Sequential, Nonzero-sum Blotto: Allocating Defensive Resources Prior to Attack*

Robert Powell

April 2008

Travers Department of Political Science
210 Barrows Hall, 1950
University of California
Berkeley, CA 94720-1950
RPowell@Berkeley.edu

* For helpful comments, criticisms and discussion, I thank Ernesto Dal Bó, Santiago Oliveros, Larry Samuelson, Jacob Shapiro, and Leo Simon.

Sequential, Nonzero-sum Blotto:
Allocating Defensive Resources Prior to Attack

Abstract

The strategic allocation of resources across multiple fronts has long been studied in the context of Blotto games in which two players simultaneously select their allocations. But many allocation problems are sequential. For example, a state trying to defend against a terrorist attack generally allocates some or all of its resources before the attacker decides where to strike. This paper studies the allocation problem confronting a defender who must decide how to distribute limited resources across multiple sites before an attacker chooses where to strike. Unlike many Blotto games which only have very complicated mixed-strategy equilibria, the sequential, nonzero-sum "Blotto" game always has a pure-strategy subgame perfect equilibrium, the defender always plays the same pure strategy in any equilibrium, and the attacker's equilibrium response is generically unique and entails no mixing. The defender minmaxes the attacker, and the attacker strikes the site among its best replies that minimizes the defender's expected losses.

Sequential, Nonzero-sum Blotto: Allocating Defensive Resources Prior to Attack

Allocating resources against a strategic adversary is an old problem in game theory and has long been studied in the context of Colonel Blotto games (e.g., Borel 1921, Tukey 1949, Blackett 1958, Shubik and Weber 1981). Blotto games are two-person, zero-sum games in which the players simultaneously decide how to allocate limited resources across N independent battlefields. How well a player does on a particular battlefield depends on how much each player allocates to that front, and each player's payoff in the game is the sum of its battlefield payoffs.

However, some allocation problems are sequential. Following the attacks of 9/11, for example, the United States undertook a massive effort to protect its critical infrastructure and key assets. More generally, defenders often have to allocate some or all of their resources before an attacker decides where to strike.

In the model below, a defender chooses an allocation prior to an attack. After observing the defender's allocation, the attacker, e.g., a terrorist group, decides where to strike. (With additional assumptions the model can also be interpreted as one in which the attacker decides how to allocate a fixed amount of resources or effort.) The more the defender allocates to a site, the "harder" that site becomes and the less likely an attack is to succeed. The defender and attacker may value the sites differently so that the game will generally not be zero-sum.[1]

This sequential, nonzero-sum "Blotto" game has a generically unique equilibrium path and there is no mixing along it. Even though the game is not zero-sum, the defender's equilibrium strategy is to minmax the attacker. The attacker, by contrast, acts in the way that is most favorable to the defender; it strikes the site among its best responses that minimizes the defender's loss.

The existence of a pure-strategy equilibrium path, much less the fact that it is the generically unique path, contrasts with the emphasis of recent work on Blotto games (e.g., Golman and Page 2006, Roberson 2006, Hart 2007). But this contrast is more

---

[1] The Minimax Theorem trivially implies that that the equilibrium paths of the static and dynamic games are identical when the game is zero-sum.

apparent than real for two reasons. First, mixing in Blotto games often arises from an additional assumption on the payoff functions, namely that a player is sure to win a battle if it allocates more to that battlefield than its adversary does. A natural motivation for this specification is that the battlefields are legislative districts, the player's are political parties, a party's vote share in a given district is the ratio of its allocation to that of its adversary, and, crucially, a party wins the district with 50-percent-plus-one of the votes (see, for example, Laslier and Picard 2002 and the discussion of Blotto games in Dekel, Jackson, and Wolinsky 2008). The discontinuity in the payoff functions ensures that these games have no pure-strategy equilibria.

It seems more natural in other substantive settings, such as the allocation of military forces, to assume that the outcome of a fight on a battlefield is stochastic. The more a player dedicates to a particular front, the higher the probability of prevailing on that front (e.g., Shubik and Weber 1981). Blotto games with continuous payoff functions may have pure-strategy equilibria (Blackett 1948, Shubik and Weber 1978, Coughlin 1992). Indeed, the defender's unique equilibrium strategy in the static version of the game studied here is pure although the attacker does mix (Powell 2007). In particular, the defender minmaxes the attacker (as it does in the sequential game) even though the game is nonzero-sum.

The second reason that the contrast is more apparent than real is that the sequential, non-zerosum Blotto game is a game of perfect information albeit it one with an infinite action space. Harris (1985) demonstrates that pure-strategy equilibria exist in a class of perfect information, infinite games which includes the model studied here. Offering a much simpler proof of existence, this paper also shows that mixing plays a negligible role in the sequential allocation game and that the generically unique equilibrium path is very simple.

The next section describes the model. The subsequent section constructs an equilibrium in which the defender minmaxes the attacker and then shows that the defender must minmax the attacker in any equilibrium. The final section characterizes the attacker's equilibrium strategy.

