CEDEX

CENTRE FOR DECISION RESEARCH & EXPERIMENTAL ECONOMICS

The University of Nottingham

UNITED KINGDOM · CHINA · MALAYSIA

Maria Montero, Alex Possajennikov, Martin Sefton and Theodore Turocy

December 2013

Majoritarian Contests with Asymmetric Battlefields: An Experiment

CEDEX

CENTRE FOR DECISION RESEARCH & EXPERIMENTAL ECONOMICS

The Centre for Decision Research and Experimental Economics was founded in 2000, and is based in the School of Economics at the University of Nottingham.

The focus for the Centre is research into individual and strategic decision-making using a combination of theoretical and experimental methods. On the theory side, members of the Centre investigate individual choice under uncertainty, cooperative and non-cooperative game theory, as well as theories of psychology, bounded rationality and evolutionary game theory. Members of the Centre have applied experimental methods in the fields of public economics, individual choice under risk and uncertainty, strategic interaction, and the performance of auctions, markets and other economic institutions. Much of the Centre's research involves collaborative projects with researchers from other departments in the UK and overseas.

Please visit http://www.nottingham.ac.uk/cedex for more information about the Centre or contact

Suzanne Robey
Centre for Decision Research and Experimental Economics
School of Economics
University of Nottingham
University Park
Nottingham
NG7 2RD
Tel: +44 (0)115 95 14763
Fax: +44 (0) 115 95 14159
suzanne.robey@nottingham.ac.uk

The full list of CeDEx Discussion Papers is available at

http://www.nottingham.ac.uk/cedex/publications/discussion-papers/index.aspx

Majoritarian Contests with Asymmetric Battlefields: An Experiment[1]

Maria Montero[2]
University of Nottingham
University of the Basque Country (FAE I)
IKERBASQUE, Basque Foundation of Science

Alex Possajennikov, Martin Sefton[3]
University of Nottingham

Theodore L. Turocy[4]
University of East Anglia

December 2013

Abstract

We investigate a version of the classic Colonel Blotto game in which individual battles may have different values. Two players allocate a fixed budget across battlefields and each battlefield is won by the player who allocates the most to that battlefield. The winner of the game is the player who wins the battlefields with highest total value. We focus on the case where there is one large and several small battlefields, such that a player wins if he wins the large and any one small battlefield, or all the small battlefields. We compute the mixed strategy equilibrium for these games and compare this with choices from a laboratory experiment. The equilibrium predicts that the large battlefield receives more than a proportional share of the resources of the players, and that most of the time resources should be spread over more battlefields than are needed to win the game. We find support for the main qualitative features of the equilibrium. In particular, strategies that spread resources widely are played frequently, and the large battlefield receives more than a proportional share in the treatment where the asymmetry between battlefields is stronger.

Keywords: Colonel Blotto, majoritarian contests, experiment

JEL Codes: C72, C91

1

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?

Little is known about this type of 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 to win a majority of votes. Young considers the case where there is 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. Young discusses two games, differing in the number of voters, and reports that equilibrium expenditure is disproportionately skewed towards the large voter.

Our experimental treatments are based on the games discussed by Young. In the experiment, two subjects compete for objects and receive points for objects won. An object is won by the player spending most on it, and some objects are worth more points than others.

The winner of the game is the subject obtaining 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 precisely the equilibrium: identifying the equilibrium is computationally challenging and presumably beyond the ability of experimental subjects (indeed, we use numerical methods to pin down the 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.

We use numerical methods to completely describe equilibrium strategies for our experimental setting. We find that there is a unique mixed strategy 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, the equilibrium also predicts that players almost always spread their budget over more objects than are needed to win the contest.

The equilibrium mixed strategy is complicated 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, firstly because there are more possible strategies, and secondly because some simple strategies are more exploitable. Nevertheless, the degree of exploitability of actual play is 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 in detail, we find evidence for some of the 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. 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) extended this analysis to the case of a discrete budget. Thomas (2013) extends the analysis to the case of asymmetric battlefields and shows that uniform marginals, where the mean expenditure on a battlefield is proportional to its value, are a sufficient condition for equilibrium.

Young (1978) extended the original majoritarian objective model of Borel (1921) to the case of asymmetric battlefields. He studies vote-buying games in which players with opposite interests allocate their budgets across voters. He considers two games, both involving one large voter and several identical small voters, where a player needs to secure the votes of the large voter and one small voter, or the votes of all small voters, to win. In both games the large voter is predicted to receive a share of the budget above its proportion of the votes. This vote-buying game is equivalent to a Colonel Blotto game with asymmetric battlefields.

A small number of studies have recently examined variants of the Colonel Blotto game experimentally.[5] 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 naive 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 the 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

---

[5] Dechenaux et al. (2012) survey the experimental literature on contests more generally.

subjects who were unfamiliar with this game to exhibit behavior consistent with equilibrium".[6] 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 have to 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, which is not possible 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. (2013) study a multiple-battlefield contest with asymmetric battlefields and an auction CSF, but an additive objective. As shown by Thomas (2013), in equilibrium each object receives a share of expenditure proportional to its value. Duffy and Matros (2013) study a multiple-battlefield contest setting with asymmetric battlefields and a majoritarian objective, but where the outcome of 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. Both Avrahami et al. (2013) and Duffy and Matros (2013) find that treatment differences conform to equilibrium predictions.

3. AN ASYMMETRIC BLOTTO GAME

The game we study is an asymmetric version of the classic Colonel Blotto game introduced by Borel (1921). Two players, A and B, simultaneously allocate identical endowments E across n battlefields. Let $N = \{1, \ldots, n\}$ denote the set of battlefields. Each battlefield has a

---

[6] 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 the observed choices as reflecting iterated reasoning in several dimensions. Note however, that their subjects play a one-shot game, with no opportunity for learning.

value which is the same for both players, but some battlefields may be worth more than others. We denote the value of battlefield i by v_i. This asymmetry across battlefields is the fundamental difference from the classical Colonel Blotto game.

Each battlefield is won by one of the players according to an auction contest success function. Let $x_i^j \geq 0$ be the amount allocated to battlefield i by player j. 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 $x_i^A = x_i^B$. We will use N_A to denote the set of battlefields won by A and N_B to denote the set won by B.

The winner of the game is the player who wins battlefields with the greatest total value. That is, we consider a majoritarian objective, where A wins if and only if $\sum_{i \in N_A} v_i > \sum_{i \in N_B} v_i$. We assume throughout that $\sum_{i \in N_A} v_i \neq \sum_{i \in N_B} v_i$ for any partition of battlefields, so that there is always a winner and a loser.

