School of Economics and Finance

An Experiment on Non-Zero Sum Colonel Blotto Games

Rafael Hortala-Vallve and Aniol Llorente-Saguer

Working Paper No. 779 December 2015 ISSN 1473-0278

Queen Mary University of London

An Experiment on Non-Zero Sum Colonel Blotto Games

Rafael Hortala-Vallve*
Government Department
London School of Economics

Aniol Llorente-Saguer†
Queen Mary
University of London

December 20, 2015

Abstract

We study a version of the Colonel Blotto game where valuations across battlefields are heterogeneous and asymmetric. These games can exhibit unique pure strategy equilibria, some of which are non-monotonic with respect to the battlefield valuations. We test our theoretical predictions in the laboratory and find low initial levels of equilibrium play but substantial learning throughout the experiment. Learning is higher for games with monotonic equilibria. Finally, we find that deviations from equilibrium predictions benefit aggregate welfare.

JEL Classification: C92, D70

Keywords: Colonel Blotto, Non-zero sum, Experiments

*r.hortala-vallve@lse.ac.uk.

†a.llorente-saguer@qmul.ac.uk.

1 Introduction

Since proposed by Borel in 1921, the Colonel Blotto game and its different variations have been studied extensively and applied to a variety of settings.[1] In the original version of the game, two colonels fight over a number of battlefields and simultaneously distribute a (discrete) number of troops among them. A battlefield is won by the colonel with most troops, and the winner is the colonel that wins the most battlefields. Despite the variety of formulations in the game (discrete or continuous troops, equal or unequal forces, equal or different valuations of battlefields), Colonel Blotto is typically a zero-sum game where the colonels equally value each battlefield. Such games generically display equilibria in mixed strategies.

However, in most situations bargaining parties might value battlefields differently. Hortala-Vallve and Llorente-Saguer (2012) shows that when battlefield valuations are asymmetric and heterogeneous, there exist pure strategy equilibria.[2,3] Moreover, the games with pure strategy equilibria have a unique equilibrium and can be conveniently characterized. In this paper, we experimentally test our previous theoretical findings.

Studying these games in controlled laboratory conditions is interesting for a number of reasons. First, unlike in standard Colonel Blotto, our game has pure strategy equilibria, which eliminates the subjects' need to randomize and substantially simplifies the nature of the strategic interaction whilst keeping the essential features of the conflict inherent in a Colonel Blotto game. Second, in contrast with the classical example, strict Pareto improvements may now exist: a colonel may accept losing a battlefield if it implies winning a battlefield that is of more value to him. By allowing differing relative intensities, we depart from the zero-sum nature of the game; thus social preferences now play a role. Third, some equilibria display non-monotonic equilibria (i.e., more troops are sometimes deployed in valuations that are less valued). Previous research has shown that under incomplete information, subjects play monotonic strategies (see Hortala-Vallve and Llorente-Saguer, 2010). Yet it remains an open question whether subjects learn to play non-monotonic strategies in line with some of our theoretical predictions when information about the

[1] Applications range from models of distributive politics and electoral competition (Myerson, 1993, Lizzeri and Persico, 2001, Laslier, 2002, Laslier and Picard, 2003, and Laslier, 2005) to mechanisms to resolve conflict (Hortala-Vallve and Llorente-Saguer, 2010), models of terrorism (Arce, Kovenock and Roberson, 2012), competition and contests (Avrahami and Kareev, 2009, and Avrahami et al., 2014) and social networks (Kohli et al., 2011).

[2] Duffy and Matros (2014) and references therein show the existence of pure strategy equilibria in stochastic blotto games.

[3] Various papers study non-zero sum versions of the Colonel Blotto games: Kvasov (2007) characterizes the equilibrium when the allocation of forces is costly and both colonels have exactly the same number of troops; Roberson and Kvasov (2012) extend the analysis to cases in which the colonels' number of troops differs. More recently, Kovencock and Roberson (2014) studies a continuous version of our game where the possibly asymmetric, budget constraints hold only on average. Schwartz et al. (2014) extends the analysis of Hortala-Vallve and Llorente-Saguer (2012) to characterize mixed strategy equilibria in games with a sufficient number of battlefields of each possible value relative to the players' resource asymmetry.

opponent's valuations is common knowledge.

In our experiment, subjects are randomly matched in pairs to play Colonel Blotto games with three battlefields, all of which display a unique equilibrium in pure strategies.[4] We observe substantial learning across periods as the frequency of equilibrium play increases from 12.92% in the first five periods to 58% in the last five periods. Learning is more pronounced in games where equilibrium strategies are monotonic (i.e., more troops are invested in battlefields that are more highly valued) and when equilibrium play yields the highest payoffs. We observe high heterogeneity in the frequency equilibrium play among subjects. Moreover, and perhaps surprisingly, subjects display little correlation in strategic sophistication across games. In terms of welfare, subjects obtain on average a higher payoff than we had predicted when the equilibrium entails monotonic strategies, yet their payoff is lower when the equilibrium entails non-monotonic strategies.

1.1 Related Literature

Although we are the first to experimentally study complete information non-zero sum Colonel Blotto games, several other experiments have studied alternative versions of the Colonel Blotto game. Avrahami and Kareev (2009), Chowdhury, Kovenock, and Sheremeta (2013) and Avrahami et al. (2014) study versions of the game in which each subject values each battlefield identically. While in Chowdhury, Kovenock, and Sheremeta (2013) the payoff is the sum of the battlefields won, in Avrahami and Kareev (2009) and Avrahami et al. (2014) participants are paid according to the outcome in a randomly selected battlefield. The main objective of these papers is to compare the effect of symmetric versus asymmetric endowments across colonels. Despite the complexity of the mixed strategy equilibria, overall behavior mimics the theoretical predictions.

