

9-2012

Weakest-Link Attacker-Defender Games With Multiple Attack Technologies

Daniel G. Arce
University of Texas - Dallas

Dan Kovenock
Chapman University, kovenock@chapman.edu

Brian Roberson
Purdue University

Follow this and additional works at: http://digitalcommons.chapman.edu/economics_articles

 Part of the [Defense and Security Studies Commons](#), [Military Studies Commons](#), [Peace and Conflict Studies Commons](#), and the [Political Theory Commons](#)

Recommended Citation

Arce, Daniel G., Dan Kovenock, and Brian Roberson. "Weakest-link attacker-defender games with multiple attack technologies." *Naval Research Logistics (NRL)* 59.6 (2012): 457-469.
DOI:10.1002/nav.21500

This Article is brought to you for free and open access by the Economics at Chapman University Digital Commons. It has been accepted for inclusion in Economics Faculty Articles and Research by an authorized administrator of Chapman University Digital Commons. For more information, please contact laughtin@chapman.edu.

Weakest-Link Attacker-Defender Games With Multiple Attack Technologies

Comments

This article was originally published in *Naval Research Logistics*, volume 59, issue 6, in 2012. DOI: [10.1002/nav.21500](https://doi.org/10.1002/nav.21500)

Copyright

Wiley-Blackwell

Weakest-Link Attacker-Defender Games with Multiple Attack Technologies

Daniel G. Arce,¹ Dan Kovenock,² Brian Roberson³

¹ *Department of Economics, University of Texas at Dallas, GR31, School of Economic, Political, and Policy Sciences, 800 W. Campbell Rd, Richardson, Texas 75080*

² *Economic Science Institute, Argyros School of Business and Economics, Chapman University, One University Drive, Orange, California 92866*

³ *Department of Economics, Purdue University, Krannert School of Management, 403 W. State Street, West Lafayette, Indiana 47907*

Received 29 September 2011; revised 14 June 2012; accepted 1 July 2012

DOI 10.1002/nav.21500

Published online 30 July 2012 in Wiley Online Library (wileyonlinelibrary.com).

Abstract: In this article, we examine a model of terrorism that focuses on the tradeoffs facing a terrorist organization that has two qualitatively different attack modes at its disposal. The terrorist organization's objective is to successfully attack at least one target. Success for the target government is defined as defending all targets from any and all attacks. In this context, we examine how terrorist entities strategically utilize an efficient but discrete attack technology — e.g., suicide attacks — when a more conventional mode of attack is available, and the optimal anti-terrorism measures. © 2012 Wiley Periodicals, Inc. *Naval Research Logistics* 59: 457–469, 2012

Keywords: conflict; suicide terrorism; weakest link; colonel Blotto game

1. INTRODUCTION

Terrorism is a form of asymmetric conflict in which terrorists utilize violent actions against (mainly civilian) noncombatants to influence a target audience beyond the immediate victims and, ultimately, to obtain ideological, political, or religious objectives. Whereas terrorism is asymmetric in that terrorist groups have a relative resource disadvantage with respect to the target government, there also exist structural asymmetries between attack and defense which terrorists can turn to their advantage through their selection of targets and tactics. In particular, governments with high-profile counterterror policies (e.g., the war on terror), or those facing a coordinated terrorist campaign (e.g., the French in Algeria) are often judged by their ability to deter or interdict all attacks. If one target is successfully attacked, then counterterror policy and the competency of the government itself can be subject to public scrutiny. For example, as is written in the *Joint House-Senate Intelligence Inquiry into September 11, 2001* (US Congress [51]), terrorists need to be successful only once to kill Americans and demonstrate the inherent vulnerabilities they face. Similarly, after bombing the Brighton hotel where Margaret Thatcher was staying in the 1980s, and

failing to kill her, the IRA issued a statement that read: “Today you have been lucky. But you have to be lucky every time. We only have to be lucky once,” (King [32]). This suggests that, as a whole, the set of targets of interest to terrorist groups may be viewed as a weakest-link network¹ from the perspective of a target government,² whose success is defined in terms of security against all possible attacks.³

¹ The term weakest link stems from Hirshleifer's [25] metaphor about the public good provided by dike builders on the perimeter of a circular island (c.f. Cornes [12], Hausken [23]). Whoever builds the lowest dike will define the entire island's level of defense against a flood.

² The weakest-link viewpoint may be due to the policymakers' (or voters') perception that successful counterterror policy involves the complete absence of incidents within a defined protectorate (e.g., Gassebner et al. [17, 18], King [32], and Rosenbaum [42]), or the target itself may be a network corresponding to a weakest-link technology, as is the case with inter-airline baggage handling or critical infrastructure. Under either interpretation there is a structural asymmetry in the terrorist's favor.

³ When the vulnerability of one target not only depends on its choice of security measures but also on the actions of others, a situation of interdependent security can arise that is consistent with a weakest link. Heal and Kunreuther [24] give the example of airline baggage screening. Specifically, the 1988 crash of Pan Am 103 over Lockerbie, Scotland was due to a bomb that was contained in a bag initially screened by Malta Airlines in Malta, thereby constituting the weak link.

Part of this work was completed while Kovenock was Visiting Professor at the Social Science Research Center Berlin (WZB).