As in the classical Colonel Blotto game, we assume that players' entire endowments must be spent on the battlefields, i.e., $\sum_{i \in N} x_i^j = E$, $j = A, B$. In the terminology of Kovenock and Roberson (2012) the contest cost function exhibits the use-it-or-lose-it technology. From the point of view of the players, there are only two possible outcomes of the game: either win or lose. Thus, assuming that the utility of a win exceeds that of a loss, each player maximizes her expected utility by maximizing her probability of winning. This has the implication that equilibria are independent of risk attitudes.

Since players are symmetric, each player wins with probability one-half in equilibrium. Moreover, except in trivial cases, the only equilibria of the game are in mixed strategies. However, beyond this little is known about equilibrium.[7] To address this question

---

[7] One thing we know is that results for the additive objective game do not carry over to majoritarian objective games. To see this suppose $E = 5$ and $v = (2, 1, 1, 1)$. 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 case (Thomas, 2013). (An example of such a strategy is $x = (4 - 4\varepsilon, 2\varepsilon, \varepsilon, 1 + \varepsilon)$ with probability 0.5 and $x = (4 - 4\varepsilon, 2\varepsilon, 1 + \varepsilon, \varepsilon)$ with probability 0.5, where ε is uniformly distributed between 0 and 1.) 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.

we follow Young (1978) and consider two games from the class of apex games.[8] In one of his games there are four battlefields with values $v = (2, 1, 1, 1)$, while in the other there are five battlefields with values $v = (3, 1, 1, 1, 1)$. Young solved these two games 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).

To identify theoretical predictions for these games we assume each player has a budget of 120 indivisible units. We restrict attention to "object-symmetric" strategies, i.e., mixed strategies that put equal weight on all possible permutations of a given allocation across symmetric battlefields. For example, in the four battlefield case, where the first battlefield is the large battlefield and the other three are symmetric small battlefields, one possible object-symmetric strategy consists of playing each of (80,40,0,0), (80,0,40,0) and (80,0,0,40) with probability one-third. Note that the equilibrium of the game with object-symmetric strategies is also an equilibrium of the original game. Even with this restriction, the number of available strategies is rather large.[9] We then calculate equilibrium strategies numerically using the Gambit package (McKelvey et al., 2013).

Although the strategy spaces are large, the fact that these are two-player constant-sum games makes the problem of computing equilibria tractable. In particular, the problem of finding a mixed strategy equilibrium (more specifically in this case a minimax strategy) can be expressed as a linear program for which practical solution algorithms are available. Moreover, since 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 verify if it is unique. Appendix A gives technical details.

In both games the equilibrium is unique, and confirms Young's results concerning the amounts allocated to the large battlefield. In equilibrium, the expected share of resources devoted to the large battlefield is around 50% in the four battlefield case and 58% in the five battlefield case. In each case the large battlefield receives a share of resources exceeding its value as a proportion of the total value. We refer to this as the *super-proportionality* property

---

[8] In an apex game with n battlefields one "large" battlefield is more important than the other $(n-1)$ "small" battlefields. In particular, to win overall a player must win either all the small battlefields or the large and one small battlefield.

[9] The number of object-symmetric strategies is approximately 52,000 in the four battlefield case and 430,000 in the five battlefield case.

of the equilibrium. The equilibrium also allows us to see which battlefields a player targets. Does a player concentrate resources on all the small battlefields, or on the large battlefield and just one small battlefield? Or does a player hedge and spread the budget over all battlefields? At first sight it seems that there is no point spending resources on more battlefields than are needed to win. However, in equilibrium a player places positive amounts on more battlefields than are needed to win the game with a probability exceeding 90%. We refer to this as the *hedging* property of equilibrium. In the next section we describe an experimental design to test these predictions.

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.[10] 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 'objects', each of which was worth a given number of points. A subject could only submit bids that added up to 120, and had 90 seconds to submit the bids.[11] A subject won an object if he outbid his opponent on that object (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 object, who won each object, and how much they earned.

---

[10] Instructions for one of the treatments can be found in Appendix B.
[11] If subjects timed out, the computer made a default decision allocating zero tokens to each object. Across all sessions only 28 out of 6,660 allocation decisions resulted in a timeout.

We ran three treatments. APEX4 and APEX5 use Young's (1978) apex games with four and five battlefields respectively. For these treatments Object A represented the large battlefield. For comparison, we also ran a treatment using a (degenerate) apex game with three symmetric battlefields (APEX3).[12] For each treatment we conducted three sessions with between 14 and 20 subjects 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.[13]

Table 1. Experimental treatments					
Treatment	Values of objects	Proportional share of expenditure on object A $(v_A/\Sigma v_i)$	Equilibrium share of expenditure on object A	Number of pairs	Number of subjects
APEX3	$v = (1, 1, 1)$	0.33	0.33	23	46
APEX4	$v = (2, 1, 1, 1)$	0.40	0.50	26	52
APEX5	$v = (3, 1, 1, 1, 1)$	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, 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). Second, 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). Also, keeping subjects in the same pairs simplifies the structure of possible dependencies between decisions

---

[12] 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).

[13] 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. For completeness we report these sessions in Appendix C.

and in particular allows us to treat each pair as an independent observation since subjects in one pair cannot influence or be influenced by the decisions of subjects in any other pair.


5. RESULTS

5.1 Predicted and observed distributions

We begin with an overview of our results and how they relate to equilibrium predictions. Figure 1 displays the equilibrium and the empirical distribution of allocations for the APEX3 treatment. In equilibrium, the players use mixed strategies where the marginal distribution of tokens on each object is approximately uniform on {0,…,80}. In contrast, there is a pronounced bi-modality in subject choices, with subjects tending to place either very small amounts or about half their budget on an object 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, and Mago and Sheremeta, 2012). Figure 1 also shows that the distributions of bids are similar across the three objects, although subjects allocated slightly more to Object A than to Object B than to Object C.[14]

Figure 1. Predicted and observed bids in APEX3
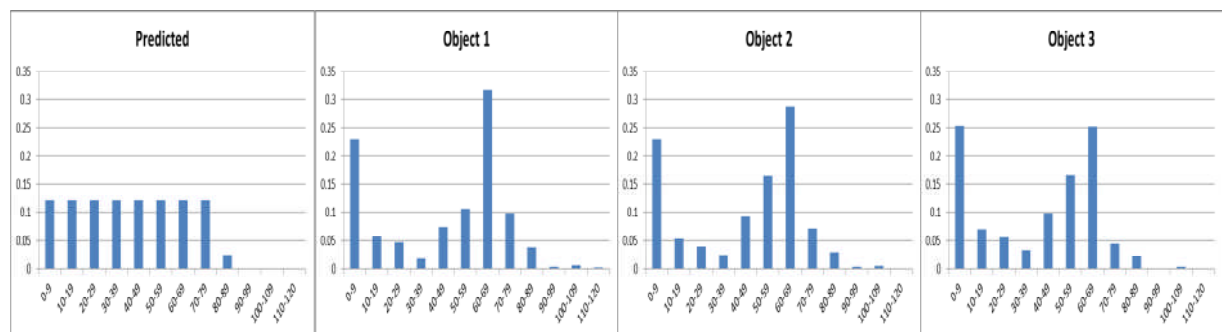
Next we turn to the treatments with asymmetric battlefields. Figure 2 displays the equilibrium marginal distribution of bids for Object A (the large battlefield) and the equilibrium marginal distribution of bids for one of the other objects (a small battlefield) for the APEX4 treatment. The figure also shows the distributions actually observed in the experiment, where for the marginal distribution on a small battlefield we pool the data from

---

[14] Again, this echoes previous experimental findings. Chowdhury et al. (2013) also observe mild positional order effects (see their table 3 and figure 3).

all the small battlefields.