Montero et al. (2014) analyze games with asymmetric but homogeneous valuations, where the goal is to maximize the probability of winning a majority of battlefields. They find that subjects do invest more in more highly valued battlefields, but the additional investment is less than predicted. In a similar study, Duffy and Matros (2014) compare theoretically and experimentally two payoff objectives in a stochastic environment: maximizing the total expected payoff versus maximizing the probability of winning a majority of all battlefields. They show that the two objectives result in qualitatively different predictions and substantially different behavior in the laboratory. The games played in Duffy and Matros (2014) allow for equilibria in pure strategies. However, since they do not allow for heterogeneous values across colonels, equilibrium strategies are always monotonic.

[4] The case of two battlefields is not particularly interesting, since moving all the troops to the most valued battlefield is a weakly dominant strategy. Hortala-Vallve and Llorente-Saguer (2010) studied the two-battlefield case in a setup with incomplete information. Subjects quickly learned to play the weakly dominant strategy even with incomplete information.

Hortala-Vallve and Llorente-Saguer (2010) is the first paper that experimentally studies a setup with asymmetric and heterogeneous battlefield valuations. As opposed to the study in this paper, their setting includes incomplete information about the colonel valuations. When information is incomplete, subjects learn to play strategically by deploying more troops in battlefields that are more highly valued.[5]

The rest of the paper is organized as follows: Section 2 presents the model and characterizes the set of equilibria. Section 3 describes the experimental design and procedures. Section 4 describes the experimental results, and Section 5 concludes.

2 The Model

We present a simple model in which we can test the theoretical predictions about pure strategy equilibria of non-zero sum Colonel Blotto games.[6] There are two colonels ($i \in \{1, 2\}$) who fight over three battlefields ($m \in \{1, 2, 3\}$). Colonel i's valuation towards battlefield n is denoted $\theta_n^i \geq 0$. Non-zero sum Colonel Blotto games are characterized as those in which the condition that both colonels equally value both territories is not imposed (i.e., it is possible that $\theta_n^i \neq \theta_n^j$). We assume that colonels attach different values to different battlefields ($\theta_n^i \neq \theta_m^i$ for $n \neq m$).

Colonels need to simultaneously divide $P \in \mathbb{N}$ (indivisible) troops across the battlefields. The winner in each battlefield is the colonel that deploys the most troops to that battlefield. Ties are broken with the toss of a fair coin. Colonel i's payoff when winning battlefield n is θ_n^i, and the payoff is 0 when the opponent wins the battlefield. The player's total payoff is the sum of the payoffs across the battlefields.

Troop deployment strategies require two necessary conditions to constitute a pure strategy Nash equilibrium: (i) a battlefield is never won by more than one troop (it is useless to have a large victory) and (ii) no troops are deployed in a battlefield that is lost (it is useless to invest resources in a defeat). Both statements imply that the colonels' deployment in a battlefield can only be one of two types: (1) one colonel wins with a single troop or (2) both colonels deploy the same number of troops. With this in mind, we need to characterize the type of preference profiles that yield a pure strategy equilibrium.

We start by looking at two diametrically opposed situations. First, consider a situation in which

[5] In a related paper, Hortala-Vallve, Llorente-Saguer and Nagel (2013) compare behavior in a Colonel Blotto game to behavior in a situation with unrestricted communication among colonels. They find that outcomes and overall welfare comparison crucially depend on the level of uncertainty on each other's valuations.

[6] See Hortala-Vallve and Llorente-Saguer (2012), for a general characterization.

colonels' preferred battlefields coincide. In this case, all conflict occurs in the most preferred battlefield, and all troops belonging to both colonels are deployed there. For this deployment of troops to constitute a Nash equilibrium, we need both colonels to prefer the top battlefield at least as much as the sum of the other two battlefields. Second, consider a situation in which the preferred battlefield of any colonel is the least preferred of the other colonel. In this case, the conflict will occur in their mid-ranked battlefield, and both colonels will win their preferred battlefield by deploying a single troop.

In Hortala-Vallve and Llorente-Saguer (2012), we show that both preference profiles above are the only non-zero sum Colonel Blotto games with pure strategy equilibrium (and indeed this equilibrium is unique). These equilibria have interesting properties: when both colonels care enough about the same battlefield all troops are deployed to that territory, but this prevents colonels benefiting from potential *gains from trade* inherent in differences in others than the most preferred battlefield. Instead, when colonels have diametrically opposed preferences they extract all *gains from trade*, yet their deployment of troops is non-monotonic, because they deploy one troop in their most preferred battlefield but five in their second preferred one.

Two pieces of notation will be needed to formally state our result. First, $n^i_{(k)}$ denotes colonel i's k-th most preferred battlefield and $\theta^i_{(k)}$ the valuation of colonel i's k-th most preferred battlefield. In other words, when there are three battlefields $n^i_{(1)}$ is the most preferred battlefield of colonel i and $\theta^i_{(1)}$ is colonel i's valuation of his/her most preferred battlefield. Second, we need to characterize the valuation profiles that sustain an equilibrium in pure strategies. We define

$$\begin{aligned} \Theta_1 &= \left\{ \left(\theta^1, \theta^2\right) \in \Theta^2 \middle| n^1_{(1)} = n^2_{(1)} \ \& \ \theta^i_{(1)} \geq \theta^i_{(2)} + \theta^i_{(3)} \quad \forall i \right\}, \\ \Theta_2 &= \left\{ \left(\theta^1, \theta^2\right) \in \Theta^2 \middle| n^1_{(2)} = n^2_{(2)} \ \& \ n^1_{(1)} = n^2_{(3)} \right\}. \end{aligned}$$

where Θ_1 is the set of valuation profiles such that both colonels most preferred battlefields coincide and they both value this battlefield more than the sum of the other two; Θ_2 is the set of valuation profiles where only the second most preferred battlefields coincide. Below we formally state our result.