Correspondence to: B. Roberson (brobers@purdue.edu)

In addition to the asymmetry that for terrorists one success is often enough to alter the political landscape, airways, and so forth,⁴ terrorists have multiple tactics at their disposal. Consider for example, suicide terrorism, which accounts for an average of 12 times more damage than conventional attacks (Sandler et al. [45]) and could be even deadlier were it not for the associated crowd blocking of casualties (Kress [37]). The modern use of this tactic dates to the 1983 Beirut bombings by Hezbollah against US and French military personnel, with the bombings being viewed as bringing about these nations' troop withdrawals from Lebanon. This tactic was subsequently adopted by the Tamil Tigers (LTTE) and Kurdistan Workers' Party (PKK), and has been indelibly ingrained into the American psyche subsequent to the mass casualty suicide attacks of September 11, 2001 (called 9/11 hereafter). However, no group that employs suicide terror does so exclusively (Crenshaw [13]). For example, the March 3, 2004 train station bombings in Madrid are associated with affiliates of al-Qaeda, but it was not a suicide operation, as the bombs were left on trains.⁵ Indeed, the use of cell phones as detonators in the Madrid bombings is one rationale for why the Aznar government initially suspected that the Basque organization *Euskadi ta Askatasuna* (ETA) was behind the attacks, as this form of detonation was a signature of past ETA attacks. Moreover, before the bombings the Aznar government was favored to easily win the elections that were scheduled three days hence. Instead, it lost; a result that is widely interpreted as stemming from electoral accountability in the aftermath of the bombings.⁶ For the al-Qaeda organization and its affiliates alone, the 1993 World Trade Center bombing; the May 29, 2004 Al-Khobar massacres in Saudi Arabia; and the June 30, 2007 discovery of explosives found in unattended cars parked at Piccadilly Circus and Trafalgar Square are additional examples of non-suicide attacks.

In this article, we examine and characterize — in the context of a weakest-link network — how terrorist entities strategically utilize two qualitatively different modes of attack: a discrete one being capable of inflicting more damage at a lower per unit cost and a more conventional one. Throughout this article, we use the convention of referring to suicide attacks as this deadlier mode of attack. The invisibility of suicide operatives is a common feature that is

applicable to a number of military technologies including chemical, biological and radio-nuclear (CBRN) weapons of mass destruction (WMD). However, tactics such as these have yet to be proven to be cost efficient for terrorists relative to suicide or conventional attacks, due to the difficulties of procurement and weaponization of CBRN and increased vulnerability to detection of WMD by intelligence services (Franck and Melese [16]).⁷ Our choice of labeling the discrete but efficient mode of attack as suicide terrorism is motivated by: (i) on average, suicide operatives produce more than they cost (Atran [3]), (ii) the utilization of suicide tactics has increased in recent years (Economist [14]), and (iii) a terrorist organization that has the ability to utilize either or both suicide terrorism tactics and conventional tactics faces a non-trivial tradeoff.⁸ In regard to this last point, Crenshaw [13] provides a review of 13 books on the subject of suicide terrorism/martyrdom, all of which were published post-9/11, and deal almost exclusively with suicide bombing from the perspective of the bomber/operative. However, both Hoffman and McCormick [26] and Crenshaw [13] recognize that suicide actions are rational for the group that operatives represent, and that explaining how suicide tactics fit into the groups' overall strategy of violence is remarkably understudied.⁹

In examining suicide terrorism, our analysis highlights two critical features: (i) weakest-link networks of targets and (ii) the availability of both conventional and suicide tactics for the attacker. We find that, in equilibrium: (i) a terrorist organization may choose with positive probability not to launch any attacks, (ii) in the case that an attack is launched, at most one target is attacked, and (iii) conditional on an attack being launched, the suicide attack technology is not utilized with probability one. Remarkably, we find that the frequency and magnitude of suicide attacks in our model depends on a simple measure that incorporates the structural asymmetry arising in the weakest-link network and the asymmetry between the characteristics of the attacker and the defender, which we term “the normalized relative strength of the attacker.” As the normalized relative strength of the attacker approaches unity the conflict becomes more symmetric, and the equilibrium frequency and magnitude of suicide

⁴ Kaplan et al. [31] identify an alternative asymmetry in “low-level” intelligence that acts in the terrorists'/insurgents' favor. This asymmetry counteracts the government's advantages in resources and force sizes. Similarly, in a related attack and defense game, Bernhardt and Polborn [6] examine a cost-based asymmetry between attack and defense. In that case, the “committed” attacker experiences no opportunity costs from allocating forces and continues attacking targets until either he runs out of targets or is defeated.

⁵ The suspects blew themselves up later to avoid capture.

⁶ Suicide attacks are, on average, far more severe than conventional attacks, and the severity of attack has been shown to increase the likelihood of cabinet changes within a government (Gassebner et al. [17, 18]).

⁷ As noted by a referee, the “lumpiness” of suicide operatives is a feature of a number of traditional military applications. For example, a naval commander must choose a discrete number of aircraft carriers to execute an attack. For an early discussion of traditional military applications of this general type of game see Blackett [8].

⁸ As discussed below, this issue is examined in Rosendorff and Sandler [43] and Feinstein and Kaplan [15].

