and numerous ones following, it is apparent that player behavior tends to be quite dichotomous, even if allocations average out to the expected value. Players allocate either a very low or high number of resources to a battlefield, essentially sacrificing certain battlefields in order to more certainly win others. Through intuition, players normally used a “Big-3” strategy, choosing three high (≥ 15) and two very low (≤ 2) numbers to distribute.

The top favorite set of allocations, irrespective of battlefield order, was, by far, $\{0, 0, 16, 17, 17\}$. Three of the top four sets in Table 3.2 are Big-3. The uniform set of $\{10, 10, 10, 10, 10\}$ is clearly inferior, but we conclude that players will try this intuitively obvious option at least once during their time of play. Other common allocations were mostly a variant of the favorites, like $\{0, 0, 16, 16, 18\}$. Big-3 strategies are not necessarily the best; their number of wins is only half a standard deviation more than the expected number. Big-2 strategy wins, on

the other hand, are more than one deviation above, although Big-2 is not used often. The most utilized specific Big-2 set is {3, 3, 4, 20, 20}, a count of twelve times. Most of the top five players by accumulated score, however, employed Big-2 strategies more frequently than the average player did. Accumulated scores are calculated as +2 for game win, +1 tie, +0 loss, over the 35 games total.

TABLE 3.2: Favorite Sets of Allocations Employed

Set of Allocations	Number of Instances Used
{0, 0, 16, 17, 17}	65
{1, 1, 16, 16, 16}	41
{10, 10, 10, 10, 10}	37
{1, 1, 15, 16, 17}	24

In the following sections, we look at detailed changes in allocation strategy across games and among players through order statistics, equilibrium behavior, and serial correlations.

3.2 Order Statistics

In order to further validate our general results, we choose to look at order statistics of player allocations for a given round: the maximum, minimum, and medium. We expected that the median would show the prevalence of Big-2 or Big-3 strategies, and the minimum and maximum would show a preference to the extremes.

Figure 3.2 exhibits the distributions of medians across all tournaments, and broken down by round. If players allocated according to the equilibrium strategy of a uniform distribution in each battlefield, we would expect that the median would follow a symmetric distribution with mean 10. What we observe, however, is a highly skewed distribution with mean 11.27. This is significantly different ($p < 0.01$) from the assumption that the mean is 10. The distribution of

medians implies that the majority of plays follow a Big-3 strategy, with a minority playing Big-2, confirming our analysis.