Proposition 1 *A valuation profile θ has a unique equilibrium in pure strategies if and only if $\theta \in \Theta_1 \cup \Theta_2$. Otherwise, θ has no equilibrium in pure strategies. If $\theta \in \Theta_1$, colonels deploy all points (P) in the preferred battlefield. If $\theta \in \Theta_2$, they deploy 1 point in the preferred battlefield and the remaining points ($P-1$) in the second preferred battlefield.*

Proof. See details in Hortala-Vallve & Llorente-Saguer (2012). ■

3 Experimental Design and Procedures

In order to test our theoretical predictions, we ran an experiment at Universitat Pompeu Fabra. Students were recruited through the online recruitment system ORSEE (Greiner, 2015) and the experiment was programmed and conducted using the software z-Tree (Fischbacher, 1999). We ran 2 sessions with 20 subjects per session. The subjects interacted for 35 rounds. In each round, subjects were randomly divided into groups of two.[7] At the beginning of each round, subjects were told their valuations, and these valuations were also made publicly known among the group. Next, the subjects had to distribute six troops across the three battlefields ($P = 6$).

Valuations were drawn uniformly and independently for each colonel from the set of four possible profiles specified in Table 1; all profiles are normalized so that they add up to 600.[8]

Profiles	θ_1	θ_2	θ_3
P1	300	250	50
P2	350	200	50
P3	400	150	50
P4	450	100	50

Table 1: Preferences' profiles (in euro cents).

The order of the valuations is permuted in order to obtain three classes of games. In the first class of games, subjects have the same ordinal preferences across battlefields –we denote these games as Same Order, SO. In the second class of games, the preferred battlefield of both subjects coincides, yet the second most preferred battlefield of one subject is the least preferred of the other one –we denote these as Same Preferred, SP. Finally, in the last class of games we have that the most valued battlefield for a subject is the least valued of his/her opponent –we denote these the Opposite Order, OO. Valuations were permuted so that there were 20%, 40% and 40% of SO, SP and OO games respectively. Participants were not aware of the random process generating valuation profiles but had complete information about preferences. That is, at the beginning of each round we gave full information about the preferences to each matched pair.

The payoff received by each subject was equal to the sum of valuations of the battlefields in which they deployed more troops than their opponent plus half the valuations of the battlefields in which both subjects had deployed the same number of troops. At the end of each period, each subject received the following information: (i) her/his valuations on each battlefield; (ii) her/his

[7] We partition subjects into 2 sets of 10 players so as to obtain 2 independent observations per session.

[8] Our experiment was framed in terms of votes (instead of troops) and issues (instead of battlefields) because of other related work we have previously published. This different framing should not change the strategic incentives of our subjects.

partner's valuations, (iii) the points that he/she had distributed over the battlefields, (iv) the points distributed by the partner, (v) the outcome in each battlefield, and (vi) the payoff.

All experimental sessions were organized along the same procedure. Subjects received detailed written instructions, which an instructor read aloud (see Online Appendix). Before the start of the session, subjects were asked to answer a set of control questions to check their full understanding of the experimental design. At the end of the experiment, subjects had to fill in a short questionnaire on the computer and were given their final payment in private. To determine payment, the computer selected three periods randomly. In total, subjects earned an average of €13.93, including a show-up fee of €3. Session length, including wait time and payment, was approximately 70 minutes.

4 Experimental Results

4.1 Aggregate Behavior

Figure 1 displays the temporal evolution of the frequency of equilibrium play in the three types of games. We can see that there is substantial learning throughout the experiment: the percentage of equilibrium play in the first five periods is 12.92%, and this increases to 58% in the last five periods of the experiment. We can also see that there are noticeable differences across games, especially in the second half of the experiment, when equilibrium play for the *SO* games takes off and reaches values around 80% of equilibrium play.

Despite the fact that the action space is large (there are 28 different ways to distribute 6 troops across 3 battlefields), deviations from the theoretical predictions always follow a similar pattern. Subjects should never invest a single troop in their least preferred battlefield, and indeed most actions in our sample follow this prediction: in 84.93% of our observations, zero troops are invested in the least preferred battlefield. When looking in detail at this data we can see that there is noticeable learning: there are 67.5% observations with zero troops in the first five periods, but this quickly increases to levels above 85% (82.5%, 88.21% and 82.26% observations with zero troops in the least preferred battlefield in SO, SP, and OO, respectively).[9]

The high incidence of investing zero troops in the least preferred battlefield seems to suggest that

[9] There is no significant difference in the frequency of deploying zero troops on the last battlefield across SO and OO games (Wilcoxon, $z = 0.000$, $p = 1.000$). This frequency is, however, significantly higher in SP games than in the other two types of games (Wilcoxon, $z = 1.841$, $p = 0.067$ against SO and Wilcoxon, $z = 1.671$, $p = 0.095$ against OO). For these and the other non-parametric tests in the rest of the paper, we use the average in a matching group as the unit of observation. All these tests are two-tailed sign-rank tests.