⁹ In her review, Crenshaw [13] concludes that there is no longer any need to introduce an analysis of suicide attacks by explaining to the uninitiated that it is not rooted in psychopathology or fanaticism or indeed in any single cause such as deprivation, religious belief, or frustration (p. 162).

attacks increases. In addition, we find that the incidence of suicide terrorism increases as the total cost of utilizing suicide operatives decreases. Given that this total cost includes the costs of recruiting and training suicide operatives as well as the final force expenditure, our model is consistent with the stylized fact that suicide terrorism is likely to arise in an environment in which a group has significant political support (i.e., lower costs of recruitment), but not the means for political expression (Hoffman and McCormick [26]).

Strategic interdependencies, or linkages, across targets is a central theme in the literature on multidimensional contests.¹⁰ Although the particular linkages arising in weakest-link networks differ from those arising in traditional models of strategic resource allocation, these two types of linkages are related. For example, in the classic Colonel Blotto game¹¹ each player has a fixed level of forces, each target is won by the player who allocates the higher level of force, and the payoff to each player is the sum of the wins across the entire set of targets. In that game, the constraint on the total force expenditure creates a cost-based linkage among the battlefields, that is, forces allocated to a specific battlefield reduce the level of force that can be allocated to the other battlefields. By contrast, in our study the linkage arises through the players' objectives. In particular, both the government's (defender's) and terrorists' (attacker's) payoffs are a function of the most successful of the attacks across targets.¹² Analyses of this type of weakest-link defense technology include Clark and Konrad [11] and Kovenock and Roberson [35];¹³ however, these models differ from our focus here in that the attacker has only the conventional allocation of homogeneous resources across targets. In our study, the attacker has two qualitatively different attack technologies, conventional and suicide tactics. This framework allows us to characterize how

and why terrorist organizations choose between conventional and suicide attacks, as those organizations that have suicide operatives at the ready do not rely exclusively on suicide attacks.

As in the multidimensional contest literature, defense, for the target government in our model, involves the hardening of targets. The macrotechnologies of conflict for target governments may include defensive and/or proactive/preemptive measures (Arce and Sandler [2]) and disruptive and/or defensive tactics (Franck and Melese [16]). In contrast to our defensive tactics, Kaplan *et al.* [31] and Jacobsen and Kaplan [30] examine the "targeted killing" of entrenched insurgent cells, which would fall under the category of proactive/preemptive measures in Arce and Sandler's [2] taxonomy of counter-error strategies and disruptive tactics under the Franck and Melese [16] system. By focusing on only defensive tactics, our model is silent about considerations such as how targeted killings may lead to a desire for terrorist retaliation as in Jacobsen and Kaplan [30], or how the costs of civilian casualties resulting from government intervention affect the government's optimal anti-terrorism policy, as in Kaplan *et al.* [31] and Jacobsen and Kaplan [30].

The issue of multiple attack technologies is also examined by Rosendorff and Sandler [43] and Feinstein and Kaplan [15]. In both of those papers, the attacker has two attack technologies, normal and spectacular. In Feinstein and Kaplan [15], a normal attack is characterized as having low fixed costs and high marginal costs. Correspondingly, a spectacular attack has high fixed costs and low marginal costs. There are a number of important distinctions between the modeling approach we use, and those examined in these papers.¹⁴ However, our attack technology generalizes the multi-attack technology utilized in these papers by modeling multiple input use explicitly, allowing for a continuous range of normal attacks and multiple (discrete) levels of suicide spectaculars. We assume that each suicide operative requires a fixed cost for recruiting and training, but that the force effectiveness in the presence of suicide operatives can also be increased at a constant marginal cost through the use of conventional resources. To simplify the discussion of the model and the results on the tradeoffs facing the terrorist organization and the optimal counter-terrorism strategies, we focus on the simplest case in which each suicide operative has the same fixed cost and the marginal costs of increasing force effectiveness through the use of conventional resources in both normal and suicide attacks are set to one. However, it is straightforward, but somewhat tedious, to extend our results to allow for differing marginal costs and for the attacker to face increasing fixed

¹⁰ For a survey see Kovenock and Roberson [36].

¹¹ This game originates with Borel [9] who examines the case of symmetric players and three homogeneous battlefields. Early extensions include Gross and Wagner [21] who allow for a finite number of symmetric battlefields, and Gross [20] who allows for heterogeneous battlefield valuations. Recent extensions include: asymmetric players (Hart [22], Macdonell and Mastronardi [39], Roberson [40], Weinstein [52]), non-constant-sum variations (Kvasov [38], Hortala-Vallve and Llorente-Saguer [27, 28], Roberson and Kvasov [41]), and alternative definitions of success (Golman and Page [19], Szentes and Rosenthal [48], [49], and Tang, Shoham, and Lin [50]).

¹² Similar objective-based linkages arise in Szentes and Rosenthal's [48] chopstick auction in which three chopsticks are being auctioned and each of two players seeks to win at least two of the auctions. In the context of politicians engaged in a campaign resource allocation game, Snyder [47] and Klumpp and Polborn [33] examine the case in which each player seeks to win a majority of the component contests. The related case in which success requires winning a super-majority is examined by Szentes and Rosenthal [49].

¹³ See Shubik and Weber [46] for an early treatment of a related game, and Kovenock and Roberson [36] for further discussion of literature utilizing the weakest-link framework.

¹⁴ Most notably, Feinstein and Kaplan [15] is a dynamic model and the authors use simulations to characterize optimal strategies, and our model examines a weakest-link network of targets rather than a single target.

costs for suicide operatives. Furthermore, our main results on the tradeoff between conventional and suicide tactics hinge on the cost-effectiveness of suicide operatives rather than differences in the marginal costs.