A Model

A defender is trying to protect N sites and must decide how to distribute R among them. The more the defender spends on a particular site, the less likely an attack on that site is to succeed. After observing the defender's allocation, an attacker decides which site to strike.

A strategy for the defender in this game is simply an allocation vector $r = (r_1, ..., r_N)$ where $r_j \geq 0$ for all j and $r_1 + \cdots + r_N \leq R$. Let Δ denote the set of all such allocations. A pure strategy for the attacker strategy specifies which site it will strike after observing any $r \in \Delta$. A mixed strategy is a function $\alpha : \Delta \rightarrow [0, 1]^N$ where $\alpha_j(r)$ is the probability that the attacker hits j after seeing allocation $r \in \Delta$ and $\alpha_1(r) + \cdots + \alpha_N(r) = 1$.

As for payoffs, suppose that the defender suffers a loss $\lambda_j > 0$ and the attacker gains $\gamma_j > 0$ if site j is successfully attacked. The probability that an attack on site j succeeds is $v_j(r_j)$ with v_j continuous and strictly decreasing as long as this site is imperfectly defended (i.e., as long as $v_j(r_j) > 0$). Then the defender's loss and the attacker's gain are $L(r, \alpha) \equiv \sum_{j=1}^{N} \lambda_j v_j(r_j)\alpha_j(r_j)$ and $G(r, \alpha) \equiv \sum_{j=1}^{N} \gamma_j v_j(r_j)\alpha_j(r_j)$ respectively.

Finally resources are assumed to be scarce in the sense that the defender cannot perfectly defend all of the sites at once. However the defender allocates R, one or more sites will be imperfectly defended. Were this not the case, the allocation problem would be trivial. The defender would defend everything perfectly and thereby hold its losses down to zero.

ASSUMPTION 1 (SCARCE RESOURCES): *Let r be any allocation of R, then $v_j(r_j) > 0$ for at least one j.*

Subject to an important qualification, the probabilities of attack α_j can also be interpreted as a resource-allocation decision. Suppose that rather than choosing which site to strike, the attacker has resources or total effort E to allocate. In this effort game, the attacker's strategy is a function $e : \Delta \rightarrow \mathbb{R}_+^N$ such that $e_j(r)$ is the effort or resources the attacker invests in striking site j and $e_1(r) + \cdots + e_N(r) \leq E$. The probability that an attack on site j succeeds is $V_j(r_j, e_j)$. V_j is assumed to be continuous as well as decreasing in r_j and increasing in e_j as long as $0 < V_j(r_j, e_j) < 1$. The defender's loss and the attacker's gain are $L_E(r,e) \equiv \sum_{j=1}^{N} \lambda_j V_j(r_j, e_j)$ and $G_E(r,e) \equiv \sum_{j=1}^{N} \gamma_j V_j(r_j, e_j)$.

It is straightforward to show that the game is zero-sum if and only if $\lambda_j/\lambda_k = \gamma_j/\gamma_k$ for all j and k. If in addition to being zero-sum, the defender and attacker also choose their allocations simultaneously, the model would be a standard Blotto game.

A key assumption makes the two sequential allocation models equivalent. Suppose that the marginal return on effort at a site is independent of the level of effort devoted to attacking that site. That is, $V_j(r_j, e_j) = v_j(r_j)e_j$ for some continuous and decreasing function v_j.

Thinking of α_j as the fraction of E dedicated to site j, i.e., $e_j(r) = \alpha_j(r)E$, it is easy to see that $(r, \alpha(r))$ is a subgame perfect equilibrium of the game in which the defender chooses where to attack if and only if $(r, e(r))$ is subgame perfect in the game where the attacker makes an allocation decision. This follows immediately from recognizing that if the attacker has to choose where to strike, it goes after the site or sites offering the highest expected payoff. If the attacker allocates effort, it invests in hitting the site or sites offering the highest marginal return. But these are exactly the same set of sites when the marginal return is independent of the effort. When an attacker chooses where to strike, its expected return to hitting j is $\gamma_j v_j(r_j)$. When the attacker allocates effort, the marginal return on investing effort in going after j is $\partial G_E/\partial e_j = \gamma_j v_j(r_j)$. As a result, any mixture over the set of sites offering the attacker its highest payoff can also be seen as an allocation of E across the sites offering the highest expected marginal return.[2]

The remainder of this paper focuses on the interpretation in which the attacker chooses which site to hit.

The Defender's Equilibrium Allocation

A very simple intuition leads to the optimal allocation. Suppose that the defender has not allocated any resources ($r = 0$). The attacker's expected payoff to hitting j is $\gamma_j v_j(0)$. Suppose that attacking k offers the highest expected payoff at $r = 0$, i.e.,

---

[2] The assumption of scarce resources also implies that there will always be a positive return to effort, so the attacker will fully allocate E.