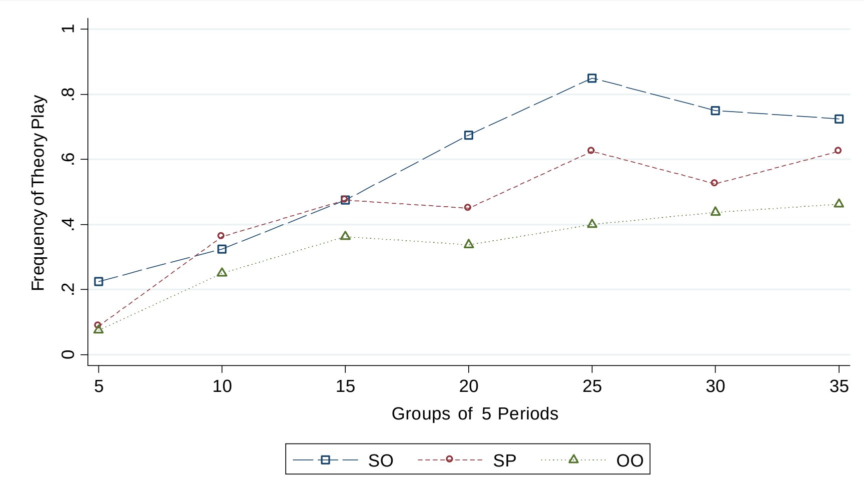

Figure 1: Frequency of theory play for each type of equilibrium in groups of 5 periods.

the main deviation with respect to the theoretical prediction is due to an improper distribution of votes across the top two battlefields. In SO and SP games, subjects invest less than the total number of votes in the battlefield with highest valuation; conversely, in OO games, they invest too many troops.

When deviating from equilibrium behavior, subjects seem to move toward a troop allocation that more closely matches their relative preference across battlefields. We analyze equilibrium play across different preference profiles in Figure 2. This Figure displays the frequency of equilibrium play by type of game and type of preference profile. Recall that there are four possible preference profiles (see Table 1), and these are uniquely characterized by the valuation of the most preferred battlefield. We restrict the analysis to the last 20 periods due to the learning we observed in the initial periods (our results remain unchanged, but become a little noisier when looking at the whole sample). We observe that the frequency of equilibrium play is positively correlated with the valuation of the most preferred battlefield when the top battlefield coincides (SO and SP games). In other words: the higher the stakes of winning the most preferred battlefield, the higher the likelihood of deploying all troops into it. Equilibrium play is usually lower in SP games than in SO games: that is, when conflict is highest, there is less risk of departing from equilibrium predictions. Instead, when there are gains from trade in the last two preferred battlefields, subjects tend to deviate and extract gains from such differences at the risk of losing their most preferred battlefield. Finally, when preferences are opposed (OO games), equilibrium play decreases with the valuation of the most preferred battlefield: investing a single vote in the most preferred battlefield is increasingly difficult for subjects who strongly prefer such a battlefield.

In Table 2 we test and validate our previous claims by regressing the frequency of playing the equilibrium predictions on period, dummies on the class of game, and the interaction of the last

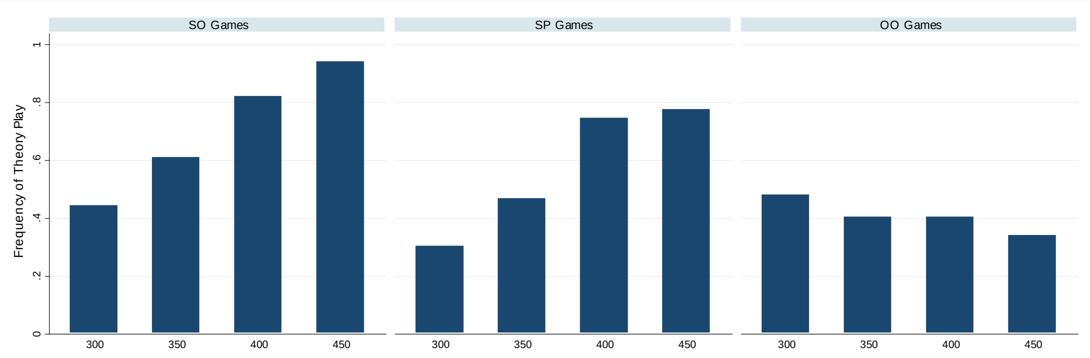

Figure 2: Frequency of theory play by type of game and by the highest valuation in the last 20 periods.

two variables. As we observed before, we find an increasing trend in playing equilibrium strategies, and this is especially pronounced in the SO games.

Are subjects more likely to play equilibrium strategies when their relative preferences over battlefield varies? To test this claim, we define a dummy that captures a 'high intensity' valuation for the most preferred battlefield for each subject[10] and an analogous variable for the opponent (high intensity of opponent, HIO). In the last three columns, we report our results disaggregated by type of game so that our findings are easier to interpret. We observe that in SP and SO games, high intensity yields more equilibrium deployment of troops (all troops on the most preferred battlefield). Conversely, in OO games, having more extreme preferences correlates negatively with playing equilibrium strategies; recall that the equilibrium action in that case has a single troop in the most preferred battlefield, but the more a subject values that battlefield, the higher the risks of leaving a unique troop on that territory.

Finally, given that our game is one with complete information, one might argue that deviations from our equilibrium predictions are due to social preferences. If subjects care about the welfare of others, they should take into account the preferences of their opponent. We introduce two dummies that capture the change in behavior by our subjects when their opponent changes her/his preferences: the first dummy, "Dummy HI=1 and HIO=1", captures the behavior of a subject with high preferences when her/his opponent's preferences change from being low to high. We observe that there is an effect in the first two games where playing equilibrium strategies is more likely. An increase in the valuation of the top battlefield of the opponent increases the likelihood that more resources are invested in that top battlefield –an indication that there is no support for other-regarding preferences such as inequality aversion or efficiency concerns. The last dummy captures the same change in the opponent's preferences when our subject does not have a strong

[10] The 'High Intensity' dummy takes value one if the most preferred battlefield is valued 400 or 450 and zero if it is valued 300 or 350. All results are robust to considering alternative specifications that capture how intensely a subject values his most preferred battlefield.