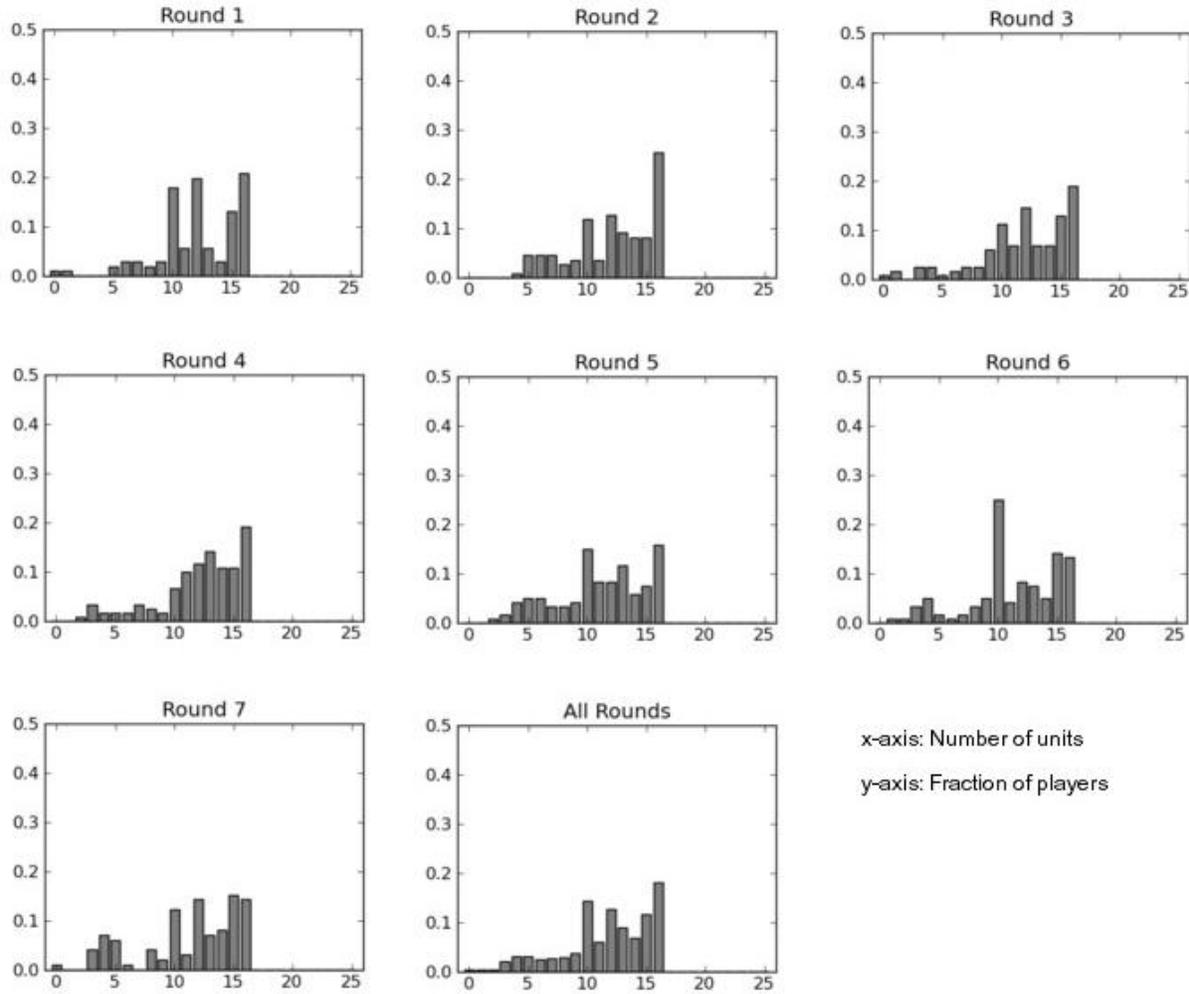


FIGURE 3.2: Distribution of the Median across Rounds. This figure illustrates how the minimum unit allocation of the player set of five allocations is distributed across the tournament.

Figure 3.3 shows the distributions of minimum allocations across the tournaments. The distribution of minimums follows roughly the expected distribution, except for the larger-than-expected presence of people playing a minimum of 10. This implies that a significant percentage of plays are $\{10, 10, 10, 10, 10\}$ (remember this was a top four favorite strategy), while under the

equilibrium strategy, that would be very unlikely to result. Additionally, the mean of the minimums is 1.91, which is significantly different ($p < 0.01$) from the expected mean of 3.33.

