

# Majoritarian Blotto Contests with Asymmetric Battlefields: An Experiment on Apex Games

Maria Montero, Alex Possajennikov, Martin Sefton, Theodore L. Turocy

This is the final version accepted by *Economic Theory*, July 2015

The final publication is available at Springer via <http://dx.doi.org/10.1007/s00199-015-0902-y>

## Abstract

We investigate a version of the classic Colonel Blotto game in which individual battlefields may have different values. Two players allocate a fixed discrete budget across battlefields. Each battlefield is won by the player who allocates the most to that battlefield. The player who wins the battlefields with highest total value receives a constant winner payoff, while the other player receives a constant loser payoff. We focus on apex games, in which there is one large and several small battlefields. A player wins if he wins the large and any one small battlefield, or all the small battlefields. For each of the games we study, we compute an equilibrium and we show that certain properties of equilibrium play are the same in any equilibrium. In particular, the expected share of the budget allocated to the large battlefield exceeds its value relative to the total value of all battlefields, and with a high probability (exceeding 90% in our treatments) resources are spread over more battlefields than are needed to win the game. In a laboratory experiment, we find that strategies that spread resources widely are played frequently, consistent with equilibrium predictions. In the treatment where the asymmetry between battlefields is strongest, we also find that the large battlefield receives on average more than a proportional share of resources. In a control treatment, all battlefields have the same value and our findings are consistent with previous experimental findings on Colonel Blotto games.

**Keywords:** Colonel Blotto, contest theory, majoritarian objective, resource allocation, experiment

**JEL Codes:** C72, C92, D72, D74

---

We thank the co-editor (Dan Kovenock), four anonymous referees, Subhasish Chowdhury, Judith Avrahami, seminar participants at New York University, Keele University, University of East Anglia, University of the Basque Country, and conference participants at the Voting Power in Practice Symposium at LSE 2011, M-BEES 2011, SING7 2011, Contest, Mechanisms and Experiments Conference at Exeter 2012, SAET 2012, GAMES 2012 and ESEM-EEA 2013. The equilibrium computations were carried out on the High Performance Computing Cluster supported by the Research and Specialist Computing Support service at the University of East Anglia.

---

M. Montero (✉)

University of Nottingham, University of the Basque County (FAE I) and Ikerbasque  
School of Economics, University of Nottingham, University Park, Nottingham NG7 2RD, United Kingdom  
Email : [maria.montero@nottingham.ac.uk](mailto:maria.montero@nottingham.ac.uk)

A. Possajennikov

School of Economics, University of Nottingham, University Park, Nottingham NG7 2RD, United Kingdom  
Email : [alex.possajennikov@nottingham.ac.uk](mailto:alex.possajennikov@nottingham.ac.uk)

M. Sefton

School of Economics, University of Nottingham, University Park, Nottingham NG7 2RD, United Kingdom  
Email : [martin.sefton@nottingham.ac.uk](mailto:martin.sefton@nottingham.ac.uk)

T.L. Turocy

School of Economics, University of East Anglia, Norwich Research Park, Norwich NR4 7TJ, United Kingdom  
Email : [T.Turocy@uea.ac.uk](mailto:T.Turocy@uea.ac.uk)

## 1. INTRODUCTION

We investigate theoretically and experimentally a multiple-battlefield conflict in which battlefields may have different values. As in the classic Colonel Blotto game, two players compete to win battlefields by allocating a fixed amount of resources across them. A battlefield is won by the player who spends the most on it, and the winner of the game is the player winning the battlefields with highest total value. In terms of the classification in Kovenock and Roberson (2012) we study a multiple battlefield conflict with auction contest success function (CSF), budget-constrained use-it-or-lose-it costs, and a weighted majority objective since for the overall win a player needs to win a majority of battlefields, weighted by their values.

It is perhaps obvious that a player should favor more important battlefields relative to less important ones. But by how much? If a large battlefield is worth twice as much as a small battlefield, should it command twice as many resources? Should a player concentrate the resources on the minimal set of battlefields necessary for the overall win or should he spread resources over all battlefields? And do human subjects behave in the way predicted by equilibrium? To address these questions, we study Blotto games with asymmetric battlefields using apex games to characterize the asymmetry. This class of games is the simplest class of majoritarian contests with asymmetric battlefields without ties (see Isbell 1959). In this class of games, there is one large and several small battlefields, and the contest can be won by winning the large and any one of the small battlefields, or by winning all small battlefields.

Little is known about this class of Blotto games, either theoretically or empirically. Theoretically, if players are symmetric each player has the same equilibrium probability of winning the contest. Further, except for trivial cases (e.g. where the value of one battlefield is greater than the combined value of all other battlefields) any equilibrium of the game must involve mixed strategies, as in the classic Colonel Blotto game. Beyond this we know of only limited results due to Young (1978). He interprets this contest as a game between lobbyists with opposing interests, competing to bribe voters that may differ in the number of votes they control. The lobbyists' aim is to win a majority of votes. Young considers two cases with one large and several small voters, such that a lobbyist wins if he gets the support of the large and any one of the small voters, or of all the small voters; both of his games are particular instances of apex games. He reports the equilibrium expected expenditures, which are

disproportionately skewed towards the large voter, but is silent on other properties of equilibrium behavior.<sup>1</sup>

For each of our treatments, we use numerical methods to compute an equilibrium under the restriction that the small battlefields are treated symmetrically. As well as predicting that the large battlefield receives on average more than a proportional share of the resources of the players, this equilibrium also predicts that with a very high probability players spread their budget over more battlefields than are needed to win the contest. We also show that these features must hold in any equilibrium of the game, whether it treats small battlefields symmetrically or not (see proposition 2 in section 3).

Our experimental treatments are based on the games discussed by Young. In the experiment, two subjects compete for battlefields and receive points for battlefields won. A battlefield is won by the player spending the most on it, and some battlefields are worth more points than others. The winner of the game is the subject obtaining the most points. Although this is a simple game to describe, it is not clear whether subject behavior will conform well to theoretical predictions. On the one hand it is extremely unlikely that behavior will match equilibrium precisely: identifying equilibrium is computationally challenging and presumably beyond the ability of experimental subjects (indeed, we use numerical methods to pin down an equilibrium). On the other hand, even simpler versions of Colonel Blotto are notoriously difficult to solve, and yet, as we discuss in the next section, experiments with these have found behavior to be qualitatively in line with key features of equilibrium.

Equilibrium mixed strategies have a complex structure, and play in the experiment does not match it in detail. In all treatments we can identify strategies that, if pitted against our subjects' strategies, would win more often than not. With more battlefields the game is more complicated, both because there are more possible strategies, and because some simple strategies are more exploitable in the sense that it is possible to obtain a higher expected payoff against those strategies. Nevertheless, the degree of exploitability of actual play is

---

<sup>1</sup> Young describes the lobbying game in which lobbyists have infinitely divisible budgets. He then states that "Instead of attempting to compute [the expected allocation of the lobbyists' budget to each voter] explicitly by integration, a useful approach is to approximate the continuous strategy spaces by a finite grid. A natural way to do this is by allocating a large but finite number  $u$  of *indivisible* units to each lobbyist that may be distributed among the voters. This leads to a two-person, zero-sum matrix game whose equilibrium solutions were used to approximate the equilibrium solution to the infinite lobbying game." He then states, without giving further details of the derivations, the resulting values.

roughly the same across treatments, indicating that even though the game with a larger number of battlefields is more complicated, the behavior of subjects is as close to equilibrium, at least by this measure.

Even though play does not match the equilibrium exactly, we find evidence for some of the qualitative features of equilibrium predictions. Strategies that spread resources over more battlefields than is necessary to win the contest are played often and their frequency increases over time in all treatments. In the game with more battlefields where the asymmetry between small and large battlefields is more pronounced, the large battlefield receives on average more than a proportional share of the total resources.

In the next section we review the related literature on Colonel Blotto games. Section 3 describes our game and its theoretical properties. Section 4 describes our experimental design and procedures. Results are presented in Section 5 and Section 6 concludes.

## **2. RELATED LITERATURE**

The Colonel Blotto game was introduced in Borel (1921), where he considers three identical battlefields, an auction contest success function, budget-constrained use-it-or-lose-it costs and a majoritarian objective. Equilibria for the game were first identified in Borel and Ville (1938). Later studies introduced the term “Colonel Blotto” and adopted an additive objective (i.e. players maximize the total value of battlefields won). For this latter formulation of the game, Roberson (2006) shows that when all battlefields are identical a player’s marginal distribution of her expenditure on a battlefield must be uniform in any equilibrium; Hart (2008) extends this analysis to the case of a discrete budget. Also using the additive objective formulation with symmetric budgets, Gross and Wagner (1950) consider the game with three asymmetric battlefields with arbitrary values, and Thomas (2013) extends the analysis to an arbitrary number of asymmetric battlefields. Thomas shows that uniform marginals, where the mean expenditure on a battlefield is proportional to its value, are a sufficient condition for equilibrium. It is not known if this condition is also necessary, or whether it is always possible to construct multivariate distributions with this property.<sup>2</sup>

Young (1978) extended the original majoritarian objective model of Borel (1921) to

---

<sup>2</sup> There are also results on the additive objective with asymmetric budgets. Gross and Wagner (1950) solve the case of two battlefields with arbitrary values. Roberson (2006) characterizes the equilibrium for an arbitrary number of symmetric battlefields.

the case of asymmetric battlefields. He studies vote-buying games in which players with opposite interests allocate their budgets across voters ('battlefields'). Young considered two particular apex games, and solved them assuming a finitely divisible budget (the size of which is not stated) and reported the expected amounts allocated to the large battlefield (but not the equilibrium strategies). These are the two games that we study theoretically and experimentally in this paper.

A small number of studies have recently examined variants of the Colonel Blotto game experimentally.<sup>3</sup> Avrahami and Kareev (2009) focus on contests between players with differing strengths. In their contests the two players have different budgets, and they find that subject behavior is sensitive to the relative budgets in the way predicted by equilibrium. They conclude: "the results indicate that naïve players can behave, intuitively, in a way that approximates the sophisticated game-theoretic solution." Chowdhury et al. (2013) also study a game between asymmetric players, and compare the auction CSF with a lottery CSF. They find that the probabilities of winning for players 1 and 2 are as predicted by equilibrium, and the bidding strategies differ across treatments in the direction predicted. They note some interesting deviations from equilibrium, but overall conclude "... it took only one hour for subjects who were unfamiliar with this game to exhibit behavior consistent with equilibrium".<sup>4</sup> In both of these studies battlefields are symmetric and players have an additive objective. In contrast, our experiment studies a game with a majoritarian objective and asymmetric battlefields.

Mago and Sheremeta (2012) study a setting with a majoritarian objective and an auction CSF, but with linear costs. Thus, in their experiment subjects decide how much of their budget to allocate to the contest, as well as how to allocate resources across battlefields. They find that subjects make higher aggregate bids than predicted; this is a common finding in contest experiments, but one which is not possible in our fixed-budget setting.<sup>5</sup>

---

<sup>3</sup> Dechenaux et al. (2014) survey the experimental literature on contests more generally.

<sup>4</sup> Arad and Rubinstein (2012) also study a game with an auction CSF and budget-constrained-use-it-or-lose-it costs using a round-robin tournament in which each subject's allocation is pitted against everybody else's. They observe significant deviations from equilibrium and interpret subjects' choices as reflecting iterated reasoning in several dimensions. Note however that their subjects play a one-shot game, with no opportunity for learning.

<sup>5</sup> Another related paper is Hortala-Vallve and Llorente-Saguer (2010). They consider a modification of the Colonel Blotto game with an additive objective in which the two players have cardinal valuations of the battlefields, and these valuations are private information. Their main focus is on the efficiency of the outcome, an issue that does not arise in our setting.

There are two recent papers with asymmetric battlefields and budget-constrained use-it-or-lose-it costs, but that differ from ours in other dimensions. Avrahami et al. (2014) study a multiple-battlefield contest with asymmetric battlefields and an auction CSF, but with an additive objective. There are eight battlefields with different values, and players may also differ in the amount of resources they command. Avrahami et al. (2014) find that average wins reflect the players' relative strength and that there is a correlation between the allocation to a battlefield and its value. Duffy and Matros (2015) study a multiple-battlefield contest setting with asymmetric battlefields and a majoritarian objective, but where the outcome for each individual battlefield is determined using a lottery rather than an auction CSF. An important implication of this assumption is that, unlike in our setting, their game has a pure-strategy equilibrium rather than a mixed-strategy equilibrium. Interestingly, in the settings considered by Young (1978), the equilibrium of the game with a lottery CSF also predicts that the large voter receives a share of the budget that exceeds its proportion of the votes. Duffy and Matros (2015) find that treatment differences conform to equilibrium predictions.