Figure 2. Predicted and observed bids in APEX4
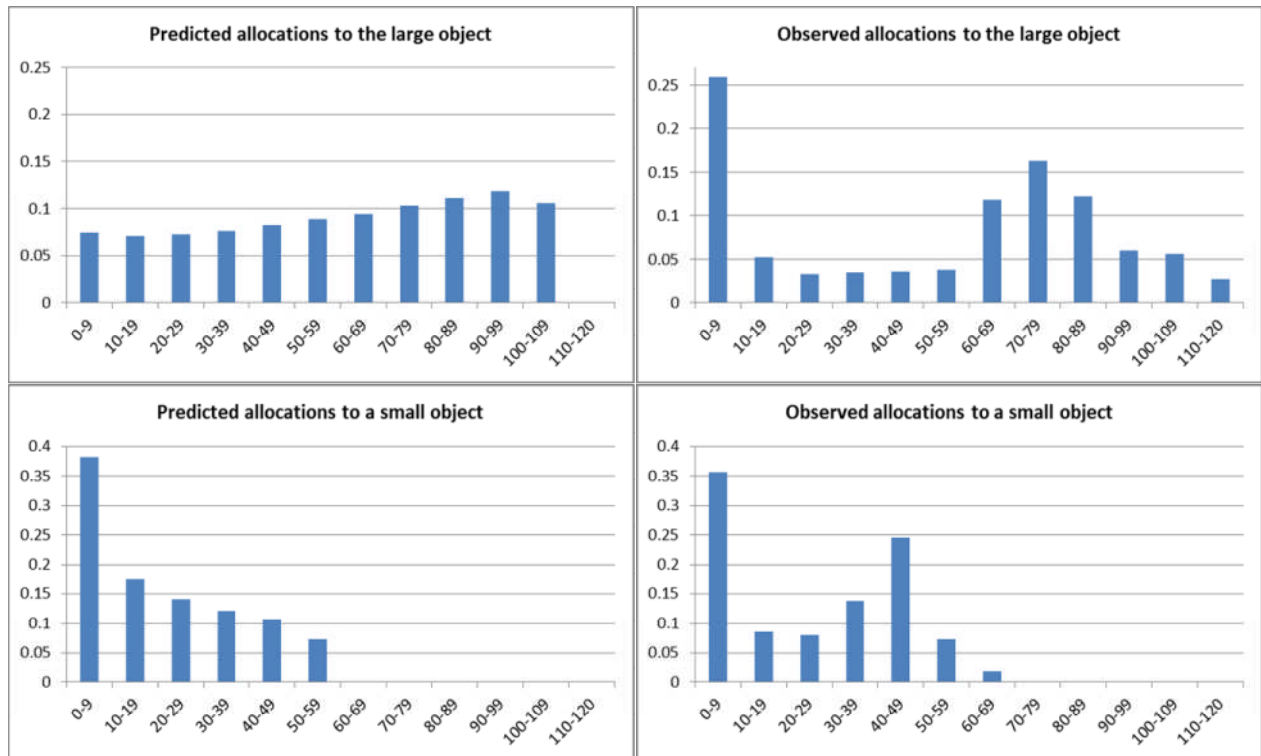
Again, the observed distributions are bi-modal, and are markedly different from the theoretical distributions. Relative to equilibrium, subjects too often place either very low amounts or around two-thirds of their budget on the large battlefield. Analogously, a small battlefield is often allocated an amount close to 0 or an amount close to one-third of the budget.

Figure 3 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.

Figure 3. Predicted and observed bids in APEX5
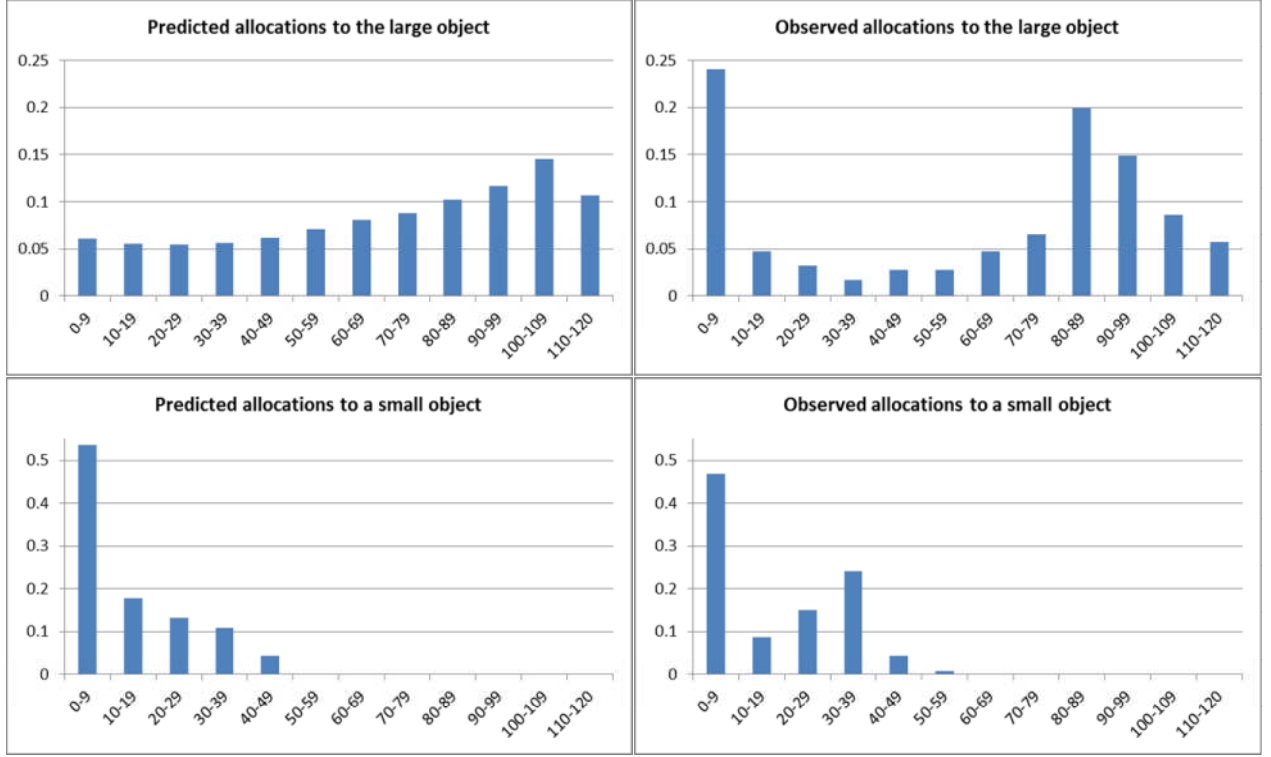
In summary, for all treatments we observe bi-modal bidding patterns, suggesting subjects sometimes give up on a battlefield and concentrate their forces on a subset of them. The positions of the modes suggest that when concentrating on a subset of battlefields subjects allocate budgets roughly in proportion to the value of objects within this subset.