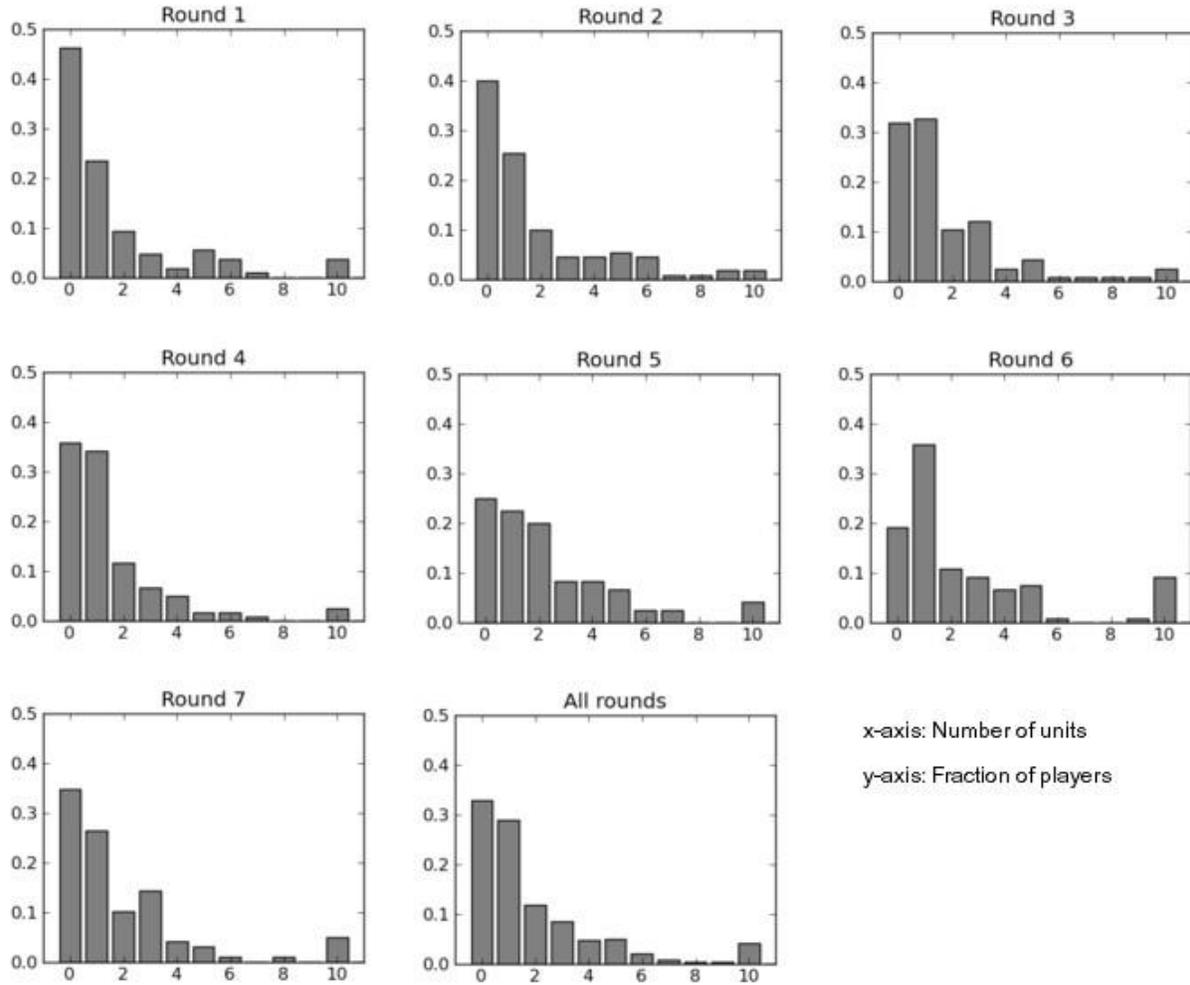


FIGURE 3.3: Distribution of the Minimum across Rounds. This figure illustrates how the minimum unit allocation of the player set of five allocations is distributed across the tournament.

Finally, Figure 3.4 exhibits the distribution of maximums. This distribution is also quite singular from equilibrium predictions since more players allocate above 20 units and fewer allocate 19 or 20 than expected. The mean of the maximums is 16.94, which is significantly

different ($p < 0.05$) from the expected mean of 16.67. The fact that the mean of the minimums is smaller than expected, and the mean of the maximums is larger, implies that players have a tendency to allocate towards the extremes more than expected by equilibrium play, again reconfirming our analysis.