### 3. AN ASYMMETRIC BLOTTO GAME

The game we study is a version of the classic Colonel Blotto game introduced by Borel (1921). The main difference from the original game is that we consider the case in which battlefields do not all have the same value. In our game  $\Gamma$ , two players,  $A$  and  $B$ , simultaneously allocate identical discrete endowments  $E \in \mathbb{Z}_{++}$  across a set  $N = \{1, \dots, n\}$  of battlefields. Each battlefield  $i$  has an associated value  $v_i > 0$ , which is the same for both players. A pure strategy for player  $j$  is an allocation of the endowment  $E$  across the battlefields; formally, it is a vector of nonnegative integers  $(x_i^j)_{i \in N}$  with  $\sum_{i \in N} x_i^j = E, j = A, B$ .

Denote the set of pure strategies as  $X$ . As in Borel's original formulation, both players' entire endowments must be spent on the battlefields. Each battlefield is won by one of the players according to the auction contest success function. Battlefield  $i$  is won by player  $A$  if  $x_i^A > x_i^B$ , by player  $B$  if  $x_i^A < x_i^B$ , and is randomly allocated with equal probability if  $x_i^A = x_i^B$ . In the event of ties on multiple battlefields, randomization is realized independently on each.

The winner of the game is the player who wins battlefields with the greatest total

value. Let  $N_j$  denote the set of battlefields won by player  $j$ . Player  $A$  wins the game if

$$\sum_{i \in N_A} v_i > \sum_{i \in N_B} v_i, \text{ and player } B \text{ wins the game if } \sum_{i \in N_A} v_i < \sum_{i \in N_B} v_i. \text{ In the cases we consider, the}$$

battlefield values are such that, for any subset  $G \subseteq N$  of battlefields,  $\sum_{i \in G} v_i \neq \sum_{i \in N \setminus G} v_i$ , so that

there is always a clear winner. Assuming that the utility of a win exceeds that of a loss, the payoff from winning can be normalized to 1 and the payoff from losing can be normalized to 0. The binary outcome structure entails that equilibrium is independent of risk attitudes. Let  $u(x, y)$  denote the probability that player  $A$  wins the game if player  $A$  chooses strategy  $x$  and player  $B$  chooses strategy  $y$ .

Since the game is finite, a Nash equilibrium always exists, and because the game is symmetric across players, each player wins with probability one-half in any equilibrium.<sup>6</sup> Moreover, except in trivial cases, the only equilibria of the game are in mixed strategies. However, beyond this little is known about equilibrium.<sup>7</sup> In this paper, we specialize to the case of apex games, and consider the two games previously discussed by Young (1978). In one of his games, which we will call APEX4, there are four battlefields with values  $\mathbf{v} = (2, 1, 1, 1)$ , while in the other, which we will call APEX5, there are five battlefields with values  $\mathbf{v} = (3, 1, 1, 1, 1)$ . Young reported the expected amounts allocated to the large battlefield as exactly one-half in APEX4, and approximately 56% in APEX5. However, Young does not report equilibrium strategies for these games, which is the problem to which we now turn.

Calculating equilibria for these games is a daunting task. The size of the strategy

---

<sup>6</sup> Constant-sum games have a well-defined value, so all equilibria of a constant-sum game yield the same payoffs to the players. Symmetric games always have at least one equilibrium that is symmetric, and a symmetric equilibrium in this game necessarily results in each player winning with probability one-half.

<sup>7</sup> One thing we know is that results for the additive objective game with a continuously divisible budget do not carry over to majoritarian objective games. Take  $\mathbf{v} = (2, 1, 1, 1)$  and suppose the budget  $E = 5$  is continuously divisible. Consider any strategy such that the amount allocated to the first battlefield is uniformly distributed between 0 and 4, and the amount allocated to each of the other battlefields is uniformly distributed between 0 and 2. Such a strategy constitutes an equilibrium in the additive objective case (Thomas 2013). (For an example of such a strategy draw  $\varepsilon$  from a uniform distribution on the interval  $[0, 1]$  and place  $4 - 4\varepsilon$  units on the large battlefield and assign  $2\varepsilon$ ,  $\varepsilon$ , and  $1 + \varepsilon$  to the small battlefields where each of the 6 permutations of assignments to small battlefields is equally likely. Note that the total placed on all battlefields is  $4 - 4\varepsilon + 2\varepsilon + \varepsilon + 1 + \varepsilon = 5$  and the marginal distribution on each battlefield is uniform.) For the majoritarian objective case any such strategy can be bettered by a strategy that puts 3 on the first battlefield and 2 on the second. This alternative strategy wins the second battlefield with probability 1, and the first battlefield with probability 3/4, hence it wins with probability 3/4 overall. Analogously, for  $\mathbf{v} = (3, 1, 1, 1, 1)$  and  $E = 7$ , any strategy such that the allocation is uniform on  $[0, 6]$  for the first battlefield and on  $[0, 2]$  for each of the others can be beaten with a probability of at least 5/6, since placing 5 on the first battlefield and 2 on the second would beat it with probability 5/6.

space of the game increases rapidly in both the budget and the number of battlefields. For an endowment  $E = 120$ , corresponding to the parameterization we use in our experiments, there are  $|X| = 302,621$  pure strategies in APEX4, and  $|X| = 9,381,251$  in APEX5. The analysis can be simplified substantially by using the fact that small battlefields are strategically equivalent, in the sense that they can be relabeled arbitrarily without changing the main strategic characteristics of the game. There are  $M = 6$  such labelings for APEX4, and  $M = 24$  for APEX5. We say that two pure strategies  $x$  and  $y$  are permutations, and write  $x \sim y$  if there exists a permutation  $\mu: N \rightarrow N$  with  $\mu(1) = 1$ , such that  $y_{\mu(i)} = x_i$  for all  $i \in N$ . For example, in APEX4 with an endowment of  $E = 120$ ,  $(80, 40, 0, 0) \sim (80, 0, 40, 0) \sim (80, 0, 0, 40)$ . The relation  $\sim$  induces a partition of the set  $X$  of pure strategies of the game. Let  $P$  denote the partition induced by  $\sim$ , where  $P(x)$  denotes the element of the partition to which strategy  $x$  belongs.

There are many interesting qualitative features of a mixed strategy over  $X$  which depend only on the distribution of play over the elements of  $P$ . For example, we will be interested in expected allocations to the large battlefield, and in the probability with which players use “minimal winning strategies” (i.e. strategies that allocate positive amounts only to a set of battlefields which are just enough to win the game if the player wins all those battles). This information can be recovered from knowledge of the distribution over  $P$ .

Consider the set of mixed strategies  $\Delta(X)$  in game  $\Gamma$  (pure strategies are included in this set as degenerate mixed strategies). We denote a generic mixed strategy of player  $j$  (i.e. an element of  $\Delta(X)$ ) by  $\pi^j$ , with  $\pi^j(x)$  denoting the probability that the pure strategy  $x$  is played in  $\pi^j$ . We use  $\pi = (\pi^A, \pi^B)$  to denote a generic strategy profile. We say a mixed strategy  $\pi^j$  in  $\Delta(X)$  is *object-symmetric* if, for all pure strategies  $x$  and  $y$  such that  $x \sim y$ ,  $\pi^j(x) = \pi^j(y)$ . Such a strategy treats objects (battlefields) that have the same value symmetrically in terms of allocation of the endowment.

For each strategy  $x \in X$ , we can construct a mixed strategy in  $\Delta(X)$  that randomizes uniformly between all the permutations of  $x$  (i.e., the elements of  $P(x)$ ) as  $\sum_{z \in P(x)} \frac{1}{|P(x)|} z$ . A mixed strategy constructed in this way is a particular type of object-symmetric strategy which



puts the entire weight on one element of the partition. We define an auxiliary game  $\hat{\Gamma}$ , in which the pure strategy set  $\hat{X}$  corresponds to the set of strategies constructed above and consequently to the partition  $P$ .<sup>8</sup> The payoffs from a strategy combination in  $\hat{\Gamma}$  will be the payoffs from the corresponding combination of object-symmetric strategies in the game  $\Gamma$ . The practical effect of this construct is that analyzing the auxiliary game is equivalent to analyzing the original game restricted to object-symmetric strategies.

Formally, there is a one-to-one mapping  $p$  that assigns to each pure strategy  $\hat{x}$  in  $\hat{X}$  an element  $p(\hat{x})$  of the partition  $P$ . We define the probability that player  $A$  wins game  $\hat{\Gamma}$  if player  $A$  plays strategy  $\hat{x}$  and player  $B$  plays strategy  $\hat{y}$  as

$$\hat{u}(\hat{x}, \hat{y}) = \sum_{x \in p(\hat{x})} \sum_{y \in p(\hat{y})} \frac{1}{|p(\hat{x})|} \frac{1}{|p(\hat{y})|} u(x, y).$$

Intuitively, each strategy  $\hat{x}$  in  $\hat{\Gamma}$  corresponds to uniform randomization over members of the element  $p(\hat{x})$  of the partition  $P$  in  $\Gamma$ .

The benefit of considering the object-symmetric restriction encoded in auxiliary game  $\hat{\Gamma}$  is that the auxiliary game has a much smaller set of pure strategies, with  $|\hat{X}| = |P| = 52,311$  for APEX4 and  $|\hat{X}| = |P| = 430,256$  for APEX5. Our first result is that, for the problem of computing one equilibrium of  $\Gamma$ , it suffices to compute an equilibrium of  $\hat{\Gamma}$ . An equilibrium of  $\hat{\Gamma}$  corresponds to an equilibrium of  $\Gamma$  in which object-symmetric strategies are played. In the latter equilibrium, the probability of each partition element  $p(\hat{x})$  is divided equally between all the strategies in this partition element.

**Proposition 1.** *Let  $(\hat{\pi}^A, \hat{\pi}^B)$  be an equilibrium in mixed strategies in  $\hat{\Gamma}$ . Then, there is a corresponding equilibrium  $(\pi^A, \pi^B)$  in  $\Gamma$ , in which  $\pi^j(x) = \frac{1}{|P(x)|} \hat{\pi}^j(\hat{x})$ , where  $p(\hat{x}) = P(x)$ , for  $j = A, B$  and for all strategies  $x$  in  $X$ .*

**Proof.** The proof is by contradiction. Suppose  $(\hat{\pi}^A, \hat{\pi}^B)$  is an equilibrium of  $\hat{\Gamma}$  but  $(\pi^A, \pi^B)$  in the statement of the proposition is not an equilibrium of  $\Gamma$ . Then for at least one

---

<sup>8</sup> Note that  $\hat{X}$  does not correspond to the set of all object-symmetric strategies in  $\Gamma$ ; this set would correspond to  $\Delta(\hat{X})$ .

player, say player  $A$ , there is a strategy  $y$  in  $X$  which gives strictly higher payoff to that player. Because  $\pi^B$  is object-symmetric by construction, any strategy  $z \sim y$  must also be a strictly better response to  $\pi^B$ . Then, all elements of  $P(y)$  are strictly better responses for  $A$  in  $\Gamma$ , and so it follows that the strategy  $\hat{y}$  with  $p(\hat{y}) = P(y)$  is a strictly better response for  $A$  in  $\hat{\Gamma}$ , which contradicts the original assumption. ■

Note that the converse statement, i.e., that an equilibrium in object-symmetric strategies of  $\Gamma$  corresponds to an equilibrium in  $\hat{\Gamma}$ , is also true. This is because the set  $\Delta(\hat{X})$  corresponds to the set of all object-symmetric strategies in  $\Delta(X)$ , and thus a subset of  $\Delta(X)$ .

An immediate consequence of proposition 1 is that, if  $\Gamma$  has a unique equilibrium, then this equilibrium must be object-symmetric; it must also be player-symmetric (i.e., both players use the same mixed strategy) since  $\Gamma$  is a symmetric game. However,  $\Gamma$  need not have a unique equilibrium, and therefore there may be equilibria of  $\Gamma$  where the strategies cannot be obtained by computing an equilibrium of  $\hat{\Gamma}$  as above.<sup>9</sup>

We will show that for any equilibrium of  $\Gamma$  there is an equilibrium of  $\hat{\Gamma}$  (which corresponds to an object-symmetric equilibrium of  $\Gamma$  by proposition 1) that preserves the distribution of play over  $P$ . We first introduce some notation. For any mixed strategy  $\pi^j$  of player  $j$  in  $\Gamma$ , we denote the induced probability distribution over  $P$  as  $Q_{\pi^j}$ , where

$$Q_{\pi^j}(P(x)) = \sum_{z \in P(x)} \pi^j(z). \text{ Since each pure strategy in } \hat{\Gamma} \text{ corresponds to an element of } P, \text{ a}$$

mixed strategy  $\hat{\pi}^j$  induces the probability distribution  $Q_{\hat{\pi}^j}$  over  $P$  as  $Q_{\hat{\pi}^j}(P(x)) = \hat{\pi}^j(\hat{x})$ ,

where  $p(\hat{x}) = P(x)$ . For a strategy profile  $\pi = (\pi^A, \pi^B)$  in  $\Gamma$ , we take the average of the distributions of the two players,  $Q_{\pi}(P(x)) = 1/2 \cdot Q_{\pi^A}(P(x)) + 1/2 \cdot Q_{\pi^B}(P(x))$ , as the induced probability distribution over  $P$ . Similarly, the distribution over  $P$  induced by strategy profile

$\hat{\pi} = (\hat{\pi}^A, \hat{\pi}^B)$  in  $\hat{\Gamma}$  is  $Q_{\hat{\pi}}(P(x)) = 1/2 \cdot Q_{\hat{\pi}^A}(P(x)) + 1/2 \cdot Q_{\hat{\pi}^B}(P(x))$ . Now we prove the result.