5.2 Minimal winning and hedging strategies

Note that in order to win a round a player only needs to capture battlefields with a combined value exceeding that of his opponent's captured battlefields. To do this in APEX3 a player simply needs to win two battlefields, thus subjects may find it natural to give up on one of the battlefields and concentrate resources on just two of them. We refer to a strategy that places zero on one battlefield as a *minimal winning strategy*. In contrast, a *hedging strategy* places a positive amount on all three battlefields. 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. Minimal winning strategies have an equilibrium probability of around 7% in APEX3, but they are observed much more

often than predicted (42% of submitted strategies are minimal winning).[15]

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 voter 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 4 presents the evolution of the proportion of hedging and minimal winning strategies over time. For APEX4 and APEX5 "MWL" denotes minimal winning strategies 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.

**Figure 4. Proportions of minimal winning and hedging strategies**

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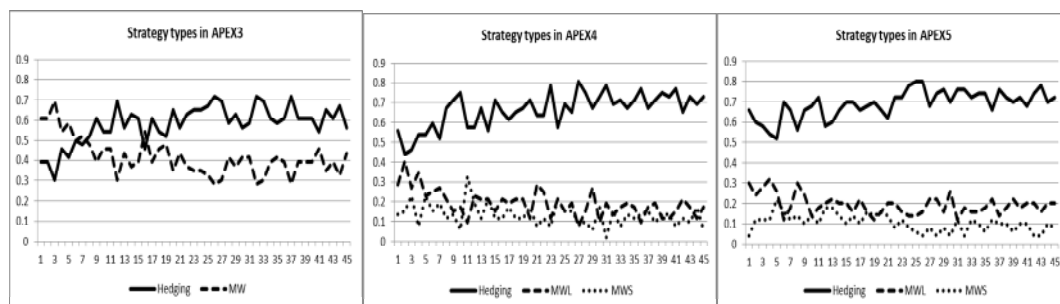
[15] A similar result is obtained by Mago and Sheremeta (2012) in a majoritarian contest with linear costs. In their experiment 35% of the time subjects bid only on two out of three objects, whereas in equilibrium they should make positive bids on all three.

5.3 Super-proportionality

Equilibrium predicts that players spend a disproportionate amount of their budgets on the large battlefield. If subjects' bids on objects were proportional to object 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 5 shows the share of expenditure on the large battlefield (Object A) relative to its proportional share, i.e. $(x_A/120) \div (v_A/\Sigma v_i)$. The theoretical prediction is that this measure is 1 for APEX3 (where the battlefields are symmetric) and increases with the number of battlefields. Even though bids on the large battlefield are not as high, on average, as predicted by equilibrium, Figure 5 suggests that Object A receives more than a proportional share of the budget. For APEX3, the battlefields are symmetric and so the small deviation from proportionality suggests a positional order effect. For APEX4 the allocation ratio is essentially the same, and so again the small departure from proportionality may be attributed to a positional effect rather than to strategic considerations. Allocations are clearly super-proportional in APEX5, although even in this case they are well below the equilibrium prediction.

Figure 5. Share of budget allocated to large battlefield relative to proportional share
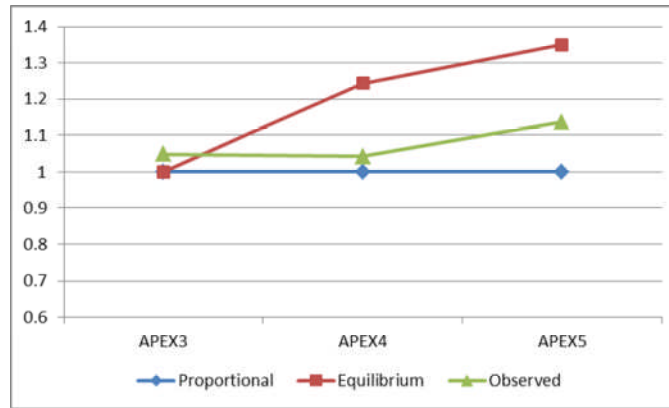
Formal statistical tests are presented in Table 2, which shows the share of the budget allocated to object A in each treatment and p-values for tests against proportionality. We use two-sided sign-rank tests treating each pair as an independent observation.[16]

---

[16] 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 Object A (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 Object A is proportional to value.

Object A has a small but significant positional advantage in APEX3, although this advantage becomes insignificant in the later rounds. In APEX4 the share allocated to the large battlefield is slightly, but insignificantly, higher than a proportional share. In APEX5 the share allocated to the large battlefield is significantly higher than proportional, whether we average across all rounds or focus on later rounds. Thus, we find significant evidence of super-proportionality in APEX5, but not in APEX4.

After using this conservative approach to test our main hypothesis, 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:

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

where $AllocA_{it}$ refers to the number of tokens allocated by subject i to object A 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 object A 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.

Table 3. Determinants of allocation to Object A			
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)
$1/t$	13.848** (6.001)	24.782*** (7.336)	27.788** (9.977)
Constant	43.990*** (1.886)	45.724*** (2.949)	58.261*** (3.691)
# Observations	2,024	2,262	2,143
# Subjects	46	52	50

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

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 when subjects play repeatedly against the same opponent, and the hot-box effect disappears. In our treatments 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. On the other hand, the significant trend means that allocations are predictable to some extent. Note however, that the variable $1/t$ models a diminishing trend, and the significant negative trend is essentially capturing a reduction in average amount allocated to Object A in initial rounds.[17]

### 5.4 Heterogeneity

For all treatments the lobbying game is a symmetric constant-sum game and so in equilibrium each lobbyist wins with probability 1/2 in any play. This means that a player expects to win

---

[17] Restricting 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.

22.5 out of the 45 games. In fact, some do considerably better than this. For example, in one of the pairs one subject won 11 rounds and the other won 34 rounds, so that the difference in wins was 23. Figure 6 shows the observed frequencies of each possible value of the difference in wins. For comparison the theoretical distribution is also shown.