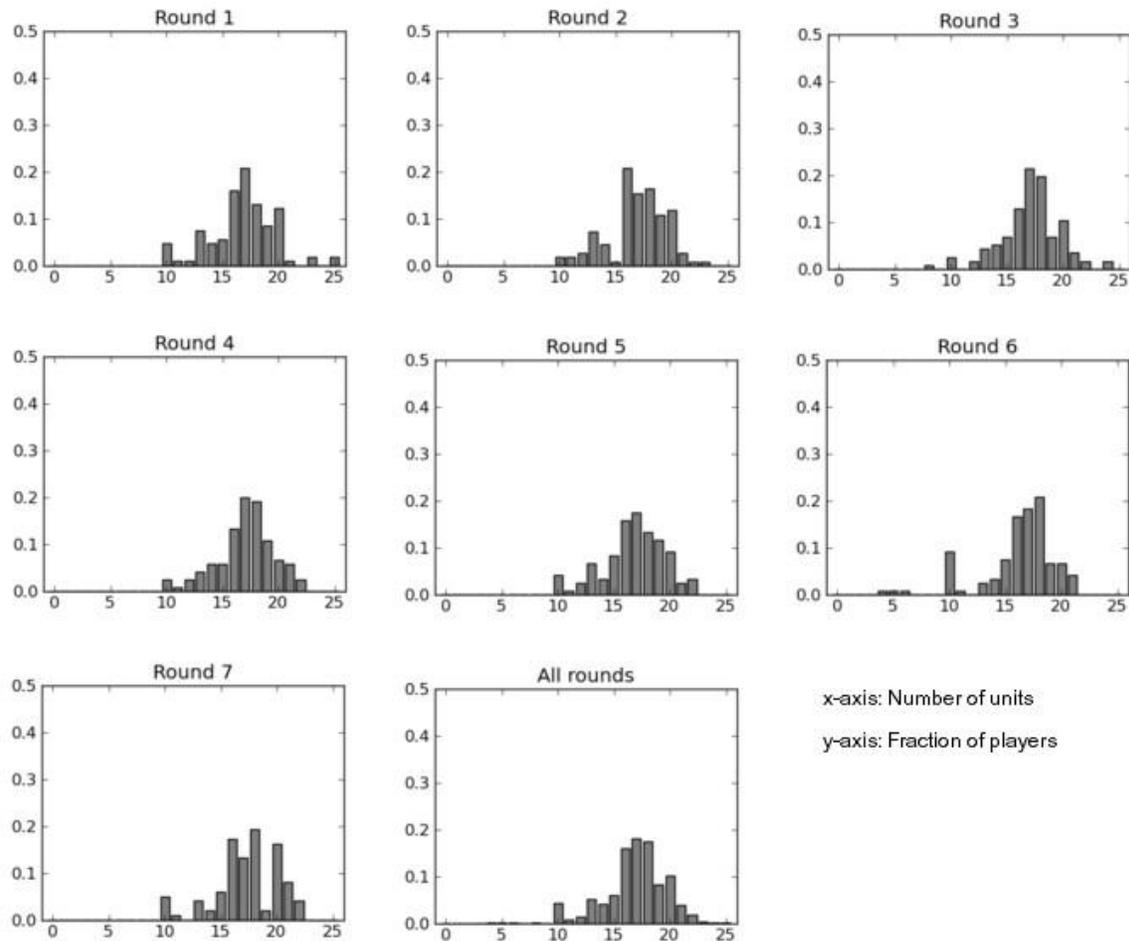


FIGURE 3.4: Distribution of the Maximum across Rounds. This figure illustrates how the maximum unit allocation of the player set of five allocations is distributed across the tournament.

3.3 Equilibrium Behavior

Since deviations from equilibrium behavior were evident, we first investigated if success in the tournament correlates with following the equilibrium strategy. We assessed the two main

aspects of equilibrium behavior that players did not follow: averaging the same number of units in each battlefield, and uniformly allocating between 0 and 20 resources. To measure how well players mix between the five different battlefields, we calculated the sum of squares of the difference between the fraction of units that they placed in the battlefield, and 0.2, the theoretical equilibrium number to place. Similarly, to measure how well players mix their allocations, we calculated the sum of squares of the difference between the fraction of times that the player chose each allocation and $1/21$, the theoretical result.

With these metrics of how well players mixed, we correlated them with the average number of points won per game (scoring as defined previously). (We used this measure, as opposed to the total-number-of-rounds-won metric, in order to get finer-grained results.) There is a correlation coefficient of -0.26 between winning and mixing between battlefields, and -0.08 between winning and mixing the allocated number of units. Neither of these values are significant at the 95% level. This has a few interpretations: that people are behaving irrationally enough that the equilibrium behavior is not optimal, or perhaps, that we are only measuring aggregate behavior, in which case, it does not tell us if people are truly mixing, only that it averages out in the end.