Variables	*(1)*	*(2)*	*(3)*	SO games *(4)*	SP games *(5)*	OO games *(6)*
Period	0.01***	0.01***	0.01***	0.02***	0.02***	0.01***
	(0.001)	(0.001)	(0.001)	(0.002)	(0.002)	(0.001)
Dummy SO game		0.24***	0.07*			
		(0.048)	(0.042)			
Dummy SP game		0.12	0.03			
		(0.087)	(0.073)			
Dummy SO game × Period			0.01***			
			(0.002)			
Dummy SP game × Period			0.00			
			(0.003)			
Highest Intensity (HI)				0.42***	0.28***	−0.15**
				(0.103)	(0.035)	(0.067)
Dummy HI=1 and HIO=1				0.12**	0.24***	−0.04
				(0.058)	(0.039)	(0.044)
Dummy HI=0 and HIO=1				0.13***	0.09*	−0.15***
				(0.034)	(0.044)	(0.030)
Constant	0.17***	0.07**	0.14***	−0.02	−0.05	0.26***
	(0.038)	(0.031)	(0.031)	(0.059)	(0.051)	(0.058)
Observations	1,400	1,400	1,400	280	560	560

Table 2: Random effects GLS regression of the frequency of voting in line with the theoretical prediction as a function of the listed variables. (4), (5) and (6) restrict the data to SO, SP and OO games respectively. Hubbard-White robust standard errors, clustered by independent groups. Standard errors in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$

preference in her/his most preferred battlefield. This is precisely where we should expect the other-regarding preferences to be stronger. Yet once again, we see that subjects are more likely to play equilibrium strategies, thus not allowing an opponent with stronger preferences to win her/his preferred battlefield. Note also that the negative coefficient in the OO games captures the fact that as the opponent values the second issue more (values the first issue less), playing equilibrium strategies is more likely. Once again, we see that an increase in the valuation of the opponent exacerbates conflict.

An alternative explanation for deviations from equilibrium play is that subjects play a 'truthful' strategy by which they distribute troops proportionally to their battlefields' valuations. We find that this is not the case: the percentage of truthful play in our sample is below 10% (6.8%, 10.2% and 8.4% in SO, SP, and OO games, respectively).

4.2 Individual Behavior

We now turn our attention to individual behavior. Figure 3 displays histograms of the frequency of equilibrium play by each subject (in the last twenty periods) for each type of game. It shows substantial heterogeneity across subjects in all types of games. A noticeable difference across the

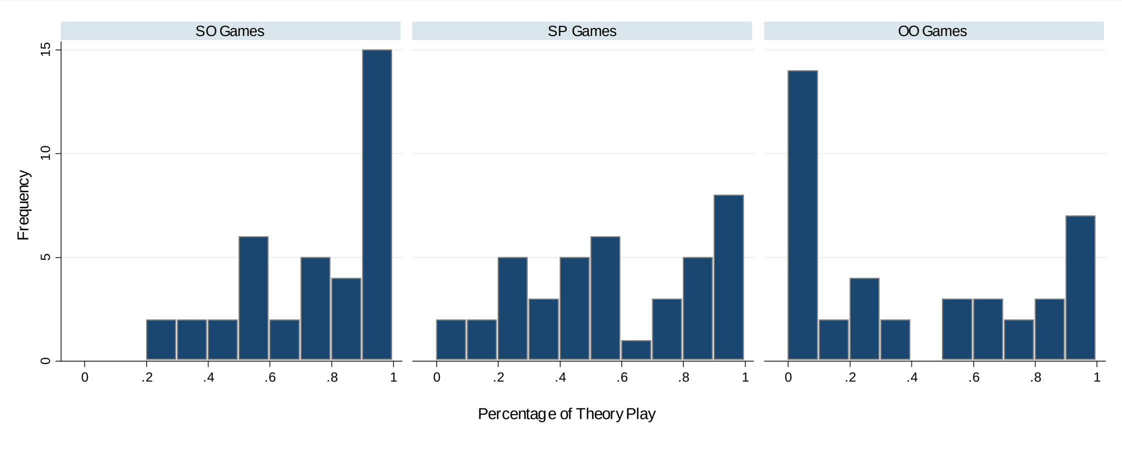

Figure 3: Histogram of individual frequencies of theory play in the last 20 periods for the three types of games.

three games is the mode of zero in OO games: while the percentage of subjects never playing the equilibrium strategy is as low as 0% and 5% in SO and SP games, this percentage is 35% in OO games. Both pairwise comparisons with OO games are statistically significant (Wilcoxon, $z = 1.826$, $p = 0.068$ in both cases).

It is an interesting endeavour to analyze whether the likelihood to follow equilibrium predictions between the three different games is correlated. Is it true that there are some individuals who are well aware of strategic considerations across the different situations? Or are there some individuals who are more likely to act as prescribed in equilibrium predictions in some particular games and less likely in others? Figure 4 displays the joint distribution of frequency of equilibrium play across any two of our games. The horizontal axis plots the frequency with which each subject played the equilibrium action in one of the games, and the vertical axis plots the frequency in a different game. We can see a strong positive correlation between playing equilibrium in SP and SO games (the correlation is 0.577). However, the other two figures show a very weak, if any, relation between both frequencies: this conforms with the findings of Georganas et al. (2015) that, in a very different setting, find no correlation in subjects' estimated sophistication across games.