Closely related is the literature on attacker–defender games, also known as interdiction models, that feature a sequential-move structure in which the defender is an exogenously imposed leader, the defender maximizes his value for the system, and the attacker’s objective is to minimize the maximum value of the system. For example, Brown et al. [10] consider a critical infrastructure model in which there are complex network effects involving weakest-link substructures. In contrast to our weakest-link system approach, in Brown et al. [10] the overall system is expected to operate in a degraded fashion following a successful attack on a weakest-link subsystem. This leads governments to identify weakest-link infrastructure (subsystems) and terrorists to probe for weakest links in the government’s payoff function(s). Our model differs from this game — and from the attacker-defender game literature in general — in terms of the structural asymmetry in the players’ objectives (weakest-link), multiple modes of attack, and simultaneous-move structure. With a sequential-move structure in which the defender is an exogenously imposed leader, the defense has no opportunity to conceal the allocation of forces and all attacker allocations can be made contingent on given defensive allocations. However, in most applications defensive resources can either be concealed or randomly allocated with sufficient speed that it is difficult to argue that attacker allocations can be made contingent on defensive allocations. This is certainly true in the case of information or transportation network defense or border defense, where either attackers must take actions before being certain of the allocation of defensive resources or where strategies like random monitoring or deployment may be used by defenders. Indeed, in these contexts, exogenous leader structures seem rather implausible. Maybe this is why the early literature on Blotto games expended substantial effort to solve the simultaneous move case, despite the simplicity of the solution of sequential games of perfect information.

The contest success function we use is a special case of $a^m/(a^m + d^m)$ where a and d denote, respectively, the attacker’s and defender’s effective expenditures, and m equals infinity. However, it is not true that the existing literature that uses this formulation is more general than our approach. In simultaneous-move games this literature generally places very severe restrictions on the exponent m to ensure the existence of a pure-strategy equilibrium. For a single contest with linear costs, (the famous Tullock rent seeking model), a pure strategy equilibrium exists only for m less than or equal to 2, and concavity in a player’s variable holds only for m less than or equal to 1. For m greater than 2, as in the $m = \infty$ case, no pure-strategy equilibria exist. Although there has not

been a complete characterization of the equilibrium set for the $m > 2$ case, except for Baye et al.’s [5] characterization for $m = \infty$, we do know that there exist equilibria in one shot contests that are payoff equivalent to the $m = \infty$ case whenever $m > 2$ (Baye et al. [4], Alcalde and Dahm [1]). Hence, as argued in Konrad and Kovenock [34] simultaneous move models with pure-strategy equilibria employing the form of $a^m/(a^m + d^m)$ are not more general than all-pay auction-based equilibria and hold only for a small range of the parameter space of m . The interpretation of models with low m is that they involve a sufficiently large amount of noise (see Konrad and Kovenock [34] for a discussion of how much noise is implied). Models with high m , with $m = \infty$ the limiting case, are models with low or no noise.

The article proceeds as follows. In Section 2, we describe a model of conflict with technologies of attack and defense in terms of the players, their strategies, and payoffs. In Section 3, the model is solved with the result being a mixed strategy Nash equilibrium. This characterization is consistent with the observation that suicide terrorism is not the exclusive *modus operandi* of the way in which terrorists broaden the impact of their actions by creating an aura of uncertainty through tactics that appear to be random. In particular, we are able to characterize the frequency of suicide attacks and the nature of terrorist “spectaculars,” whether of the suicide or conventional variety. The final section contains brief concluding remarks.

2. THE MODEL

2.1. Players and Strategies

We examine a complete-information, simultaneous-move, one-shot game in which two players, an attacker, A , and a defender, D , allocate their forces across a weakest-link network consisting of a finite number, $n \geq 2$, of homogeneous targets. The defender chooses a level of a continuous (conventional) one-dimensional defensive force for each of the targets. The attacker can also choose a level of the conventional (non-suicide) force for each target. For both players conventional forces have a unit cost equal to one, and the level of conventional force allocated to each target must be nonnegative. In addition to a conventional attack, for each target i the attacker has the opportunity to send any discrete number of suicide operatives denoted by $s_i \in \{0, 1, 2, \dots\}$ (where $s_i = 0$ denotes no suicide attack) at cost c for each operative which provides an effective force allocation of S for each operative. Note that our focus is on suicide operatives, an inherently discrete resource for the attacker. Let $\mathbf{s} = (s_1, s_2, \dots, s_n)$ denote the n -tuple of the attacker’s allocation of suicide operatives across the n targets. Observe that to get the same effect, with conventional forces, as a suicide attack on target i with s_i suicide operatives, the attacker

would have to allocate $s_i \cdot S$ units of the conventional forces to target i . In addition, the attacker has the ability to continuously increase the force effectiveness of a suicide attack by using conventional forces at a constant marginal unit cost equal to one. That is, to launch a suicide attack on target i with s_i operatives the attacker incurs a fixed cost of $s_i c$ and faces a constant marginal cost of one per unit of additional force effectiveness beyond $s_i S$. Note that our assumption that conventional forces can be used to augment the effectiveness of a suicide attack does not imply that conventional forces are utilized in the same manner in both types of attacks.