Next, we analyzed whether player behaviors approach equilibrium as the tournament progressed and subjects played more games. Figure 3.5 depicts the fraction of their units that players place in the five battlefields, and how this allocation varies over the rounds. The first five games for a player, Round 1, exhibit a strong aversion to the last battlefield, averaging only 6.5 resource units instead of the expected 10 units. This is concurrent with previously discussed general results, but more interestingly, this is also significantly less than what players allocate to Battlefield 5 in the rest of the rounds, an average of 8.1 units ($p < 0.05$). No other battlefields

demonstrated any significant differences across the seven rounds at the 95% level. Therefore, our reasoning is that players soon notice the disfavored Battlefield 5 and learn to place more resources there, but through the length of this tournament, still do not entirely overcome the bias. This demonstrates the player ability to perceive weaknesses and modify their strategies accordingly.

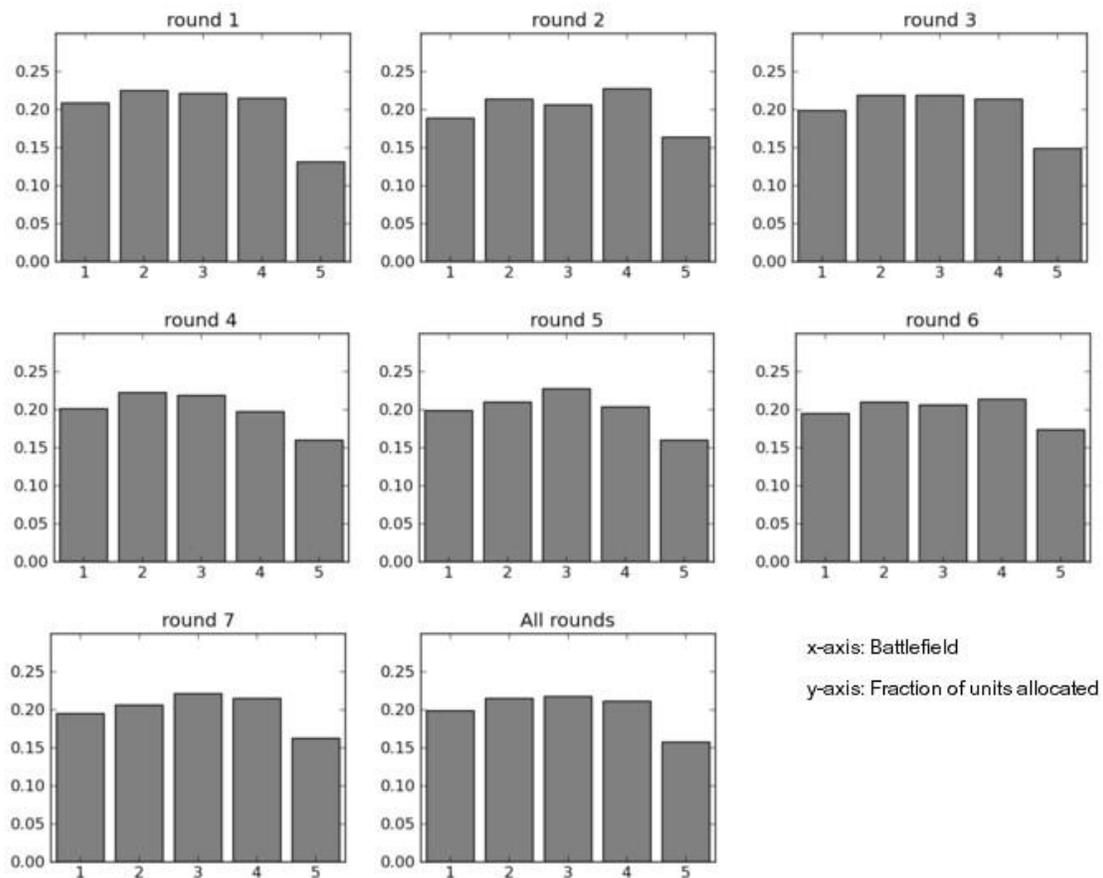


FIGURE 3.5: Distribution of Units among Battlefields across each Round. This figure illustrates how resource units are distributed among battlefields across each round.

