Colonel Blotto with Imperfect Targeting

Deborah Fletcher[a,*], Adam Narkiewicz[b], Steven Slutsky[b]

[a]**Department of Economics, Miami University, Oxford, OH 45056**
[b]**Department of Economics, University of Florida, Matherly Hall, Gainesville, FL 32611**

November 2014

Abstract: Colonel Blotto games have been applied in a variety of contests where players allocate resources across multiple battlefields, and a battlefield is won by the player with the most resources there. One drawback of this standard model is the assumption that players can perfectly target their efforts toward different battlefields. In many situations, however, players can only imperfectly target different battlefields, with an allocation affecting multiple battlefields. We develop an extension of Blotto incorporating this type of interrelation. We use a novel numerical method, as well as standard analytic techniques, to explore equilibrium behavior. We find that when the players' resources are equal, at least one player tends to allocate resources asymmetrically across battlefields. When resources are unequal, however, even by an arbitrarily small amount, resource allocations may be much symmetric. We also find that in many cases, a player can reduce total resources without lowering his expected payoff.

JEL code: C72
Keywords: Colonel Blotto; mixed strategies

*Corresponding author.
Email addresses: fletchd@miamioh.edu, anarkiewicz@ufl.edu, slutsky@ufl.edu

(I) Introduction

A variety of contests exist in which the opponents compete by expending resources. In many games of this type, the players benefit by allocating resources differently from their opponents; the game is thus one of "strategic allocative mismatch."[1] Prominent examples of this are military contests and political campaigns. One approach to modeling such situations is the classic Colonel Blotto game.[2] In its simplest version, the players have equal resources that they allocate across different battlefields of equal size. The winner in a battlefield is determined by which player puts in more resources, that is "gets there first with the most."[3] This creates discontinuities in the payoffs of the players. The overall goal of the players is to maximize the expected number of battlefields they win. Extensions of this model allow for the players to have unequal resources, for battlefields to be of different sizes, and for one of the players to have an advantage on some battlefields so as to be able to win them with fewer resources (perhaps, having gotten there first, within limits, a player does not need the most).[4, 5]

Although the Colonel Blotto game can be straightforward to specify, equilibrium behavior can be complicated. Pure strategy equilibria exist only in special cases, such as when only two battlefields exist or when one battlefield is much larger than all the others

[1] See Golman and Page (2009).

[2] Borel (1921) first proposed this game, with further analysis in Borel and Ville (1938). Modern analysis began with Tukey (1949), with follow-ups by Blackett (1954, 1958) and Bellman (1969). Only recently has a general solution to the continuous game been provided (Roberson (2006)); see also Weinstein (2012).

[3] In response to a question about who would win a battle, Gen. Nathan Bedford Forest of the Confederate Army replied, "The one who gets there first with the most men." This quotation is often used in discussions of strategy in a variety of contexts.

[4] Kovenock and Roberson (2008) consider a model of two-party competition where parties attract voters by offering redistributive policies when voters have heterogeneous loyalties to the different parties. There is a formal overlap between getting votes through redistributive policies and attracting them through campaign expenditures, as analyzed here.

[5] Colonel Blotto models of advertising begin with Friedman (1958). Colonel Blotto models of elections go back at least to Sankoff and Mellos (1972). Recently, Szentes and Rosenthal (2003) study a type of all-pay auction of which a Colonel Blotto model of Electoral College competition is an example. Laslier and Picard (2002) consider a model of distributive politics that reduces to a Colonel Blotto game.

combined. In the latter case, the pure strategy equilibrium has both players spending all resources in the larger battlefield. Pure strategy equilibria also exist if one player has such an overwhelming resource advantage that it can win all battlefields, regardless of the other players' allocations.

Typically, when there are more than two battlefields, only mixed strategy equilibria exist, and can be quite complex. Consider a player maximizing over n battlefields. The marginal distributions of the mixed strategy give the probability of allocating different amounts to any battlefield, while the joint distributions give the probability of allocating different n-tuples across the battlefields. Necessary and sufficient conditions on the marginal distributions for a mixed strategy equilibrium have been specified (see Roberson (2006)). Consistent with any set of marginal distributions, multiple joint distributions can exist.

Characterizing the equilibrium joint distributions has been more difficult. Examples of these go back to Borel and Ville (1938) with other possibilities shown only recently by Roberson (2006) and Weinstein (2012) for Blotto games with more than two battlefields, and Macdonell and Mastronardi (2014) for Blotto games with exactly two battlefields. The equilibrium joint distributions can vary widely in their nature. Some are asymmetric and have supports that are of lower dimension than the n-1 dimensional space of allocations. There are also equilibria with full support such as the Hex equilibrium given by Borel and Ville; see Weinstein for a discussion of this equilibrium. In an equilibrium with full support, any neighborhood in the space of allocations, including any containing the point at which resources are allocated equally to all battlefields, has positive measure.

One assumption in all these variants is that players can perfectly target their resources. That is, resources can be allocated to a single battlefield with no spillover effects to other battlefields. In many contexts, however, spillovers do exist. For example, consider advertising campaigns by firms or political candidates. A battlefield might be considered to be some group defined by gender, ethnicity, age, socioeconomic status, geographic location, or preferences. In a perfect world for the players, they would be able to target expenditures in the form of ads with different messages very narrowly aimed at each demographic group, with only that group seeing the ad designed for it. While firms or candidates do try to carry out such targeting through such actions as the placement of different ads on specific radio or television programs, perfect targeting rarely is possible. Ads that are aimed at one group will often be observed by members of other groups.[6] This is only beginning to change with the advent of targeted Internet advertisements.

Fletcher and Slutsky (2011) develop a structure that allows for imperfect targeting in the probabilistic voting context. Their structure has two media markets where ads can be purchased. In each market, the ads are viewed by three groups of voters: those initially supporting each candidate and those indifferent between them. The partisan types are assumed to have an intensity of preference toward their preferred candidates. Ads move this intensity in favor of the candidate running the ad. Under probabilistic voting, the support each candidate receives from a type varies continuously with its post-campaign intensity.

[6] Golman and Page (2009) also offer extensions of the Colonel Blotto game that involve externalities across theaters. However, their externalities are in the payoffs, with players valuing combinations of battlefields beyond simply summing the values of each battlefield. They do not consider having the outcomes on multiple theaters depend on the same expenditures as we do here.

In this paper, we replace probabilistic voting with a Blotto-style assumption that all members of each group in a market vote for the candidate the group prefers after all ads have been shown, no matter how small the group's preference for that candidate. In addition, we apply the model to situations beyond the political context.[7] We consider a situation in which multiple battlefields are grouped together because they cannot be targeted separately, and the players allocate resources to each grouping. For convenience, we will call a grouping a theater and each component of the grouping a battlefield. We will analyze allocations across two theaters, each of which has three battlefields.[8] The two theaters overall may be of unequal importance, and the different battlefields can also vary in importance. One of the players may have an advantage on some battlefield so can win there, even though she allocates fewer resources to that battlefield than the other player. The effectiveness of allocations can differ across theaters or even across battlefields. The effectiveness function is, however, the same for both players.[9] Unlike the typical Colonel Blotto game or the model in Fletcher and Slutsky (2011), we also assume that allocations can have diminishing rather than constant marginal effectiveness.

Our results under imperfect targeting are not only consistent with the basic nature of standard Blotto, but strengthen those results. First, pure strategy equilibria are even less likely to exist with imperfect targeting. Unlike the standard results, pure strategy

[7] Imperfect targeting has been analyzed in some contexts, such as development economics (see Bibi and Duclos (2007)). However, it has not been explored in the contest literature.

[8] While it would be interesting to analyze the case where there is a continuum of battlefields, such an approach would no longer be in the spirit of a Blotto model. In a continuum model, the amount a player wins would vary continuously with the resources each side expends, while in a Blotto model, there are significant discontinuities in the amounts won as functions of the resources expended.

[9] This generalization has an important implication in the political context. Since the effectiveness function may differ across markets, it incorporates the possibility of ad price differences across markets. Thus, the decision variable could be either the number of ads played in the market or the total advertising expenditure.

equilibria may not exist even when there are only two theaters of unequal size. As follows from Theorem 1 and Lemma 2 below, in some cases it is even true that the neutral battlefield in one theater must be larger than the entire other theater for a pure strategy equilibrium to exist. Of course, then the equilibrium has all resources expended in the larger theater.

Second, consider situations where pure strategy equilibria do not exist. As we show in Theorem 2, when the two players have equal resources that lie in some range determined by the size of the advantages the players have in the non-neutral battlefields, the equilibrium joint distributions cannot have full support for both players. At least one player puts no probability weight in a neighborhood around equal allocation across the two theaters.[10]

To more fully analyze mixed strategy equilibria, we convert our two-person constant-sum games into a set of linear programming problems. We use a novel method to develop necessary conditions about the equilibrium strategies. To do this, we numerically solve more than 500,000such problems over a range of parameter values, which allows us to more fully understand what happens in the interval around equal allocation across the two theaters. The results for exactly equal resources turn out to be knife-edge. When resources differ, even by an arbitrarily small amount, it is possible for both players to have full support around equal distribution across the two theaters.

Theorem 3 and related numerical observations give some insight into the shape of the players' probability distributions in the neighborhood around equal allocations across the two theaters. We find evidence that in the neighborhood around equal allocation, the

---

[10] Note that although imperfect targeting introduces an element of having more than two theaters, an important element of just two theaters remains. Taking into account budget balance, each player has only a one-dimensional allocation decision.

two players tilt their probability weight toward different theaters. One interesting insight is that in many circumstances, there are situations of "slackness," where adding to a player's resources does not increase that player's expected payoff. This would imply that in the political context, a rent-seeking candidate would be able to extract resources from the campaign without lowering his expected vote or probability of winning.

The formal model is described in Section II. Pure strategy equilibria are considered in Section III. The results on mixed strategy equilibria are presented in Section IV, and conclusions are discussed in Section V. All proofs and discussion of the numerical analysis are contained in two Appendices.

(II) The Model

1. Specification

Consider a contest between two players A and B, each of whom divides resources between two theaters denoted m and n. Each theater has three battlefields denoted 1, 2 and 3. A pair (i, j) denotes battlefield i of theater j. Each battlefield possesses two characteristics: importance and advantage. Importance may be the size of the battlefield, but could more generally relate to its significance to the players. The importance of battlefields are denoted by θ_{ij}, i = 1, 2, 3, j = m, n. The overall importance of a theater is assumed to be the sum of the importance of the battlefields, with $M = \Sigma\, \theta_{im}$ and $N = \Sigma\, \theta_{in}$. The two players have the same evaluation of the importance of battlefields.

Advantage relates to the extent to which the characteristics of a battlefield favor one of the players over the other, and is denoted α_{ij}. Battlefield 1 in each theater favors B, while battlefield 3 favors A. For convenience, advantages are from the point of view of A, so advantages in battlefields 1 will have negative values, while those in battlefields 3 will be positive. Battlefields 2 are neutral, with neither player having an advantage before the battle begins.

The total resources available to the two players are R_A and R_B, which we assume are exogenously determined. The players can allocate their resources across the two theaters with x_m and x_n denoting the allocations of player A, and y_m and y_n denoting the allocations of player B. Thus, the two players face resource constraints

$$x_m + x_n = R_A \text{ and } y_m + y_n = R_B \tag{1}$$

An allocation to a theater goes to all the battlefields in that theater in the same amount. Resources allocated to a theater by a player will shift the advantage toward that player. The effectiveness of resources in shifting the advantage are given by the functions h_{ij} in battlefield i of theater j, where

$$h'_{ij} > 0, h''_{ij} \leq 0, h_{ij}(0) = 0 \tag{2}$$

For example, in battlefield (1, n), the post-battle advantage is $\alpha_{1n} + h_{1n}(x_n) - h_{1n}(y_n)$. Figure 1 shows the post-battle situation in the two theaters.

This specification generalizes the standard Colonel Blotto assumption that effectiveness is linear in expenditures. Here, we allow for the possibility of diminishing marginal effectiveness. In addition, specifying different functions for each battlefield allows the effectiveness of allocations to differ across battlefields. In the military context, this could be due to differences between the battlefields in aspects such as size and terrain. We continue to assume that the effectiveness functions are the same for both players and are independent of the level of the other player's allocation.

Any battlefield is won by the player with the post-battle advantage. Let I denote the post-battle advantage. If neither has an advantage, each player wins the battlefield with probability 0.5. Then the outcome for player A in any battlefield is specified as:

$$V(I) = \begin{cases} 1 & iff \quad I > 0 \\ .5 & iff \quad I = 0 \\ 0 & iff \quad I < 0 \end{cases} \tag{3}$$

The overall payoff to Player A is

$$f(x_m, y_m) = \theta_{1m} V\big(\alpha_{1m} + h_{1m}(x_m) - h_{1m}(y_m)\big) + \theta_{2m} V\big(h_{2m}(x_m) - h_{2m}(y_m)\big) + \theta_{3m} V\big(\alpha_{3m} + h_{3m}(x_m) - h_{3m}(y_m)\big) + \theta_{1n} V\big(\alpha_{1n} + h_{1n}(R_A - x_m) - h_{1m}(R_B - y_m)\big) + \theta_{2n} V\big(h_{2n}(R_A - x_m) - h_{2n}(R_B - y_m)\big) + \theta_{3n} V\big(\alpha_{3n} + h_{3n}(R_A - x_m) - h_{3n}(R_B - y_m)\big) \tag{4}$$

The players simultaneously make their allocation decisions, with A choosing x_m to maximize f and B choosing y_m to minimize f, given the resource constraints (1).

Within this general structure, we make some notational conventions and some technical assumptions to assure that the analysis is interesting. Without loss of generality, we rename the players so A's resources are at least as great as B's (i.e. $R_A \geq R_B$). Similarly, we rename the theaters so that overall, n is at least as important as m ($N \geq M$).