$\gamma_k v_k(0) = \max\{\gamma_j v_j(0) : j = 1, ..., N\}$.[3] Clearly, the defender should begin by allocating resources to k. But the more the defender spends on k, the harder and less attractive that site becomes. Eventually k is no more attractive than what was initially (i.e., at allocation $r = 0$) the second most attractive target. Call this site k'. If the defender still has resources to spend on defense, it must now allocate them to both k and k' so as to keep the attacker's payoffs to going after these sites equal.[4]

As the defender spends more on k and k', they become less attractive targets and eventually are no more attractive than what was initially the third most attractive target, say, k''. At this point the defender must allocate resources to these three sites so as to keep the attacker's expected payoffs to hitting these sites equal to each other. And, of course, the more the defender devotes to these three sites, the less attractive they become. Eventually they are no more attractive than what was originally the fourth most attractive site, and so on. The defender continues to allocate its resources in this way, having to spread them across more and more sites, until if fully allocates R. Allocating resources in this way minimizes the attacker's maximum payoff. That is the optimal allocation r^* solves $\min_{r \in \Delta}\{\max\{\gamma_1 v_1(r_1), ..., \gamma_N v_N(r_N)\}\}$.

This section proves that the minmax allocation r^* is unique and that the defender plays r^* in all subgame perfect equilibria. As for the attacker, striking the site among the best replies to r^* which minimizes the defender's loss is always an equilibrium response to r^* and it is the generically unique response to r^*.

To ease the notational burden, let $M_A(r) \equiv \max\{\gamma_1 v_1(r_1), ..., \gamma_N v_N(r_N)\}$. Then r^* is a minmax allocation if and only if $r^* \in \arg\min\{M_A(r) : r \in \Delta\}$. Lemma 1 shows that there is only one minmax allocation.

LEMMA 1: *The minmax allocation r^* is unique.*

Proof: At least one minmax allocation is sure to exist because the $v_j(r_j)$ are continuous in r and the set of possible allocations is compact. Arguing by contradiction to show

---

[3] In other words, attacking k is a best reply to allocation $r = 0$. To ease the intuitive exposition, assume this is the only best reply.

[4] Suppose the payoffs are unequal with $G_{k'} v_{k'}(r_{k'}) > G_k v_k(r_k)$ for $r_{k'} > 0$, $r_k > 0$, and $r_j = 0$ for $j \neq k, k'$. Then defender could have reduced its expected loss $L_{k'} v_{k'}(r_{k'})$ by allocating less to k and more to k'.

that only one minmax allocation exists, assume that $r^* \neq r'$ both minmax the attacker. Then $M_A(r') = M_A(r^*) > 0$ where the inequality follows from the fact that resources are scarce.

That $r^* \neq r'$ implies $r_j^* \neq r_j'$ for at least one j. Without loss of generality suppose $r_j^* < r_j'$. If $v_j(r_j^*) > 0$, then v_j is imperfectly defended and greater than $v_j(r_j')$. This leaves $M_A(r') = M_A(r^*) \geq \gamma_j v_j(r_j^*) > \gamma_j v_j(r_j')$. Continuity now ensures that there is an $\varepsilon > 0$ such that $M_A(r') > \gamma_j v_j(r_j' - \varepsilon)$. This ε of resources can be distributed across the sites $k \neq j$ to form the allocation $\widehat{r}$ where $\widehat{r}_j = r_j' - \varepsilon$ and $\widehat{r}_k = r_k' + \varepsilon/(N-1)$ for all $k \neq j$. But this yields the contradiction that $M_A(\widehat{r}) < M_A(r')$.

If $v_j(r_j^*) = 0$, then there exists an $\varepsilon > 0$ such that $r_j^* > r_j' - \varepsilon$ which implies $v_j(r_j^*) = v_j(r_j' - \varepsilon) = 0$. Once again ε resources can be distributed across the sites $k \neq j$ to form the allocation $\widehat{r}$ where $\widehat{r}_j = r_j' - \varepsilon$ and $\widehat{r}_k = r_k' + \varepsilon/(N-1)$ for all $k \neq j$. This too yields the contradiction that $M_A(\widehat{r}) < M_A(r')$. ■

Now define $\text{br}_A(r)$ to be the set of pure best replies to r: $\text{br}_A(r) \equiv \{j : \gamma_j v_j(r_j) = M_A(r)\}$. In a subgame perfect equilibrium, the attacker has to strike sites in $\text{br}_A(r)$ after observing r. It will be useful to establish three related facts which follow from the intuition that minmaxing the attacker entails initially allocating resources to the most attractive target and then spreading them out over more sites as those sites become best replies for the attacker. Allocating resources in this way means, first, that the defender only invests in protecting sites the attacker might actually hit, i.e., $\sum_{j \in \text{br}_A(r^*)} r_j^* = R$. Second, if r' differs from the minmax allocation, then the set of best replies to r' is contained in the set of best replies to the minmax allocation r^*. Consequently, the amount of resources spent on the best replies to r' is less than R. Formally,

Lemma 2: *Let r^* be the unique minmax allocation and suppose $r' \neq r^*$. Then:*
(i) $r_j^* = 0$ *for all* $j \notin \text{br}_A(r^*)$ *and* $\sum_{j \in \text{br}_A(r^*)} r_j^* = R$*;*
(ii) $\text{br}_A(r') \subseteq \text{br}_A(r^*)$*;*
(iii) $\sum_{j \in \text{br}_A(r')} r_j' < R$.