Consider next the other aspect of equilibrium strategy: choosing evenly between 0 and 20 resources. Figure 3.6 illustrate how players change their resource allocations as the tournament

progresses. There is no clear evidence of convergence on equilibrium behavior, but there are a few interesting points to observe. Note the pronounced spike in Round 2 at 16 units, which we hypothesize is due to the success of Big-3 strategies in the first round, and a sudden increase at 10 units in Round 6. The usage of allocating 1 unit also gradually increased synchronously with the decline for 0 units. A probably motivation for this pattern would be that, after seeing many 0 unit battlefields, players learn to allocate 1 unit in a battlefield to win if the opponent's corresponding battlefield held zero.

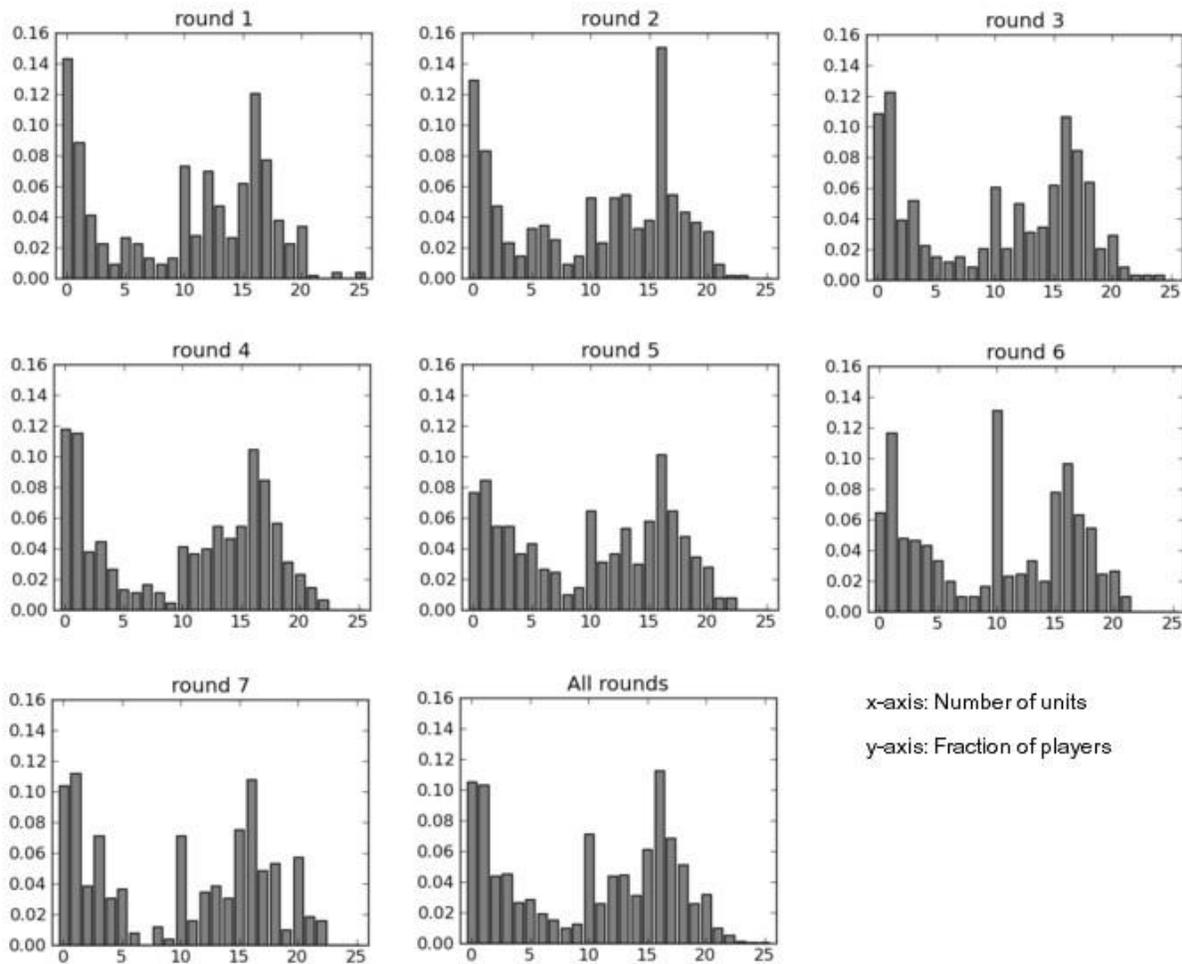


FIGURE 3.6: Distribution of Units across Rounds. This figure illustrates how players distribute their resource units across the seven rounds.