We focus on the case in which the suicide attack is strongly efficient, $c < S$. For example, the improvised explosive devices worn or carried by a suicide bomber can cost less than \$150 to produce and the bombers themselves are regarded as expendable assets from the organizational perspective (Hoffman and McCormack [26]). Furthermore, the requirements to be a successful suicide operative are not trivial, involving a level of intelligence that exceeds what is required of operatives in a conventional attack.¹⁵ The suicide attack technology captures the notion that a tactic such as a suicide attack is a discrete decision that, although cost effective, entails costs — including recruitment, training, and the final force expenditure. As Iannaccone [29] observes, the number of “martyrs” is very small relative to the total number of the members in the groups that use suicide terrorism.

2.2. Payoffs

Our focus is on a weakest-link network of targets, and the players have asymmetric payoff functions reflecting the structural asymmetry arising in the weakest-link network. For each target, the player that allocates the higher level of force wins that target. In the case that the players allocate the same level of force to a target, the defender wins the target. For the defender success consists of allocating at least as high a level of force to all targets within the network. Conversely, an attacker is successful if he allocates a higher level of force to at least one target in the network.

For example, using Memorial Institute for the Prevention of Terrorism (MIPT) data, Gassebner et al. [17, 18] find statistically significant evidence of a “one strike and you’re out” phenomenon whereby the presence of at least one terror event increases the likelihood of a cabinet change within a target government, with the likelihood of a change increasing with the severity of attack. Furthermore, some targets themselves are, by definition, weakest links. The luggage transfer of the suitcase bomb that downed Pam Am flight 103 is an

¹⁵ For example, Sageman [44] finds that the suicide operatives of the global Salafist movement (which includes al-Qaeda) were far more educated than the average person worldwide, with 60% having college degrees.

example, as is the interdiction of a twin car bombing plot against Saudi Arabia’s main oil processing facility (Economist [14]). Similarly, pipeline attacks in Nigeria have had a significant impact on Nigerian oil production as well as on crude prices internationally.

For the defender, let $\mathbf{d} = (d_1, d_2, \dots, d_n)$ denote an n -tuple of forces across the n targets. Similarly, let $\mathbf{a} = (a_1, a_2, \dots, a_n)$ denote an n -tuple of the attacker’s conventional forces, where a_i denotes the attacker’s allocation of conventional force to target i . Recall that $\mathbf{s} = (s_1, s_2, \dots, s_n)$ denotes the attacker’s n -tuple of suicide operatives. Given that for each target i the attacker may utilize either conventional or suicide tactics, we define the attacker’s effective force allocation as follows.

DEFINITION 1: The attacker’s effective force allocation for target i is equal to the allocation of the continuous conventional resource to target i plus any and all suicide operatives allocated to target i :

$$\hat{a}_i = a_i + S s_i.$$

The n -tuple of the attacker’s effective force allocations is denoted by $\hat{\mathbf{a}}$.

Observe that if the attacker has chosen an effective force level of \hat{a}_i for target i such that $S \leq \hat{a}_i < 2S$, it is clearly cost minimizing for the attacker to set $s_i = 1$. Similarly, given an effective force level of \hat{a}_i for target i such that $\lambda S \leq \hat{a}_i < (\lambda + 1)S$ for some integer λ , an optimizing attacker has implicitly chosen $s_i = \lambda$. Note that any effective force level of \hat{a}_i such that $\lambda S + c < \hat{a}_i < (\lambda + 1)S$ is provided at the lowest cost by using λ suicide operatives combined with an investment of $\hat{a}_i - \lambda S$ additional units of force effectiveness.¹⁶ However, the lowest cost of this effective force level is $\hat{a}_i - \lambda(S - c)$, which is greater than the cost of using $\lambda + 1$ suicide operatives, attaining an effective force level of $(\lambda + 1)S > \hat{a}_i$. Consequently, no effective force allocation \hat{a}_i such that $\lambda S + c < \hat{a}_i < (\lambda + 1)S$ will be optimally used by the attacker.

Success for the attacker is formally defined as follows.

DEFINITION 2: The weakest-link indicator function, denoted by ι^{WL} , takes a value of one if there exists a target i for which the attacker’s effective force (as defined in Definition 1) exceeds the defensive forces allocated to that target and takes a value of zero otherwise.

$$\iota^{WL} = \begin{cases} 1 & \text{if } \exists i \mid \hat{a}_i > d_i \\ 0 & \text{otherwise} \end{cases}.$$

¹⁶ If $\lambda = 0$, then these are conventional forces, but if $\lambda > 0$, then this is an investment in continuously increasing the suicide attack force effectiveness.

In the event that all targets are successfully defended, the weakest-link indicator function takes a value of zero, but if any single attack is successful this indicator takes a value of one. Again, this corresponds to a “one strike and you’re out” implication for an incumbent target government (Gassebner et al. [17, 18] and Rosenbaum [42]). It also refers to the terrorist’s need to only be lucky once to highlight the government’s vulnerability (King [32]). Alternatively, a collection of specific targets (e.g., critical infrastructure) may intrinsically exhibit a weakest-link network structure.