Proof: To see that (i) holds, suppose that $r_j^* > 0$ for some $j \notin \text{br}_A(r^*)$. Now redistribute $\varepsilon > 0$ from r_j^* across all of the sites in $\text{br}_A(r^*)$. That is, define $\widehat{r}$ as $\widehat{r}_k = r_k^* + \varepsilon / \|\text{br}_A(r^*)\|$ for all $k \in \text{br}_A(r^*)$, $\widehat{r}_k = r_k^*$ for all $k \notin \text{br}_A(r^*)$ and $k \neq j$, and $\widehat{r}_j = r_j^* - \varepsilon$. Since

$M_A(r^*) = \gamma_k v_k(r_k^*) > \gamma_j v_j(r_j^*)$ for all $k \in \text{br}_A(r^*)$, continuity ensures we can take ε sufficiently small that $\gamma_k v_k(r_k^* + \varepsilon / \|\text{br}_A(r^*)\|) > \gamma_j v_j(r_j^* - \varepsilon)$ for all $k \in \text{br}_A(r^*)$. This leaves $M_A(\widehat{r}) < M_A(r^*)$ and contradicts the assumption that r^* minimizes $M_A(r)$.

Turning to (ii), let $j \notin \text{br}_A(r^*)$. Then (i) implies $r_j^* = 0$. Now observe that $\gamma_j v_j(r_j') \leq \gamma_j v_j(0) = \gamma_j v_j(r_j^*) < M_A(r^*) < M_A(r')$ where the last inequality holds because r^* is the unique minmax allocation. But $\gamma_j v_j(r_j') < M_A(r')$ means $j \notin \text{br}_A(r')$. Hence, $j \notin \text{br}_A(r^*) \Rightarrow j \notin \text{br}_A(r')$. The is equivalent to $j \in \text{br}_A(r') \Rightarrow j \in \text{br}_A(r^*)$ or, alternatively, $\text{br}_A(r') \subseteq \text{br}_A(r^*)$.

As for (iii), $\gamma_j v_j(r_j') = M_A(r') > M_A(r^*)$ for all $j \in \text{br}_A(r')$ where the strict inequality follows from the fact that r^* is the unique minmax allocation. Further, $M_A(r^*) = \gamma_j v_j(r_j^*)$ for all $j \in \text{br}_A(r')$ as $\text{br}_A(r') \subseteq \text{br}_A(r^*)$. This gives $\gamma_j v_j(r_j') > \gamma_j v_j(r_j^*)$ or $r_j' < r_j^*$ for all $j \in \text{br}_A(r')$. Summing over $\text{br}_A(r')$ yields $\sum_{j \in \text{br}_A(r')} r_j' < \sum_{j \in \text{br}_A(r')} r_j^* \leq R$. ■

Turning to the equilibria of the game, a subgame perfect equilibrium of is a strategy profile $(\widehat{r}, \widehat{\alpha}(r))$ satisfying two conditions: (i) $\widehat{r}$ minimizes the defender's loss given that the attacker plays according to $\widehat{\alpha}(r)$, and (ii) playing according to $\widehat{\alpha}(r)$ maximizes the attacker's expected payoff at every $r \in \Delta$, i.e., $\widehat{\alpha}_j(r) > 0$ only if $j \in \text{br}_A(r)$. It follows that:

PROPOSITION 1: *Subgame perfect equilibria exist, and the defender plays the unique minmax allocation r^* in all of them.*

We prove the proposition by establishing two lemmas. The first ensures existence by constructing a pure-strategy, subgame perfect equilibrium in which the defender's allocation is r^*. Strictly speaking, existence is not an issue as Harris (1985) demonstrates. But the proof of Lemma 1 is much simpler than his analysis and provides some intuition for subsequent results. The second lemma shows that if the defender's allocation differs from r^*, then it cannot be part of any subgame perfect equilibrium.

The key to constructing the equilibrium is recognizing that if $\widehat{r}$ is an equilibrium allocation then even though the attacker is indifferent to hitting any of the sites in $\text{br}_A(\widehat{r})$, the defender may not be indifferent when the game is nonzero-sum. Indeed, $\lambda_i v_i(\widehat{r}_i)$ will generally not equal $\lambda_j v_j(\widehat{r}_j)$ for $i, j \in \text{br}_A(\widehat{r})$ in a nonzero-sum game. Let k be the site in $\text{br}_A(\widehat{r})$ the defender prefers the attacker to hit, i.e., $\lambda_k v_k(\widehat{r}_k) = \min\{\lambda_j v_j(\widehat{r}_j) : j \in \text{br}_A(\widehat{r})\}$.

(To ease the exposition assume for the moment this site is unique.) Then save for some nongeneric circumstances, the attacker must strike k with probability one in response to $\widehat{r}$.