Lastly, we studied the distribution as players competed more against a single partner. In Figure 3.7, we observe how well people mixed their resource allocations for the first game that they played against each opponent, the second game, etc. There is a clear trend that people mix more evenly as they play more games against a single opponent. This implies that the dominant learning effect is that people learn how their opponent plays rather than how the set of their opponents play in aggregate, to the point that people default back to different strategies when playing a new opponent for the first time.

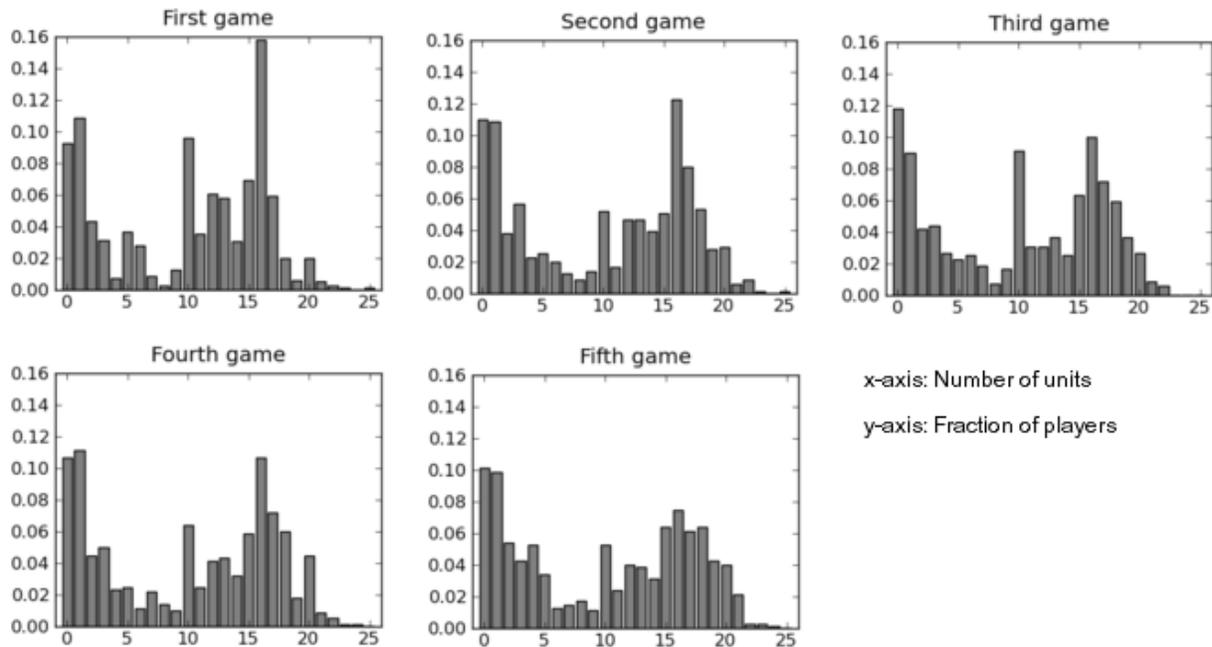


FIGURE 3.7: Distribution of Units across all Games in the Series. This figure illustrates how players distribute their resource units across all the first games in their rounds, all second games, etc.

3.4 Serial Correlation

Supporting our initial hypotheses, we find evidence of serial correlation for players between games. A correlation value of 0.072 for a battlefield's unit allocation between games is

statistically significant at a p-value of 0.000005. This denotes that battlefield allocations are ‘sticky’ – players tend to play similarly between games even when with a repeated opponent.