Although we focus on the case in which conventional forces and suicide operatives are linearly additive, in practice a terrorist organization may choose to send a suicide operative to a target and then follow that up with another attack on the first responders. If such a technology dominates the pure suicide-bomber technology, then in our model the combination of a suicide attack and follow up force could be treated as a single unit which, given its cost structure and effective force implications, would fit directly into our framework.¹⁷

The attacker’s (terrorist’s) payoff function is given by

$$\pi_A(\mathbf{a}, \mathbf{s}, \mathbf{d}) = v_A \iota^{WL} - c \sum_{i=1}^n s_i - \sum_{i=1}^n a_i.$$

When any target is successfully attacked, so that $\iota^{WL} = 1$, the terrorist receives the value of a successful attack, v_A , less the total cost of all suicide attacks (if any) and the cost of all conventional attacks (if any). If no target is successfully attacked these costs are still born by the terrorist. Also, we have normalized the per unit cost of conventional forces to one, and c represents the fixed cost per suicide operative.

The defender’s payoff function is given by

$$\pi_D(\mathbf{a}, \mathbf{s}, \mathbf{d}) = v_D(1 - \iota^{WL}) - \sum_{i=1}^n d_i.$$

As in an insurance policy, the defender always pays the cost of defense, $\sum_{i=1}^n d_i$. This is augmented by the value of a successful defense, v_D , when every target is successfully defended, thereby reflecting a weakest-link vulnerability.

Given that terrorism is a form of asymmetric conflict with respect to both the resource disparity and the structural externalities arising in the weakest-link network of targets, it will be useful to introduce a simple summary statistic which captures both of these forms of asymmetry. Recall that the floor function $\lfloor x \rfloor$ gives the largest integer less than or equal to

x , and observe that $\lfloor \frac{v_A}{c} \rfloor$ is the maximum number of suicide operatives that the terrorist organization can profitably use.

Note that the maximum profitable expenditure for the attacker (defender) is v_A (v_D), which if used solely by conventional means, translates into a maximal effective force of v_A (v_D). To capture the notion that terrorist organizations have a relative resource disadvantage with respect to the target government, we focus on the case that $v_D > v_A$. However, terrorist organizations also have the ability to utilize suicide operatives. An allocation \mathbf{s} of suicide operatives across the targets increases the effective force by $\sum_i s_i S$ at a cost of $\sum_i s_i c$, implying that the maximal effective force that can be allocated at a cost of v_A is $v_A + \lfloor \frac{v_A}{c} \rfloor (S - c)$.

DEFINITION 3: The normalized relative strength of the attacker, denoted by α , is the ratio of n times the attacker’s maximal effective force allocation to the defender’s maximal effective force allocation,

$$\alpha = \frac{n(v_A + \lfloor \frac{v_A}{c} \rfloor (S - c))}{v_D},$$

where $\alpha < 1$ implies that the attacker is relatively disadvantaged and $\alpha > 1$ implies that the attacker is relatively advantaged. As $\alpha \rightarrow 1$ the situation becomes (relatively) symmetric.

The coefficient n applies because under a weakest-link structure the target government must successfully defend all possible targets from potential attacks. Hence, for the defender the maximum profitable expenditure that may be allocated equally to all targets is v_D/n . If $\alpha < 1$, the attacker is relatively disadvantaged, and the defender has the ability to profitably apply to all n targets a level of force that is greater than the level the attacker can profitably apply to a single target. Conversely, if $\alpha > 1$, then the attacker is relatively advantaged, and the defender does not have the ability to profitably apply to all n targets a level of force that is greater than what the attacker can profitably apply to a single target.

Given the structural asymmetries arising in the weakest-link network of targets, the normalized relative strength of the attacker identifies whether or not the defender has the ability to allocate more defensive forces to all n targets than the amount of effective force the attacker can allocate to any one target ($\alpha < 1$ and $\alpha > 1$, respectively). Furthermore, as the normalized relative strength of the attacker approaches unity we will refer to the conflict as being more symmetric, where this symmetry takes into account both the resource and structural asymmetries.

In the next section, we provide an equilibrium in our model in which the attacker creates an aura of uncertainty over the mode of attack, conventional and/or suicide, as well as the identity of the target to be attacked. Hence, the defender faces

¹⁷ We thank an anonymous referee for this example. Our assumption that conventional forces and suicide operatives are linearly additive, and neither experience decreasing returns to scale, is clearly restrictive. However, we consider it a reasonable starting point for examining the problem. An obvious direction of extension is to more general technologies.

strategic uncertainty over both the method of attack and the identity of the target to be attacked.

3. EQUILIBRIUM AND CHARACTERIZATION

Note that in our formulation: (i) force expenditures are sunk, (ii) force expenditures have a positive opportunity cost and (iii) the player who allocates the higher level of force to a target wins that target with certainty.¹⁸ Consequently, if one player wins with certainty, then the other player's best response is the strategy vector $\mathbf{0}$, which minimizes cost in a losing effort. Then, the winner will reduce the winning force arbitrarily close to zero to reduce cost as well. But then, $\mathbf{0}$ is no longer a best reply to this strategy. It clearly follows that there is no pure strategy equilibrium for this class of games.

Let \mathbf{x} denote a generic n -tuple of (effective) forces. For the defender, a mixed strategy (which we term a distribution of force for the defender) is an n -variate distribution function $P_D : \mathbb{R}_+^n \rightarrow [0, 1]$, where $P_D(\mathbf{x}) = \Pr\{d_i \leq x_i \text{ for all } i\}$ denotes the probability that each d_i in a random n -tuple \mathbf{d} drawn from the n -variate distribution function P_D is less than or equal to the corresponding x_i in the n -tuple $\mathbf{x} \in \mathbb{R}_+^n$. Note that the univariate marginal distribution of P_D for the i th target, $F_D^i(x_i) = \Pr\{d_i \leq x_i\}$, denotes the probability that at target i the level of force d_i is less than or equal to x_i .