To see why, suppose the contrary, i.e., the attacker hits $j \neq k$ where $j, k \in \mathrm{br}_A(\widehat{r})$ and $\lambda_k v_k(\widehat{r}_k) < \lambda_j v_j(\widehat{r}_j)$. Then the defender could deviate from $\widehat{r}$ by spending slightly less on site k, say $\widehat{r}_k - \varepsilon$. This makes k a more attractive target than any other site in $\mathrm{br}_A(\widehat{r})$. Indeed, the attacker now strictly prefers going after k to striking any other site. This leaves the defender with a loss of $\lambda_k v_k(\widehat{r}_k - \varepsilon)$. Taking ε small enough ensures that $\lambda_k v_k(\widehat{r}_k - \varepsilon) < \lambda_j v_j(\widehat{r}_j)$ and therefore that this is a profitable deviation. This, however, is a contradiction as the defender cannot profitably deviate from an equilibrium allocation, and this contradiction means that the attacker must strike k in equilibrium.

Formalizing this requires some additional notation. Let $\mu_D(r)$ be the defender's minimum loss given allocation r if the attacker strikes a site in $\mathrm{br}_A(r)$: $m_D(r) \equiv \min\{\lambda_j v_j(r_j) : j \in \mathrm{br}_A(r)\}$. Now take $\mu(r)$ to be the sites in $\mathrm{br}_A(r)$ at which the defender's losses are lowest: $\mu(r) \equiv \{j \in \mathrm{br}_A(r) : \lambda_j v_j(r_j) = m_D(r)\}$.

Finally, define the attacker's strategy $\alpha^*(r)$ to be $\alpha^*_k(r) = 1$ if $k = \min \mu(r)$ and $\alpha^*_k(r) = 0$ otherwise. In words, the attacker plays a best response to r, i.e., hits a site in $\mathrm{br}_A(r)$. If $\mathrm{br}_A(r)$ contains two or more sites, the attacker strikes the site in $\mathrm{br}_A(r)$ that minimizes the defender's expected loss, i.e., $j \in \mu(r)$. If two or more sites in $\mathrm{br}_A(r)$ minimize the defender's expected loss, i.e., if $\mu(r)$ has two or more elements, the attacker goes after the site with the smallest index among the sites in $\mu(r)$. Note that $\alpha^*(r)$ entails no mixing and is a pure strategy.[5]

It is now straightforward to show that $(r^*, \alpha^*(r))$ is a subgame perfect.

LEMMA 3: *The profile* $(r^*, \alpha^*(r))$ *is a pure-strategy subgame perfect equilibrium.*

Proof: Clearly $\alpha^*(r)$ maximizes the attacker's payoff following any r as the attacker strikes an element of $\mathrm{br}_A(r)$ and therefore maximizes its payoff after r. Consequently, it suffices to show that r^* is a best reply to α^*.

---

[5] For the purposes of constructing an equilibrium, α^* could be defined as any mixed strategy with support in $\mu(r)$. As shown below, $\mu(r)$ is generically a singleton.

To see that the defender cannot profitably deviate to any $r' \neq r^*$, note that the attacker strikes site $k^* = \min \mu(r^*)$ with probability one after r^* and $k' = \min \mu(r')$ with probability one after r'. Hence, the defender's payoff is $\lambda_{k^*} v_{k^*}(r^*_{k^*})$ to r^* and $\lambda_{k'} v'_k(r'_{k'})$ to r'.

Lemma 2(ii) ensures $k' \in \text{br}_A(r') \subseteq \text{br}_A(r^*)$. That $k' \in \text{br}_A(r^*)$ implies $\gamma_{k'} v_{k'}(r^*_{k'}) = M_A(r^*)$. The fact that r^* is the unique minmax allocation also means $M_A(r^*) < M_A(r')$. Moreover, $M_A(r') = \gamma_{k'} v_{k'}(r'_{k'})$ because $k' \in \text{br}_A(r')$. Combining these relations gives $\gamma_{k'} v_{k'}(r^*_{k'}) = M_A(r^*) < M_A(r') = \gamma_{k'} v_{k'}(r'_{k'})$ which yields $v_{k'}(r^*_{k'}) < v_{k'}(r'_{k'})$.

This inequality leads to $\lambda_{k^*} v_{k^*}(r^*_{k^*}) \leq \lambda_{k'} v_{k'}(r^*_{k'}) < \lambda_{k'} v_{k'}(r'_{k'})$ where the weak inequality follows from the fact that k^* is a site in $\text{br}_A(r^*)$ at which the defender's loss is minimized. The strict inequality shows that the defender's expected loss to playing r^* is less than that of playing r'. Hence, r^* does better against $\alpha^*(r)$ than any $r' \neq r^*$ and so is a best reply to $\alpha^*(r)$. ■

To demonstrate that the defender can never play any $r' \neq r^*$ in any subgame perfect equilibrium, consider any strategy profile $(r', \alpha'(r))$ in which $\alpha'_j(r) > 0$ only if $j \in \text{br}_A(r)$. It suffices to show that the defender can profitably deviate from r' to some other allocation thereby establishing that $(r', \alpha'(r))$ is not an equilibrium.