Next, we investigate if player strategy remains consistent between rounds by examining similarity in placement when ordered by size. For example, {0, 5, 10, 15, 20} and {19, 16, 10, 4, 1} appear very different when comparing them on a battlefield level, but are similar under the Big-3 strategy category. When we compare in this fashion, self-correlations are statistically significant at every point, much like what we expected from the background discussion. Even more importantly, we tested if the self-correlations were larger or smaller if the previous game was a win or loss. From largest to smallest allocations, self-correlations were {0.157, 0.298, 0.261, 0.154, 0.08} when the previous game was a win and {0.197, 0.382, 0.558, 0.103, 0.298} when the previous game was a loss. As evidenced, the correlations are almost strictly higher when the previous game is a loss. This implies that players who lose remain more consistent in their allocations.

4. Conclusion

We devised an experimental study to investigate the classic Colonel Blotto game class created by Borel in 1921. In our research project, we examined equilibrium strategy properties and further focused on ways in which human play differed from the theoretical equilibrium. Subjects competed against one another in a classic Colonel Blotto game with five battlefields and 50 resources each. A theoretical equilibrium is characterized by a uniform distribution of unit allocations between zero and 20 for each battlefield.

Results of the experimental play diverged from the equilibrium in several aspects. There was a significant reduction in resource allocation to the fifth, rightmost battlefield. The distribution of overall allocations was bimodal, rather than uniform, concentrating on very low

levels near zero and highs around 16. Majority of players used Big-3 strategies, although Big-2 worked very well. In addition to a different distribution, we also discovered correlations of a player's allocations between games. This was both with individual battlefields as well as with size of the allocations. With the latter, we find increased correlation when a player lost their previous game.

There are a few limitations of this study to keep in mind. More data would greatly improve the quality of the results, visibility of some player strategies, and our ability to rigorously analyze statistical significance. The removal of bad data also slightly decreased our beginning sample size. Faulty data occurred when, for example, players submitted no armies in a game, usually a result of a player computer crash or glitch in the tournament software. Another concern is that we cannot assume subjects in the experiment game always act rationally; their end goal might not always be to maximize the expected number of wins. We are unable to identify the presence of irrational actions and therefore, noise is introduced into the data.

This experimental project is unique among the published Colonel Blotto research and helps us further understand the distinctions between human play and expected equilibria in constant sum games. Humans have consistent biases in many of these games, such as rock-paper-scissors, and consequently, these biases can be exploited in manners that theoretical studies never predicted. The Colonel Blotto class of games has many applications to real world problems including business investment under competition and strategy computer games. There is a great deal of work left for illustrating and predicting human biases and deciding optimal play in games where humans are involved.

References

1. Bellman, R. (1969). On Colonel Blotto and Analogous Games. *Siam Review*, 11, 66–68.
2. Borel, E. (1921). La theorie du jeu les equations integrales a noyau symetrique. *Comptes Rendus del Academie*. 173, 1304–1308; English translation by Savage, L. (1953). The Theory of Play and Integral Equations with Skew Symmetric Kernels. *Econometrica*, 21, 97–100.
3. Borel, E., & Ville, J. (1938). Application de la theorie des probabilities aux jeux de hasard. Paris: Gauthier-Villars, 1938; reprinted in Borel E., Cheron, A.: *Theorie mathematique du bridge a la portee de tous*. Paris: Editions Jacques Gabay, 1991.
4. Chowdhury, Subhasish M., Kovenock, Dan and Sheremeta, Roman M. (June 2009). An Experimental Investigation of Colonel Blotto Game. CESifo Working Paper Series No. 2688.
5. Clark, D.J., & Konrad, K.A. (2007). Asymmetric Conflict: Weakest Link against Best Shot. *Journal of Conflict Resolution*, 51, 457-469.
6. Friedman, L. (1958). Game-theory Models in the Allocation of Advertising Expenditure, *Operations Research*, 6, 699-709.
7. Gross, O., & Wagner, R. (1950). A Continuous Colonel Blotto game, Unpublished article, RAND Corporation RM-408.
8. Laslier, J.F., Picard, N. (2002). Distributive Politics and Electoral Competition. *Journal of Economic Theory*, 103, 106-130.
9. Roberson, B. (2006). The Colonel Blotto game. *Economic Theory* 29 (1), 1-24.
10. Shubik, M., & Weber, R.J. (1981). Systems Defense Games: Colonel Blotto, Command and Control. *Naval Research Logistics Quarterly*, 28, 281–287.