4.3 Best response to observed behavior

Deviations from equilibrium predictions could be explained by the presence of naive or noisy subjects. Due to their presence, it might not be the best response to act as prescribed by equilibrium predictions. In order to study whether deviations from equilibrium can be due to non-equilibrium behavior of other subjects, we compute the best response to observed behavior. For SO and SP games, the best response coincides precisely with our theoretical prediction: invest all troops in the

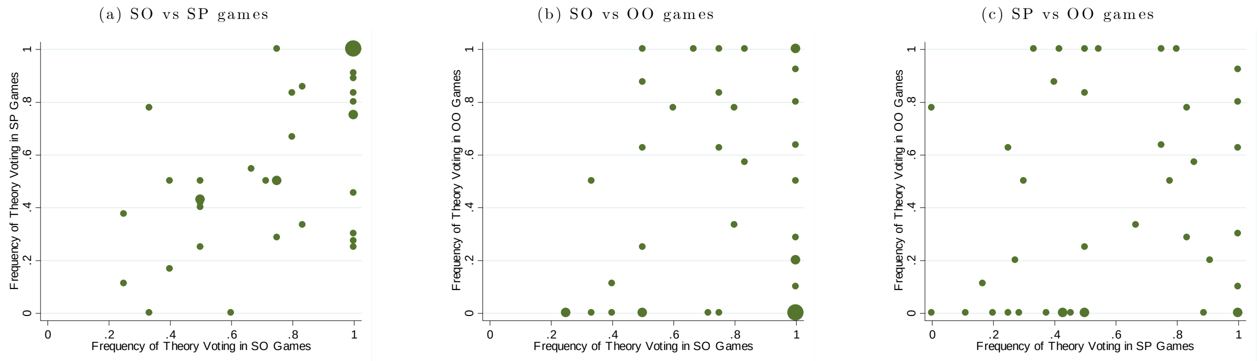

Figure 4: Joint distribution of frequencies of equilibrium play in the last 20 periods (in pairs of types of games). Panels (a), (b) and (c) display the joint distribution of frequencies of equilibrium play in So vs SP, SO vs OO and SP vs OO games respectively.

most valued battlefield. For *OO* games, the best response is to invest two troops in the preferred battlefield and four in the second most preferred.[11] This is largely due to the variance of actions in the initial periods when subjects are learning to play. When we restrict the analysis to the last twenty periods, our equilibrium prediction is the best response to observed behavior if colonels are risk neutral or not too risk averse.[12] In all, deviations from equilibrium predictions cannot be explained by best-responding actual behavior.

4.4 Costs of not playing equilibrium behavior

Another aspect worth studying is the costs of not playing equilibrium. Given the previous analysis, we know that a subject that does not play equilibrium strategies should (on average) be damaged and obtain a lower payoff than a subject that played equilibrium strategies. However, can we learn anything from analyzing the individual payoffs of subjects who are not following equilibrium strategies? We look at the payoff difference between not playing equilibrium strategies and doing so. We restrict our analysis to those subjects who did not play equilibrium strategies and partition our data according to the following: (1) observations when subjects obtain a strictly smaller payoff than that which they would have obtained if they had played equilibrium strategies; (2) observations when subjects obtain exactly the same payoff; and (3) observations when subjects obtain a strictly higher payoff. In Table 3 we report the average difference in payoffs in these three cases for each of the three games our subjects played. We also report the percentage difference and the frequency with which each cell occurs (i.e., among the subjects that did not play equilibrium strategies in a particular game, we find the percentage that obtained a higher, smaller, or equal

[11] The equilibrium strategy where subjects invest 1 troop in the most preferred battlefield and 5 in the second most preferred yields an almost identical but strictly lower payoff.

[12] The difference in expected payoff between playing $(1, 5, 0)$ and $(2, 4, 0)$ decreases in $\theta^i_{(1)}$. If $\theta^i_{(1)} = 300$ the payoff diference is 57.42 while if $\theta^i_{(1)} = 450$ the payoff diference is 4.92.

		Lower payoff	Equal	Higher payoff	Total
SO games	Payoff difference	-102.25	0	65.38	-69.33
	% difference	-64.47	0	13.55	-19.86
	Frequency	74.79	14.29	10.92	100
SP games	Payoff difference	-117.50	0	115.63	-76.54
	% difference	-57.23	0	22.29	-22.93
	Frequency	77.92	9.09	12.99	100
OO games	Payoff difference	-103.97	0	212.50	-39.37
	% difference	-27.94	0	43.96	-8.33
	Frequency	67.38	18.18	14.44	100

Table 3: Average cost and average percentual loss of payoff of not playing equilibrium separated by the sign of it and the class of game, and frequency of each of the cases.

payoff). Looking first at the aggregate for each game (last column), we see that the overall impact of not playing equilibrium strategies is negative: in other words, on average subjects would have been better off by playing equilibrium strategies. This is especially so when there is most conflict –i.e., when both subjects prefer the same battlefield, cases SO and SP. This mirrors what we observed in the previous section –that equilibrium play is the best response to observed behavior. The smallest percentage cost is when both subjects have opposed preferences.

When looking at the disaggregated data we see that not playing equilibrium strategies will most likely yield a lower payoff: 74.79%, 77.92%, 67.38% of our observations in SO, SP, and OO games, respectively. Conditional on obtaining a higher payoff, the payoff differences increase (in absolute and relative terms) with the misalignment of subjects' preferences. Finally, conditional on obtaining a lower payoff, the expected cost is significantly higher in SO and SP games than in OO games (Wilcoxon, $z = 1.826$, $p = 0.068$ in both cases). There are no significant differences between SO and SP games (Wilcoxon, $z = 1.095$, $p = 0.273$).

4.5 Welfare

Unlike the traditional formulation of the Colonel Blotto game, where all battlefields are equally valued by both colonels, our non-zero sum games allow for "gains of trade" –both subjects can improve their utility if battlefield victories are appropriately assigned. In the analysis that follows, we use two different measures to compare the aggregate welfare of our subjects with that predicted by theory: the average payoffs and the percentage of Pareto efficient allocations.