---

<sup>9</sup> For example, in the game APEX3 (a degenerate apex game where  $v = (1, 1, 1)$ ) with an endowment  $E = 5$  analyzed in the appendix of the working paper Montero et al. (2013), consider both players randomizing equally between the three pure strategies (3,2,0), (0,3,2), and (2,0,3). This strategy combination is an equilibrium of the game but is not object-symmetric, since there are three other permutations, namely (3,0,2), (0,2,3) and (2,3,0), that are not being used. This equilibrium is in fact part of a continuum of equilibria which place probability 1 on the permutations of (3,2,0), and the equilibrium set does include the object-symmetric case where all six permutations are given equal probability.

**Proposition 2.** *Let  $(\pi^A, \pi^B)$  be an equilibrium mixed strategy profile in  $\Gamma$ , and consider the associated probability distribution of play over the partition  $P$  of  $X$ . Then there exists an equilibrium  $(\hat{\pi}^A, \hat{\pi}^B)$  of the auxiliary game  $\hat{\Gamma}$  that induces the same probability distribution of play over  $P$ .*

**Proof.** Consider all  $M$  permutations of the small battlefields and construct a set of  $M$  mixed strategy profiles  $\{(\pi_m^A, \pi_m^B)\}_{m=1, \dots, M}$ , each strategy profile corresponding to a renaming of the small battlefields. Each mixed strategy profile in this set is an equilibrium in  $\Gamma$  and induces the same probability distribution over  $P$ . Because the set of Nash equilibria of a two-person constant-sum game is convex, the convex combination of these  $M$  equilibria  $\sum_{m=1}^M \frac{1}{M} (\pi_m^A, \pi_m^B)$  must also be an equilibrium in  $\Gamma$ . This latter equilibrium is by construction object-symmetric and hence corresponds to an equilibrium  $(\hat{\pi}^A, \hat{\pi}^B)$  of  $\hat{\Gamma}$ . For each player  $j$ , the strategy  $\hat{\pi}^j$  in this equilibrium induces the same distribution over  $P$  as the strategy  $\pi^j$  in the original equilibrium  $(\pi^A, \pi^B)$ , thus the equilibrium induces the same distribution over  $P$  as the original equilibrium. ■

As a final step to reduce computational complexity, we compute only player-symmetric equilibria in  $\hat{\Gamma}$ . If the game  $\hat{\Gamma}$  has a unique player-symmetric equilibrium, the analysis comes at no cost in terms of the distribution of play over  $P$ .

**Proposition 3.** *Suppose  $\hat{\pi} = (\hat{\pi}^A, \hat{\pi}^B)$  with  $\hat{\pi}^A = \hat{\pi}^B$  is the unique player-symmetric equilibrium of  $\hat{\Gamma}$ . Any equilibrium  $(\pi^A, \pi^B)$  of  $\Gamma$  induces the same probability distribution over  $P$  as  $\hat{\pi}$ .*

**Proof.** Suppose there are two equilibria of  $\Gamma$  with different distributions over  $P$ . By proposition 2, each of the two distributions can be realized by an equilibrium in  $\hat{\Gamma}$ . Therefore, there are two equilibria in  $\hat{\Gamma}$ . If these two equilibria are player-symmetric this would be a contradiction, hence at least one of them, say  $(\hat{\rho}^A, \hat{\rho}^B)$ , has  $\hat{\rho}^A \neq \hat{\rho}^B$ . It is a property of zero-sum games that  $(\hat{\rho}^A, \hat{\rho}^A)$  and  $(\hat{\rho}^B, \hat{\rho}^B)$  are also equilibria of  $\hat{\Gamma}$ , contradicting the assumption that  $\hat{\Gamma}$  has a unique player-symmetric equilibrium. ■

In the case of  $\hat{\Gamma}$  having multiple player-symmetric equilibria, we can formalize a

bound on how inaccurate our characterization is of the properties of distributions over the partition  $P$  arising from equilibrium play in  $\Gamma$ . We say that a set  $\hat{\Psi}$  of mixed strategy profiles in  $\hat{\Gamma}$  is bounded by  $\varepsilon$  if the maximum difference in probabilities that any two strategy profiles in the set assign to any pure strategy is less than  $\varepsilon$ , i.e.

$\max_{\hat{\pi}, \hat{\rho} \in \hat{\Psi}} \max_{\hat{x} \in \hat{X}} \left| 1/2 \cdot (\hat{\pi}^A(\hat{x}) + \hat{\pi}^B(\hat{x})) - 1/2 \cdot (\hat{\rho}^A(\hat{x}) + \hat{\rho}^B(\hat{x})) \right| \leq \varepsilon$ . We say that a set  $\Theta$  of probability distributions over  $P$  is bounded by  $\varepsilon$  if  $\max_{Q_\pi, Q_\rho \in \Theta} \max_{p \in P} |Q_\pi(p) - Q_\rho(p)| \leq \varepsilon$ .

Let  $\Pi$  denote the set of equilibria in  $\Gamma$  and  $\hat{\Pi}_{ps}$  denote the set of player-symmetric equilibria in  $\hat{\Gamma}$ .

**Proposition 4.** *If the set of player-symmetric equilibria of  $\hat{\Gamma}$  is bounded by  $\varepsilon > 0$ , then the set of equilibrium probability distributions over  $P$  in  $\Gamma$  is also bounded by  $\varepsilon$ .*

**Proof.** Consider any two equilibria  $(\pi^A, \pi^B)$  and  $(\rho^A, \rho^B)$  of the game  $\Gamma$  with probability distributions  $Q_\pi = 1/2 \cdot Q_{\pi^A} + 1/2 \cdot Q_{\pi^B}$  and  $Q_\rho = 1/2 \cdot Q_{\rho^A} + 1/2 \cdot Q_{\rho^B}$  over  $P$ . By proposition 2, there exist equilibria  $(\hat{\pi}^A, \hat{\pi}^B)$  and  $(\hat{\rho}^A, \hat{\rho}^B)$  with probability distributions  $Q_{\hat{\pi}} = Q_\pi$  and  $Q_{\hat{\rho}} = Q_\rho$ . Since the game  $\hat{\Gamma}$  is zero-sum,  $(\hat{\pi}^A, \hat{\pi}^A), (\hat{\pi}^B, \hat{\pi}^B)$  and  $(\hat{\rho}^A, \hat{\rho}^A), (\hat{\rho}^B, \hat{\rho}^B)$  are also equilibria of  $\hat{\Gamma}$ . They are player-symmetric, and by the reasoning in the proof of proposition 2, they induce distributions  $Q_{\pi^A}, Q_{\pi^B}$  and  $Q_{\rho^A}, Q_{\rho^B}$  over  $P$ . Since each  $\hat{x}$  maps to exactly one partition element  $p(\hat{x})$  of  $P$ , the bound on the set  $\hat{\Pi}_{ps}$  of player-symmetric equilibria of  $\hat{\Gamma}$  implies  $\max_{\hat{\pi}, \hat{\rho} \in \hat{\Pi}_{ps}} \max_{p \in P} |Q_{\hat{\pi}}(p) - Q_{\hat{\rho}}(p)| \leq \varepsilon$ . Therefore  $\max_{p \in P} |Q_{\pi^A}(p) - Q_{\rho^A}(p)| \leq \varepsilon$  and  $\max_{p \in P} |Q_{\pi^B}(p) - Q_{\rho^B}(p)| \leq \varepsilon$ . Consider  $\max_{p \in P} |Q_\pi(p) - Q_\rho(p)| = \max_{p \in P} \left| 1/2 \cdot (Q_{\pi^A}(p) + Q_{\pi^B}(p)) - 1/2 \cdot (Q_{\rho^A}(p) + Q_{\rho^B}(p)) \right| = 1/2 \cdot \max_{p \in P} |Q_{\pi^A}(p) - Q_{\rho^A}(p) + Q_{\pi^B}(p) - Q_{\rho^B}(p)| \leq 1/2 \cdot \left( \max_{p \in P} |Q_{\pi^A}(p) - Q_{\rho^A}(p)| + \max_{p \in P} |Q_{\pi^B}(p) - Q_{\rho^B}(p)| \right) \leq \varepsilon$ , where the last step follows from the bounds on the set  $\hat{\Pi}_{ps}$ . Thus  $\max_{Q_\pi, Q_\rho \in Q_\Pi} \max_{p \in P} |Q_\pi(p) - Q_\rho(p)| \leq \varepsilon$ , where  $Q_\Pi$  is the set of probability distributions over  $P$  induced by equilibria in  $\Gamma$  (player-symmetric or

not, in object-symmetric strategies or not), which means that this set is also bounded by  $\varepsilon$ . ■

Appendix A gives technical details of our computational approach, which we implemented with the support of the Gambit package (McKelvey et al. 2013). For the endowment  $E = 120$ , for APEX4 we compute an equilibrium in which only 1,479 out of 52,311 strategies in  $\hat{X}$  are played (and therefore only that many elements of  $P$  have a positive probability in the induced distribution), while for APEX5 the computed equilibrium involves randomization over 821 elements out of 430,256. The fact that equilibrium supports are small makes the computation of equilibrium substantially more feasible.

We have analyzed numerically the equilibria of games APEX4 and APEX5 with endowments varying up to 120. Because we use floating-point arithmetic in computing equilibria, the possibility of round-off error means we cannot conclusively show uniqueness. However, for larger values of the endowment, the range of strategy probabilities and other properties of behavior which are consistent with equilibrium is small. For most values of the endowment, we are able to bound these ranges to be at most  $10^{-6}$ , which, in light of proposition 4, means we can conclude that the equilibrium predictions we assert below are for all practical purposes pinned down. We conjecture that for many values of the endowment, the uncertainty is indeed rounding error and the equilibrium in object-symmetric strategies is unique, which would yield a unique equilibrium distribution over  $P$  in  $\Gamma$ . We show below the properties of equilibria in the game with  $E = 120$ ; our computations indicate these predictions are not especially sensitive to the size of the endowment.

Although only a small fraction of the possible strategies appear in equilibrium, nevertheless the equilibria are sufficiently complex that it is not plausible that participants would follow them exactly. Instead, we next report aggregated statistics about equilibrium behavior, based on the equilibrium distributions over  $P$ , which we will use as the foundation for the analysis of the data from our experiments.


We begin by calculating the expected proportion of resources allocated to the large battlefield. For APEX4, the predicted allocation to the large battlefield is 50% (to 2 d.p.), as claimed by Young (1978). However, for APEX5, we compute an equilibrium prediction of 58%, greater than the 56% claimed by Young. In each case the large battlefield receives a share of resources exceeding its value as a proportion of the total value, and this is true of any

equilibrium of the game.

**Equilibrium Prediction 1.** *The large battlefield receives, on average, a super-proportional share of resources.*

We next turn to distributional predictions on allocations to the large and small battlefields generated by equilibrium strategies. Figure 1 displays the equilibrium distributions of allocations in both APEX4 and APEX5 for a budget of  $E = 120$ . We note that the marginal cumulative distribution function on the large battlefield is convex, and that on the small battlefields is concave. This implies that for the large battlefield the marginal probability mass function is increasing and larger allocations on its support are more likely than smaller ones, whereas for the small battlefields the marginal probability mass function is decreasing and smaller allocations on the support are more likely than larger ones. For both games, the maximum allocation to the large battlefield is strictly less than the entire endowment; in APEX4 the maximum allocation is about 90% of the endowment, whereas in APEX5 it is about 95%.