For the attacker, a pure strategy is a $2n$ -tuple consisting of the n -tuple of the attacker's allocation of the continuous resource across the n targets and the n -tuple of the attacker's allocation of suicide operatives across the n targets. It follows directly that our focus on the attacker's effective force allocation does not place any restrictions on the correlation structures available to the attacker. To simplify the following expressions we will focus on the attacker's effective force allocation.

A mixed strategy for the attacker (which we term a distribution of effective force for the attacker) may thus be written as an n -variate distribution function $\hat{P}_A : \mathbb{R}_+^n \rightarrow [0, 1]$, where $\hat{P}_A(\mathbf{x}) = \Pr\{\hat{a}_i \leq x_i \text{ for all } i\}$ denotes the probability that the n -tuple of forces $\mathbf{x} \in \mathbb{R}_+^n$ successfully defends each and every target i from attack given that the attacker's effective allocation of force across the n targets, $\hat{\mathbf{a}}$, is a random n -tuple drawn from the n -variate distribution function \hat{P}_A .

Below, we examine an equilibrium for all parameter configurations in which neither the suicide attack technology is prohibitively costly for the attacker ($v_A \leq c$) nor the defender is so weak that suicide tactics are always suboptimal for the attacker ($\frac{v_D}{n} < c$). The remaining cases, as well as the proof of our main theorem, are included in appendix.

Recall that if there exists an integer λ such that $\lambda S \leq \hat{a}_i < (\lambda + 1)S$, then an optimizing attacker has implicitly chosen $s_i = \lambda$. In the analysis that follows it will also be helpful to define the following two functions, for $\mathbf{x}, \hat{\mathbf{x}} \in [0, (\lfloor \frac{v_A}{c} \rfloor + 1)S]^n$ and $\lambda = 0, \dots, \lfloor \frac{v_A}{c} \rfloor$

$$g(x_i) = \begin{cases} x_i - \lambda(S - c) & \text{if } \lambda S \leq x_i < \lambda S + c \\ (\lambda + 1)c & \text{if } \lambda S + c \leq x_i < (\lambda + 1)S \end{cases}$$

and

$$h(\hat{x}_i) = \begin{cases} \hat{x}_i & \text{if } \lambda S \leq \hat{x}_i < \lambda S + c \\ \lambda S + c & \text{if } \lambda S + c \leq \hat{x}_i < (\lambda + 1)S \end{cases}$$

To interpret $g(x_i)$ and $h(\hat{x}_i)$ note that it is suboptimal for a cost-minimizing attacker to allocate an effective force of $\hat{x}_i \in (\lambda S + c, (\lambda + 1)S)$ for any integer $\lambda = 0, \dots, \lfloor \frac{v_A}{c} \rfloor$ [doing so is strictly dominated by the effective force $\hat{x}_i = (\lambda + 1)S$]. For the defender, this suboptimal region corresponds to force allocations $x_i \in (\lambda S + c, (\lambda + 1)S)$ for any integer $\lambda = 0, \dots, \lfloor \frac{v_A}{c} \rfloor$. Over the set of attacker's cost-minimizing effective force levels and the corresponding defensive force levels [i.e., $\hat{x}_i, x_i \in [\lambda S, \lambda S + c]$ for any integer $\lambda = 0, \dots, \lfloor \frac{v_A}{c} \rfloor$], the functions $g(x_i)$ and $h(\hat{x}_i)$ identify the attacker's minimal cost for allocating an effective force equal to x_i units of defensive force and the defender's minimal cost for allocating force equal to \hat{x}_i units of effective attack force, respectively. For $\hat{x}_i, x_i \in (\lambda S + c, (\lambda + 1)S)$, the function $g(x_i)$ is completed by inserting the attacker's cost of effective force allocation at the upper endpoint of the interval, where $\lambda + 1$ suicide operatives are used, and the function $h(\hat{x}_i)$ is completed by inserting the defender's cost of force allocation at the lower endpoint of the interval.

THEOREM 1: A Nash equilibrium of the model of terrorism with suicide attack is for each player to allocate his forces as follows.

- a. If $\alpha < 1$, then for player D and $\mathbf{x} \in [0, v_A + \lfloor \frac{v_A}{c} \rfloor (S - c)]^n$,

$$P_D(\mathbf{x}) = \frac{\min_i \{g(x_i)\}}{v_A} \tag{1}$$

Similarly for player A and $\hat{\mathbf{x}} \in [0, v_A + \lfloor \frac{v_A}{c} \rfloor (S - c)]^n$,

$$\hat{P}_A(\hat{\mathbf{x}}) = 1 - \alpha + \frac{\sum_i h(\hat{x}_i)}{v_D} \tag{2}$$

The expected payoff for player A is 0, and the expected payoff for player D is $v_D(1 - \alpha)$.

- b. For $\alpha \geq 1$ and $c \leq (v_D/n)$, let $\bar{\lambda}$ be the largest nonnegative integer such that $\bar{\lambda}S < \frac{v_D}{n}$.