Figure 6. Theoretical and Observed Distributions of Differences in Wins
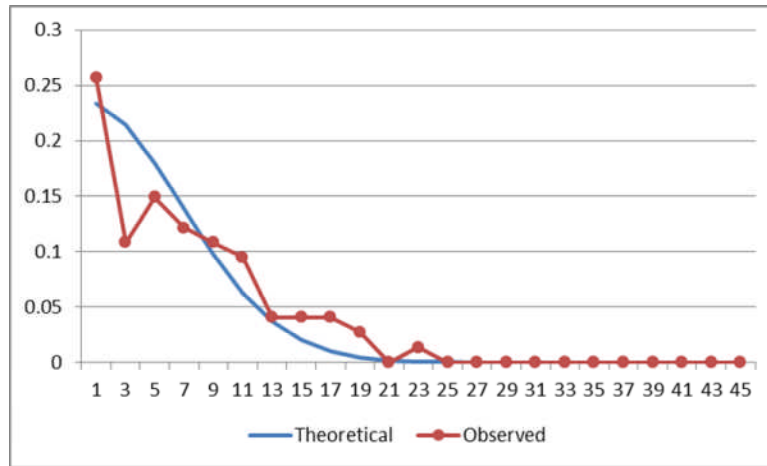
The figure shows that fewer than expected pairs have a small difference in wins and more than expected have a large difference in wins. Theoretically, the expected difference in wins is 5.38, while in the data the average difference in wins is 6.84. This difference is significant (Chi-square test p = 0.028). The obvious interpretation is that some subjects are better than others at playing the lobbying game.

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

$$\Pr\{Win_{it} = 1\} = \Phi(\beta_0 + \beta_1 AllocA_{it} + \beta_2 MADTemp_{it} + \beta_3 MADAcr_{it} + \beta_4 MWL_{it} + \beta_5 MWS_{it} + u_i + \varepsilon_{it}).$$

Here Φ 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. $MADTemp_{it}$ is the mean absolute deviation of allocations across rounds and measures the variability of a subject's allocation between rounds, while $MADAcr_{it}$ measures the variability of the

allocation across objects in a given round.[18] MWL_{it} and MWS_{it} are indicator variables for subject i playing different types of minimal winning strategy. In APEX3, there is only one type of minimal winning strategy and so we replace the MWL_{it} and MWS_{it} variables with a single variable MW_{it}. The regression includes subject level random effects. Results are reported in Table 4.

Table 4. Determinants of winning			
Treatments	APEX3	APEX4	APEX5
Dependent variable: Win_{it}			
$AllocA_{it}$	-0.002** (0.001)	-0.003* (0.001)	0.002 (0.002)
$MADTemp_{it}$	0.005*** (0.002)	0.009*** (0.002)	0.020*** (0.002)
$MADAcr_{it}$	0.007 (0.005)	0.007 (0.005)	-0.009 (0.009)
MW_{it}	-0.138** (0.066)		
MWL_{it}		0.115 (0.076)	0.106 (0.082)
MWS_{it}		-0.166* (0.097)	-0.005 (0.112)
Constant	-0.189* (0.137)	-0.279*** (0.092)	-0.367*** (0.104)
# Observations	2,024	2,254	2,124
# Subjects	46	52	50

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

Consistent with Chowdhury et al. (2013) we find that $MADTemp_{it}$ is highly significant, hence unpredictability of the allocation is one of the main determinants of success; also consistent with their results, the variability of the allocation across objects does not have a significant effect on the probability of winning. In APEX3, subjects tend to place too much of their budgets on Object A, and to play minimal winning strategies too often.

---

[18] If y_{ijt} is the allocation of subject i to object j at time t and n is the number of objects in the treatment,

$$MADTemp_{it} = \sum_{j=1}^{n} | y_{ijt} - y_{ij(t-1)} | / n \text{, and } MADAcr_{it} = \sum_{j=1}^{n} | y_{ijt} - 120/n | / n .$$

Against this actual distribution of play subjects who placed less on Object A and used minimal winning strategies less often were more likely to win. The picture in APEX4 and APEX5 is a little different. Recall that, relative to equilibrium, subjects play minimal winning strategies too frequently and place too little of their budgets on Object A. Thus, one might expect that minimal winning strategies will win less often and strategies placing more on Object A will win more often in these treatments. In fact this is not always the case. In APEX4 minimal winning strategies that focus on the small battlefields are less likely to win, but strategies (not necessarily minimal winning) that place more on Object A are also less likely to win (both effects are marginally significant). In APEX5 we do not observe any significant effect of the use of minimal winning strategies or the amount placed on Object A on the probability of winning. 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 the next section we look more closely at which strategies perform best against the empirical distribution.

5.5 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 5 displays this measure of exploitability, for some selected strategies. We consider a "uniform" strategy and a "minimal winning" strategy. The uniform strategy is a mixed strategy that induces a uniform marginal distribution of tokens on each battlefield, with the expected allocation to a battlefield being proportional to its valuation. 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).[19] Finally 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.

Table 5. 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	0.750	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	0.833	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 also be beaten 75% of the time in APEX4, and does even worse than the minimal winning strategy in APEX5, where it can be beaten 83% 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

---

[19] These minimal winning strategies are mixtures of particular MWS and object-symmetric MWL strategies. Different weights on the MWS and MWL components increase the exploitability of the strategy.

consistent with the finding from the previous section that subjects spending less on Object A 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 and less exploitable in the less complex APEX3 than in the more complex APEX4, but this pattern breaks down once we consider 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 APEX3, suggesting that in the game with more battles subjects "raise their game" and find ways to protect themselves from being exploited.

6. CONCLUDING REMARKS

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, 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. If one battlefield has a value that is twice as much as that of another battlefield, should it get twice as many resources? The answer crucially depends on how battlefield values contribute to final victory or defeat. For the case we consider, where 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. First, 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. Second, and related to the first point, 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. Third, 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. For example, 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

Computing an equilibrium and verifying uniqueness

The games APEX3, APEX4, and APEX5 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). We first develop some generic notation. Write the pure strategies available to each player as {1, 2, ..., M}. Let u_{ij} denote the payoff to a player, if the player chooses pure strategy i while the other player chooses pure strategy j. Let π denote a mixed strategy, where π_i indicates the probability assigned in the mixed strategy to pure strategy i. Then, the payoff to a player of choosing strategy i if his opponent plays a given mixed strategy π is $\sum_{j=1}^{M} u_{ij}\pi_j$. Two-player, constant-sum games have a value, which we write as ω. We must therefore have that no pure strategy for a player can give him a payoff greater than the value, that is,

$$\sum_{j=1}^{M} u_{ij}\pi_j \le \omega \quad \forall i \in 1,2,\ldots,M, \tag{1}$$

and that the mixed strategy is a proper distribution,

$$\pi_i \ge 0 \quad \forall i \in 1,2,\ldots,M, \tag{2}$$

$$\sum_{j=1}^{M} \pi_j = 1. \tag{3}$$

Because of the symmetry of the games APEX3, APEX4, and APEX5, and the payoff structure, the value of the game is known in advance to be $\omega = 1/2$.[20] Therefore, finding a minimax strategy reduces to finding a mixed strategy π which satisfies the constraints (1), (2), and (3). This can be embedded into a linear program, where the choice of the objective function is arbitrary. Any feasible solution is a minimax equilibrium of the game.