If the advantages in battlefields 1 and 3 were so large as to be insurmountable by any allocation of resources by the disadvantaged candidate, then the game would reduce to a single theater, and the extension would not be interesting. Thus, for all battlefields to be in play, we assume each player has sufficient resources to potentially win those battlefields in which the other player has an advantage:[11]

$-\alpha_{1j} < h_{1j}(R_A)$ and $\alpha_{3j} < h_{3j}(R_B)$, j = m, n (5)

If the resource differences were very large, one player could overwhelm the other on all battlefields. To rule this out, we assume the resource difference between the players is not very large relative to the battlefield advantages:

$-\alpha_{1j} > h_{1j}(R_A – R_B)$ and $\alpha_{3j} > h_{3j}(R_A – R_B)$, j = m, n (6)

In essence, allocating just the difference in resources available to the players to a theater is not enough to sway the outcome of any non-neutral battlefield.[12] Note that (5) and (6) with respect to α_{3j} imply that $R_A < 2R_B$.

Next, if there is an exact balancing of the importance of certain battlefields, then knife-edge equilibria may result. For example, if $\theta_{2m} = \theta_{2n}$, multiple equilibria may exist, because the gains and losses from small changes in allocations that only change the outcomes on the neutral battlefields cancel out. Such equilibria are non-generic, arising only on measure zero sets in the parameter space. The following condition rules out this type of equilibrium:

$\sum_{i=1}^{3}(t_i\theta_{in} + s_i\theta_{im}) \neq 0$ when t_i and s_i each take any values from the set (7)

(0, 1, -1, ½, -½), but are not all zero.

The expression in (7) is a general formula for the change in *f* due to changes in a player's strategy. In a battlefield, either a strategy change does not change the outcome, or it causes victory to switch from one player to the other, or it creates or breaks a tie. One

[11] Note that other battlefields could exist with very high advantages, for example, with $\alpha_{3j} > h_{3j}(R_B)$, but those battlefields will always be won by the favored player. Thus, players need not consider them when setting strategies.

[12] Note that the second condition in (6) is imposed to ensure that there are not extreme asymmetries in the two players' advantages.

implication of this assumption is that $M \neq N$, so that the two theaters cannot be exactly equal in importance. Therefore, given the naming convention above, $M < N$ must hold. Another implication is that $\theta_{2m} \neq \theta_{2n}$, so the importance of the initially neutral battlefields cannot be exactly equal.

2. Examples

Consider the following settings that fit within the context of this model.

Military campaigns

First, to extend the classic Colonel Blotto warfare model to incorporate imperfect targeting, consider a military campaign between armies A and B. A allocates x_m soldiers to theater m and x_n soldiers to theater n, while B allocates y_m and y_n soldiers across the theaters. In each theater, the armies face off across a linear battlefront, with the terrain giving advantage to one army or the other at different points along the battlefront. Battlefields 1 and 3 denote points along the front that are advantageous to B and A, respectively, while battlefields 2 are neutral, favoring neither army. For example, one location in theater m might have boulders on B's side of the front, giving that army's soldiers convenient cover. The magnitude by which B gains from this cover is represented by $-\alpha_{1m}$. The importance parameter θ_{1m} could be based on the size of this location or on its strategic value to the armies. Such points of advantage need not be geographically contiguous, so that it would be difficult for an army to send soldiers only to neutral points or to those points where that army has an advantage. Thus, each army is unable to perfectly target its troops to areas with a given advantage. The effectiveness of troops, represented by the h_{ij} functions, could differ across battlefields because of differences in such factors as ease in communications or resupply.

Advertising campaigns

A second application of the model is in advertising campaigns, either between two firms competing for sales, or two political candidates competing for votes. In the product advertising case, the firms do not compete by changing the characteristics of the products, but through advertising. Consumers differ in which product they prefer, with some being indifferent. For example, a few years ago the University of Florida switched from selling only Coke products to selling only Pepsi products on campus. Some at the University were ecstatic about the change, others (including one of the authors) were outraged, and some were completely indifferent. The Coca-Cola Company and PepsiCo do not attempt to win consumers over by changing the formulas of Coke and Pepsi, but instead have advertising wars. As in Golman and Page (2009), the product advertising case is winner-takes-all if a store has shelf space for only one firm's product. [13] Another possibility is advertising to institutional users that will sign exclusive contracts to use one firm's good, such as hospitals buying all their cleaning supplies from a single supplier, or schools buying all their milk from a single producer because of advertising.

In the political advertising example, the two candidates' platforms are fixed prior to the beginning of the campaigns. Each campaign seeks to alter voter preferences in favor of its candidate. Since no voter is likely to prefer a single candidate's position on

[13] If some customers use market ranking as a signal of quality, then a discontinuous payoff would result similar to the winner-takes-all setting. A relatively small increase in advertising might increase the firm's market ranking, then induce a large increase in sales through the quality perception effect. This is seen in Amazon, which uses a complex sales volume rating, and recommendations to customers are based on these ratings.

every issue, ads might be designed to change voter preferences in other ways. Candidates might attempt to change the saliency of different issues in the minds of the voters, or focus on non-issue factors such as character or competence.[14]

The theaters in either advertising case would be media markets m and n. In the political case, these are different media markets in the same electoral district, such as two cities in the same state for a Senate race. Market segments 1 represent consumers with a preference for product B or partisans for candidate B, while segments 3 prefer product or candidate A. Segments 2 are the neutral consumers or voters. Firms or candidates allocate advertising dollars across markets, but cannot perfectly target who will see a given ad within a market. Since all sales and votes are equally important,[15] the importance parameters θ_{ij} simply represent numbers of consumers or voters. α_{1j} and α_{3j} are the preference intensities of the partisans for B and A, respectively. The h_{ij} are the advertising effectiveness functions in the various sub-markets. Ads might have different effectiveness in different media markets because of different television viewing habits or differences in ad prices between large and small markets. Additionally, individuals with initially strong preferences toward a product or candidate might respond differently to ads than do initially neutral individuals.

Campaign promises by political candidates

[14] In the political context, maximizing the probability of winning rather than expected plurality is considered preferable, but is often not used as it is significantly more complicated to analyze. A number of papers including Aranson, Hinich and Ordeshook (1974), Ledyard (1984), Snyder (1989), Duggan (2000), and Patty (2007) have considered the relation between outcomes under the two objectives in a spatial voting context. In some, perhaps special, circumstances, they are equivalent. It should be noted that expected plurality maximization is not only more convenient but in some circumstances can be justified as more realistic. Candidates with a low probability of winning may desire to lose by as small a margin as possible.

[15] In the political setting, this would be true in a statewide or district-wide race. Obviously, all votes are not created equal in an Electoral College setting.

In the previous example, candidate competition occurs with fixed platforms. However, in order to win votes, candidates may promise to spend government resources in a way that will benefit voters in specific geographic areas, such as the choice of location for a new military base. The benefits of the project will accrue to all voters in the geographic area, not just to a candidate's partisans or neutral voters. The interpretations of markets' importance and intensities would be the same as in the political advertising case. The effectiveness functions h_{ij} would now relate to the amount of benefits a project in an area would give to a particular block of voters.[16]

Quality competition among providers of education

In the final example, two private school providers compete for students by varying the quality of education in two districts, m and n. A is a religious organization and B is a secular firm. Quality differences are achieved by A via allocations x_m and x_n, and by B via y_m and y_n. The providers wish to maximize the total number of pupils they serve because of economies of scale arising from such factors as reduced textbook prices for large orders or fixed administrative costs.

Parents differ in their preferences for secular or religious schools. These underlying preferences α_{ij} could be overcome with large enough quality differentials. Quality differentials are difficult to target toward pupils of one type and not another. 1, 3 and 2 are the parent types preferring secular or religious education and the neutrals, respectively. h_{ij}, the effectiveness of dollars spent on quality, might differ across communities because dollars buy different amounts of quality due to differences in land

[16] This is the setting analyzed by Lindbeck and Weibull (1987), but their model is based on probabilistic voting.

costs and prevailing wage rates, or due to differently sized physical plants. All of these factors would likely be correlated with community size. The effectiveness functions might also differ across parental types, as these may have different sensitivities to changes in quality. The importance parameters would generally relate to the numbers of the different types in each community but might also incorporate desires by the schools to have students with certain skills (academic or athletic).

In three of these four examples, each contestant would have his own independent budget to allocate, and there is no reason to think the two contestants would have identical budgets. In the campaign promises example, though, candidates are making promises about future allocations of the overall government budget. Thus, the two candidates have the exact same amount of resources to allocate. This difference is significant since, as shown below, the results for equal and unequal resources are different in crucial ways. These examples illustrate that both cases need to be considered.

(III) Pure Strategy Equilibria

In a Blotto game with perfect targeting, a pure strategy equilibrium generally exists when there are only two battlefields, with both contestants expending all their resources in the larger battlefield. With imperfect targeting, even though the players have only two alternatives on which to spend resources, there are really more than two battlefields and pure strategy equilibria typically do not exist. To analyze this, we begin with Lemma 1, which rules out interior pure strategy equilibria.

Lemma 1: Given parameter restrictions (5) – (7), there do not exist any pure strategy equilibria other than $x_m = y_m = 0$ or $x_n = y_n = 0$.

To specify the conditions for pure strategy equilibria, it is useful to see the graphical representation of the payoff functions (4), shown in Figure 2. Define cutoff values $\bar{x}_1$, $\bar{x}_2$, $\bar{x}_3$, and $\bar{x}_4$ by $h_{3m}(\bar{x}_1) - h_{3m}(R_B) + \alpha_{3m} = 0$, $h_{1n}(R_A - \bar{x}_2) + \alpha_{1n} = 0$, $h_{1m}(\bar{x}_3) + \alpha_{1m} = 0$, and $h_{3n}(R_A - \bar{x}_4) - h_{3n}(R_B) + \alpha_{3n} = 0$ Similarly, let $\bar{y}_1$, $\bar{y}_2$, $\bar{y}_3$, and $\bar{y}_4$ be defined by $h_{1m}(R_A) - h_{1m}(\bar{y}_1) + \alpha_{1m} = 0$, $-h_{3n}(R_B - \bar{y}_2) + \alpha_{3n} = 0$, $-h_{3m}(\bar{y}_3) + \alpha_{3m} = 0$, and $h_{1n}(R_A) - h_{1n}(R_B - \bar{y}_4) + \alpha_{1n} = 0$,

The curve connecting $\bar{x}_1$ and $\bar{y}_3$ shows all allocation pairs for A and B that cause a tie in battlefield (3, m). For allocation pairs below that curve, A wins (3, m), while B wins for allocations above the curve. Points on the curve connecting $\bar{x}_2$ and $\bar{y}_4$ yield a tie in (1, n), while A wins (1, n) above the curve, and B wins below it. Similarly, points on the curves connecting $\bar{x}_3$ and $\bar{y}_1$ and $\bar{x}_4$ and $\bar{y}_2$ indicate ties in battlefields (1, m) and (3, n), respectively. A wins (1, m) below the curve at $\bar{x}_3$ and wins (1, n) above the curve at $\bar{x}_4$. The curve through (0, 0) in the figure shows the allocations at which battlefield (2, m) is tied, while the final curve through (R_A, R_B) shows the allocations leading to a tie in (2, n). Payoffs to the two players in the spaces between any two curves are constant. That is, in such a region, a small change in allocation by either player does not change the payoffs.

The exact relation of the $\bar{x}_i$ and $\bar{y}_j$ to each other can vary, and Figure 2 only shows a representative case. It is worth noting that, following from (2), (5) and (6), $0 < \bar{x}_1, \bar{x}_2, \bar{x}_3, \bar{x}_4 < R_A$ and $0 < \bar{y}_1, \bar{y}_2, \bar{y}_3, \bar{y}_4 < R_B$ must hold.[17]

Theorem 1: Under (1) – (7) and the convention that the theaters are named so that $M < N$, the only possible pure strategy equilibrium is $x_m = y_m = 0$. This equilibrium exists if and only if $R_A = R_B$, $\theta_{2n} > \theta_{2m}$, and both of the following conditions hold:

(A) Either $\bar{y}_3 < \bar{y}_4$ and $\frac{1}{2}\theta_{2n} > \frac{1}{2}\theta_{2m} + \theta_{3m}$ or $\bar{y}_4 \leq \bar{y}_3$ and $\theta_{1n} + \frac{1}{2}\theta_{2n} > \frac{1}{2}\theta_{2m} + \theta_{3m}$

(B) Either $\bar{x}_3 < \bar{x}_4$ and $\frac{1}{2}\theta_{2n} > \theta_{1m} + \frac{1}{2}\theta_{2m}$ or $\bar{x}_4 \leq \bar{x}_3$ and $\frac{1}{2}\theta_{2n} + \theta_{3n} > \theta_{1m} + \frac{1}{2}\theta_{2m}$

From Theorem 1, when each player has sufficient resources which differ from those of the other player but not by too much, then no pure strategy equilibrium exists. When $R_A = R_B$, an equilibrium exists only under further restrictions on the parameters of the model. To gain some intuition behind this result, recognize that given the budget constraint, each player has a one-dimensional strategy space: how much is allocated to theater m. Fixing the opponent's allocation, a player's payoff is nonmonotonic over this strategy space. As a player allocates more resources toward m, battlefields in m will switch allegiance to that player, while those in n switch away at different allocation

[17] To see these, note that the left-hand side of each equality defining the $\bar{x}_i$s and $\bar{y}_i$s is a continuous and strictly monotonic function of $\bar{x}_i$ or $\bar{y}_i$, respectively. Using (2). (5), and (6), it can be easily shown that at 0, value of each function has opposite sign to its value at the corresponding upper bound. Thus, there must be a value strictly between zero and the upper bound at which the value of the function is zero.

levels. Because of this non-monotonicity, at least one player will have an improving deviation over any pair of pure strategy allocations, one for each player.