**Fig. 1** CDFs of equilibrium marginal distributions for  $E = 120$   
Left panel: APEX4. Right panel: APEX5

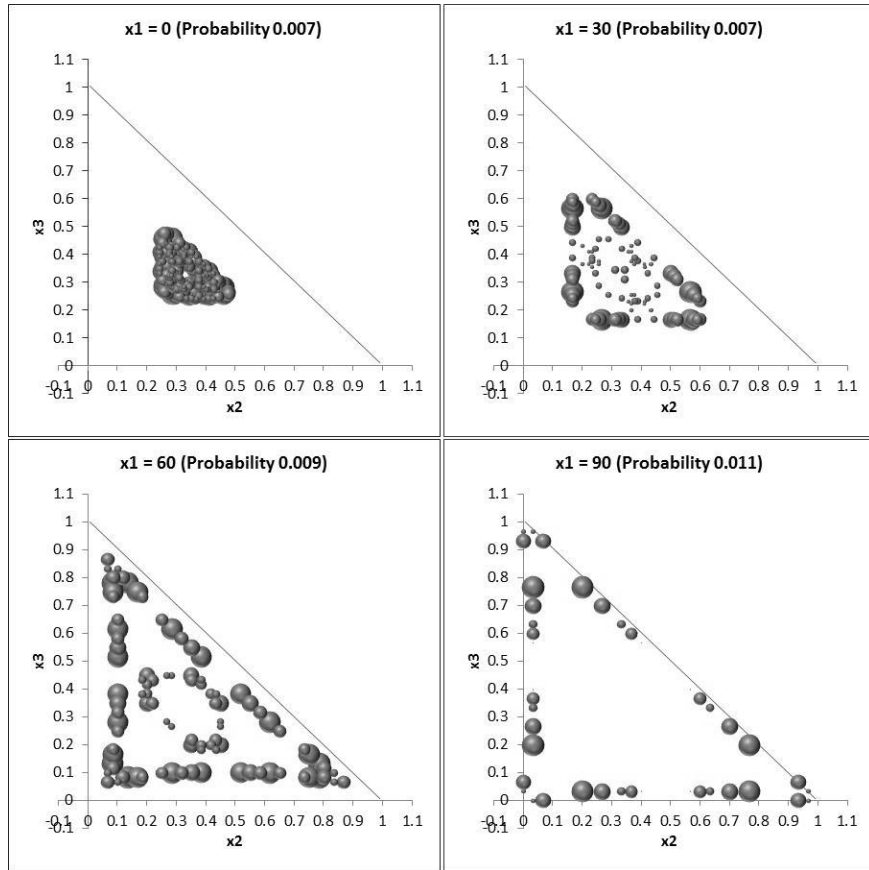


**Equilibrium Prediction 2.** *For the large battlefield the relative frequency of an allocation is increasing in the allocation (up to an upper bound); the opposite is true for a small battlefield.*

Because of the majoritarian objective of the game, the marginal distributions on large and small battlefields alone are not enough to characterize optimal play; the correlation structure among the components of the allocation is also important. One way to describe the

equilibrium distribution is to ask, conditional on a given allocation to the large battlefield, how the remaining resources are distributed over the small ones. Figure 2 plots, for APEX4 with an endowment of 120, the support of the joint distribution over small battlefields, conditional on four selected levels of allocation to the large one (the equilibrium probability of each of the four selected levels is reported in parentheses).

**Fig. 2** Equilibrium conditional distributions over small battlefields in APEX4



The triangles in figure 2 are to be interpreted as follows. For the specified value of  $x_1$ , the corresponding triangle shows  $x_2$  and  $x_3$  as proportions of the remainder  $120 - x_1$ ; the value of  $x_4$  can be read as the residual. The size of the dots is proportional to their probability mass. In equilibrium, when allocating zero or small amounts to the large battlefield, resources are allocated roughly equally across the small ones. As the allocation to the large battlefield grows, two of the remaining three battlefields are favored. The same qualitative pattern is also found in the equilibrium of APEX5, although the extra dimension in the allocation does not lend itself as easily to visualization in plots.

We can distill some information about joint distributions of allocations over battlefields in an intuitive and economically interesting way by considering what proportion of battlefields receive nonzero allocations. On the one hand, in apex games there are two types of sets of battlefields that are minimally necessary for winning: one can win the large battlefield and any one small battlefield, or one can win all the small battlefields. However, targeting resources only to such sets is risky, in that losing any single battlefield can result in a loss overall. For this reason, strategies that allocate resources to larger sets could be useful.

We say that a (pure) strategy has a target value of  $k$  if it allocates nonzero resources to battlefields with a total value  $k$ . Therefore, in APEX4, targeting a minimal winning set of battlefields means using a strategy with a target value of  $k = 3$ ; in APEX5, the minimal winning set of battlefields has a target value of  $k = 4$ . We call any strategy with a target value strictly above those values a *hedging* strategy. For any given mixed strategy, we can compute the proportion of play of pure strategies with a particular target value.

In APEX4, the equilibrium probability of strategies targeting a minimal set of battlefields (i.e., strategies with  $k = 3$ ) is just 6% (1% target the three small battlefields and 5% target the large and one small battlefield). In equilibrium it is more likely that a hedging strategy is used: the probability of a strategy targeting the large battlefield and two small (i.e., strategies with  $k = 4$ ) is 18%, and that of targeting all battlefields ( $k = 5$ ) is 76%. For APEX5, the values are 8% for a minimal winning set of battlefields (i.e.  $k = 4$ ) (1% target the four small battlefields and 7% target the large and one small battlefield), 28% for the large battlefield and two small battlefields ( $k = 5$ ), and 64% for all battlefields ( $k = 7$ ) (strategies with  $k = 6$  are never played in equilibrium). The qualitative relationship between these probabilities is robust to the value of the endowment.

**Equilibrium Prediction 3.** *Strategies targeting all battlefields have the highest probability, followed by those targeting the large battlefield and two small battlefields, followed by those targeting only a minimal winning set of battlefields.*

In the next section we describe an experimental design to see if the behavior of human subjects reflects the observations about equilibrium play.

#### 4. EXPERIMENTAL DESIGN AND PROCEDURES

The experiment was conducted at the University of Nottingham with 148 subjects recruited



from a university-wide pool of undergraduate students using ORSEE (Greiner 2004). The experiment consisted of nine computerized sessions, with no subject participating in more than one session. The experiment was programmed in z-Tree (Fischbacher 2007).

All sessions used an identical protocol. Upon arrival, subjects were given a written set of instructions that the experimenter read aloud.<sup>10</sup> Subjects were then randomly paired and played a sequence of 45 rounds of a game against the same opponent. Subjects were not told who of the other people in the room was paired with them, but they knew that they were playing the same subject throughout. Subjects were not allowed to communicate with one another throughout the session. In each round a subject won either £0.50 or nothing and at the end of the session subjects were paid their accumulated earnings from all 45 rounds.

In each round subjects were given a budget of 120 tokens and used these to bid for battlefields, each of which was worth a given number of points.<sup>11</sup> A subject could only submit bids that added up to 120, and had 90 seconds to submit the bids.<sup>12</sup> A subject won a battlefield if he outbid his opponent on that battlefield (or, in the case of a tie, if he won a random computer draw). The subject that won the most points in a given round earned £0.50. At the end of each round subjects were informed of how much they and their opponent bid for each battlefield, who won each battlefield, and how much they earned.

We conducted three treatments. APEX4 and APEX5 use Young's (1978) apex games with four and five battlefields respectively. For these treatments battlefields and values were listed on subjects' computer screens with the large battlefield always listed first. For comparison, we also ran a treatment using a game (APEX3) with three symmetric battlefields.<sup>13</sup> APEX3 can be viewed as a degenerate apex game in which any of the battlefields can arbitrarily be designated as large, and the other two as small, but we will refer to the battlefield listed first on-screen as the "large battlefield". The difference between bids on the "large" and "small" battlefields in APEX3 can be viewed as a measure of a positional order effect. For each treatment we conducted three sessions with between 14 and 20 subjects

---

<sup>10</sup> Instructions for one of the treatments can be found in Appendix B.

<sup>11</sup> In the instructions and on screen we did not use the "battlefield" terminology, but rather referred to "objects".

<sup>12</sup> If a subject timed out, the computer made a default decision allocating zero tokens to each battlefield. Across all sessions only 28 out of 6,660 allocation decisions resulted in a timeout.

<sup>13</sup> The APEX3 game is isomorphic to the Colonel Blotto game with additive objective studied in Hart (2008). For this game equilibrium marginal distributions are approximately uniform, with different weights placed on odd and even allocations (see Hart 2008).

in a session. Each session took approximately 1.5 hours and subjects earned on average £11.25 (about \$17 at the time of the experiment). Table 1 summarizes the experimental design.<sup>14</sup>

**Table 1** Experimental treatments

Treatment	Values of battlefields	Proportional value of large battlefield ( $v_l/\sum v_i$ )	Equilibrium share of expenditure on large battlefield	Number of pairs	Number of subjects
APEX3	$\mathbf{v} = (1, 1, 1)$	0.33	0.33	23	46
APEX4	$\mathbf{v} = (2, 1, 1, 1)$	0.40	0.50	26	52
APEX5	$\mathbf{v} = (3, 1, 1, 1, 1)$	0.43	0.58	25	50

Note that, in contrast to the theoretical analysis of a one-shot game discussed in the previous section, in our experiment subjects play a repeated game. This motivates several remarks. First, because the game is constant sum, any equilibrium of the repeated game involves playing an equilibrium of the stage game in each period. Second, even though subjects play repeatedly, since a subject either wins £0.50 or nothing in each round, equilibrium strategies are independent of risk preferences (Wooders and Shachat 2001). Third, use of a repeated play design requires a choice of how subjects will be matched across plays: most experiments use either a random matching protocol in which subjects are randomly re-matched from round to round or a fixed matching protocol where subjects are kept in the same pairs. An advantage of the fixed-pair protocol is that it gives subjects a greater incentive to be unpredictable (Chowdhury et al. 2013). Keeping subjects in the same pairs simplifies the structure of possible dependencies between decisions. In particular, this allows us to treat each pair as an independent observation.

## 5. RESULTS


We organize our results to examine the three predictions about the equilibria presented in section 3. We then turn to a further analysis of behavioral patterns in the data and try to explain possible reasons for deviations from these predictions.

<sup>14</sup> At an early stage of our research we also ran some sessions with a budget of 5 indivisible units. This permitted identification of equilibrium benchmarks without resorting to numerical methods. See Montero et al. (2013) for details.

## 5.1 Super-proportionality

Our first theoretical prediction was that players spend a disproportionate amount of their budgets on the large battlefield. If subjects' bids were proportional to values they would place 40% and 43% of their budgets on the large battlefield in APEX4 and APEX5 respectively, while in equilibrium expected bids are 50% and 58% of budgets. Figure 3 shows the share of expenditure on the large battlefield relative to its proportional share,  $(x_1 / 120) \div (v_1 / \sum v_i)$ , aggregated over all rounds and subjects. The theoretical prediction is that this measure is 1 for APEX3 (where the battlefields are symmetric) and increases with the number of battlefields. Figure 3 shows that the large battlefield receives more than a proportional share of the budget in all treatments. For APEX3 the battlefields are symmetric and so the small deviation from proportionality suggests a positional order effect, i.e. the effect of the ordering of the battlefields in the on-screen presentation. We also examined the data from the last 15 rounds separately and find that this deviation from proportionality disappears with experience. For APEX4 the allocation ratio is close to that observed in APEX3. Allocations are clearly super-proportional in APEX5, although even in this case they are well below the equilibrium prediction.

**Fig. 3** Share of budget allocated to large battlefield relative to proportional share



Formal statistical tests are presented in table 2, which shows the share of the budget allocated to the large battlefield in each treatment and p-values for tests against proportionality. We use two-sided sign-rank tests treating each pair as an independent observation.<sup>15</sup>

<sup>15</sup> This is equivalent to a test against equilibrium for APEX3. For tests against equilibrium in APEX4 and APEX5 all p-values are less than 0.0005.

**Table 2** Budget share allocated to the large battlefield

Treatment	Proportional Share (%)	Equilibrium Share (%)	Number of pairs	All rounds		Last 15 rounds	
				Share (%)	p-value*	Share (%)	p-value*
APEX3	33	33	23	35	0.0177	34	0.7380
APEX4	40	50	26	42	0.3037	42	0.3740
APEX5	43	58	25	49	0.0001	47	0.0094


\* p-values based on two-sided sign-rank test that mean allocation to large battlefield is proportional to value.

**Observation 1.** *We find significant evidence of super-proportionality in APEX5, but not in APEX4. In both games the average allocation levels to the large battlefield are below equilibrium predictions.*

## 5.2 Predicted and observed distributions

Next we look in more detail at subjects' strategies and how they relate to equilibrium predictions. Figure 4 displays the empirical distribution of allocations for the APEX3 game, aggregated over all subjects and rounds. Equilibrium strategies imply that the marginal distribution of tokens on each battlefield is approximately uniform on  $\{0, \dots, 80\}$ , represented by the horizontal probability density line in the figure. In contrast, there is a pronounced bimodality in subject choices, with subjects tending to place either very small amounts or about half their budget on a battlefield too often. This is similar to what is observed in previous experiments with Colonel Blotto or related games (see Avrahami and Kareev 2009; Chowdhury et al. 2013; Mago and Sheremeta 2012). Figure 4 also shows that the distributions of bids are similar across the battlefields, although subjects allocated slightly more to the first battlefield than the second, and slightly more to the second than the third.<sup>16</sup>


**Fig. 4** Observed relative frequencies of bids on each battlefield in APEX3 and the equilibrium prediction (horizontal line in each panel)



<sup>16</sup> Again, this echoes previous experimental findings. Chowdhury et al. (2013) also observe mild positional order effects (see their table 3 and figure 3).