This formulation also permits verification of the uniqueness of the equilibrium. It is easy to see that the set of equilibria is convex in this setting, as the set of π satisfying the constraints (1), (2), and (3) is convex. Let Φ denote the set of equilibria. The set of equilibria is a singleton if and only if $\min_{\pi \in \Phi} \pi_i = \max_{\pi \in \Phi} \pi_i$ for all strategies $i \in \{1,2,\ldots,M\}$.

---

[20] Knowing the value in advance allows us to simplify the construction of the linear program. Dantzig's original construction computes the value of the game as part of its output.

Therefore, uniqueness can be verified as follows. Pick a strategy i. Let $\underline{\pi_i}$ be the optimal value of the linear program $min_\pi \pi_i$ subject to the constraints (1)-(3), and let $\overline{\pi_i}$ be the optimal value of the linear program $max_\pi \pi_i$ subject to the constraints (1)-(3). If $\underline{\pi_i} = \overline{\pi_i}$ for all $i \in \{1,2,\ldots,M\}$, then the equilibrium is unique.

The step of verifying uniqueness can be done relatively efficiently once one equilibrium has been found, as this equilibrium is a feasible solution; the linear programming algorithm can be started at this feasible solution. Note also that if one is interested in knowing simply whether Φ is a singleton, it is enough to consider only the strategies i which have value one-half at the first known equilibrium; strategies which are strictly inferior need not be checked. The main computational challenge is in finding the first feasible solution.

Improving 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 battlefield-symmetric strategies, APEX3 has 1261 strategies, APEX4 has 52311, and APEX5 has 430256. However, preliminary explorations with smaller budgets led us to conjecture that equilibria in these games would have small support. We therefore used an iterative method to identify the set of strategies.

Iteration on supports

Consider the game APEXk with a budget of E tokens. Let S be the set of pure strategy token allocations. We construct an increasing sequence of supports, $S_0 \subseteq S_1 \subseteq \ldots \subseteq S_n$, such that S_n is the support of an equilibrium of the game. Pick some initial guess at the support of the equilibrium, and call it $S_0 \subseteq S$. (The correctness of the construction does not depend on the value of the initial guess S_0; for this approach to work efficiently, it should be small in size.)

At each step i of the algorithm, we consider the restriction of APEXk in which players can choose only strategies in S_i. This induces a well-defined two-player constant-sum game, which can be solved for some equilibrium π_i; insofar as $|S_i| << |S|$ solving the restricted game should be much faster than solving the full game. Then, given π_i, we consider all the

24

strategies which were deleted from the restricted game, $S \setminus S_i$, and order them in decreasing order by their payoff relative to the candidate equilibrium π_i. If there are no strategies which attain a payoff greater than the equilibrium payoff of one-half, then the algorithm terminates, and π_i is an equilibrium of the full game with strategy set S. If not, then we construct S_{i+1} by adding the top t strategies from $S \setminus S_i$ to those in S_i, and iterating.

The number of strategies t added at every step is arbitrary; we obtained sufficiently good performance with $t = 25$. The tradeoff is that if t 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 t is too large the algorithm will consider many strategies which are not in the equilibrium, 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 S_{i+1} is always a strict superset of S_i, and because all the supports are bounded above (in the sense of set inclusion) by the whole strategy set S, this iterative process is guaranteed to terminate in a finite number of steps.

Iteration on budgets

The support iteration approach is most effective if the initial guess S_0 on the support of pure strategies used in equilibria is accurate. Working on a conjecture that the equilibrium has a similar qualitative structure for various budget sizes, we considered a sequence of games with token budgets $E = 15; 30; 60; 120$. Given an equilibrium $\pi^*(E)$ for the game with budget E, we then constructed an initial guess for the support iteration for the game with budget $2E$ by doubling all the token allocations as follows. A pure strategy in APEXk can be written as a vector of token allocations $(a_b)_{b=1}^{k}$. The initial guess for the support of the equilibrium in the game with budget $2E$ is then the set of allocations such that the allocation $(a_b / 2)_{b=1}^{k}$ is played with positive probability in the equilibrium $\pi^*(E)$ of the game with budget E.

Using this approach, we verified that the main qualitative properties of the equilibrium referenced in the main body of the paper regarding the marginal distributions of token allocations on each contest and the prevalence of supermajority strategies also hold with other budget sizes. We also found that the equilibrium is unique for most - but not all - choices of the budget E.

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.

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

26

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):

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)? _____.

We also run experiments using a coarse budget of 5 indivisible units. The motivation for the coarse budget is that there are relatively few strategies and this makes the game much easier to analyze (and perhaps to play). Having a coarser budget has little effect on the equilibrium expected expenditures on objects. Experimental results show that object A's share of expenditure is higher than proportional in both asymmetric battlefield treatments and is very close to the equilibrium share in the APEX5 treatment. Section C.1 contains the equilibrium prediction for all three apex games with a coarse budget. Section C.2 contains the experimental results.

C.1 Equilibrium predictions for the coarse budget

The game APEX3

We restrict attention to object-symmetric strategies. There are 5 possible object-symmetric strategies in this game, and the following table gives player 1's expected payoff in the normal-form game with these five strategies (for example, "410" represents the object-symmetric mixed strategy that puts equal probability on each of the pure strategies (4,1,0), (4,0,1), (1,4,0), (0,4,1), (1,0,4), (0,1,4)). Recall that, because there are only two possible outcomes (winning and losing), risk attitudes are irrelevant under expected utility theory and a player's payoff can be identified with the probability of winning. Recall also that any equilibrium of this game is also an equilibrium of the original game.

	500	**410**	**320**	**311**	**221**
500	0.5	1/3	1/3	0	0
410	2/3	0.5	0.5	1/3	1/6
320	2/3	0.5	0.5	0.5	2/3
311	1	2/3	0.5	0.5	1/3
221	1	5/6	1/3	2/3	0.5

There is a continuum of equilibria in the game, described by strategy combinations $(\lambda_1\, 320 + (1 - \lambda_1)\, 311, \lambda_2\, 320 + (1 - \lambda_2)\, 311)$ for $1/2 \leq \lambda_1, \lambda_2 \leq 1$. This includes a "pure" strategy equilibrium (320, 320), but recall that 320 is actually a particular mixed strategy. Note also that 320 is the only strategy in the normal form game above that survives the iterated elimination of weakly dominated strategies.