The following special case helps in understanding the conditions determining the relations between $\bar{y}_3$ and $\bar{y}_4$, and between $\bar{x}_3$ and $\bar{x}_4$. Assume that effectiveness is the same in the battlefields of each theater, with $h_{im}(z) = h_m(z)$ and $h_{in}(z) = h_n(z)$, all i and z. Further, assume that theater m is less important than n because it is smaller. This could lead allocations to be more effective there with $h_m(z) \geq h_n(z)$, all z, because if advertising prices are related to market population, a dollar spent in a smaller market would buy more ads.

Lemma 2: Assume (1) – (7), $R_A = R_B$, and $h_m(z) \geq h_n(z)$. If $\alpha_{3m} < -\alpha_{1n}$, or if $\alpha_{3m} = -\alpha_{1n}$ and $h_j'' < 0$, then $\bar{y}_3 < \bar{y}_4$. If $\alpha_{3n} > -\alpha_{1m}$, or if $\alpha_{3n} = -\alpha_{1m}$ and $h_j'' < 0$, then $\bar{x}_3 < \bar{x}_4$.

Therefore, if the more important theater n also has larger advantages ($\alpha_{3m} < -\alpha_{1n}$ and $\alpha_{3n} > -\alpha_{1m}$), then from Lemma 2 and conditions (A) and (B) of Theorem 1, $\frac{1}{2}\theta_{2n} > \frac{1}{2}\theta_{2m} + \theta_{3m}$ and $\frac{1}{2}\theta_{2n} > \theta_{1m} + \frac{1}{2}\theta_{2m}$ must hold at a pure strategy equilibrium. Together these imply that $\theta_{2n} > M$. Not only is theater n more important overall, its neutral battlefield must be more important than all of theater m. In this case, both players put all their resources in the more important theater and devote no resources to the less important one. If the advantage relations in Lemma 2 do not hold, the conditions on θ_{2n} relative to the importance of battlefields in m in Theorem 1 are still sufficient but are not necessary for the pure strategy equilibrium to exist.

Pure strategy equilibria may exist if one player's resources are significantly larger than the other with assumption (6) violated. For example, with enough resources, player A could set $x_m = x_n = \frac{1}{2}R_A$ and, no matter what B did, win all battlefields in theater m provided $h_{1m}\left(\frac{1}{2}R_A\right) - h_{1m}(R_B) + \alpha_{1m} > 0$, and all battlefields in theater n if $h_{1n}\left(\frac{1}{2}R_A\right) - h_{1n}(R_B) + \alpha_{1n} > 0$.

Although the formal results in Theorem 1 apply to a situation with only two theaters, the intuition that pure strategy equilibria are unlikely to exist carries over to more than two theaters. For a pure strategy equilibrium to exist in a setting of K theaters, the equilibrium strategies restricted to any pair must be an equilibrium in that pair holding fixed allocations to all the other theaters. Let R_A^{ij} and R_B^{ij} be the sum of resources allocated to theaters i and j by the two players. If the players have total resources which do not differ much, then, on at least one pair of theaters i and j, the resources R_A^{ij} and R_B^{ij} will satisfy assumptions (5) and (6) with an equilibrium then unlikely on the pair.

(IV) Mixed Strategy Equilibria

In most circumstances, no pure strategy equilibrium exists and the players use mixed strategies. Finding the equilibrium mixed strategies in general is difficult. Some notational conventions and definitions are useful for the analysis. Let $E(x_m)$ and $F(y_m)$ denote the cumulative distribution functions for the mixed strategies of A and B, respectively. Let Ω_A and Ω_B denote the supports of the mixed strategies of A and B, respectively, defined as sets of values at which the corresponding CDF increases. That is

$\Omega_A = \{x: \forall\varepsilon > 0, E(x-\varepsilon) < E(x)\}$ and $\Omega_B = \{y: \forall\varepsilon > 0, F(y-\varepsilon) < F(y)\}$.[18] We will derive results under the following restriction on the $\bar{x}_i$ and $\bar{y}_j$:

$$\max(\bar{x}_1, \bar{x}_2) < \min(\bar{x}_3, \bar{x}_4) \text{ and } \max(\bar{y}_1, \bar{y}_2) < \min(\bar{y}_3, \bar{y}_4) \qquad (8)$$

This condition will be satisfied when advantages are quite high. Situations in which this will be likely to hold include political contests and other contexts in which individuals have strong preferences. Given (8), the open intervals $\Gamma_A \equiv (\max(\bar{x}_1, \bar{x}_2), \min(\bar{x}_3, \bar{x}_4))$ and $\Gamma_B \equiv (\max(\bar{y}_1, \bar{y}_2), \min(\bar{y}_3, \bar{y}_4))$ are nonempty. Also, let $\Gamma \equiv \Gamma_A \cap \Gamma_B$. Note that if the resource difference between the players becomes too large, Γ will be empty.

If player i chooses an allocation in the interval Γ_i, then, no matter what its opponent does, that player cannot win either of the battlefields favorable to the opponent, but also cannot lose either battlefield in which that player is favored. Only the neutral battlefields are in play. To see this, consider any x_m such that $\bar{x}_3 > x_m > \bar{x}_2$ in Figure 2. No matter what value of y_m that B chooses, the pair of allocations is always below the curves through $\bar{x}_1$ and $\bar{x}_2$. Therefore, A always wins battlefield (3, m) but loses (1, n). On the other hand, if an allocation is chosen outside the interval, there are potential gains and losses in these battlefields. For example, if x_m is below $\bar{x}_2$, for different values of y_m, A may win or lose (1, n).

From these insights, we can derive bounds on the expected equilibrium payoffs of the players.

[18] Assume that $E(x_m)$ is constant below some z and then increases above that. If $E(x_m)$ is continuous at z then $z \notin \Omega_A$ since $E(z-\varepsilon) = E(z)$ would hold for small ε. If $E(x_m)$ is discontinuous at z with a mass point there, then $z \in \Omega_A$ since $E(z-\varepsilon) < E(z)$ for any ε. The same property holds for $f(y_m)$ and Ω_B.

Proposition 1. As long as Γ is non-empty, a player's expected equilibrium payoff is greater than the sum of the importance of the less important of the neutral battlefields and the two battlefields that favor that player, and is smaller than sum of the importance of the more important neutral battlefield and the two battlefields favoring that player.

We begin with the players having the same resources, since this in an important benchmark case. Surprisingly, however, this turns out to be a knife-edge case. We then turn to exploring situations when the candidates have different resources.

1. The players have equal resources

Let $R \equiv R_A = R_B$. We assume that dividing the resources equally across theaters belongs to the intervals Γ_A and Γ_B:

$$\frac{R}{2} \in \Gamma \tag{9}$$

Given a value of R, this assumption in effect puts lower bounds on the magnitudes of the α_{ij}. These advantages cannot be too large relative to R from assumption (5). From (8), they also cannot be too close to 0. Note that (9) in at least some circumstances strengthens (6). If the h_{ij} are linear, then (9) implies $\min\left(-\frac{\alpha_{1n}}{c_{1n}}, \frac{\alpha_{3m}}{c_{3m}}\right) > \frac{R}{2}$ and $\min\left(-\frac{\alpha_{1m}}{c_{1m}}, \frac{\alpha_{3n}}{c_{3n}}\right) > \frac{R}{2}$ for some positive constants c_{ij}. Hence, for $R_A = R_B$, (6) only implies $|\alpha_{ij}| > 0$ while (9) imposes the more restrictive $|\alpha_{ij}| > \frac{c_{ij}R}{2}$.

Theorem 2: Assume parameter restrictions (5) – (9) and $R_A = R_B$. Either $\Omega_A \cap \Gamma = \emptyset$ or $\Omega_B \cap \Gamma = \emptyset$, so that at least one player never expends resources in the open interval Γ.

Although multiple mixed strategy equilibria may exist and it is difficult analytically to find the exact equilibria, Theorem 2 gives an important characterization of any mixed strategy equilibrium in circumstances when the resources available to the two players are the same. There is no equilibrium in which both players have any probability of dividing their resources nearly equally between the two theaters. This follows since the interval Γ contains $\frac{R}{2}$, from assumption 9. The interval in which at least one player puts probability weight may be larger than Γ. Note that the Theorem does not imply that even one of the players will put probability weight in Γ. In fact, if the game is symmetric, then neither player will ever have an approximately equal division.

Corollary 1: If $-\alpha_{1j} = \alpha_{3j}$, $\theta_{1j} = \theta_{3j}$, and $h_{1j} \equiv h_{3j}$ for $j = m, n$. Then in any mixed strategy equilibrium $\Gamma \cap (\Omega_A \cup \Omega_B) = \emptyset$.

Theorem 2 and Corollary 1[19] still leave open several questions: First, does any player ever put weight in Γ, and if so, which player does so? Second, how does the size of the interval in which at most one player puts probability weight compare to Γ? To answer these, we numerically solve discrete approximations of the game. The analyses are performed by choosing reasonable values for the parameters and linear or square root

[19] Corollary 1 openly violates assumption (7). However, the proof of Theorem 2 requires only $\theta_{2m} \neq \theta_{2n}$ which is not violated by Corollary 1. Thus, relaxing (7) and keeping $\theta_{2m} \neq \theta_{2n}$ guarantees that the insight provided by Corollary 1 is correct.

functional forms for h_{ij}. Our results relate to statements about the full set of equilibria, not just a single one. Since it is difficult to find the full set of equilibria, we instead find necessary conditions on the set of equilibria. We find a single equilibrium, which tells us the players' payoffs in all equilibria, since this is a two-person constant sum game. We then perturb the game slightly by adding constraints on how much or how little probability weight a player can put on a particular strategy, and see whether these constraints change the value of the game in the new equilibrium. A full description of the methodology for numerically solving and details of the solutions can be found in the Appendix. From the numerical analyses, we derive several numerical observations.

Observation 1: The fraction of times a player puts weight in Γ increases when a player's expected equilibrium payoffs are closer to the bounds in Proposition 1.

Observation 2: In over 95% of the cases in which one player puts weight in Γ, it is the player whose expected payoff from the game is lower than the payoff yielded by $(x_m = 0, y_m = 0)$.

Observation 3: In over 90% of cases, the interval around equal allocation in which a player puts no probability weight is larger than the Γ interval.

Observation 4: In over 75% of the cases, both players have a mass point allocating all resources to the theater with the more valuable neutral battlefield. In no case do both players have a mass point allocating all resources to the theater with the less valuable neutral battlefield.

Note that once we know a mixed strategy of a player who puts weight in the Γ interval, it is possible to create an uncountable number of different mixed strategies for that player that can be a part of a Nash Equilibrium. Since the other player distributes weight only outside the Γ interval, the player in question is indifferent between any two points inside the Γ interval. Thus, the distribution inside the Γ interval can be altered in any way as long as the two following conditions are satisfied: (a) the weight placed in this interval stays the same, and (b) the expected payoff from playing in Γ for the other player remains below or at the expected payoff from the entire game. An illustration of a mixed strategy having this property is shown in Figure 3.

The significance of these results depends upon the size of the interval Γ. For $\inf \Gamma$ and $\sup \Gamma$ near their respectively upper and lower bound of $\frac{1}{2}R$, the interval around $\frac{1}{2}R$ is small. Only allocations with almost exactly equal spending in the two theaters are ruled out. For Γ near its upper bound $(0, R)$, this interval is almost the entire set of allocations. Anything except almost complete asymmetry in expenditures is ruled out. The size of this interval thus depends upon the magnitude of the initial advantages relative to the resources available to the players. The less resources players have relative to these advantages, the more asymmetric must be the allocations of a player.

Throughout the analysis, we have assumed that the resources are exogenous and uncorrelated to the advantages α_{ij}. However, there may be reasons for them to be positively or negatively correlated in some of the examples of the model. Consider the case of political campaigns. When the advantages are greater voters are less persuadable, so money will be less effective and donors may prefer to donate to other candidates in other races, yielding a negative correlation. On the other hand, in larger theaters, more money is raised but ads may be more expensive, which is similar in effect to having larger advantages, creating a positive correlation.[20] Assume that $\sup \Gamma = a + bR$ and $\inf \Gamma = R - \sup \Gamma$ where the sign of b determines whether $\sup \Gamma$ and R are positively or negatively correlated. Then the interval becomes $((1 - b)R - a, bR + a)$ with the size of the interval $2a + (2b - 1)R$. The size of this interval as a fraction of R declines in R. Only if b is near $\frac{1}{2}$ and a is near 0 will the interval be very small relative to R for all values of R. Otherwise, for at least some R, it should be non-negligible in size.

2. The players have unequal resources

An important question is whether Theorem 2 can be extended to games with small differences in resources. In order to investigate this issue, we run further numerical analyses. An example of an equilibrium mixed strategy for unequal resources is given in Figure 4.

Observation 5: When resources differ by even a small amount, both players can put weight near equal division. Thus, the result in Theorem 2 appears to be knife-edge.

Figure 2 gives some insight into this knife-edge behavior. When $R_A = R_B$, the curves through 0, 0 and R_A, R_B coincide, so the region between them disappears, but exists when R_A exceeds R_B by even the tiniest amount. When the region exists, payoffs within it are constant. Thus, when the region exists, no matter how small, it affects play because in many situations, the resource-advantaged player gains by moving into that region from nearby regions.

---

[20] See Stratmann (2005) for a survey of the recent campaign contributions and campaign spending literatures, where he discusses both of these effects. Stratmann notes that when voter preferences are strong toward one candidate or the other, contributors are less likely to donate since campaign spending would have little effect in those markets. He also discusses the fact that spending is likely to buy very different amounts of advertising in different markets.