The intuition guiding the proof begins with part (iii) of Lemma 2 which ensures that if $r' \neq r^*$, then $\sum_{j \in \text{br}_A(r')} r'_j < R$. This implies that r' allocates resources to one or more sites outside $\text{br}_A(r')$. These resources can be redistributed across the sites in $\text{br}_A(r')$ to create a profitable deviation from r' and thereby establish that $(r', \alpha'(r))$ is not an equilibrium allocation.

LEMMA 4: *Suppose $r' \neq r^*$ and take $(r', \alpha'(r))$ to be any strategy profile such that $\alpha'(r)$ maximizes the attacker's payoff for all $r \in \Delta$, i.e., $\alpha'_j(r) > 0$ only if $j \in \text{br}_A(r)$. Then the defender can profitably deviate from r'.*

Proof: Let $k = \min \mu(r')$. Then the defender's loss to playing r' is bounded below by $\lambda_k v_k(r'_k)$. Lemma 2(iii) ensures $\sum_{j \in \text{br}_A(r')} r'_j < R$ since $r' \neq r^*$. Hence there exists a site $n \notin \text{br}_A(r')$ such that $r'_n > 0$. Observe that the following holds for all $j \in \text{br}_A(r')$:

$$\gamma_j v_j(r'_j) > \max\{\gamma_s v_s(r'_s) : s \notin \text{br}_A(r')\} \geq \gamma_n v_n(r'_n). \tag{1}$$

Now construct the allocation $\widehat{r}$ by shifting a small amount of resources from site n and distributing them across all of the sites in $\mathrm{br}_A(r')$ so that attacker's payoffs to striking every site in $\mathrm{br}_A(r')$ declines slightly but that the payoff to striking k declines the least. Constructing the deviation in this way will ensure that attacking site k is the attacker's unique best response to $\widehat{r}$, i.e., $\mathrm{br}_A(\widehat{r}) = \{k\}$.

Formally, inequalities (1) along with continuity guarantee that we can take $\varepsilon > 0$ small enough that

$$\gamma_n v_n(r'_n - \varepsilon) < \frac{\max_{j\in \mathrm{br}_A(r')}\{\gamma_j v_j(r'_j)\} + \max_{s\notin \mathrm{br}_A(r')}\{\gamma_s v_s(r'_s)\}}{2} \equiv B$$

Now let $\varepsilon_j \equiv \varepsilon / \|\mathrm{br}_A(r')\|$ and $\widehat{r}_j \equiv r'_j + \varepsilon_j$ for all $j \neq k$, $j \in \mathrm{br}_A(r')$. Let $\widehat{r}_j \equiv r'_j$ for all $j \neq n$, $j \notin \mathrm{br}_A(r')$. Choose $\delta_k > 0$ so that $\gamma_k v_k(r'_k + \delta_k) > \max_{j\neq k, j\in \mathrm{br}_A(r')}\{\gamma_j v_j(r'_j + \varepsilon_j), B\}$. The fact that $\gamma_k v_k(r'_k) = \gamma_j v_j(r'_j)$ for all $j \in \mathrm{br}_A(r')$ ensures that such an δ_k exists. Define $\varepsilon_k \equiv \min\{\delta_k, \varepsilon / \|\mathrm{br}_A(r')\|\}$ and let $\widehat{r}_k \equiv r'_k + \varepsilon_k$. Finally, set $\varepsilon_n \equiv \varepsilon_k + \sum_{j\neq k, j\in \mathrm{br}_A(r')} \varepsilon_j \leq \varepsilon$ and $\widehat{r}_n \equiv r'_n - \varepsilon_n$.

The attacker's unique best response to $\widehat{r}$ is to strike site k. To see why, note that $\gamma_k v_k(\widehat{r}_k) = \gamma_k v_k(r'_k + \varepsilon_k) \geq \gamma_k v_k(r'_k + \delta_k)$ and, by construction,

$$\gamma_k v_k(r'_k + \delta_k) > \max_{j\neq k, j\in \mathrm{br}_A(r')}\{\gamma_j v_j(r'_j + \varepsilon_j), B\} = \max_{j\neq k, j\in \mathrm{br}_A(r')}\{\gamma_j v_j(\widehat{r}_j), B\}.$$

These inequalities and the definition of B imply that the attacker prefers hitting k to any $j \neq n$. That $\gamma_k v_k(\widehat{r}_k) > B > \gamma_n v_n(r'_n - \varepsilon) \geq \gamma_n v_n(r'_n - \varepsilon_n)$ ensures that the attacker prefers hitting k to n. Hence $\mathrm{br}_A(\widehat{r}) = \{k\}$.