Next we turn to the treatments with asymmetric battlefields. Figure 5 displays the equilibrium marginal distribution of bids for the large battlefield and the equilibrium marginal distribution of bids to any one of the small battlefields for the APEX4 treatment. The figure also shows the distributions actually observed in the experiment, aggregated over all subjects and rounds, where for the marginal distribution on a small battlefield we pool the data from all the small battlefields.

**Fig. 5** Equilibrium probabilities (left panels) and observed relative frequencies of bids on large and small battlefields in APEX4



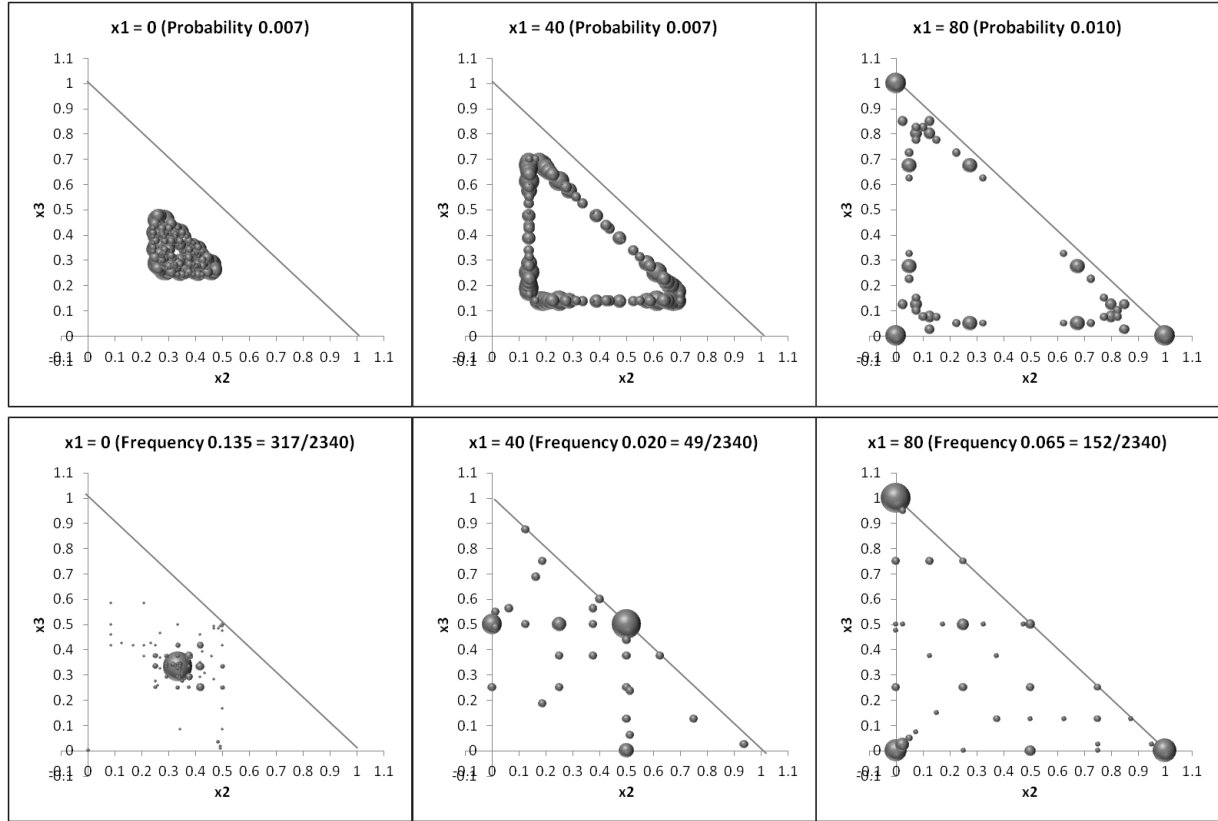
Again, the observed distributions are bi-modal, and are markedly different from the theoretical distributions. Relative to equilibrium, subjects too often place either very low amounts or around two-thirds of their budget on the large battlefield. Analogously, a small battlefield is often allocated an amount close to 0 or an amount close to one-third of the budget.

The univariate marginal distributions do not characterize the equilibrium in APEX4 and APEX5. In order to provide a better idea of what the multidimensional distributions look like in APEX4, figure 6 displays a panel of distributions over small battlefields for some commonly used values of the allocation to the large battlefield (the frequencies of these values are also reported in the figure). For comparison we also include the corresponding

panels for the equilibrium distribution in the top row of the figure. The triangles in figure 6 are to be interpreted in the same way as those in figure 2. The size of the dots is proportional to their probability mass (in the top row) or to their frequency (in the bottom row).

**Fig. 6** Conditional distributions on small battlefields in APEX4

Top row: predicted. Bottom row: observed




If no resources are allocated to the large battlefield, the equilibrium prediction is that the distribution of resources to the remaining three battlefields is not too unequal. The data reproduces this qualitative feature but in a more extreme manner: most of the mass is at the center of the triangle, corresponding to allocation (0,40,40,40). For  $x_1 = 40$ , all small battlefields receive a positive allocation of resources in equilibrium, but the distribution of resources between the battlefields is relatively unequal. The corresponding conditional distribution in the data is quite different, with notable concentrations of the data at (40,0,40,40), (40,40,0,40) and especially (40,40,40,0). Finally, for  $x_1 = 80$ , the prediction is that minimal winning strategies (and strategies that are close to minimal winning strategies) are frequent, and the data replicate this prediction in a more extreme manner: most of the

mass is at the corners, corresponding to the three minimal winning sets of the large battlefield and one small battlefield.

Figure 7 compares the equilibrium predictions with the marginal distributions actually observed in the APEX5 treatment. Similarly to APEX4, the observed distribution for the large battlefield is bi-modal. Again there is a concentration of negligible bids and a second concentration of higher bids. For this treatment the second mode is around three quarters of the budget. Analogously, the distribution on a small battlefield is bi-modal with one mode close to 0 and another around one-quarter of the budget. We omit depictions of conditional distributions in APEX5 due to the higher dimensionality, but the qualitative features of the distributions are the same, both for the theoretical and empirical distributions, as in APEX4.

**Fig. 7** Equilibrium probabilities (left panels) and observed relative frequencies of bids on large and small battlefields in APEX5



**Observation 2.** *For all treatments we observe bi-modal bidding patterns, in contrast with the shape in equilibrium prediction 2.*

### 5.3 Minimal winning and hedging strategies

Note that in order to win a round a subject only needs to win battlefields with a combined value exceeding that of the battlefields won by his opponent. Subjects may find it natural to

concentrate resources on a set of battlefields that is just large enough to win and give up on the rest; we refer to such strategies as *minimal winning*. Any strategy that allocates a positive value to a larger set of battlefields is referred to as a *hedging* strategy.

Hedging strategies can be optimal because of the uncertainty about the opponent's strategy. For example, consider APEX3 and suppose one player randomizes equally between bidding (60,60,0), (60,0,60) and (0,60,60). A best response to this is (61,58,1) which wins with probability  $2/3$ . In APEX3, minimal winning strategies have an equilibrium probability of around 7%<sup>17</sup>, but they are observed much more often than predicted (42% of submitted strategies are minimal winning).<sup>18</sup>

In APEX4 and APEX5 there are two types of minimal winning strategy. One targets all the small battlefields, giving up on the large battlefield (e.g. the strategy (0,40,40,40) in APEX4). The other type of minimal winning strategy targets the large and one small battlefield, giving up on the other small battlefields (e.g. the strategy (80,40,0,0) in APEX4). For both treatments the support of the equilibrium includes minimal winning strategies, but according to equilibrium they should be rarely played – less than 10% of the time. Instead, more than 90% of the time a player should use a hedging strategy, placing a positive amount on the large and at least two small battlefields. Again, we observe that minimal winning strategies are played much more often than is predicted by equilibrium. Minimal winning strategies that place zero on the large battlefield are predicted to be used less than 1% of the time in either treatment, but are used 13% of the time in APEX4 and 11% of the time in APEX5. Minimal winning strategies that target the large and one small battlefield are predicted to be used 5% of the time in APEX4 and 7% of the time in APEX5, but are actually used about 19% of the time in APEX4 and 20% of the time in APEX5.

Figure 8 presents the evolution of the proportion of hedging and minimal winning strategies over time. For APEX4 and APEX5 “MWL” denotes minimal winning strategies

---


<sup>17</sup> Hart (2008, p.454, Case 2.1) provides an equilibrium strategy for this game with the additive objective; due to the equivalence of the additive and majoritarian objectives in this case, the strategy is also an equilibrium in our setting. The strategy involves uniform marginals on even numbers only; in our case on  $\{0,2,4,\dots,80\}$ . The probability of a minimal winning strategy is then  $3/41 \approx 7\%$ , since each of the battlefields receives 0 with probability  $1/41$  and the strategy never puts 0 on more than one battlefield. Other approximately uniform marginal distributions would have a smaller probability of allocating 0 to a battlefield.

<sup>18</sup> A similar result is obtained by Mago and Sheremeta (2012) in a majoritarian contest with linear costs. In their experiment 35% of the time subjects bid only on two out of three objects, whereas in equilibrium they should make positive bids on all three.




that focus on the large and one small battlefield, while “MWS” denotes minimal winning strategies that focus on the small battlefields. In all three treatments the proportion of hedging strategies observed increases over time. Thus, there is some evidence that subjects learn from experience to switch from using minimal winning to using hedging strategies.

**Fig. 8** Proportions of minimal winning and hedging strategies over time



For APEX3 there is only one type of hedging strategy (namely a positive allocation to all battlefields), whereas for APEX4 and APEX5 there are strategies with different degrees of hedging. Recall the classification of strategies according to their target value  $k$  introduced in section 3. Strategies that allocate a positive value to all battlefields (i.e.  $k = 5$  in APEX4 and  $k = 7$  in APEX5) should be the most common in equilibrium, followed by strategies that allocate a positive value to the large and two small battlefields ( $k = 4$  in APEX4 and  $k = 5$  in APEX5), followed by minimal winning strategies ( $k = 3$  for APEX4 and  $k = 4$  for APEX5).

**Fig. 9** Proportions of strategies by target value over time



In both panels of figure 9, the solid line represents the probability of allocating a positive amount to all battlefields, while the dashed line shows the probability of targeting a minimal winning set of battlefields. Over the course of the experiment, the probability of targeting all battlefields increases and approaches its equilibrium value (76% for APEX4 and

64% for APEX5). Minimal winning strategies become less frequent over time but stabilize as the second most popular type. Hedging strategies that do not target all battlefields do not become more frequent; if anything their proportion diminishes.

**Observation 3.** *Strategies that target all battlefields become more frequent over time and their proportion reaches a value close to the equilibrium prediction in APEX4 and APEX5. Other types of hedging strategies are played less often than predicted in APEX4 and APEX5. In all treatments, minimal winning strategies remain more popular than equilibrium predicts.*

#### 5.4 Determinants of the allocation to the large battlefield

In this section we analyze individual allocations to the large battlefield using multivariate analysis. Following Chowdhury et al. (2013), we estimate a separate regression for each treatment taking the form:

$$\begin{aligned} AllocA_{it} = & \beta_0 + \beta_1 AllocA_{i(t-1)} + \beta_2 OppAllocA_{i(t-1)} + \beta_3 (AllocA_{i(t-1)} \times winA_{i(t-1)}) \\ & + \beta_4 (1/t) + \sum_s \beta_{5s} D_s + u_i + \varepsilon_{it}, \end{aligned}$$

where  $AllocA_{it}$  refers to the proportion of the budget allocated by subject  $i$  to the large battlefield in round  $t$ ,  $AllocA_{i(t-1)}$  is the same variable lagged,  $OppAllocA_{i(t-1)}$  is the corresponding lagged variable for the opponent of subject  $i$ , and  $winA_{i(t-1)}$  indicates whether  $i$  won the large battlefield in the previous round. The regressions also include session dummies, a reciprocal time trend and individual random effects. We exclude any observations in which a subject timed out either in the current round or in the previous one, and in which the subject's opponent timed out in the previous round (observations in which the opponent timed out in the current round are not affected). Table 3 reports the coefficients and robust standard errors from the random effects regressions. Session dummies are insignificant and are not reported.