The numerical analyses give insight into the shape of the cdfs E(x_m) and F(y_m) when both players have weight in Γ.

Observation 6: In over 95% of the cases where both players put probability weight in Γ, the weight within Γ increases toward the theater with the more valuable neutral battlefield for resource-advantaged player and toward the other theater for the resource-disadvantaged player. That is, the CDF for one player in this range tends to be approximately concave, while it is approximately convex for the opponent.

Let $\Delta \equiv R_A - R_B$. The following observation provides insights into shape of equilibrium mixed strategies.

Observation 7: There appear to be two types of patterns for the CDFs when both players put weight in Γ. In approximately 80% of these cases, the support must be disconnected in all equilibria, with mass points or separated sets smaller than Δ. In the remaining 20%, the union of the supports over all mixed strategies covers an interval greater than Δ. Some of these equilibria have connected supports over an interval of length Δ.

For the latter cases described by Observation 7, we are able to formalize some necessary conditions on the shape of equilibrium mixed strategies for both players.

Theorem 3: Assume that the open set $\Phi \subset \Gamma \cap \Omega_A \cap \Omega_B$ is a nonempty interval with length greater than Δ. Then for almost all z such that $z \in \Phi$ and $(z + \Delta) \in \Phi$,

$\frac{E(z+2\Delta)-E(z+\Delta)}{E(z+\Delta)-E(z)} = \frac{\theta_{2m}}{\theta_{2n}}$ and $\frac{F(z+\Delta)-F(z)}{F(z)-F(z-\Delta)} = \frac{\theta_{2n}}{\theta_{2m}}$.

Within the interval Φ there is a tendency for the pdf of A's mixed strategy to decrease exponentially, at least over changes of size Δ, while B's pdf increases exponentially over changes of size Δ. This pattern holds for $\theta_{2m} < \theta_{2n}$. If this inequality reverses, the pattern reverses as well. Note that the result does not assert full monotonicity. In fact, in many of the numerical analyses there is not global monotonicity. For changes smaller than Δ, A's pdf could increase and B's could decrease. An example of this can be seen in Figure 4. If, instead of a pdf, a histogram of a discrete approximation were to be graphed and the bin size for the discrete approximation were chosen to be Δ, then the histogram for A would be strictly decreasing over the interval Φ, while B's would be strictly increasing.

The unexpected difference between $E(x_m)$ and $F(y_m)$ stem from the fact that within Γ, players compete only for battlefields of importance θ_{2m} and θ_{2n}. The resource advantaged player concentrates on winning the battlefield with higher importance. On the other hand, the resource disadvantaged player concentrates on assuring that the opponent will not get both battlefields, hence is less willing to compete for the more important battlefield.

We can gain insight into our results by comparing them to those of Macdonell and Mastronardi (2014), who provide a complete characterization of mixed strategy equilibria in Blotto games consisting of exactly two battlefields, when players have unequal resources. Our model essentially reduces to theirs when the advantages in the non-neutral battlefields are either zero or become so large that a player can never win any battlefield initially favoring its opponent. Any parameter combination in our model that induces players to fight over only the neutral battlefields would be essentially equivalent to the model of Macdonell and Mastronardi. The comparison of results between our model and theirs is complicated somewhat because of different tiebreaking rules. We assume in expression (3) that if after play, neither player has an advantage in a battlefield, each has an equal probability of winning it. Macdonell and Mastronardi, on the other hand, assume that such ties are broken in favor of the player with more resources overall.

Macdonell and Mastronardi show that the support of the probability distribution of the equilibrium mixed strategies can be divided into equally spaced sets. The weights attached to the sets are exponentially increasing towards the more valuable battlefield for the player with more resources, and exponentially decreasing towards the more valuable battlefield for the player with less resources. They also derive rules about how probability should be distributed within each set.

Consider the numerical analyses discussed in Observation 7. Those with a disconnected support conform to the distributions derived by Macdonnell and Mastronardi, in the sense that for at least one player, the weight placed in Γ is distributed over equally spaced intervals in an exponentially decreasing or increasing fashion. The cases with disconnected support could differ from the Macdonnell and Mastronardi distributions for two possible reasons: we have a different tiebreaking rules, and we assume the existence of non-neutral battlefields. We run numerical analyses on their two-battlefield model with our tiebreaking rule, and these yield at least one player with a disconnected support. See Figure 5 for an example of this.[21] This strongly suggests that the existence of the non-neutral battlefields is creating the difference in results.

[21] In the two-battlefield case, the different tiebreaking rule changes which player has the disconnected support. In their model, the player with more resources will have a set of mass points and no continuous distribution. Meanwhile, the player with fewer resources can distribute probability over almost all points of the potential support. When the tie rule changes, the situation of the players reverses.

One important result implicit in Macdonell and Mastronardi is that the resource constraint in the two-battlefield Blotto game is hardly ever binding. In other words, as the resources of one player increase, the expected payoff to this player remains constant except for a countable set of points when it jumps up.[22] An important application of this "slackness" is the discussion in the public choice literature of rent-seeking candidates, who extract resources for personal benefit at the cost of lowering their expected vote or probability of winning. In the two-battlefield game described here, one candidate can, in many cases, extract resources without such a cost.

Additional numerical analyses explore whether this important property continues to hold in the presence of additional non-neutral battlefields. We conduct experiments in which the payoff function has two variants: a "regular variant" as depicted in Figure 2, and a "simple variant" in which advantages are increased so much so that the function is constant above the $x_m = y_m$ line and below the $x_m = y_m + \Delta$ line. That is, in the simple variant, battlefields 1 and 3 drop out, and the model is equivalent to that of Macdonell and Mastronardi, with our tiebreaking rule.[23] See Figure 6 for an illustration of the differences between these two variants.

Observation 8. In more than 75% of the experiments, there is an interval whose length is at least the distance between grid points over which A's expected payoff does not change with changes in A's resources.

Observation 8 suggests that slackness is very common. That is, for most possible sets of parameters there are regions in which increasing the resources for one player does

[22] Our simulations show this largely remains true for the alternative tie rule. The only difference is that the jump happens earlier by one point using Macdonell and Mastronardi's tie rule. That is, the expected payoff is a right continuous function of resources using their tie rule and left continuous using ours.

[23] Alternative ways of obtaining the simple variant are to set advantages $\alpha_{ij} = 0$ or $\theta_{1j} = \theta_{3j} = 0$.

not benefit that player. Figure 7 shows the expected payoffs of one player as a function of that player's available resources. The figure superimposes five variants which differ only in the size of the advantages. There are two simple variants. In the first, advantages are so large that the players are effectively battling only for the neutral battlefields. In the second, advantages are zero, so that players are battling for entire theaters. The three regular cases include one that satisfies assumptions (1)-(7), as well as two other cases in which advantages are made too close to zero to satisfy assumptions (1)-(7). As the Figure shows, there is a flat interval in each variant.

Finally, most other Blotto models consider only linear effectiveness functions, while we allow for nonlinearity. In particular, we consider both linear and square root *h* functions in our numerical analyses. It is interesting to compare the results under the two specifications.

Observation 9. The results in the two specifications are roughly similar, with two important differences. Having a connected support when both players have weight in Γ is less likely in the nonlinear specification, and slackness occurs 100% of the time in the square root and only about 60% of the time in the linear specification.

Observation 9 suggests that the choice of functional form for effectiveness can be significant. In fact, slackness seems to be related to the linearity of the h function. See Figure 2 for the intuition behind this idea. In the simple game, the curves connecting the cutoffs $\bar{x}_1$ to $\bar{x}_4$ and $\bar{y}_1$ to $\bar{y}_4$ in the Figure do not exist, so the players are unconcerned about what happens in the non-neutral battlefields. How concerned they are about these battlefields increases with the area between those curves and the corners. If the h functions are concave, then the curves will bow toward the corners, and there will be less area under them than in the linear case. Thus, the non-neutral battlefields will have less impact on the equilibria. When the h functions are convex, the curves bow away from the corners, making the non-neutral battlefields more important.

(V) Conclusions

Standard Colonel Blotto games assume perfect targeting of theaters. Our extension of the model adds the realistic feature that players can target battlefields only imperfectly, so that actions targeted toward one battlefield can impact others. In particular, we study a polar case where some battlefields are grouped together, and cannot be targeted separately. Each of them receives the same allocation, although they may be impacted differently by it. The battlefields in a grouping may differ in their importance and in whether one of the players has an advantage over the other in them.

We show that pure strategy equilibria will exist only in extreme circumstances and when they do, the players will choose very asymmetric allocations across groupings. We then derive a necessary condition for mixed strategy equilibria when the players have the same amount of resources to allocate, and this amount lies between upper and lower bounds that depend upon the magnitude of the advantages that one player has over the other. This condition shows that at least one player must utilize asymmetric strategies, putting no probability weight in an interval around equal allocations to the two theaters. This does not occur under probabilistic voting, even when players have equal resources, as shown in Fletcher and Slutsky (2011).

However, when resources differ, even by an arbitrarily small amount, the players can be much more symmetric in their allocations across the two theaters. Some of our results about the nature of the players' probability distributions around equal allocation may be empirically testable. For example, the resource-disadvantaged candidate tends to tilt his allocation toward the smaller neutral battlefield, while the resource-advantaged candidate tends to tilt her allocation toward the larger neutral battlefield. In addition, players can often decrease the resources allocated to a theater without decreasing their expected payoffs. This "slackness" property would imply that the function relating expenditures in political campaigns to the outcome may have flat regions.

Recently, Mastronardi and Macdonnell (2014) analyze results for a Blotto game with two completely neutral battlefields. In their simpler context, the probability distributions of the mixed strategies are placed in discrete, equally spaced sets. Our results show that when imperfectly targeted battlefields exist where one player or another has a pre-battle advantage, the support can become more continuous.

Appendix 1: Proofs

Proof of Lemma 1: Assume that $y_m \neq x_m \leq \frac{R_A}{2} < R_B$. Note that consistent with (4) and (7) we have $f(x_m, y_m) \neq f(x_m, x_m)$ and $f(x_m, y_m) \neq f(y_m, y_m)$. Moreover, if (x_m, y_m) is a Nash Equilibrium, than $f(x_m, y_m) < f(x_m, x_m)$ and $f(x_m, y_m) > f(y_m, y_m)$. Note that both at (x_m, x_m) and at (y_m, y_m) player A wins battlefield (1,m), player B wins battlefield (3,m) and there is a tie in battlefield (2,m). Player B wins battlefield (1,n) if and only if $h_{1n}(x_n) - h_{1n}(y_n) < -\alpha_{1n}$. Note however that by (2) we have $h_{1n}(x_n) - h_{1n}(y_n) \leq h_{1n}(x_n - y_n) = h_{1n}(R_A - x_m - R_B + y_m)$. Given that we are considering players playing the same strategy in theater m, by (6) we have $h_{1n}(x_n) - h_{1n}(y_n) \leq h_{1n}(R_A - R_B) < -\alpha_{1n}$. Therefore Player B wins battlefield (1,n). Player A wins battlefield (2,n) if $R_A > R_B$ or ties otherwise. Player A wins battlefield (3,n) because $h_{3n}(x_n) - h_{3n}(y_n) \geq 0$ and $\alpha_{3n} > 0$. Therefore both at (x_m, x_m) and at (y_m, y_m) players win or tie at the same battlefields. Thus $f(y_m, y_m) = f(x_m, x_m)$ which contradicts $f(x_m, x_m) > f(y_m, y_m)$.

For $y_n \neq x_n \leq \frac{R_A}{2}$, a comparison of $f(R_A - x_n, R_B - y_n)$ to $f(R_A - x_n, R_B - x_n)$ and $f(R_A - y_n, R_B - y_n)$ leads to an analogous contradiction.

If $y_m = x_m \neq 0$ or $y_n = x_n \neq 0$, then there is a tie in battlefield (2,m) or (2,n) respectively. There exists $\varepsilon > 0$ such that player A playing $x_m - \varepsilon$ and $x_m + \varepsilon$ results in resolving the tie in two opposite ways without changing the winner of the other battlefields. Since (7) ensures strict monotonicity of the sequence $f(x_m - \varepsilon, y_m)$, $f(x_m, y_m)$, and $f(x_m + \varepsilon, y_m)$, x_m is not the best response to y_m. Hence (x_m, y_m) is not a Nash Equilibrium.

The only cases not considered above are $x_m = y_m = 0$ and $x_n = y_n = 0$; all other possible combinations of x_m and y_m do not lead to a Nash Equilibrium. Q.E.D.

Proof of Theorem 1: From Lemma 1, a pure strategy equilibrium must have $x_m = y_m = 0$ or $x_n = y_n = 0$. Note that if $R_A > R_B$ neither $x_m = y_m = 0$ nor $x_n = y_n = 0$ is a Nash Equilibrium because there exists $\varepsilon > 0$ such that by playing $x_m \pm \varepsilon$ a tie on battlefield (2,m) or (2,n) is resolved in A's favor, while on other battlefields the winner does not change. That is $f(\varepsilon, 0) > f(0,0)$ or $f(R_A - \varepsilon, R_B) > f(R_A, R_B)$. Lemma 1 rules out other possibilities, hence there is no Pure Strategy Nash Equilibrium if $R_A > R_B$.