Table 4 summarizes both measures for each type of game. Both measures offer a similar conclusion: welfare is higher than predicted by our game theoretical model for SO and SP games, and lower than predicted for OO games. We find that the improvement over the equilibrium prediction in

	Payoff			% Pareto Efficiency	
	Realized	Equilibrium	Utilitarian Optimum	Realized	Equilibrium
SO games	311	300	329	86.43	67.14
SP games	334	300	380	57.86	0
OO games	442	462	479	82.50	100

Table 4: Realized and equilibrium average payoffs and percentage of Pareto efficiency for each type of game.

SO and *SP* games is around 40% to the maximum that could have been achieved. When the game is almost zero-sum (*OO* game), the possible improvements over the equilibrium predictions are very small, and the realized payoff in this case is actually slightly below our game theoretical predictions.

Deviations from equilibrium behavior can explain these results. We previously showed that such deviations harm the deviator but benefit her/his opponent. This is because in *SO* and *SP* games, subjects deviate more from equilibrium behavior the more they value their second most preferred battlefield. In other words, the higher the valuation for the most preferred battlefield is, the more likely it is that a colonel will invest all troops in this battlefield, and the more likely it is that he/she will win her/his preferred battlefield (as opposed to obtaining a tie in all three battlefields). This implies that the pair is more likely to achieve the utilitarian optimum behavior when subjects deviate from equilibrium behavior as opposed to when subjects follow equilibrium prescriptions. In *OO* games the same channel operates, yet the uncommon event in which some subjects lose their preferred battlefield (in the hands of a colonel who invests non-zero troops in their least preferred battlefield) greatly harms average welfare and explains the small losses in our data. Analogous arguments explain the differences between the predicted and realized frequencies of Pareto efficient situations.

5 Conclusion

Colonel Blotto games date back to the early twentieth century, but the difficulties in solving them have prevented their popularity. Recent interest in the application of such games to fights over limited resources, votes on multiple issues, or the capture of undecided voters has sparked a renewed interest in these games. We have now a good understanding of the colonels' equilibrium behavior as well as growing experimental evidence in controlled laboratory conditions.

In this paper, we experimentally test the theoretical predictions in a particular type of Colonel Blotto game –the non-zero sum game that has pure strategy equilibria– and find that our theoretical predictions are effective in explaining subjects' behavior. Deviations from equilibrium predictions always occur in the same direction: subjects seem to move toward matching their troop deployment to their relative valuations of battlefields. These deviations improve aggregate welfare by allowing the victories of colonels with strongly held preferences.

One avenue for future research would look at how deviations from equilibrium behavior harm or benefit aggregate welfare in different mechanisms. When thinking about the actual implementation of various mechanisms, their robustness might be viewed in terms of whether deviations from equilibrium behavior actually benefit society at large.

References

[1] Arad, A. and A. Rubinstein (2012) "Multi-Dimensional Iterative Reasoning in Action: The Case of the Colonel Blotto Game", *Journal of Economic Behavior and Organization*, 84, 571–585.

[2] Arce, D.G., Kovenock, D., and Roberson, B. (2012), "Weakest link attacker defender games with multiple attack technologies", *Naval Research Logistics*, 59, 457–469.

[3] Avrahami, J. and Y. Kareev (2009) "Do the Weak Stand a Chance? Distribution of Resources in a Competitive Environment", *Cognitive Science*, 33, 940–950.

[4] Avrahami, J., Kareev, Y., Todd, P., Silverman, B. (2014) "Allocation of Resources in Asymmetric Competitions: How do the Weak Maintain a Chance of Winning?", *Journal of Economic Psychology*, 42, 161–174.

[5] Borel, E. (1921), "La theorie du jeu les equations integrales a noyau symetrique", *Comptes Rendus de l'Academie* 173; English translation by Savage, L (1953), "The theory of play and integral equations with skew symmetric kernels", *Econometrica*, 21, 97–100.

[6] Chowdhury, S., D. Kovenock, and R. Sheremeta (2013), "An experimental investigation of Colonel Blotto games", *Economic Theory*, 52, 833–861.

[7] Duffy, J. and A. Matros (2013) "Stochastic Asymmetric Blotto Games: Theory and Experimental Evidence", *mimeo.*

[8] Fischbacher, U (2007), "z-Tree - Zurich Toolbox for Readymade Economic Experiments - Experimenter's Manual", *Experimental Economics*, 10, 171–178.

[9] Georganas, S., P.J. Healy and R.A. Weber (2015), "On the Persistence of Strategic Sophistication", *Journal of Economic Theory*, 159, 369–400.

[10] Greiner, B. (2015), "Subject pool recruitment procedures: organizing experiments with ORSEE", *Journal of the Economic Science Association*, 1(1), 114-125.

[11] Hortala-Vallve, R. (2012), "Qualitative Voting", *Journal of Theoretical Politics*, 24, 526–554.

[12] Hortala-Vallve, R and A. Llorente-Saguer (2010). "A Simple Mechanism for Resolving Conflict", *Games and Economic Behavior*, 70, 375–391.

[13] Hortala-Vallve, R and A. Llorente-Saguer (2012). "Pure strategy Nash equilibria in non-zero sum Colonel Blotto games", *International Journal of Game Theory*, 41, 331–343.

[14] Hortala-Vallve, R, A. Llorente-Saguer and R. Nagel (2013). "The Role of Information in Different Bargaining Protocols", *Experimental Economics*, 16, 88–113.