The attacker therefore hits k with probability one in any subgame perfect equilibrium which leaves the defender with a loss of $\lambda_k v_k(\widehat{r}_k)$. By construction, $\widehat{r}_k > r'_k$, so $\lambda_k v_k(\widehat{r}_k) < \lambda_k v_k(r'_k)$. But $\lambda_k v_k(r'_k)$ is a lower bound on the defender's payoff to r'. Hence, $\widehat{r}$ is a profitable deviation from r'. ■

Lemmas 3 and 4 establish Proposition 1. A pure strategy subgame perfect equilibrium exists, and the defender minmaxes the attacker by playing r^* in every subgame perfect equilibrium.

The Attacker's Equilibrium Strategy

It remains to be shown than there is a generically unique equilibrium path and that it entails pure strategies. The definition of subgame perfection implies that the support of the attacker's response to r^* must be contained in $\mathrm{br}_A(r)$. Proposition 2 shows that generically the attacker's response to r^* must be contained in $\mu(r^*)$. More specifically, the attacker can only attack sites in $\mu(r^*)$ after observing r^* for all but finitely many values of R. It is easy to see that $\mu(r^*)$ is generically a singleton. Thus the generically unique equilibrium path is for the defender to minmax the attacker with r^* and for the attacker to reply by striking the generically solitary site in $\mu(r^*)$.

The intuition for why the attacker must go after a site in $\mu(r^*)$ after observing r^* is that if it did not, then the defender could profitably deviate from r^* by allocating slightly less to a $k \in \mu(r^*)$. The attacker's unique best response to this deviation is to strike k for sure. This deviation from r^* is profitably as it gives the defender a payoff arbitrarily close to what it would have gotten had the defender hit k in response to r^*. If however $r_k^* = 0$, then the defender cannot deviate from r^* by allocating slightly less to r_k^*. However the condition $r_k^* = 0$ is nongeneric and can hold for at most finitely many values of R as long as resources are scarce.

The assumption that resources are scarce may put an upper bound on R. Suppose that the defender could defend every site perfectly if it had enough resources. That is, there exists an R' and an allocation r' of R' such that $v_j(r_j') = 0$ for all j. Let $\overline{R}$ be the smallest allocation with which it is possible to protect every site perfectly. (Continuity ensures that a minimum exists.) If it is impossible to defend everything perfectly no matter how much the defender has (e.g., $v_j(r_j)$ is bounded away from zero or goes to zero only in the limit), take $\overline{R}$ to be ∞. Then resources are scarce when $R \in (0, \overline{R})$.

Lemma 1 ensures that the for every $R \in (0, \overline{R})$ the corresponding minmax allocation of R, $r^*(R)$, is unique. Define B_j to be the allocation R that satisfies $\gamma_j v_j(0) = M_A(B_j)$ if such an an allocation exists for an $E_j \in (0, \overline{R})$. At most, only one such allocation exists since M_A is strictly decreasing over $(0, \overline{R})$. Let $B = \bigcup_{j=1}^{N} B_i$. Then:

LEMMA 5: *If $R \in (0, \overline{R})\backslash B$, then $r_j^*(R) > 0$ for all $j \in \mathrm{br}_A(r^*(R))$.*

Proof: Suppose not. Then there exists some $R' \in [0, \overline{R})\backslash B$ and $i \in \mathrm{br}_A(r^*(R'))$ such that $r_i^*(R') = 0$. But $i \in \mathrm{br}_A(r^*(R'))$ implies $\gamma_i v_i(r_i^*(R')) = M_A(r^*(R'))$. This leaves $\gamma_i v_i(0) = M_A(r^*(R'))$ which means $R' = E_i$ and this contradicts the assumption that $R' \notin B$. ■

It follows that if $R \in (0, \overline{R})\backslash B$, then the support of the attacker's equilibrium response to $r^*(R)$ must be contained in $\mu(r^*(R))$.

PROPOSITION 2: *Let $(r^*(R), \widehat{\alpha}(r))$ be any subgame perfect equilibrium for any $R \in [0, \overline{R})\backslash B$. Then $j \in \mu(r^*(R))$ whenever $\widehat{\alpha}_j(r^*(R)) > 0$.*

Proof: Arguing by contradiction, assume there exists an $R' \in (0, \overline{R})\backslash B$ and a subgame perfect equilibrium $(r^*(R'), \alpha'(r))$ such that $\alpha_i'(r^*(R')) > 0$ for some $i \notin \mu(r^*(R'))$. Dropping the argument R' to ease the notational burden, let k be any element in $\mu(r^*)$. Then $k \in \mathrm{br}_A(r^*)$ and $\lambda_k v_k(r_k^*) < \lambda_i v_i(r_i^*)$.

Lemma 5 ensures $r_k^* > 0$. We can therefore construct the allocation $\widehat{r}$ such that $\widehat{r}_k = r_k^* - \varepsilon$ and $\widehat{r}_j = r_j^*$ for all $j \neq k$. It follows that $\mathrm{br}_A(\widehat{r})) = \{k\}$ since $\gamma_k v_k(r_k^* - \varepsilon) > M_A(r^*) = \max\{\gamma_j v_j(r_j^*) : j \neq k\}$. The defender's loss to playing $\widehat{r}$ is therefore $\lambda_k v_k(r_k^* - \varepsilon)$.