Assume $R_A = R_B = R$ and $x_n = y_n = 0$ is a Nash Equilibrium. (7) ensures that $f(0, R) < f(R, R)$ and $f(R, 0) > f(R, R)$. Note that these inequalities imply respectively that $\theta_{1n} + \theta_{1m} + \frac{1}{2}\theta_{2n} + \frac{1}{2}\theta_{2m} > \theta_{1n} + \theta_{2n} + \theta_{3n} \Rightarrow \theta_{1m} + \frac{1}{2}\theta_{2m} > \frac{1}{2}\theta_{2n} + \theta_{3n}$, and $\theta_{1m} + \theta_{2m} + \theta_{3m} > \theta_{1n} + \theta_{1m} + \frac{1}{2}\theta_{2n} + \frac{1}{2}\theta_{2m} \Rightarrow \frac{1}{2}\theta_{2m} + \theta_{3m} > \theta_{1n} + \frac{1}{2}\theta_{2n}$. After adding these inequalities, we obtain $\theta_{1m} + \theta_{2m} + \theta_{3m} > \theta_{1n} + \theta_{2n} + \theta_{3n}$ which contradicts our naming convention under which $M < N$. Hence $x_n = y_n = 0$ cannot be a Nash Equilibrium.

Consider $x_m = y_m = 0$. $x_m = 0$ is the best response to $y_m = 0$ if and only if $f(0,0) > f(x, 0)$ for any $0 < x \leq R$ (strict inequality is ensured by (7)). Note that $f(x, 0)$ is constant with respect to x within the intervals $(0, \min(\bar{x}_3, \bar{x}_4))$, $(\min(\bar{x}_3, \bar{x}_4), \max(\bar{x}_3, \bar{x}_4))$, and $(\max(\bar{x}_3, \bar{x}_4), R]$. The values of *x* for which there is a tie are not interesting because the function f takes on a higher value in their neighborhood. Thus, we need to consider only four inequalities. For $x < \min(\bar{x}_3, \bar{x}_4)$ we have (a) $\frac{1}{2}\theta_{2m} + \frac{1}{2}\theta_{2n} + \theta_{3m} + \theta_{3n} > \theta_{2m} + \theta_{3m} + \theta_{3n} \Rightarrow \theta_{2n} > \theta_{2m}$. For $x >$

$\max(\bar{x}_3, \bar{x}_4)$ we have (b) $\frac{1}{2}\theta_{2m} + \frac{1}{2}\theta_{2n} + \theta_{3m} + \theta_{3n} > \theta_{1m} + \theta_{2m} + \theta_{3m} \Rightarrow \frac{1}{2}\theta_{2n} + \theta_{3n} > \theta_{1m} + \frac{1}{2}\theta_{2m}$. If $\bar{x}_3 < \bar{x}_4$ then for $x \in (\min(\bar{x}_3, \bar{x}_4), \max(\bar{x}_3, \bar{x}_4))$ we have (c) $\frac{1}{2}\theta_{2m} + \frac{1}{2}\theta_{2n} + \theta_{3m} + \theta_{3n} > \theta_{1m} + \theta_{2m} + \theta_{3m} + \theta_{3n} \Rightarrow \frac{1}{2}\theta_{2n} > \theta_{1m} + \frac{1}{2}\theta_{2m}$. If $\bar{x}_4 < \bar{x}_3$ then for $x \in (\min(\bar{x}_3, \bar{x}_4), \max(\bar{x}_3, \bar{x}_4))$ we have (d) $\frac{1}{2}\theta_{2m} + \frac{1}{2}\theta_{2n} + \theta_{3m} + \theta_{3n} > \theta_{2m} + \theta_{3m}$ which is always true if (a) is true. Note that (c) implies (a) and (b). Therefore $x_m = 0$ can be the best response only if $\bar{x}_3 < \bar{x}_4$ and (c) or $\bar{x}_4 \leq \bar{x}_3$ and (a) and (b).

Analogously, $y_m = 0$ is the best response to $x_m = 0$ if and only if $f(0,0) < f(0, y)$ for any $0 < y \leq R$. For $y < \min(\bar{y}_3, \bar{y}_4)$ we have (e) $\frac{1}{2}\theta_{2m} + \frac{1}{2}\theta_{2n} + \theta_{3m} + \theta_{3n} < \theta_{2n} + \theta_{3m} + \theta_{3n} \Rightarrow \theta_{2m} < \theta_{2n}$. For $y > \max(\bar{y}_3, \bar{y}_4)$ we have (f) $\frac{1}{2}\theta_{2m} + \frac{1}{2}\theta_{2n} + \theta_{3m} + \theta_{3n} < \theta_{1n} + \theta_{2n} + \theta_{3n} \Rightarrow \frac{1}{2}\theta_{2m} + \theta_{3m} < \theta_{1n} + \frac{1}{2}\theta_{2n}$. If $\bar{y}_3 < \bar{y}_4$ then for $y \in (\min(\bar{y}_3, \bar{y}_4), \max(\bar{y}_3, \bar{y}_4))$ we have (g) $\frac{1}{2}\theta_{2m} + \frac{1}{2}\theta_{2n} + \theta_{3m} + \theta_{3n} < \theta_{2n} + \theta_{3n} \Rightarrow \frac{1}{2}\theta_{2m} + \theta_{3m} < \frac{1}{2}\theta_{2n}$. If $\bar{y}_4 < \bar{y}_3$ then for $y \in (\min(\bar{y}_3, \bar{y}_4), \max(\bar{y}_3, \bar{y}_4))$ we have (h) $\frac{1}{2}\theta_{2m} + \frac{1}{2}\theta_{2n} + \theta_{3m} + \theta_{3n} < \theta_{1n} + \theta_{2n} + \theta_{3m} + \theta_{3n} \Rightarrow \frac{1}{2}\theta_{2m} < \theta_{1n} + \frac{1}{2}\theta_{2n}$ which is always true if (e) is true. Note that (g) implies (e) and (h). Therefore $y_m = 0$ is the best response only if $\bar{y}_3 < \bar{y}_4$ and (g) or $\bar{y}_4 \leq \bar{y}_3$ and (e) and (f). Q.E.D.

Proof of Lemma 2: Let $R_A = R_B = R$. By definition $h_m(\bar{y}_3) = \alpha_{3m}$ and $h_n(R) - h_n(R - \bar{y}_4) = -\alpha_{1n}$, so $\alpha_{3m} < -\alpha_{1n}$ implies $h_m(\bar{y}_3) < h_n(R) - h_n(R - \bar{y}_4)$. By assumption $h_m(\bar{y}_3) \geq h_n(\bar{y}_3)$. Assume $\bar{y}_3 \geq \bar{y}_4$. Then, by concavity of h we have $h_n(\bar{y}_3) \geq h_n(\bar{y}_4) - h_n(0) \geq h_n(R) - h_n(R - \bar{y}_4) > h_m(\bar{y}_3) \geq h_n(\bar{y}_3)$ which is a contradiction. Thus $\bar{y}_3 < \bar{y}_4$. If $\alpha_{3m} = -\alpha_{1n}$ and $h_j'' < 0$, assuming that $\bar{y}_3 \geq \bar{y}_4$ leads to $h_n(\bar{y}_3) \geq h_n(\bar{y}_4) - h_n(0) > h_n(R) - h_n(R - \bar{y}_4) = h_m(\bar{y}_3) \geq h_n(\bar{y}_3)$ which is also a contradiction. Thus $\bar{y}_3 < \bar{y}_4$. An identical argument holds showing $\bar{x}_3 < \bar{x}_4$. Q.E.D.

Proof of Proposition 1: Playing within the interval Γ guarantees that a player wins the battlefields that favor him, in addition to winning at least the value of the less important neutral battlefield. Since this is a constant sum game, the lower and upper bounds in the proposition must hold as a necessary condition for equilibrium. Assume that the boundaries are attained and player i has an expected value from the game equal to the sum of the battlefields that favor her and the more important neutral battlefield, and the other player, j, receives a payoff from the game equal to the sum of the battlefields that favor him and the less important neutral battlefield. Then, every strategy in the Γ interval must give player j a payoff equal to his lower bound.

Assume that $\theta_{2n} > \theta_{2m}$. Then, player i can never expend resources in the Γ interval nor at or above $\sup \Gamma$, because this gives player j an opportunity to win battlefield θ_{2n} (the more important neutral battlefield) with a positive probability while playing in the Γ interval. Thus, player i plays strategies only at or below $\inf \Gamma$. This, however, allows player j to play strategy 0, which wins both favored battlefields, and ties or wins θ_{2n} with positive probability. Thus, it produces an expected payoff greater than the lower bound. This contradicts that the lower boundary is attained. The fact that the lower bound is never attained implies that the upper bound is never attained, since it is a zero sum game and the two bounds sum to the total value of the game. Similar reasoning holds for $\theta_{2n} > \theta_{2m}$. Q.E.D.

Proof of Theorem 2: From condition (7), $\theta_{2n} = \theta_{2m}$ is ruled out. Consider $\theta_{2n} > \theta_{2m}$ with a similar argument following if $\theta_{2n} < \theta_{2m}$. Let $I_A = \Gamma \cap \Omega_A$ and $I_B = \Gamma \cap \Omega_B$ denote the points in Γ at which A and B respectively have positive weight. Assume that both players A and B put positive weight within Γ, so that $I_A \neq \emptyset$ and $I_B \neq \emptyset$. Let $b = \inf I_B$ be the lower limit of I_B. From footnote 17 in the text, $b \in I_B$ with I_B closed from the left if $F(y_m)$ is discontinuous at b and $b \notin I_B$ with I_B open from the left if $F(y_m)$ is continuous at b. If $b \notin I_B$, then player A's expected payoff strictly decreases with $x_m > b$ in the neighborhood of b and weakly decreases as x_m increases further as long as $x_m \in \Gamma$. If $b \in I_B$, playing $x_m = b - \varepsilon$ (such that $b - \varepsilon \in \Gamma$) makes player A better off than playing any $x_m > b$. These two conditions on A's expected payoffs follow from the fact that the lower is the x_m chosen by player A, the higher are A's chances of winning θ_{2n} relative to θ_{2m}. Since A cannot put positive probability on x_m's with lower expected payoffs, each of these two conditions implies that, within Γ, A will not put positive probability at values above b. Hence, $\sup I_A \leq \inf I_B$. Similarly, for $a = \inf I_A$, it follows that player B is better off by playing values that are just below a or equal to a, if possible, as opposed to playing any $y_m > a$ for $y_m \in \Gamma$. This implies that $\sup I_B \leq \inf I_A$. These two inequalities on the infimums and supremums of I_A and I_B imply in turn that $I_A = I_B = \{c\}$ where c is a single mass point such that $c \in \Gamma$. Note, however, that if player A were to play $x_m = c$, then player B would have gained by playing a y_m smaller than c since this would increase B's chances of winning (n, 2). Hence, both of them playing c cannot be a mixed strategy equilibrium and both players putting weight in Γ is ruled out. Q.E.D.

Proof of Corollary 1: Assume there exists a mixed strategy equilibrium in which one of the players, say A, has positive probability of choosing an x_m, inside the interval Γ, so that $\Gamma \cap \Omega_A \neq \emptyset$. Denote this equilibrium strategy for A as S^A. Then, because the game is symmetric, there must exist another equilibrium with B having $\Gamma \cap \Omega_B \neq \emptyset$. Call this strategy S^B. Since this is a two-person constant sum game, if a strategy for one player is ever part of an equilibrium, then it forms an equilibrium with any strategy ever an equilibrium for the other player. Hence, the pair (S^A, S^B) would form an equilibrium with both players being in the interval Γ with positive probability, contradicting Theorem 2. Q.E.D.

Proof of Theorem 3: Choose any point z s.t. $z, z + \Delta \in \Phi$. Since Player B plays both of these strategies, they must yield the same expected value. Using Lebesgue integrals:

$$\int_0^{R_A} f(x,z)dE(x) = \int_0^{R_A} f(x,z+\Delta)dE(x)$$

$$\int_0^{R_A} [f(x,z) - f(x,z+\Delta)]dE(x) = 0$$

Note that

$$f(x,z) - f(x,z+\Delta) = \begin{cases} 0 & if \quad x < z \\ \frac{\theta_{2m}}{2} & if \quad x = z \\ \theta_{2m} & if \quad z < x < z+\Delta \\ \frac{\theta_{2m} - \theta_{2n}}{2} & if \quad x = z+\Delta \\ -\theta_{2n} & if \quad z+\Delta < x < z+2\Delta \\ -\frac{\theta_{2n}}{2} & if \quad x = z+2\Delta \\ 0 & if \quad x > z+2\Delta \end{cases}$$

By the choice of z and $z+\Delta$, only battlefields (2,m) and (2,n) come into play as x changes. Hence, if $x<z$, $f(x,z)=\theta_{3m}+\theta_{2n}+\theta_{3n}$. If $x=z$, $f(x,z)=\frac{\theta_{2m}}{2}+\theta_{3m}+\theta_{2n}+\theta_{3n}$. If $x\in(z,z+\Delta)$, $f(x,z)=\theta_{2m}+\theta_{3m}+\theta_{2n}+\theta_{3n}$. If $x=z+\Delta$, $f(x,z)=\theta_{2m}+\theta_{3m}+\frac{\theta_{2n}}{2}+\theta_{3n}$. If $x>z+\Delta$, $f(x,z)=\theta_{2m}+\theta_{3m}+\theta_{3n}$. $f(x,z+\Delta)$ is defined analogously by substituting $z+\Delta$ for z. Since $z+\Delta< R_B$ and $R_A=R_B+\Delta$ it follows that $z+2\Delta< R_A$, so the integrated function is well defined for any $z, z+\Delta\in\Phi$ by the formula above.

The set of mass points has measure zero. Therefore, for almost all z there will be no mass points at $z, z+\Delta$, and $z+2\Delta$. Hence, for almost all z

$$\int_0^{R_A}[f(x,z)-f(x,z+\Delta)]dE(x)$$

$$=\theta_{2m}[E(z+\Delta)-E(z)]-\theta_{2n}[E(z+2\Delta)-E(z+\Delta)]=0$$

$$\theta_{2m}[E(z+\Delta)-E(z)]=\theta_{2n}[E(z+2\Delta)-E(z+\Delta)]$$

$$\frac{E(z+2\Delta)-E(z+\Delta)}{E(z+\Delta)-E(z)}=\frac{\theta_{2m}}{\theta_{2n}}<1$$

The inequality for the mixed strategy of player B can be proved analogously. The only difficulty can come from the fact that in general $z-\Delta$ can be negative. The remedy is to simply assume $z-\Delta<0\Longrightarrow F(z-\Delta)=0$. Q.E.D.