[15] Kohli, P., Y. Bachrach, T. Graepel, G. Smyth, M. Armstrong, D. Stillwell, R. Herbrich and M. Kearns (2011), "Behavioral Game Theory on Online Social Networks: Colonel Blotto is on Facebook", *mimeo.*

[16] Kovenock, D. and B. Roberson (2014), "Generalizations on the Colonel Blotto Game", *mimeo.*

[17] Kvasov, D. (2007), "Contests with limited resources", *Journal of Economic Theory*, 136, 738–748.

[18] Laslier, J.F. (2002), "How two-party competition treats minorities", *Review of Economic Design*, 7, 297–307.

[19] Laslier, J.-F. and N. Picard (2002), "Distributive politics and electoral competition", *Journal of Economic Theory*, 103, 106–130.

[20] Laslier, J.-F. (2005), "Party objectives in the 'divide a dollar' electoral competition" in: Social Choice and Strategic Decisions, Essays in Honor of Jeff Banks, edited by D. Austen–Smith and J. Duggan, Springer, 113–130.

[21] Lizzeri, A., and N. Persico (2001). "The provision of public goods under alternative electoral incentives", *American Economic Review*, 91, 225–239.

[22] Montero, M., A. Possajennikov, M. Sefton and T. L. Turocy (2014), "Majoritarian Blotto Contests with Asymmetric Battlefields: An Experiment on Apex Games", *mimeo.*

[23] Myerson R. (1993), "Incentives to cultivate favored minorities under alternative electoral systems", *American Political Science Review*, 87, 856–869.

[24] Roberson, B. and Kvasov, D. (2012), "The Non-Constant-Sum Colonel Blotto Game", *Economic Theory*, 51, 397–433.

[25] Schwartz, G. P. Loiseau and S. S. Sastry (2014), "The Heterogeneous Colonel Blotto Game", *mimeo.*

Online Appendix: Instructions

We are grateful for your participation and for your contribution to an Economics Department project. The sum of money you will earn during the session will be given privately to you at the end of the experiment. From now on (and until the end of the experiment) you cannot talk to any other participant. If you have a question, please raise your hand and one of the instructors will answer your questions privately. Please do not ask anything aloud!

These experiments consist of 35 periods. The rules are the same for all participants and for all periods. At the beginning of each period you will be randomly assigned to another participant with whom you will interact. None of you will know who the other participant is.

You and the other participant will simultaneously vote over a group of three issues. Each issue has three possible results: 1) you win and s/he loses it; 2) you lose and s/he wins it; and, finally 3) ties occur. These results will determine the profits that yourself and the other participant will have in each period. Remember that the participant with whom you are interacting in each period is selected randomly in each period.

Information at the beginning of each period. At the beginning of each period you will be told your 'valuations' and the valuations of the person you are paired with for each issue. The valuation of each issue specifies how much you earn when you win that issue. These valuations are expressed in terms of cents of Euro. Valuations are multiples of 50 and add up to 600.

Voting procedure. In each period you will have six votes that you will have to distribute among the different issues. After doing so you should press the 'OK' key. The participant with whom you are matched at each period has the same number of votes.

Voting result. The result on each voting procedure will be resolved by the following rule: if the number of votes you have assigned to an issue is

- ... higher than the number of votes of the other participant, you win the issue
- ... smaller than the number of votes of the other participant, you lose the issue
- ... equal to the number of votes of the other participant, ties occur

For instance, if you vote in the following way:

	Issue 1	Issue 2	Issue 3
Your votes	3	1	2
Her/his votes	0	4	2

You win Issue 1 given that you have assigned more votes (3) than him (0) in that issue; you lose issue 2 given that you have invested less votes (1) than him (4); and you tie issue 3 given that you have both invested the same number of votes (2).

Profits in each period. In each period your profits will be equal to the sum of the valuations of all issue you win plus half the valuation of the issue you tie. For instance, if in the previous example your valuations are 100, 50 and 450, according to the assignation of votes in Section 3 your benefits will be: 100 in issue 1, 0 in issue 2 and 225 (half of 450) in issue 3. As you know the preferences of your partner, you can also compute his/her profits.

Information at the end of each period. At the end of each period, as you can see in the previous screenshot, you will receive the following information:

- Your valuations in each issue
- The valuations of the participant you have interacted with
- Your votes in each issue
- The votes of the participant you have interacted with
- The issues you win, lose and tie
- Your profits

Final payment. At the end of the last period, the computer will randomly select 3 periods and you will earn the sum of the profits on those periods. Additionally you will be paid three euros for having taken part in the experiment.

Control Questions

1. Circle the correct answer. When you have to vote. . . .

- ¿Do you know your valuations? YES/NO
- ¿Do you know the valuations of the participant you are matched with? YES/NO
- ¿Do you know the identity of the participant you will interact with? YES/NO
- ¿Are you always matched to the same person? YES/NO
- ¿Are you informed about the identity of the participant you are matched with? YES/NO

2. Imagine you have the following valuations and that you and the participant with whom you are matched vote in the way specified below

	Issue 1	Issue 2	Issue 3
Your valuations	400	150	50
Valuations of the other participant	450	50	100
Your votes	3	2	1
His/her votes	4	1	1

- ¿Who wins issue 1? YOU/HIM/TIE
- ¿Who wins issue 2? YOU/HIM/TIE
- ¿Who wins issue 3? YOU/HIM/TIE
- ¿How much do you win in issue 1? ____________
- ¿How much do you win in issue 2? ____________
- ¿How much do you win in issue 3? ____________

3. ¿Which is your profit in this period? __________

4. ¿How many periods will determine your final payment? ___

Queen Mary University of London

This working paper has been produced by
the School of Economics and Finance at
Queen Mary University of London

School of Economics and Finance
Queen Mary University of London
Mile End Road
London E1 4NS
Tel: +44 (0)20 7882 7356
Fax: +44 (0)20 8983 3580
Web: www.econ.qmul.ac.uk/research/workingpapers/