The game APEX4

There are 16 object-symmetric strategies. The table below is the resulting normal form game (entries in the table correspond to the probability that the row player wins).

	5000	4100	3200	3110	2300	2210	2111	1400	1310	1220	1211	0500	0410	0320	0311	0221
5000	0.5	3/4	3/4	0.5	3/4	0.5	0	3/4	0.5	0.5	0	3/4	0.5	0.5	0	0
4100	1/4	0.5	11/12	5/6	11/12	3/4	0.5	11/12	3/4	2/3	1/3	11/12	3/4	2/3	1/3	1/6
3200	1/4	1/12	0.5	0.5	11/12	11/12	1	11/12	5/6	5/6	5/6	11/12	5/6	3/4	2/3	2/3
3110	0.5	1/6	0.5	0.5	1	11/12	3/4	1	11/12	5/6	7/12	1	11/12	5/6	7/12	1/3
2300	1/4	1/12	1/12	0	0.5	0.5	0.5	11/12	11/12	1	1	11/12	5/6	11/12	5/6	1
2210	0.5	1/4	1/12	1/12	0.5	0.5	0.5	1	23/24	11/12	11/12	1	23/24	7/8	5/6	3/4
2111	1	0.5	0	1/4	0.5	0.5	0.5	1	1	1	3/4	1	1	1	3/4	0.5
1400	1/4	1/12	1/12	0	1/12	0	0	0.5	0.5	0.5	0.5	11/12	11/12	1	1	1
1310	0.5	1/4	1/6	1/12	1/12	1/24	0	0.5	0.5	0.5	0.5	1	23/24	47/48	11/12	1
1220	0.5	1/3	1/6	1/6	0	1/12	0	0.5	0.5	0.5	0.5	1	1	11/12	1	11/12
1211	1	2/3	1/6	5/12	0	1/12	1/4	0.5	0.5	0.5	0.5	1	1	1	11/12	5/6
0500	1/4	1/12	1/12	0	1/12	0	0	1/12	0	0	0	0.5	0.5	0.5	0.5	0.5
0410	0.5	1/4	1/6	1/12	1/6	1/24	0	1/12	1/24	0	0	0.5	0.5	0.5	0.5	0.5
0320	0.5	1/3	1/4	1/6	1/12	1/8	0	0	1/48	1/12	0	0.5	0.5	0.5	0.5	0.5
0311	1	2/3	1/3	5/12	1/6	1/6	1/4	0	1/12	0	1/12	0.5	0.5	0.5	0.5	0.5
0221	1	5/6	1/3	2/3	0	1/4	0.5	0	0	1/12	1/6	0.5	0.5	0.5	0.5	0.5

Using the Gambit software (McKelvey et al., 2013) we found a unique equilibrium of the normal-form game above, with probabilities 30/77 on 4100, 12/77 on 3200, 8/77 on 2111, 24/77 on 1211 and 3/77 on 0221. In this equilibrium, the expected share of the total budget allocated to the large object is $\frac{1}{5}\left(\frac{30}{77}\times 4+\frac{12}{77}\times 3+\frac{8}{77}\times 2+\frac{24}{77}\right)=\frac{28}{55}\approx 0.51$. The expected

share of each small object is $\frac{1}{3}\left(1-\frac{28}{55}\right)=\frac{9}{55}\approx 0.16$. Note that the equilibrium places positive probability on the hedging strategies 2111 and 1211.

The game APEX5

Again we restrict attention to object-symmetric strategies. We also discard some strategies that look implausible (obtained equilibria are later checked against invasion by those strategies). We discarded strategy 14000 and strategies that allocate the budget to a losing subset of small objects, leaving 12 object-symmetric strategies:

	50000	41000	32000	31100	23000	22100	21110	13100	12200	12110	11111	02111
50000	0.5	7/8	7/8	3/4	7/8	3/4	0.5	3/4	3/4	0.5	0	0
41000	1/8	0.5	31/32	15/16	31/32	29/32	13/16	29/32	7/8	3/4	0.5	3/8
32000	1/8	1/32	0.5	0.5	31/32	31/32	1	15/16	15/16	15/16	1	7/8
31100	1/4	1/16	0.5	0.5	1	47/48	15/16	47/48	23/24	43/48	3/4	5/8
23000	1/8	1/32	1/32	0	0.5	0.5	0.5	31/32	1	1	1	1
22100	1/4	3/32	1/32	1/48	0.5	0.5	0.5	95/96	47/48	47/48	1	15/16
21110	0.5	3/16	0	1/16	0.5	0.5	0.5	1	1	31/32	7/8	25/32
13100	1/4	3/32	1/16	1/48	1/32	1/96	0	0.5	0.5	0.5	0.5	1
12200	1/4	1/8	1/16	1/24	0	1/48	0	0.5	0.5	0.5	0.5	1
12110	0.5	1/4	1/16	5/48	0	1/48	1/32	0.5	0.5	0.5	0.5	31/32
11111	1	0.5	0	1/4	0	0	1/8	0.5	0.5	0.5	0.5	7/8
02111	1	5/8	1/8	3/8	0	1/16	7/32	0	0	1/32	1/8	0.5

Using the Gambit software, we found a continuum of equilibria described by strategy combinations $(\lambda_1\,41000 + (1 - \lambda_1)\,11111, \lambda_2\,41000 + (1 - \lambda_2)\,11111)$ for $16/29 \le \lambda_1, \lambda_2 \le 4/7$. The expected share allocated to the large object is between $77/145 \approx 0.53$ and $19/35 \approx 0.54$.

C.2 Experimental results for the coarse budget

The coarse budget experiment had one APEX3 session, one APEX4 session, and two APEX5 sessions. Instructions are identical to the 120 token budget treatment (see Appendix B), except for references to the number of tokens. Table C1 summarizes the experimental design and the average share of budget allocated to Object A in the experiment.

Table C1. Experimental treatments					
Treatment	Number of pairs	Values of objects	$v_A/\Sigma v_i$	Equilibrium share of expenditure on Object A	Observed share of expenditure on Object A
APEX3	8	$v = (1, 1, 1)$	0.33	0.33	0.35
APEX4	10	$v = (2, 1, 1, 1)$	0.40	0.51	0.45
APEX5	15	$v = (3, 1, 1, 1, 1)$	0.43	0.53-0.54	0.57

Observed allocations in APEX3

The frequencies of strategy types, averaged over all 45 rounds, are shown in Table C2 below. With a slight abuse of notation, we use 320 to denote both the object-symmetric strategy in which each of the six possible permutations has equal probability (in the equilibrium analysis), and the strategy type, i.e. the set of permutations not necessarily with the same frequency (in the analysis of the data). Note that the equilibrium strategy type 320 has the highest proportion in all pairs. However, most pairs played non-equilibrium strategies more than 20% of the time. Across all pairs 35% of tokens were allocated to Object A. However, the positional advantage of Object A is statistically insignificant (sign-rank test two-sided p-value = 0.4833).