Appendix 2: Details of the numerical analyses

Methodology

To augment the analytical results, we perform numerical analyses of discrete approximations of the game. The original game is characterized by a payoff function $f(x_m, y_m)$, specified in equation (4). This function is mapped into a payoff matrix $F_{Res_B \times Res_A}$, where Res_B and Res_A are the number of rows and columns in the discrete approximation of the game, and thus describe the fineness of the matrix. Two types of mappings have been used to check for robustness: (a) a "boundary mapping" for which $F_{ij} = f\left(\frac{(j-1)R_A}{Res_A-1}, \frac{(i-1)R_B}{Res_B-1}\right)$ and (b) a "middle mapping" for which $F_{ij} = f\left(\frac{(j-0.5)R_A}{Res_A}, \frac{(i-0.5)R_B}{Res_B}\right)$, $1 \leq i \leq Res_B$, $1 \leq j \leq Res_A$. Thus, the boundary mapping includes the extreme cases of all resources going to a single theater, while for the middle mapping, each theater gets at least $\frac{R_K}{Res_K}$.

Since this is a two-player zero-sum game, we solve using the mini-max theorem. To find the mixed strategy equilibrium for Player B, probabilities are assigned to each row of the matrix F so that the largest dot product of the probability vector and any column vector is minimized. Similarly, for Player A, probabilities are assigned to each column so that the smallest dot product of the probability vector and any row vector is maximized.

Multiple equilibria often exist in this game, and it would be computationally intensive to explicitly solve for all of them. Instead, the following procedure allows us to explore the structure of the full set of equilibrium mixed strategies. Since all equilibrium mixed strategies must yield the same payoff, we begin by finding a single equilibrium, which gives us the expected payoff for any equilibrium. Then, we impose additional constraints to the linear programming problem in different ways to see if the expected payoff changes under those constraints. We use this procedure to investigate several important properties of the set of equilibrium mixed strategies:

(a) Can a particular row or column be in the support? To answer this question, we impose the constraint that the probability for that row or column is no smaller than ε, which is usually set at 0.00001. If the solution under the additional constraint yields a different value of the payoff than in the unconstrained problem, then no equilibrium mixed strategy can put positive probability on that i or j. If the payoff does not change, there exists at least one equilibrium mixed strategy with weight at that i or j.

(b) Must a particular row or column be in the support? To answer this, we impose that the probability weight at i or j must be zero. If this constraint changes the value of the objective function, the answer is yes. Otherwise, there are equilibria with zero weight at that i or j.

(c) What is the maximum possible probability weight that can be placed in a given interval? We consider the intervals [0.45, 0.55], [0.4,0.6], and [0.3,0.7] for symmetric allocations, where the bounds of the intervals represent the fraction of the total resources available to the player, and [0,0.1], [0.9,1] for asymmetric allocations. For each of these intervals, the solver is run iteratively. Each iteration requires that the sum of probabilities in the interval be at least 0.01 greater than in the previous iteration. Once this condition changes the value of the objective function, the iterative process stops, and the previous attained value is the maximum possible weight in the interval.

(d) Must probability weight fall in the intervals [0.45, 0.55] and [0.4, 0.6]? To impose this, we impose zero weight in the interval. If the value of the objective function changes under this constraint, the answer is yes, and no otherwise.

We implement this process using the Gurobi 4.51 linear optimization software and a Visual C++ program. The program accepts data about the game, including the importance and advantages of battlefields, the nature of the h functions, and the amount of resources for each player. It also accepts simulation details, including whether the boundary or middle mapping is used, the dimensions of the F matrix, the value of ε, which variable to change and by how much, and the number of changes to make. Thus, with a single line of input, the program allows us to observe how equilibrium mixed strategies change as one or more parameters are changed. Input can be organized into a CSV file so that multiple experiments can be run without supervision overnight.

During this process, the program creates linear optimization problems in a text file and iteratively calls Gurobi to solve them. In cases where Gurobi fails or produces inconsistent results, this information is included in the final report. In our numerical analyses, Gurobi very rarely failed for dimensions of the F matrix not exceeding 100. For dimensions exceeding 100, failures were more common, and increased quickly with increased resolution, rendering it impractical to use dimensions of 200 or higher. Based on the information given by Gurobi, the program constructs graphs depicting sample equilibrium mixed strategies as well as payoff matrices, and assembles them into a report that includes information on parameters, Gurobi failures, and simulation results for easy comparison.

In our numerical analyses, solving a single optimization problem took an average of 0.7 seconds. Complete characterization of the equilibrium mixed strategies with typical settings required 700-1100 problems to solve. Overall, the results in this paper required solving approximately 500,000 optimization problems.

Selected results of our numerical analyses using Gurobi as the solver have been compared to results obtained with different linear optimization software, especially Excel Solver, to check for robustness. All source codes as well as all data used for experiments and reports generated by the program are available upon request.

Results

We perform three separate experiments. In each of these, 200 sets of parameters are used, with $h(x) = x$ for half and $h(x) = \sqrt{x}$ for the other half. Each of these two groups is further divided into two subgroups of 50, with one subgroup using the boundary mapping and the other using the middle mapping.

Experiment 1: Equal Resources

Numerical Observations 1 – 4 come from the same experiment. The boundary mapping has a resolution of 51x51, while the middle mapping has a resolution of 50x50. The parameters are generated according to the following rules:

$R_A = R_B = 1, \theta_{ij} \sim U[0.1,1)$, $\alpha_{ij} \sim U[0.6,1)$ for $h(x) = x$ and $\alpha_{ij} \sim U\left[\sqrt{0.6}, 1\right)$ for $h(x) = \sqrt{x}$. That is, the importance and advantage parameters are drawn from a uniform distribution over the relevant interval. The condition on advantages ensures that $\Gamma \supset (0.4,0.6)$.

To analyze circumstances when neither player puts weight in the interval Γ, we construct two variables: (A) an indicator variable equaling 1 if there are equilibrium strategies with weight in the Γ interval for the given set of parameters, and 0 otherwise, and (B) $\frac{2|v-\bar{v}|}{b-a}$, where a is the sum of the importance assigned to the battlefields favoring player A and the less valuable neutral battlefield, b is the sum of the importance assigned to the battlefields favoring player A and the more valuable neutral battlefield, $\bar{v} = \frac{a+b}{2} = f(0,0)$ and v is the expected payoff for player A. As shown in Proposition 1, a and b are lower and upper bounds, respectively, on the payoffs players can receive. (B), therefore, has values between 0 and 1, with large values when payoffs are near either the upper or lower bound. As shown in the Table below, the correlation between (A) and (B) over all experiments is 0.64. This large, positive correlation is the basis for Observation 1.

Table A1 shows some of the results of our Experiment 1 simulations. Observation 1 follows from rows 3 – 7. Observation 2 follows from rows 8 and 9, while Observation 3 is from rows 10 and 11, and Observation 4 is from rows 12 and 13. For Observation 4, we consider equilibria in which both players have significant mass points (exceeding 0.11) at the same end of the probability distribution. Note that the results for the boundary and middle mappings are very similar for the linear case. They are also similar for the square root case, with the exception of rows 2, 6, 7, 10 and 12.

Experiment 2: Unequal Resources

Observations 5 – 7 come from the same experiment. Here, parameters are generated according to the following rules: R_A is randomly drawn from the set {1.01, 1.02, 1.03, …, 1.20} and $R_B = 1$. The conditions on advantages are designed such that

$\Gamma_B \supset (0.4,0.6)$ and $\Gamma_A \supset \left(\frac{R_A}{2} - 0.1, \frac{R_A}{2} + 0.1\right)$. This requires $\bar{x}_1, \bar{x}_2 \leq \frac{R_A}{2} - 0.1, \bar{x}_3, \bar{x}_4 \geq \frac{R_A}{2} + 0.1, \bar{y}_1, \bar{y}_2 \leq 0.4, \bar{y}_3, \bar{y}_4 \geq 0.6$. This in turn translates into the following bounds on the α_{ij}.

For $h(x) = x$:

$$\alpha_{3m} \geq R_B - \frac{R_A}{2} + 0.1, -\alpha_{1n} \geq \frac{R_A}{2} + 0.1, -\alpha_{1m} \geq \frac{R_A}{2} + 0.1, \alpha_{3n} \geq R_B - \frac{R_A}{2} + 0.1, -\alpha_{1m} \geq R_A - \frac{R_B}{2} + 0.1, \alpha_{3n} \geq \frac{R_B}{2} + 0.1, \alpha_{3m} \geq \frac{R_B}{2} + 0.1, -\alpha_{1n} \geq R_A - \frac{R_B}{2} + 0.1$$

For $h(x) = \sqrt{x}$:

$$\alpha_{3m} \geq \sqrt{R_B} - \sqrt{\frac{R_A}{2} - 0.1}, -\alpha_{1n} \geq \sqrt{\frac{R_A}{2} + 0.1}, -\alpha_{1m} \geq \sqrt{\frac{R_A}{2} + 0.1}, \alpha_{3n} \geq \sqrt{R_B} - \sqrt{\frac{R_A}{2} - 0.1}, -\alpha_{1m} \geq \sqrt{R_A} - \sqrt{\frac{R_B}{2} - 0.1}, \alpha_{3n} \geq \sqrt{\frac{R_B}{2} + 0.1}, \alpha_{3m} \geq \sqrt{\frac{R_B}{2} + 0.1}, -\alpha_{1n} \geq \sqrt{R_A} - \sqrt{\frac{R_B}{2} - 0.1}.$$

Each parameter α_{ij} is drawn for each instance from a uniform distribution bounded from below using the above formulas and bounded from above with 0.99. As previously, $\theta_{ij} \sim U[0.1,1)$. Here, the resolutions are set to $100R_A \times 100$ for the middle setting and $[100R_A + 1] \times 101$ for the boundary setting.

Table A2 shows some of the results of our Experiment 2 analyses. Observation 5 follows from rows 4 – 6. Observation 6 follows from row 7, while Observation 7 is from row 8.

Experiment 3: Slackness

To show Observation 8, we must calculate the equilibrium payoffs for different resource levels. In order to speed up the calculations in this experiment, the output files include only the calculation of the value of the objective function. The parameters are the same as in Experiment 1, and each iteration increases R_A by 0.02, and the resolution for Player A by one. We perform 51 iterations, so that for each experiment we start with $R_A = 1$ and end with $R_A = 2$. We compute the expected payoffs from the game for each of the 200 sets of parameters to see how they change as we increase resources for player A. The results are summarized Table A3.

Observation 9 follows from comparison of the identity and square root columns in the tables from all three experiments.

Notes on Figures 3 -7

Figures 3 – 6 are all from Experiment 3. Figures 3 - 5 depict simulated probability density functions for sample equilibrium strategies for Players A and B, while Figure 6 depicts the payoff function. Figure 7 is constructed separately and is not based on the results obtained with Experiment 3. The parameters used for construction of the figures, as well as h functions, mappings, and dimensions of the F matrix are given in Table A4.

Figure 3

The top line below the horizontal axis indicates which parts of the [0,1] interval may be in the support, while the bottom line indicates which parts of the [0,1] interval must be in the support. Γ is the interval between the points named as $\inf \Gamma$ and $\sup \Gamma$. The graph has been truncated from the top for clarity of exposition.

Figure 4

Γ is the interval between the points named as $\inf \Gamma$ and $\sup \Gamma$, in which density is periodically decreasing for player A – that is, decreasing over wide ranges but possibly increasing locally – and periodically increasing for player B. Note that $\theta_{2m} < \theta_{2n}$; thus, the resource advantaged player (A) is going after (2,n), the more valuable neutral battlefield.

Figure 5

The line below the horizontal axis indicates which parts of the [0, R_i] interval may be in the actual support.

Figure 6

For the simple variant, all advantages are increased to 10. Light shades of gray indicate low values of the payoff function, and increases in the expected payoff are shown with increasingly dark shades of gray. The shades of gray are adjusted automatically within each picture so that the full spectrum is utilized. Thus, inter-figure comparisons using colors may be misleading.

Figure 7

This Figure does not come from Experiment 3. The variable on the horizontal axis is R_A, and the variable on the vertical axis is the expected payoff for player A. In the series denoted by “1,” the advantages listed in Table A4 are used. In the series denoted by “0” all advantages are set to 0. In the series denoted by “0.1” the α_{ij} shown in Table A4 are reduced by 90%. In the series denoted by “0.2” the α_{ij} shown in Table A4 are reduced by 80%. In the series denoted by “10” all advantages are set in magnitude to 10, so that the game involves only two neutral battlefields. In neither variant does Player A have enough resources to win all battlefields. The total value of all battlefields is approximately 4.816. However, in all variants there are intervals between 1 and 2 in which changing R_A does not affect the expected payoff from the game.