Chowdhury et al. (2013) find that when subjects play against randomly changing opponents, their strategies are serially correlated. Specifically, they find that the lagged allocation variables and the interaction variable are useful predictors of a player's allocation to a battlefield (the significance of the latter variable they interpret as a "hot box effect"). They also find that the serial correlation is considerably reduced and the hot-box effect

disappears when subjects play repeatedly against the same opponent. Consistent with this latter finding, in our treatments (which all use fixed partner matching) we find very little evidence of serial correlation and no evidence of a hot-box effect. The opponent's lagged allocation is marginally significant only in one of the treatments. We do find a significant positive coefficient on the trend variable  $1/t$ , indicating decreasing allocations to the large battlefield across rounds. This significant coefficient is essentially picking up an effect that takes place over the early rounds. If we restrict the regression to the last 15 rounds, we find that the trend variable is not significant anymore (whereas the significance of other variables remains broadly unchanged).<sup>19</sup>

**Table 3** Determinants of allocation to large battlefield

Treatments	APEX3	APEX4	APEX5
Dependent variable: $AllocA_{it}$			
$AllocA_{i(t-1)}$	-0.037 (0.038)	-0.007 (0.048)	0.014 (0.044)
$OppAllocA_{i(t-1)}$	-0.045* (0.027)	0.034 (0.032)	-0.043 (0.035)
$AllocA_{i(t-1)} \times winA_{i(t-1)}$	-0.005 (0.030)	0.014 (0.041)	0.009 (0.035)
$1/t$	0.115** (0.050)	0.207*** (0.061)	0.232*** (0.083)
<i>Constant</i>	0.374*** (0.016)	0.377*** (0.031)	0.486*** (0.032)
# Observations	2,024	2,262	2,143
# Subjects	46	52	50

\*significant at 10%, \*\*significant at 5%, \*\*\*significant at 1%.

**Observation 4.** *There is no evidence of serial correlation in the subjects' strategies.*

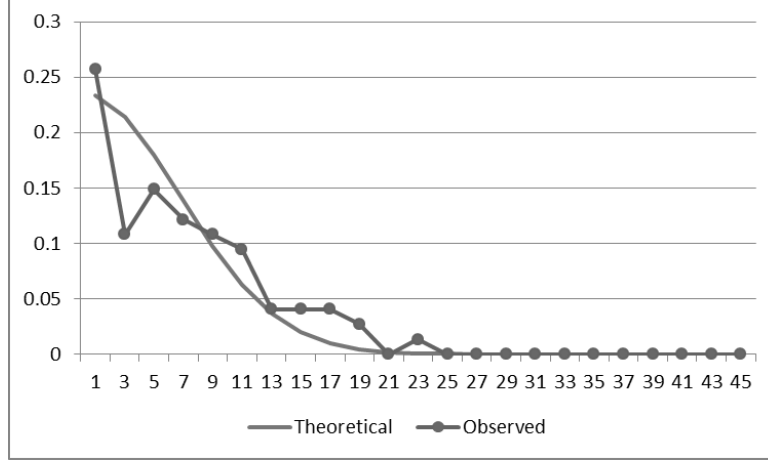
## 5.5 Determinants of success

For all treatments the game is a symmetric constant-sum game and so in equilibrium each player wins with probability 1/2 in any play. This means that a player expects to win 22.5 out of the 45 games. In fact, some do considerably better than this. For example, in one of the

<sup>19</sup> We also examined the effect of replacing the regressor  $1/t$  with  $t$ . The results are essentially the same: a regression using all the data indicates a significant decrease in the amount allocated to the large battlefield over time, but the trend variable is insignificant when we restrict the regression to the last 15 rounds.

pairs one subject won 11 rounds and the other won 34 rounds, so that the difference in wins was 23. Figure 10 shows the observed frequencies of each possible value of the difference in wins. For comparison the theoretical distribution is also shown.

**Fig. 10** Theoretical and observed distributions of differences in wins



The figure shows that fewer than expected pairs have a small difference in wins and more than expected have a large difference in wins. Theoretically, the expected difference in wins is 5.38, while in the data the average difference in wins is 6.84. This difference is significant (Chi-square test  $p = 0.028$ ). The obvious interpretation is that some subjects are better than others at playing the game.

To uncover the determinants of success in our experiment we run a probit regression of the following form:

$$\Pr\{Win_{it} = 1\} = \Phi(\beta_0 + \beta_1 AllocA_{it} + \beta_2 (AllocA_{it})^2 + \beta_3 MADTemp_{it} + \beta_4 (MADTemp_{it})^2 + \beta_5 MADAc_{it} + \beta_6 (MADAc_{it})^2 + \beta_7 MWL_{it} + \beta_8 MWS_{it} + u_i + \varepsilon_{it}).$$

Here  $\Phi$  is the c.d.f. of the standard normal distribution, and the dependent variable  $Win_{it}$  is a binary variable that takes value 1 if subject  $i$  won in round  $t$  and 0 otherwise.  $AllocA_{it}$  refers to the proportion of the budget allocated by subject  $i$  to the large battlefield in round  $t$ .

$MADTemp_{it}$  measures the variability of a subject's allocation across rounds and is defined as

$$MADTemp_{it} = \frac{1}{240} \sum_{j=1}^n |y_{ijt} - y_{ij(t-1)}|, \text{ where } y_{ijt} \text{ is the allocation of subject } i \text{ to battlefield } j \text{ at}$$

time  $t$  and  $n$  is the number of battlefields in the treatment. If a subject makes the same choice as in the previous round  $MADTemp = 0$ , while the maximum possible value is 1 if all tokens

are placed on a battlefield that received zero in the previous round.  $MADAc_{it}$  measures the variability of the allocation across battlefields in a given round and is defined as

$$MADAc_{it} = \frac{n}{240(n-1)} \sum_{j=1}^n |y_{ijt} - 120/n|. \text{ This variable also ranges from 0 to 1. It takes a}$$

minimum value of 0 when the budget is split equally across battlefields and the scaling term ensures that it takes a maximum value of 1 (when the budget is allocated to a single battlefield). To allow for non-monotonic effects of these variables we also include their squared terms in the regression. We also include indicator variables for using minimal winning strategies. In APEX3  $MW_{it} = 1$  if subject  $i$  places positive amounts on just two of the three battlefields in round  $t$  (and zero otherwise). In the other treatments we have an indicator variable for each type of minimal winning strategy.  $MWL_{it} = 1$  if positive amounts are placed on the large and just one other battlefield, and  $MWS_{it} = 1$  if nothing is placed on the large battlefield and positive amounts are placed on all small battlefields. The regression includes subject level random effects. Results are reported in table 4.

Consistent with Chowdhury et al. (2013) we find that  $MADTemp$  is highly significant, hence variability of the allocation over time is one of the main determinants of success. Note that in APEX5 the squared term is significantly negative indicating a concave relation between this measure of variability over time and the probability of winning. However, in this case the implied maximum is when  $MADTemp$  is around 0.9 and over 95% of the observations are below this, so the data are predominantly spread over the increasing range. Thus, in all treatments subjects with more variable allocations across rounds (according to this metric) win more often.

Other variables do not have a consistently significant effect on the probability of winning. Recall that, relative to equilibrium, subjects tend to play minimal winning strategies too often in all treatments. Thus it is perhaps unsurprising that, given the actual distribution of play in APEX3, subjects who used minimal winning strategies more often were less likely to win. However, in the other treatments there is no significant relationship between the use of minimal winning strategies and the chances of success. Similarly, relative to equilibrium, subjects place too little on the large battlefield in APEX4 and APEX5 and so one might expect that a subject placing more on the large battlefield would be more likely to win.

However, there is no significant effect in APEX5 and in APEX4 there is, if anything, a tendency for those allocating more to the large battlefield to lose.<sup>20</sup>

**Table 4** Determinants of winning

Treatments	APEX3		APEX4		APEX5	
Dependent variable: $Win_{it}$						
$AllocA_{it}$	-0.423 (0.503)	-0.236* (0.123)	0.434 (0.448)	-0.133 (0.110)	-0.296 (0.543)	-0.014 (0.107)
$(AllocA_{it})^2$	0.299 (0.806)		-1.155* (0.630)		0.427 (0.773)	
$MADTemp_{it}$	1.103** (0.477)	0.457*** (0.136)	1.182*** (0.424)	0.583*** (0.115)	1.945*** (0.456)	0.968*** (0.107)
$(MADTemp_{it})^2$	-1.108 (0.743)		-0.747 (0.478)		-1.083** (0.492)	
$MADAc_{it}$	1.543** (0.755)		0.281 (0.615)		-1.701** (0.809)	
$(MADAc_{it})^2$	-1.488 (0.957)		0.491 (0.743)		1.498 (0.986)	
$MW_{it}$	-0.130** (0.066)	-0.104* (0.060)				
$MWL_{it}$			0.056 (0.080)	0.140* (0.074)	0.089 (0.088)	0.086 (0.079)
$MWS_{it}$			-0.060 (0.106)	-0.132 (0.094)	-0.025 (0.120)	-0.037 (0.108)
<i>Constant</i>	-0.436** (0.201)	-0.039 (0.079)	-0.494*** (0.150)	-0.207*** (0.076)	-0.186 (0.165)	-0.425*** (0.084)
# Observations	2,024		2,254		2,124	
# Subjects	46		52		50	

\* significant at 10%, \*\* significant at 5%, \*\*\* significant at 1%.

Note that several of the right hand side variables are highly correlated. In particular the variables *AllocA* and *MADAc* are highly correlated in APEX4 ( $\rho = 0.79$ ) and APEX5 ( $\rho = 0.93$ ), as are the linear and squared terms in the quadratic specifications ( $\rho > 0.93$  in all cases). For this reason we also report results from running the regression but dropping

<sup>20</sup> The implied maximum of the concave relation between *AllocA* and the probability of winning is when 20% of the budget is allocated to the large battlefield, and two-thirds of the observations allocate more than this to the large battlefield.

*MADAc*r and the non-linear terms. The story is essentially the same. *MADTemp* is still highly significant in all treatments. Subjects who use minimal winning strategies less often are more likely to win in APEX3, but in the other treatments there is either no significant effect, or, in the case of *MWL* in APEX4 the effect goes in the opposite direction. In APEX3, consistent with subjects slightly overplaying the large battlefield in early rounds, there is a significant advantage to reducing the amount allocated to the large battlefield. However, in the other treatments where the amount placed on the large battlefield is substantially lower than the equilibrium prediction there is no significant relation between *AllocA* and the probability of success. These results for APEX4 and APEX5 may simply reflect that given the out-of-equilibrium behavior of subjects, strategies that are closer to equilibrium in terms of these specific metrics do not necessarily do better than strategies that are further away. In section 5.7 we look more closely at which strategies perform best against the empirical distribution.

**Observation 5.** *On average, subjects whose allocations to the battlefields exhibit high variability over time are more likely to win.*

## 5.6 Iterative reasoning: popular strategies and unit digits

Standard  $k$ -level reasoning (see Stahl and Wilson 1995) is impracticable in our setup, but as in Arad and Rubinstein (2012) it is possible to interpret some of the patterns observed as consistent with a form of iterative reasoning. Two caveats are in order: first, we are pooling data over 45 rounds, and iterative reasoning may be more natural for games that are played for the first time; second, the complicated structure of the game does not allow for a one-to-one identification between strategies and levels of reasoning: the same strategy may be a very naïve choice or a best response found after several levels of iteration.

We consider all permutations of a particular pure strategy as one strategy type. For example, 80-40-0-0 denotes the set of strategies  $\{(80,40,0,0), (80,0,40,0), (80,0,0,40)\}$ . This is similar to the concept of object-symmetric strategies, except that there is no presumption that different permutations have the same frequency. The following tables show the most popular strategy types in each treatment, pooling across all subjects and rounds.

In APEX3, the simple strategy 40-40-40 can be seen as a natural starting point, analogous to level 0. It can be beaten by many others, the most salient being strategy type 60-60-0, which was also the most popular strategy type in round 1, played 15 out of 46 times.

Strategy type 61-58-1 appears to be specifically designed to be played against strategy type 60-60-0, and wins with probability  $2/3$  against it.

**Table 5** Popular strategy types in APEX3

Rank	Strategy type	Frequency (n = 2070)	Average response time (sec.)
1	60-60-0	0.2599	8.59
2	70-50-0	0.0729	9.77
3	61-58-1	0.0652	12.72
4	70-40-10	0.0478	13.09
5	80-40-0	0.0425	10.78
(8)	40-40-40	0.0227	7.66

Notice that 40-40-40, 60-60-0 and 70-50-0 have shorter response times than other strategies such as 61-58-1 and 70-40-10. This is consistent with these strategies being associated with lower levels of reasoning, and in particular with 40-40-40 being the instinctive choice.

Interpretation of the popular choices as evidencing iterated reasoning is more difficult in APEX4. The equal allocation 30-30-30-30 is a less obvious starting point here, but it has the shortest response time in table 6.

**Table 6** Popular strategy types in APEX4

Rank	Strategy type	Frequency (n = 2340)	Average response time (sec.)
1	0-40-40-40	0.0645	10.90
2	80-40-0-0	0.0363	15.94
3	70-50-0-0	0.0325	13.51
4	60-60-0-0	0.0235	15.22
5	90-30-0-0	0.0192	13.96
(6)	30-30-30-30	0.0154	9.83

The next step would be to concentrate forces on a minimal winning set and beat 30-30-30-30. Strategy 0-40-40-40 can be seen as one of the simplest best responses to this strategy, and it is the most popular strategy overall and also in round 1 (6 out of 52 times). The next three most popular strategy types are also minimal winning strategies that best respond to 30-30-30-30.