Table C2. Observed allocations in APEX3									
Type\Pair	1	2	3	4	5	6	7	8	All
500	0.00	0.00	0.00	0.00	0.00	0.00	0.01	0.00	0.00
410	0.00	0.22	0.02	0.32	0.01	0.07	0.26	0.00	0.11
320	0.56	0.68	0.77	0.52	0.67	0.69	0.43	0.72	0.63
311	0.08	0.02	0.01	0.08	0.19	0.20	0.06	0.19	0.10
221	0.37	0.08	0.20	0.08	0.13	0.04	0.24	0.09	0.15

Observed allocations in APEX4

Table C3 compares the frequencies of strategy types with equilibrium in APEX4, pooling across pairs. Behavior is quite far from equilibrium. The only two hedging strategies predicted in equilibrium are 2111 (with 10% probability) and 1211 (with 31% probability). These are observed only 4 and 5% of the time respectively in the experiment. Although we observed hedging strategies quite frequently (e.g. 3110 was observed 10% of the time), they are not played as frequently as in equilibrium, and they are not the hedging strategies that should be played in equilibrium. Across all pairs, 45% of the budget is allocated to Object A.

This is significantly higher than proportional (sign-rank test two-sided p-value = 0.0217), but also significantly lower than the equilibrium prediction (sign-rank test two-sided p-value = 0.0050).

Table C3. Observed allocations in APEX4								
Strategy Type	4100	3200	2111	1211	0221	3110	2210	1220	Other
Predicted frequency	0.39	0.16	0.10	0.31	0.04	-	-	-	-
Observed frequency	0.20	0.20	0.04	0.05	0.22	0.10	0.05	0.04	0.09

Observed allocations in APEX5

Table C3 compares the frequencies of strategy types with equilibrium in APEX5, pooling across pairs. (There is a small interval of equilibria, with varying weights on 41000 and 11111. Table C4 reports the midpoint). As in APEX4, the proportion of hedging strategies observed is lower than predicted, 34% overall, and these often correspond to strategies that should not be played in equilibrium. Across all pairs, 57% of the budget is allocated to Object A. This is significantly higher than proportional (sign-rank test two-sided p-value = 0.0007), and we cannot reject the hypothesis that the share is different from 19/35 (the upper bound of the equilibrium prediction) (sign-rank test two-sided p-value = 0.1728).

Table C4. Observed allocations in APEX5								
Strategy Type	41000	11111	50000	02111	31100	32000	21110	Other
Predicted frequency	0.56	0.44	-	-	-	-	-	-
Observed frequency	0.27	0.10	0.17	0.12	0.12	0.07	0.07	0.08

Summary

Overall the results for the coarse budget are qualitatively similar to the results with the fine budget. Over all rounds, object A is allocated a slightly more than proportional share in APEX4 and a clearly super-proportional share in APEX5; hedging is frequently observed though not as frequently as equilibrium theory would predict. The main difference between the coarse budget and the fine budget experiments is that, with the coarse budget, the observed share for the large object in APEX4 is significantly higher than proportional and the observed share in APEX5 is substantially closer to equilibrium.

REFERENCES

Arad, A. and A. Rubinstein (2012) "Multi-Dimensional Iterative Reasoning in Action: The Case of the Colonel Blotto Game", *Journal of Economic Behavior and Organization* 84, 475-712.

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

Avrahami, J., Kareev, Y., Todd, P. and B. Silverman (2013) "Allocation of Resources in Asymmetric Competitions: How do the Weak Maintain a Chance of Winning?", *Journal of Economic Psychology*, in press.

Borel, E. (1921[1953]) "La théorie du jeu et les équations intégrales à noyau symétrique", *Comptes Rendus de l'Académie des Sciences* 173, 1304-1308; English translation: Savage, L. (1953) "The Theory of Play and Integral Equations with Skew Symmetric Kernels", *Econometrica* 21, 97-100.

Chowdhury, S.M., Kovenock, D. and R.M. Sheremeta (2013) "An Experimental Investigation of Colonel Blotto Games", *Economic Theory* 52, 833-861.

Danzig, G. (1951) "A proof of the equivalence of the programming problem and the game problem", in Koopmans (ed), *Activity Analysis of Production and Allocation*. Wiley, pp. 330-335.

Dechenaux, E., Kovenock, D. and R.M. Sheremeta (2012) "A Survey of Experimental Research on Contests, All-pay Auctions and Tournaments", Working Paper.

Duffy, J. and A. Matros (2013) "Stochastic Asymmetric Blotto Games: Theory and Experimental Evidence", Working Paper.

Fischbacher, U. (2007) "z-Tree: Zurich Toolbox for Ready-Made Economic Experiments", *Experimental Economics* 10, 171-178.

Greiner, B. (2004) "An Online Recruitment System for Economic Experiments", in: Kremer, K., Macho, V. (Eds.), *Forschung und Wissenschaftliches Rechnen* 2003. GWDG Bericht 63, Göttingen: Ges. für Wiss. Datenverarbeitung, pp. 79-93.

Hart, S. (2008) "Discrete Colonel Blotto and General Lotto Games", *International Journal of Game Theory* 36, 441-460.

Kovenock, D. and B. Roberson (2012) "Conflicts with Multiple Battlefields", in *Oxford Handbook of the Economics of Peace and Conflict* (ed. by M. R. Garfinkel and S. Skaperdas). Oxford University Press, pp. 503-531.

McKelvey, R., McLennan, A. and T. Turocy (2013) Gambit: Software Tools for Game Theory, Version 13.1.0, http://www.gambit-project.org.

Mago, S.D. and R.M. Sheremeta (2012) "Multi-Battle Contests: An Experimental Study", Working Paper.

Roberson, B. (2006) "The Colonel Blotto Game", *Economic Theory* 29, 1-24.

Thomas, C. (2013) "N-Dimensional Blotto Game with Asymmetric Battlefield Values", Working Paper.

Walker, M. and J. Wooders (2001) "Minimax Play at Wimbledon", *American Economic Review* 91, 1521-1538.

Wooders, J. and J. Shachat (2001) "On the Irrelevance of Risk Attitudes in Repeated Two-Outcome Games", *Games and Economic Behavior* 34, 342–363

Young, H. P. (1978) "A Tactical Lobbying Game", RR-78-10, International Institute for Applied Systems Analysis, Laxenburg, Austria.