Table A1: Basic Statistics for Experiment 1

	Quantity	Boundary, identity	Middle, identity	Boundary, sqrt	Middle, sqrt	Total
1	Observations	50	50	50	50	200
2	Seemingly pure strategy equilibrium	12	10	1	15	38
3	Player A puts weight in Γ	15	15	14	12	56
4	Player B puts weight in Γ	15	16	13	15	59
5	Neither player puts weight in Γ	20	19	23	23	85
6	Neither player puts weight in Γ, MSNE	8	9	22	8	47
7	Correlation between (A) and (B)	0.58	0.59	0.57	0.81	0.64
8	Player with weight in Γ is the player winning less than (0,0)	28	30	25	27	110
9	Player with weight in Γ is the player winning more than (0,0)	2	1	2	0	5
10	Non-PSNE cases where interval that must not contain weight is larger than Γ	33	34	49	35	151
11	Non-PSNE cases where interval that must not contain weight is approximately Γ	5	6	0	0	11
12	Both players have a mass point assigning a weight to a theater with more valuable neutral battlefield	40	36	32	45	153
13	Both players have a mass point assigning a weight to a theater with less valuable neutral battlefield	0	0	0	0	0

Table A2: Basic Statistics for Experiment 2

	Quantity	Boundary, identity	Middle, identity	Boundary, sqrt	Middle, sqrt	Total
1	Observations	50	50	50	50	200
2	Solver failures	9	8	7	6	30
3	Seemingly PSNE	0	0	0	0	0
4	Both players have weight in Γ	36	34	39	41	150
5	Exactly one player has weight in Γ	3	2	4	3	12
6	Neither player has weight in Γ	2	6	0	0	8
7	Player A clearly tends to put more resources towards theater with more valuable battlefield inside Γ and Player B does the opposite	36	31	39	40	146
8	Support may be connected for both players in considerable parts of Γ	13	10	5	3	31

Table A3: Basic Statistics for Experiment 3

Quantity	Boundary, identity	Middle, identity	Boundary, sqrt	Middle, sqrt	Total
Experiments	50	50	50	50	200
Observations	2550	2550	2550	2550	10200
Solver failures	58	56	30	20	164
Observations with the same expected value as the previous observation	96	98	809	823	1826
Experiments with at least one instance of “slackness”	29	30	50	50	159

Table A4: Parameters, h functions, mappings and dimensions of the F matrix for Figures 3 - 7

	Figure 3	Figure 4	Figure 5	Figure 6	Figure 7
θ_{1m}	0.215	0.264	1	0.955	0.611
θ_{2m}	0.782	0.561	0.56	0.633	0.894
θ_{3m}	0.150	0.395	1	0.661	0.850
θ_{1n}	0.548	0.669	1	0.51	0.857
θ_{2n}	0.631	0.800	0.52	0.746	0.954
θ_{3n}	0.541	0.621	1	0.505	0.631
α_{1m}	-0.726	-0.897	-10	-0.829	-0.940
α_{1n}	-0.733	-0.839	-10	-0.946	-0.830
α_{3m}	0.748	0.667	10	0.886	0.879
α_{3n}	0.679	0.871	10	0.814	0.869
R_A	1	1.06	1.08	1.23	n/a
R_B	1	1	1	1	n/a
$h_{ij}(x)$	x	x	$\sqrt{x}$	x	$\sqrt{x}$
Mapping	middle	middle	boundary	middle	boundary
Dimensions of F matrix	50×50	106×100	109×101	123×100	$(100R_A+1)\times101$

Note: The values for θ_{ij} and α_{ij} for Figures 3, 4 ,6 and 7 are reported to the first three digits. The advantages given in this Table for Figure 7 are for the series denoted by "1" in the Figure. The values for other series in Figure 7 can be found in the Notes for Figure 7.

References

P. Aranson, M. Hinich, P. Ordeshook, Election goals and strategies: equivalent and nonequivalent candidate objectives. Am Pol Sci Rev 68:4 (1974), 135 - 152

R. Bellman, On Colonel Blotto and analogous games. Siam Rev 11 (1969), 66–68

Bibi, S. and J. Duclos, Equity and policy effectiveness with imperfect targeting. J Development Economics 83 (2007), 109 - 140

D.W. Blackett, Some Blotto games. Nav Res Log Quart 1 (1954), 55–60

D.W. Blackett, Pure strategy solutions to Blotto games. Nav Res Log Quart **5** (1958), 107–109

E. Borel, La Theorie du Jeu et les Equations Integrales a Noyan Symetrique. Comptes Rendus de l'Academic des Sciences 173 (1921), 1304-1308; English translation by L. Savage, "The Theory of Play and Integral Equations with Skew Symmetric Kernels," Econometrica 21 (1953), 97-100.

E. Borel and J. Ville, Application de la Théorie des Probabilités aux Jeux de Hasard. Gauthier-Villars, Paris 1938. [Reprinted at the end of E. Borel and A. Cheron, "Théorie mathématique du bridge à la portée de tous," Editions Jacques Gabay, Paris, 1991]

J. Duggan, Equilibrium equivalence under expected plurality and probability of winning maximization. University of Rochester working paper (2000)

D. Fletcher, S. Slutsky, Campaign Allocations under Probabilistic Voting, Public Choice, 146:3-4 (2011), 469 - 499.

L. Friedman, Game-Theory Models in the Allocation of Advertising Expenditures. Operations Research 6:5 (1958), 699 - 709

R. Golman, S.E. Page, General Blotto: Games of Strategic Allocative Mismatch. Public Choice 138 (2009), 279–299

D. Kovenock, B. Roberson, Electoral Poaching and Party Identification. J Theoretical Politics 20:3 (2008), 275–302

J. Laslier, N. Picard, Distributive Politics and Electoral Competition. J Ec Theory, Vol. 103:1 (2002), 106 - 130

J. Ledyard, The Pure Theory of Large Two-Candidate Elections. Public Choice 44 (1984), 7 - 41

A. Lindbeck, J.W. Weibull, Balanced-Budget Redistribution as the Outcome of Political Competition. Public Choice 52:3 (1987), 273 - 297

S. Macdonell, N. Mastronardi, Waging simple wars: a complete characterization of two-battlefield Blotto equilibrium. Economic Theory (2014), 1-34.

J. Patty, Generic Difference of Expected Vote Share and Probability of Victory Maximization in Simple Plurality Elections with Probabilistic Voters. Social Choice and Welfare 29(1) (2007), 149-173

B. Roberson, The Colonel Blotto Game. Ec Theory, 29:1 (2006), 1 - 24

D. Sankoff, K. Mellos, The Swing Ratio and Game Theory. Am Pol Sci Review, 66:2 (1972), 551 - 554

J. Snyder, Election goals and the allocation of campaign resources. Econometrica 57:3 (1989), 637 – 660

T. Stratmann, Some Talk: Money in Politics. A (Partial) Review of the Literature. Public Choice 124 (2005), 135 - 156

B. Szentes, R.W. Rosenthal, Beyond Chopsticks: Symmetric Equilibria in Majority Auction Games. Games and Ec Beh 45:2 (2003), 278 - 295

J.W. Tukey, A problem of strategy. Econometrica 1**7** (1949), 73

J. Weinstein, Two Notes on the Blotto Game. BE Journal of Theoretical Economics (Contributions) 12:1 (2012), Article 7

Figure 1: Post-campaign advantages in each theater

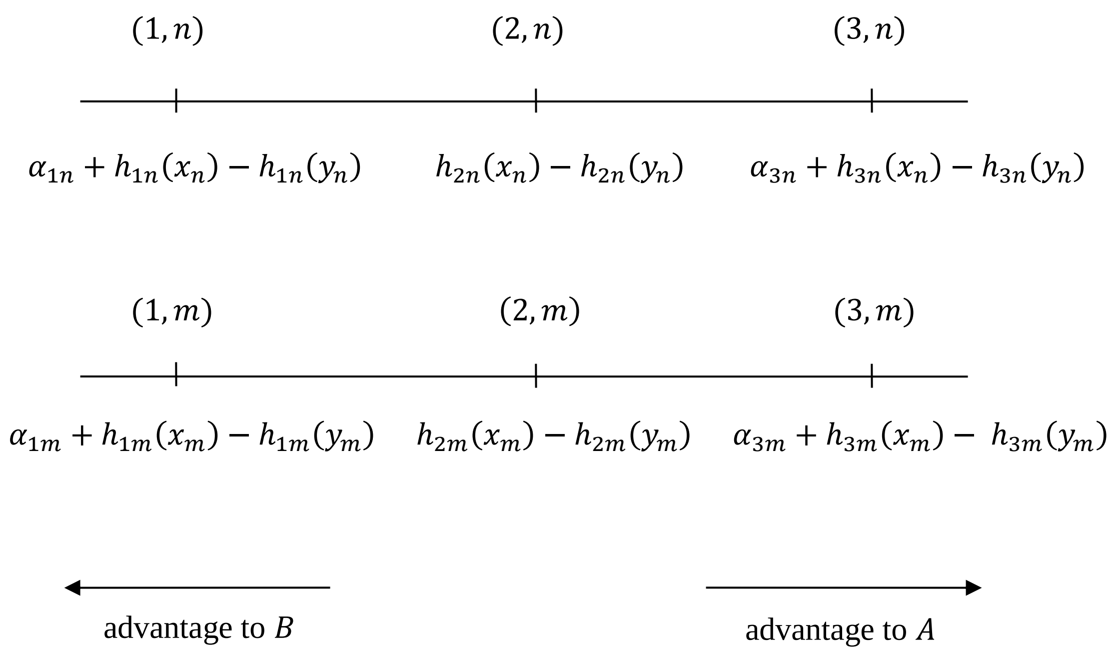

Figure 2: Payoff function

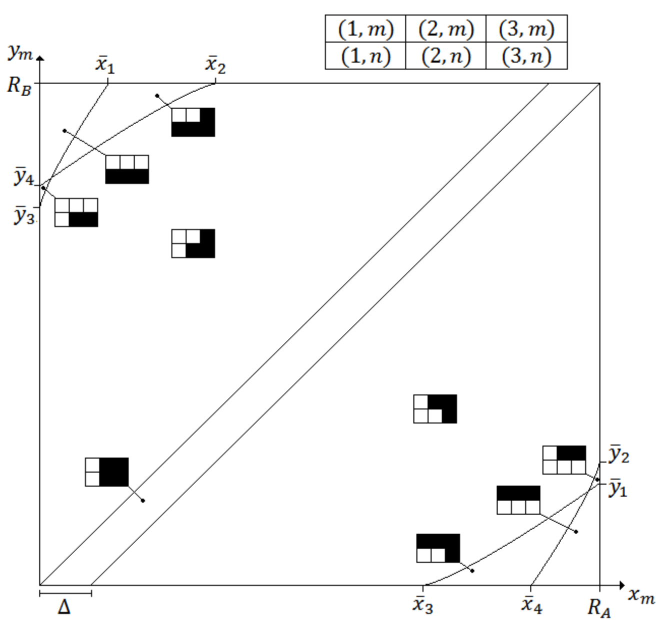

Note: schematic representation of the $f(x_m, y_m)$ function. Flat areas are separated by lines at which there are ties. Squares filled with black indicate which battlefields are won by player A according to the table at the top of the picture.

Figure 3. Example mixed strategies for the case with equal resources.

Player A:

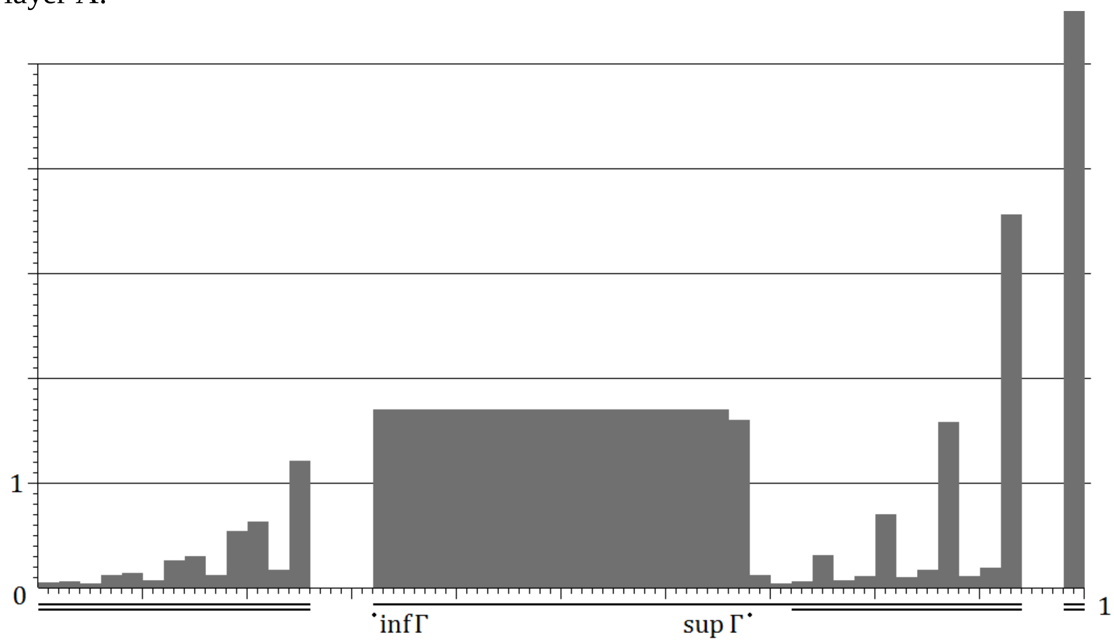

Player B:

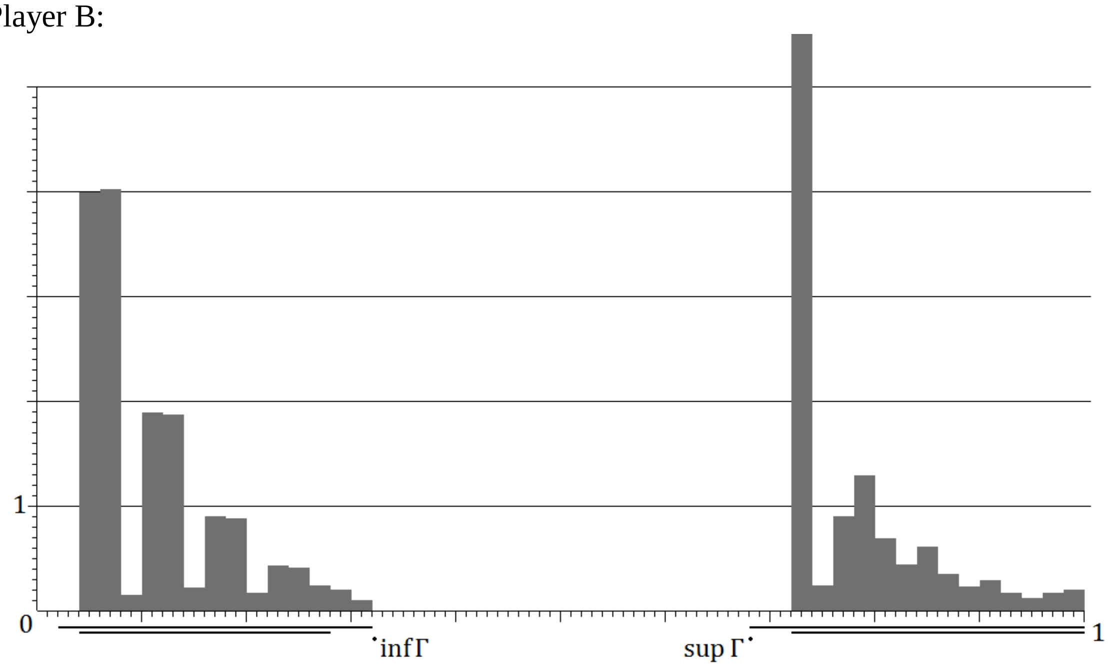

See Appendix 2 for a description of parameters and mappings used to create this Figure.