A possible further step is that, if the other player is likely to focus on a minimal winning set of battlefields, a player may be better-off by putting some tokens on all battlefields. For example, strategy 117-1-1-1 (the 11<sup>th</sup> most popular strategy type in our data) is a best response against some popular strategy types like 90-30-0-0, 80-40-0-0 and 60-60-0-0 (though it of course loses against seemingly more naïve strategies like 0-40-40-40 and 30-30-30-30). Hedging strategies that do not distribute tokens equally can be interpreted as the result of at least two steps of reasoning, since they appear to be designed to be played against minimal winning strategies.


**Table 7** Popular strategy types in APEX5

Rank	Strategy type	Frequency (n = 2250)	Average response time (sec.)
1	0-30-30-30-30	0.0484	17.26
2	80-40-0-0-0	0.0316	17.55
3	90-30-0-0-0	0.0293	22.56
4	1-30-30-30-29	0.0236	18.98
5	85-35-0-0-0	0.0182	16.39
(6)	100-5-5-5-5	0.0178	10.05
(16)	24-24-24-24-24	0.0071	13.06

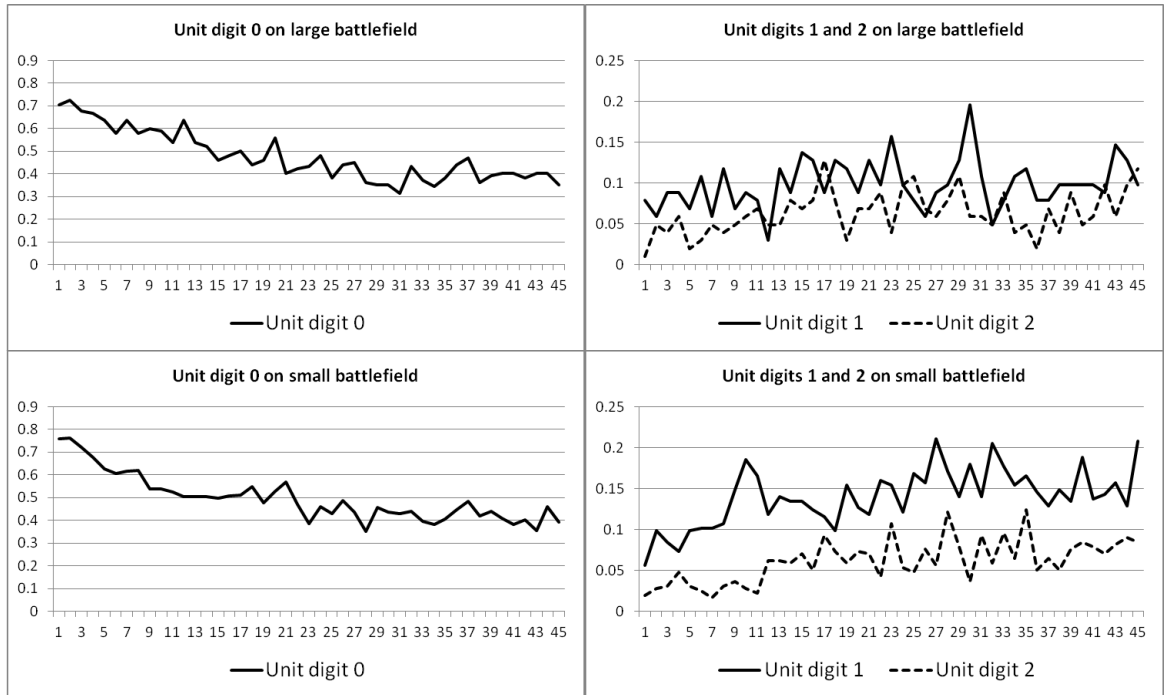
In APEX5, analogously to the previous case, 24-24-24-24-24 has one of the shortest response times and can be taken as the starting point. Other popular strategy types like 0-30-30-30-30 and 80-40-0-0-0 are best replies to this strategy that divide the budget across a minimal winning set of battlefields. A further possible step of reasoning is to spread the budget over the entire set of battlefields. Strategy type 100-5-5-5-5 is a best reply against some popular minimal winning strategies in table 7 (all except 0-30-30-30-30).

Another indication of iterative reasoning is the evolution of unit digits over time. Many of the observed strategies used round numbers (unit digit 0), but the frequency of unit digit 0 decreases over time. The next figures show the proportion of unit digits 0, 1 and 2 over time. Figure 11 refers to treatment APEX3 pooling all battlefields together. Figure 12 pools treatments APEX4 and APEX5 together (the figures for each of the treatments are very similar) but separates the large and small battlefields. The proportion of unit digit 0 decreases over time in all treatments, but this is more pronounced in APEX4 and APEX5.

**Fig. 11** Proportion of unit digits over time in APEX3



**Fig. 12** Proportion of unit digits over time across APEX4 and APEX5



**Observation 6.** Popular strategy types are consistent with one or two steps of iterative reasoning. The evolution of the frequency of unit digits is consistent with an increasing level of sophistication over time.

### 5.7 A measure of deviation from equilibrium: Exploitability

To measure how far observed play is from equilibrium we take advantage of the fact that equilibrium mixed strategies restrict an opponent's probability of winning to one-half. Any mixed strategy that can be beaten with probability exceeding one-half cannot be an equilibrium strategy. Intuitively, the greater the expected payoff one can obtain against a

mixed strategy, the further that strategy is from equilibrium play. Thus, to measure the extent to which a strategy deviates from equilibrium we take the expected payoff from the best response to this strategy.

Table 8 displays this measure of exploitability for some selected strategies. We consider two specific benchmarks for comparison, denoted “uniform” and “minimal winning”. The particular uniform strategies we consider are mixed strategies that induce a marginal distribution of tokens on each battlefield that approximates a uniform distribution with the expected allocation to a battlefield being proportional to its valuation. In other words, they approximate strategies that are equilibria in the game with an additive objective and a continuously-divisible budget. For APEX3, we use the strategy provided by Hart (2008, p. 454, Case 2.1), discussed in footnote 17; the strategy is also an equilibrium strategy in our case and thus its exploitability is exactly 0.5. For APEX4, we use the object-symmetric version of  $(96-4\varepsilon, 2\varepsilon, \varepsilon, 24+\varepsilon)$ , where  $\varepsilon$  is uniformly distributed over the integer values in  $[0,24]$ . For APEX5, we use the object-symmetric version of  $(103-6\varepsilon, 2\varepsilon, 2\varepsilon, \varepsilon, 17+\varepsilon)$ , where  $\varepsilon$  is uniformly distributed over the integer values in  $[0,17]$ . These strategies are parametrized by a single variable  $\varepsilon$ , which makes them relatively simple. Their exploitability is similar to what we know to be a lower bound on the exploitability of a uniform strategy in the continuously-divisible budget case.<sup>21</sup> The particular minimal winning strategy we consider models naïve behavior: it allocates the budget to a minimal winning set of battlefields and, within this set, each battlefield receives an allocation proportional to its value. For APEX3 the minimal winning strategy randomizes equally between  $(60,60,0)$ ,  $(60,0,60)$  and  $(0,60,60)$ ; for APEX4 it randomizes equally between  $(80,40,0,0)$ ,  $(80,0,40,0)$ ,  $(80,0,0,40)$ , and  $(0,40,40,40)$ ; for APEX5 it randomizes equally between  $(90,30,0,0,0)$ ,  $(90,0,30,0,0)$ ,  $(90,0,0,30,0)$ ,  $(90,0,0,0,30)$ , and  $(0,30,30,30,30)$ . In the remaining columns we report the exploitability of (and below it the best response to) a mixed strategy defined over allocations used in the experiment, where the probability of each allocation is equal to its empirical relative frequency.

---

<sup>21</sup> Footnote 7 provides, for the continuous case, strategies that achieve winning probability  $3/4$  for APEX4 and  $5/6$  for APEX5 against *any* strategy with uniform marginals on the appropriate intervals. For each strategy with uniform marginals there may exist strategies that achieve even higher winning probability;  $3/4$  for APEX4 and  $5/6$  for APEX5 are thus lower bounds on the exploitability of strategies with uniform marginals in the continuous case. Due to the discreteness of the budget, exploitability values for our games can be below these bounds.

**Table 8** Exploitability of strategies

	Uniform	Minimal Winning	Rounds 1-15	Rounds 16-30	Rounds 31-45	All Rounds
APEX3	0.500	0.667	0.681 2-56-62	0.670 71-6-43	0.664 7-71-42	0.664 8-71-41
APEX4	0.733	0.750	0.707 6-42-41-31	0.696 76-1-1-42	0.673 71-2-2-45	0.674 71-2-2-45
APEX5	0.814	0.800	0.682 5-31-31-32-21	0.648 21-41-11-11-36	0.677 25-37-11-36-11	0.655 11-21-21-35-32

Note that as the number of battlefields increases the game gets more complex. With more battlefields not only are there more strategies, but sub-optimal strategies are more exploitable. While the minimal winning strategy can be beaten two-thirds of the time in APEX3, it can be beaten 75% of the time in APEX4 and 80% of the time in APEX5. The uniform strategy, which is an equilibrium strategy in APEX3, can be beaten 73% of the time in APEX4, and does even worse than the minimal winning strategy in APEX5, where it can be beaten 81% of the time.

Turning to the data from the experiment, for any treatment we can find strategies that beat the empirical distribution more than 60% of the time. In all cases the best responses to the empirical distribution hedge by placing at least a small amount on all battlefields, beating the modes in the data at zero. Indeed, more generally the best responses score highly by beating modes in the data. Note also that although subjects in APEX4 and APEX5 place too little on the large battlefield, relative to equilibrium, the best response to the empirical distribution sometimes involves placing a small amount on the large battlefield. This is consistent with our previous finding that subjects spending less on the large battlefield sometimes win more often. This also underscores how challenging the super-proportionality hypothesis is in these games. Not only is it difficult to identify equilibrium strategies, but if the data is out-of-equilibrium sophisticated subjects might exploit this by choosing allocations that are even further from equilibrium.

Notably, we find no clear pattern in the degree of exploitability of our subjects across treatments. Subjects are less exploitable with experience in APEX3 and APEX4, but not in APEX5. Overall, we find that for all treatments the empirical distribution can be beaten around two-thirds of the time, with very little variation across treatments. Thus, although simple strategies are more exploitable in more complex games, subject behavior is no more

exploitable in APEX5 than in APEX3.

**Observation 7.** *Naïve strategies are more exploitable in APEX5 than in APEX4 than in APEX3, but the exploitability of actual play does not change much across treatments.*

## 6. CONCLUDING REMARKS

Our extension of classic Colonel Blotto games to allow for asymmetries in the values of battlefields captures an important feature of real-world conflicts, where not all battles carry equal value. Even restricting attention to the simplest asymmetries, using apex games, analytic theoretical solutions are unavailable. By studying a game that restricts strategies to be object-symmetric and using numerical methods, we are able to identify equilibrium predictions for these games and derive some interesting qualitative properties. Moreover, we show that these properties are shared by any equilibrium in the original, unrestricted, game.

Nevertheless, whether these theoretical features are evident in the play of human subjects remains an important question. As Walker and Wooders (2001) have remarked in the context of other games with mixed strategy equilibria, games requiring unpredictable play are often easy to play, but difficult to play well. The Colonel Blotto game provides a good example. Although it is easy to describe to subjects, and subjects have no trouble understanding the rules, sophisticated play is very demanding. For our simplest treatment (APEX3), corresponding to the classic Colonel Blotto game, it is perhaps obvious that a sophisticated player should not favor one battlefield, as such favoritism can be exploited by an opponent. It is perhaps equally obvious that in the more complex version with asymmetric battlefields a player should favor the more important battlefield. But it is not clear by how much it should be favored. The answer crucially depends on how battlefield values contribute to final victory or defeat. For the case we consider, in which the player who captures battlefields with the greatest total value wins, equilibrium requires players to allocate super-proportional amounts to the large battlefield. We would argue that the precise amount to be placed (in expectation) on the large battlefield, and the precise equilibrium strategies, are far from obvious. Indeed, we obtain equilibrium predictions using numerical methods that are beyond the cognitive capabilities of our subjects.

Perhaps unsurprisingly, behavior in our experiment deviates from these equilibrium predictions. We observe bi-modal distributions, where subjects choose either to spend very

little or a substantial proportion of their budget on a battlefield, with too little weight placed on intermediate allocations. This is also observed in related Colonel Blotto experiments. We also observe that subjects too often submit allocations that concentrate their resources on a minimal winning set of battlefields, placing zero on other battlefields, whereas equilibrium behavior requires hedging more than 90% of the time. Finally, in games with asymmetric battlefields subjects spend too little on the large battlefield. An implication of these deviations is that it is possible to find strategies that beat the empirical distribution of allocations more than 60% of the time.