The defender's payoff to offering r^* is $\sum_{j \in \mathrm{br}_A(r^*)} \lambda_j v_j(r_j^*)\alpha_j'(r^*) \geq \lambda_i v_i(r_i^*)\alpha_i'(r^*) + [1 - \alpha_i'(r^*)]\lambda_k v_k(r_k^*) > \lambda_k v_k(r_k^*)$ where the strict inequality follows from the fact that $\alpha_i'(r^*) > 0$. The strict inequality along with continuity guarantee that we can choose ε so that $\sum_{j \in \mathrm{br}_A(r^*)} \lambda_j v_j(r_j^*)\alpha_j'(r^*) > \lambda_k v_k(r_k^* - \varepsilon) = \lambda_k v_k(\widehat{r}_k)$. Hence the defender strictly prefers deviating to $\widehat{r}$. This contradicts the assumption that $(r^*, \alpha'(r))$ is subgame perfect and establishes the claim. ■

It is also trivial to see that $\mu(r^*)$ is generically a singleton. Suppose $j \neq k$ are both in $\mu(r^*)$. Since $\mu(r^*) \subseteq \mathrm{br}_A(r^*)$, $\gamma_j v_j(r_j^*) = \gamma_k v_k(r_k^*)$. That j and k are in $\mu(r^*)$ also means that they minimize the defender's losses over the sites in $\mathrm{br}_A(r^*)$. Consequently, $\lambda_j v_j(r_j^*) = \lambda_k v_k(r_k^*)$. These equalities imply that $\lambda_j/\lambda_k = \gamma_j/\gamma_k$ which is clearly non-generic. Hence, the defender's unique equilibrium strategy is to minmax the attacker who replies by attacking the generically unique site among its best replies that minimizes the defender's losses.

Conclusion

The strategic allocation of resources across multiple fronts has long been studied in the context of Blotto games. In these two-person, zero-sum games, the actors simultaneously decide how to distribute their resources. However many allocation problems are fundamentally dynamic. An attacker, for example, often decides where and how to strike after observing (perhaps imperfectly) how the defender has allocated its defensive resources. In the present analysis, a defender decides how to distribute limited resources across multiple sites. The more a defender dedicates to a specific site, the less likely an attack on that site is to succeed. After observing how the defender has allocated its resources, the attacker decides where to strike.

Unlike many Blotto games which only have very complicated mixed-strategy equilibria, the sequential, nonzero-sum Blotto game always has a subgame perfect equilibrium, the defender always plays the same pure strategy in any equilibrium, and the attacker's equilibrium response is generically unique and entails no mixing. The defender minmaxes the attacker, and the attacker strikes the site among its best replies that minimizes the defender's expected losses.

References

Blackett, D.W. 1958. "Pure Strategy Solutions to Blotto Games," *Naval Research Logistics Quarterly* 5:107-109.

Borel, Emil. 1921. "La Théorie du Jeu les Equations Integrales à Noyau Symétrique," *Comptes Rendus de l'Académie* 173:1304-308. English translation by Leonard Savage, "The Theory of Play and Integral Equations with Skew Symmetirc Kernels," *Econometrica* 21(1953):97-100.

Coughlin, Peter. 1992. "Pure Strategy Equilibria in a Class of Systems Defense Games," *International Journal of Game Theory* 20:195-210.

Dekel, Eddie, Matthew Jackson, Asher Wolinsky. 2008. "Vote Buying: General Elections," *Journal of Political Economy* (forthcoming).

Golman, Russell and Scott E. Page. 2006. "General Blotto: Games of Allocative Strategic Mismatch," Manuscript, Center for Complex Systems, University of Michigan. Accessed November 3, 2006 at http://www.cscs.umich.edu/~spage/papers_files/blotto.pdf.

Harris, Christopher. 1985. "Existence and Characterization of Perfect Equilibrium in Games of Perfect Information." *Econometrica* 55:613-28.

Hart, Sergiu. 2007. "Discrete Blotto Games and General Lotto Games." Manuscript, Deparatment of Economics, Hebrew University.

Laslier, Jean-Francois and Nathalie Picard. 2002. "Distributive Politics and Electoral Competition," *Journal of Economic Theory* 103: 1-6-130.

Powell, Robert. 2007. "Defending Against Terrorist Attacks with Limited Resources," *American Political Science Review* 101(August):527-41.

Roberson, Brian. 2006. "The Colonel Blotto Game," *Economic Theory* 29:1-24.

Shubik, Martin and Robert James Weber. 1978. "Competitive Valuation of Cooperative Games." Discussion Paper No. 482, Cowles Foundation, Yale University.

Shubik, Martin and Robert James Weber. 1981. "Systems Defense Games: Colonel Blotto, Command and Control," *Naval Research Logistics Quarterly* 28(2): 281-87.

Tukey, John W. 1949. "A Problem of Strategy," *Econometrica* 17: 73.