Figure 4. Example mixed equilibrium strategies when resources are unequal.

Player A:

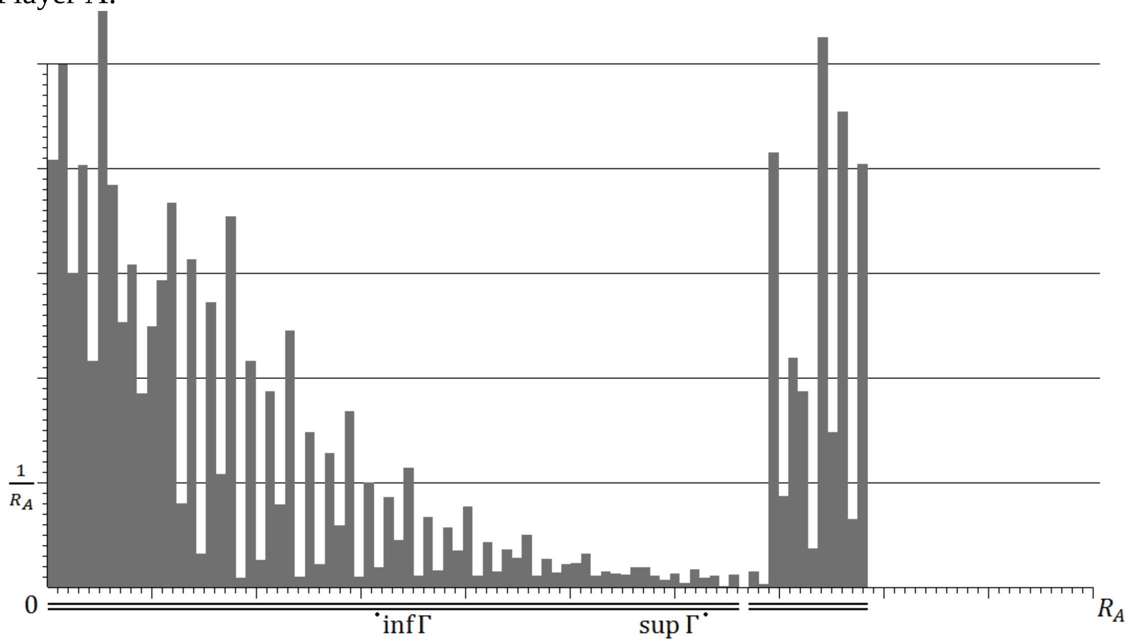

Player B:

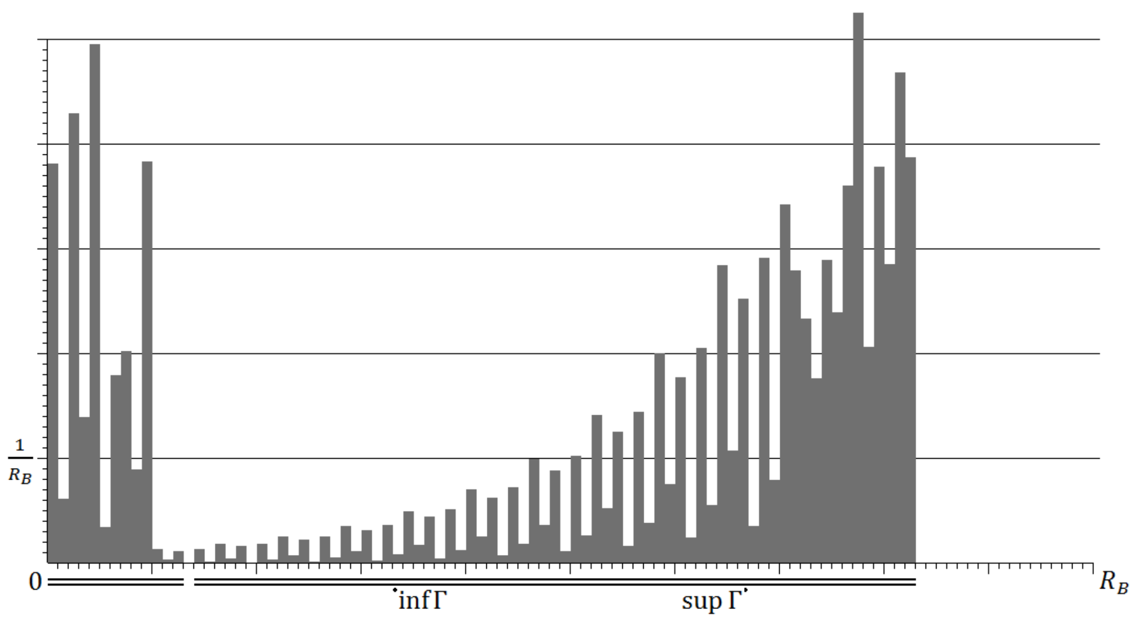

See Appendix 2 for a description of parameters and mappings used to create this Figure.

Figure 5. Example mixed strategy for simple Blotto game with unequal resources

Player A:

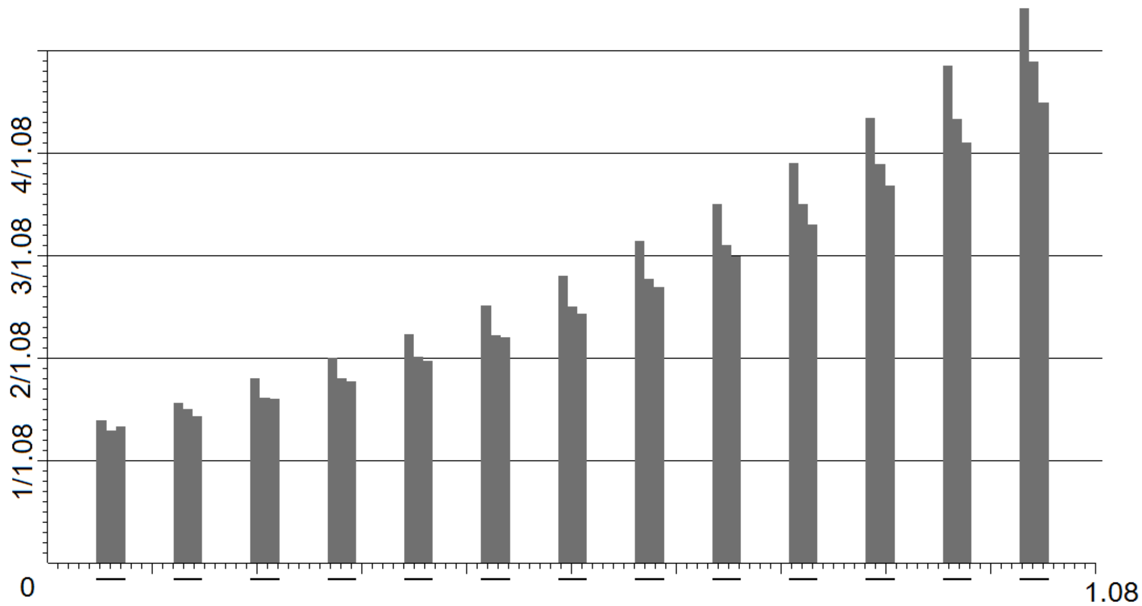

Player B:

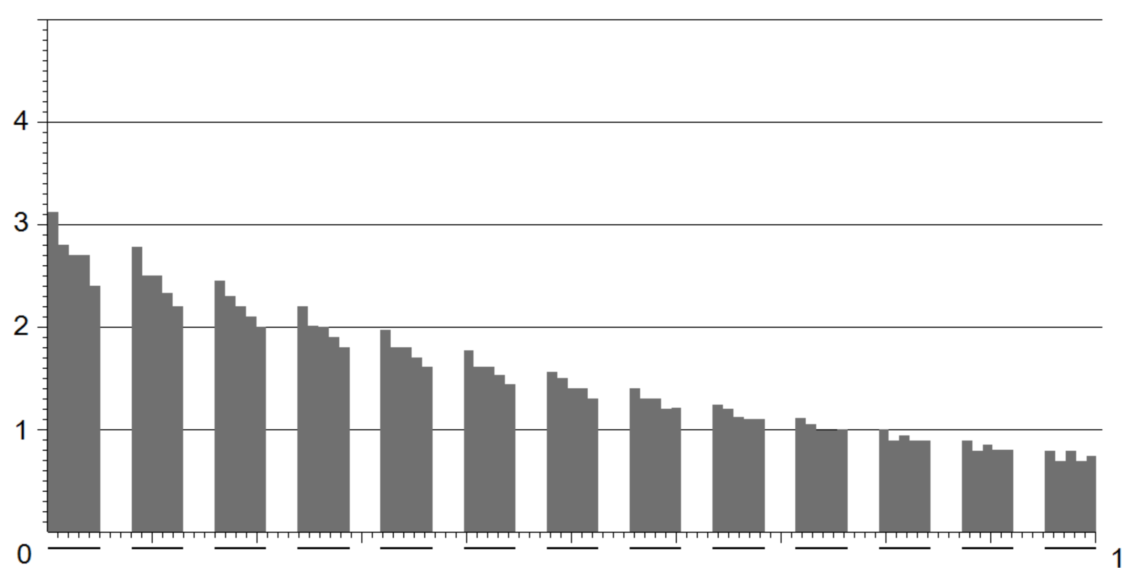

See Appendix 2 for a description of parameters and mappings used to create this Figure.

Figure 6: Sample payoff functions for the “simple variant” and the “regular variant”

Simple variant:

Regular variant:

See Appendix 2 for a description of parameters and mappings used to create this Figure.

Figure 7: Change in expected payoff with respect to difference in resources

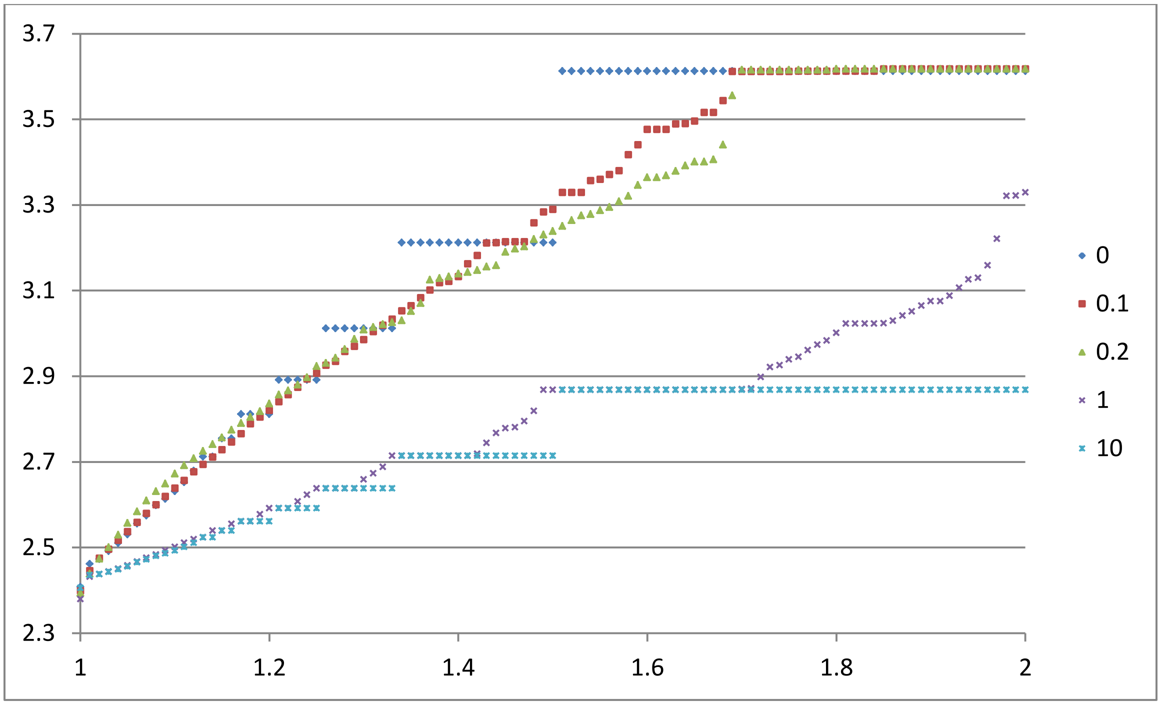

See Appendix 2 for a description of parameters and mappings used to create this Figure.