On the other hand, in spite of these deviations from equilibrium, we find evidence of strategic sophistication. Analysis of popular strategy types reveals patterns consistent with several steps of iterated reasoning and the depth of this reasoning appears to increase over time. Furthermore, while simple strategies are more exploitable in games with more battlefields, we find that the exploitability of our subjects is quite similar across treatments. Our experiment also finds support for some of the qualitative features of equilibrium predictions. Although our subjects play hedging strategies only around 60% of the time, this proportion increases over the course of the session. Similarly, in games with asymmetric battlefields, subjects place less on the large battlefield than predicted, but in APEX5, where the asymmetry between battlefields is more pronounced, they place a more than proportional amount on the large battlefield.

## APPENDIX A: COMPUTATIONAL TECHNIQUES

The games APEX4 and APEX5, and the corresponding auxiliary games restricting to object-symmetric strategies, are two-player, constant-sum, symmetric games. As two-player, constant-sum games, a minimax strategy, and therefore a Nash equilibrium, can be written as the solution to a linear program (Dantzig 1951). In practice, it is possible to solve linear programs of substantial size. Moreover, because the set of mixed strategy Nash equilibria for two-player constant-sum games is convex, once an equilibrium has been found it is straightforward to explore the whole set of equilibria, and to place tight bounds on the size of this set. In particular, it is in principle possible to verify uniqueness, if one uses exact rational arithmetic. However, doing so is infeasible in our games with large endowments. We instead use floating-point arithmetic for our calculations. For many values of the endowment, we bound all strategy probabilities to within  $10^{-6}$ , which suggests the object-symmetric equilibrium is in fact unique.

### *Improving computational efficiency*

Turning specifically to the games studied in this paper, with a budget of  $E = 120$  tokens, the strategy spaces of these games are quite large. Even restricting attention to object-symmetric strategies, APEX4 has 52,311, and APEX5 has 430,256 strategies. However, preliminary explorations with smaller budgets led us to conjecture that equilibria in these games would have small supports. We therefore used an iterative method to identify the set of equilibrium strategies.

Consider the game with a budget of  $E$  tokens. Let  $X$  be the set of pure strategy token allocations. We construct an increasing sequence of supports,  $X_0 \subseteq X_1 \subseteq \dots \subseteq X_T$ , such that  $X_T$  is the support of an equilibrium of the game. Pick some initial guess at the support of the equilibrium, and call it  $X_0 \subseteq X$ . (The correctness of the construction does not depend on the value of the initial guess  $X_0$ ; for this approach to work efficiently, it should be small in size.)

At each step  $\tau$  of the algorithm, we consider the restriction of the game in which players can choose only strategies in  $X_\tau$ . This induces a well-defined two-player constant-sum game, which can be solved for some equilibrium  $\pi_\tau$ ; insofar as  $|X_\tau| \ll |X|$ , solving the

restricted game should be much faster than solving the full game. Then, given  $\pi_\tau$ , we consider all the strategies which were deleted from the restricted game,  $X \setminus X_\tau$ , and order them in decreasing order by their payoff relative to the candidate equilibrium  $\pi_\tau$ . If there are no strategies which attain a payoff greater than the equilibrium payoff of one-half, then the algorithm terminates, and  $\pi_\tau$  is an equilibrium of the game with strategy set  $X$ . If not, then we construct  $X_{\tau+1}$  by adding the top  $h$  strategies from  $X \setminus X_\tau$  to those in  $X_\tau$ , and iterating.

The number of strategies  $h$  added at every step is arbitrary; we obtained sufficiently good performance with  $h = 25$ . The tradeoff is that if  $h$  is too small, then the algorithm will require many iterations, and therefore many calls to the linear program solver, to find the equilibrium, while if  $h$  is too large the algorithm will consider many strategies which are not in the equilibrium support, slowing down individual runs of the linear programming solver.

In any event, the correctness of the approach does not depend on the scheme used for adding strategies. Because  $X_{\tau+1}$  is always a strict superset of  $X_\tau$ , and because all the supports are bounded above (in the sense of set inclusion) by the strategy set  $X$ , this iterative process is guaranteed to terminate in a finite number of steps.



## APPENDIX B: INSTRUCTIONS

### General rules

Welcome! This session is part of an experiment in the economics of decision making. If you follow the instructions carefully and make good decisions, you can earn a considerable amount of money.

In this session you will be competing with one other person, randomly selected from the people in this room, over the course of forty-five rounds. Throughout the session your competitor will be the same but you will not learn whom of the people in this room you are competing with. The amount of money you earn will depend on your decisions and your competitor's decisions.

It is important that you do not talk to any of the other people in the room until the session is over. If you have any questions raise your hand and a monitor will come to your desk to answer it.

### Description of a round

Each of the forty-five rounds is identical. At the beginning of each round your computer screen will look like the one below.

Round

1 out 45

Remaining time [sec]: 79

You have a budget of 120 tokens and you need to distribute it across 4 objects in whole amounts.

Object	Points	Your Bid
A	2	<input type="text"/>
B	1	<input type="text"/>
C	1	<input type="text"/>
D	1	<input type="text"/>

Submit

Help

The table gives the labels of the objects and how many points each is worth. In the rightmost column, enter 4 whole numbers (possibly including zeros) that add up to 120 and press the "Submit" button. If the bids you submit do not add up to 120 the computer will indicate by how many tokens the bid needs to be corrected.

You have 120 tokens. You must use these to bid on 4 objects labelled A, B, C and D. You get points for winning objects – object A is worth 2 points and the other objects are worth 1 point each. For each object you can bid any whole number of tokens (including zero), but the total bid for all objects must add up to 120 tokens. You bid by entering numbers in the boxes, and then clicking on the “Submit” button. If the bids you submit do not add up to 120 the computer will indicate by how many tokens the bid needs to be corrected. If you do not

submit a valid bid within 90 seconds the computer will bid for you and will place zero tokens on each object.

When everyone in the room has submitted their bids, the computer will compare your bids with those of your opponent. Your computer screen will look like the one below (the bids in the figure have been chosen for illustrative purposes only):

Round

1 out 45

Remaining time [sec]: 24

You and your opponent submitted the following bids this round:

Object	Points	Your Bid	Opponent's Bid	Points won by you
A	2	60	30	2
B	1	30	30	0*
C	1	0	0	1*
D	1	30	60	0

\* - random decision

You won 3 points of the available 5.  
You won more points than the opponent.  
Your earnings this round: 50 pence.

Continue

Help  
Press the "Continue" button when you are ready to continue the session.

You win an object if you bid more for it than your opponent. (If you and your opponent bid the same amount the computer will randomly decide whether you or your opponent wins the object, with you and your opponent having an equal chance of winning the object. In this case the computer screen will indicate with an asterisk that the object was awarded randomly).

**The winner of the round is the person who gets the most points.**

The winner of the round earns 50 pence, the other person earns zero.

### Ending the Session

At the end of the session you will be paid the amount you have earned from all forty-five rounds. You will be paid in private and in cash.

Now, please complete the quiz. If you have any questions, please raise your hand. The session will continue when everybody in the room has completed the quiz correctly.

## Quiz

1. Suppose your bids and your competitor's bids were as follows:

Object	Points	Your Bid	Opponent's Bid
A	2	30	60
B	1	30	60
C	1	30	0
D	1	30	0

How many points would you receive? \_\_\_\_\_.

How many points would your opponent receive? \_\_\_\_\_.

What would your earnings from this round be (in pence)? \_\_\_\_\_.

What would your opponent's earnings from this round be (in pence)? \_\_\_\_\_.

2. Suppose your bids and your competitor's bids were as follows:

Object	Points	Your Bid	Opponent's Bid
A	2	30	10
B	1	30	30
C	1	30	40
D	1	30	40

Who wins object A? Me / My Opponent / Randomly Determined (Circle One)

Who wins object B? Me / My Opponent / Randomly Determined (Circle One)

For the remaining questions suppose the computer awards object B to your opponent:

How many points would you receive? \_\_\_\_\_.

How many points would your opponent receive? \_\_\_\_\_.

What would your earnings from this round be (in pence)? \_\_\_\_\_.

What would your opponent's earnings from this round be (in pence)? \_\_\_\_\_.

## REFERENCES

- Arad, A., Rubinstein, A.: Multi-dimensional iterative reasoning in action: the case of the Colonel Blotto game. *J. Econ. Behav. Organ.* 84, 475-712 (2012)
- Avrahami, J., Kareev, Y.: Do the weak stand a chance? Distribution of resources in a competitive environment. *Cognitive Sci.* 33, 940-950 (2009)
- Avrahami, J., Kareev, Y., Todd, P., Silverman, B.: Allocation of resources in asymmetric competitions: how do the weak maintain a chance of winning? *J. Econ. Psychol.* 42, 161-174 (2014)
- Borel, E.: La théorie du jeu et les équations intégrales à noyau symétrique. *Comptes Rendus de l'Académie des Sciences* 173, 1304-1308 (1921); English translation: Savage, L.: The theory of play and integral equations with skew symmetric kernels. *Econometrica* 21, 97-100 (1953)
- Borel, E., Ville, J.: Application de la théorie des probabilités aux jeux de hasard. Gauthier-Villars, Paris (1938); reprinted in Borel E., Cheron, A.: *Théorie mathématique du bridge à la portée de tous*. Editions Jacques Gabay, Paris (1991)
- Chowdhury, S.M., Kovenock, D., Sheremeta, R.M.: An experimental investigation of Colonel Blotto games. *Econ. Theory* 52, 833-861 (2013)
- Dantzig, G.: A proof of the equivalence of the programming problem and the game problem. In: Koopmans, T.C. (ed.) *Activity Analysis of Production and Allocation*, pp. 330-335. Wiley, New York (1951)
- Dechenaux, E., Kovenock, D., Sheremeta, R.M.: A survey of experimental research on contests, all-pay auctions and tournaments. *Forthcoming in Exp. Econ.* (2014). doi: 10.1007/s10683-014-9421-0
- Duffy, J., Matros, A.: Stochastic asymmetric Blotto games: theory and evidence. Typescript (2015)
- Fischbacher, U.: z-Tree: Zurich toolbox for ready-made economic experiments. *Exp. Econ.* 10, 171-178 (2007)
- Greiner, B.: An online recruitment system for economic experiments. In: Kremer, K., Macho, V. (eds.) *Forschung und Wissenschaftliches Rechnen 2003*, pp. 79-93. GWDG Bericht 63, Göttingen: Ges. für Wiss. Datenverarbeitung (2004)

- Gross, O., Wagner, R.: A continuous Colonel Blotto game. RAND Corporation RM-408 (1950)
- Hart, S.: Discrete Colonel Blotto and General Lotto games. *Int. J. Game Theory* 36, 441-460 (2008)
- Hortala-Vallve, R., Llorente-Saguer, A.: A simple mechanism for resolving conflict. *Game. Econ. Behav.* 70, 375-391 (2010)
- Isbell, J.R.: On the enumeration of majority games. *Mathematical Tables and Other Aids to Computation* 13, 21-28 (1959)
- Kovenock, D., Roberson, B.: Conflicts with multiple battlefields. In Garfinkel, M.R., Skaperdas, S. (eds.) *Oxford Handbook of the Economics of Peace and Conflict*, pp. 503-531. Oxford University Press, Oxford (2012)
- McKelvey, R., McLennan, A., Turocy, T.: Gambit: software tools for game theory, Version 13.1.0. <http://www.gambit-project.org> (2013)
- Mago, S.D., Sheremeta, R.M.: Multi-battle contests: an experimental study. Available at <http://ssrn.com/abstract=2027172> (2012)
- Montero, M., Possajennikov, A., Sefton, M., Turocy, T.: Majoritarian contests with asymmetric battlefields: an experiment. Available at <http://dx.doi.org/10.2139/ssrn.2369056> (2013)
- Roberson, B.: The Colonel Blotto game. *Econ. Theory* 29, 1-24 (2006)
- Stahl, D.O., Wilson, P.W.: On players' models of other players: theory and experimental evidence. *Game. Econ. Behav.* 10, 218-254 (1996)
- Thomas, C.: N-dimensional Blotto Game with asymmetric battlefield values. Typescript (2013)
- Walker, M., Wooders, J.: Minimax play at Wimbledon. *Am. Econ. Rev.* 91, 1521-1538 (2001)
- Wooders, J., Shachat, J.: On the irrelevance of risk attitudes in repeated two-outcome games. *Game. Econ. Behav.* 34, 342-363 (2001)
- Young, H. P.: A tactical lobbying game. RR-78-10, International Institute for Applied Systems Analysis, Laxenburg, Austria (1978)