¹⁸ More formally, the conflict at each target utilizes the deterministic auction contest success function. See Baye et al. [5].

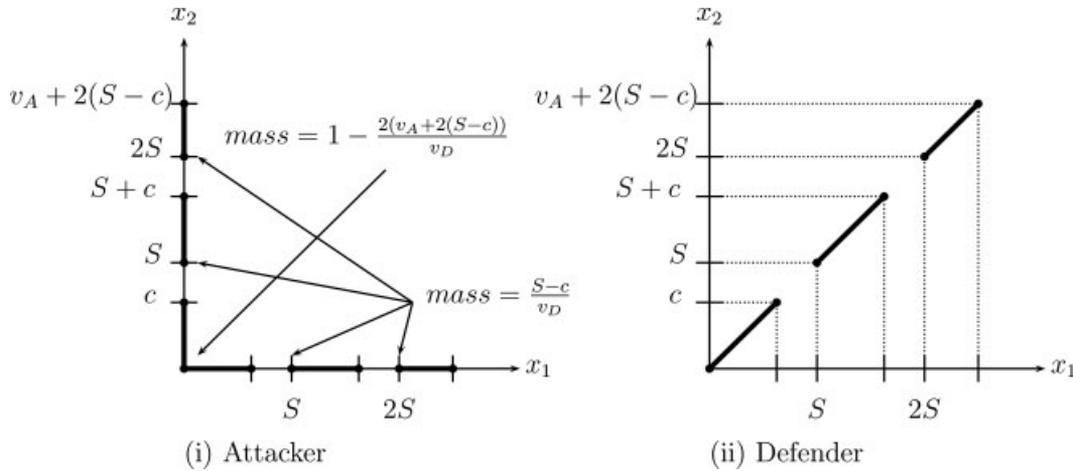


Figure 1. Supports of case (a) equilibrium joint distributions with $\lfloor \frac{v_A}{c} \rfloor$ and $n = 2$.

i. If $\bar{\lambda}S < \frac{v_D}{n} < \bar{\lambda}S + c$, then for player D and $\mathbf{x} \in [0, (v_D/n)]^n$,

$$P_D(\mathbf{x}) = 1 - \frac{(v_D/n) - \bar{\lambda}(S - c)}{v_A} + \frac{\min_i \{g(x_i)\}}{v_A} \quad (3)$$

Similarly for player A and $\hat{\mathbf{x}} \in [0, (v_D/n)]^n$,

$$\hat{P}_A(\hat{\mathbf{x}}) = \frac{\sum_i h(\hat{x}_i)}{v_D} \quad (4)$$

The expected payoff for player A is $v_A - (v_D/n) + \bar{\lambda}(S - c)$, and the expected payoff for player D is 0.

ii. If $\bar{\lambda}S + c \leq \frac{v_D}{n} \leq (\bar{\lambda} + 1)S$, then for player D and $\mathbf{x} \in [0, \bar{\lambda}S + c]^n$,

$$P_D(\mathbf{x}) = 1 - \frac{(\bar{\lambda} + 1)c}{v_A} + \frac{\min_i \{g(x_i)\}}{v_A} \quad (5)$$

Similarly for player A and $\hat{\mathbf{x}} \in [0, (\bar{\lambda} + 1)S]^n$,

$$\hat{P}_A(\hat{\mathbf{x}}) = \frac{\sum_i [\min \{h(\hat{x}_i), \frac{v_D}{n}\}]}{v_D} \quad (6)$$

The expected payoff for player A is $v_A - (\bar{\lambda} + 1)c$, and the expected payoff for player D is 0.

Figure 1 provides the supports of the equilibrium distributions of effective force for case (a) of Theorem 1 with $\lfloor \frac{v_A}{c} \rfloor = 2$ and only two targets in the weakest-link network

($n = 2$).¹⁹ In case (a), as in all cases, the attacker launches an attack on at most one target. Note also that in case (a), as in all cases, the defender's allocation of force has perfect positive correlation. One property of this correlation structure is that for any given level of force the probability that the attacker destroys at least one target is maximized if the attack is on a single target. As a result the attacker launches an attack on at most one target and may simultaneously use more than one suicide operative in the attack. When as in case (a) the normalized relative strength of the attacker is less than one, the attacker launches at most one attack and launches no attacks with probability $1 - \alpha$. Figure 2 provides the supports of the equilibrium distributions of effective force for subcase (i) of case (b) of Theorem 1 with $\lfloor \frac{v_A}{c} \rfloor = 2$ and only two targets in the weakest-link network ($n = 2$). In this case, the attacker launches exactly one attack with certainty. The equilibrium number of attacks is summarized in corollary 1.

COROLLARY 1: In cases (a) and (b) of Theorem 1, for any realization of his equilibrium strategy, the attacker attacks at most one target. In any case (b) realization the attacker launches an attack on exactly one target. In case (a), the attacker's equilibrium strategy attacks a single target with probability α , and launches no attacks with the remaining probability.

The proof of corollary 1 is contained in the proof of Theorem 1 given in appendix. It is important to note that our formulation of attack and defense features endogenous entry and force expenditure decisions and allows for the players to use general correlation structures for force expenditures

¹⁹ Recall that the support of an n -variate distribution function P , is the complement of the union of all open sets of \mathbb{R}^n with P -volume zero.

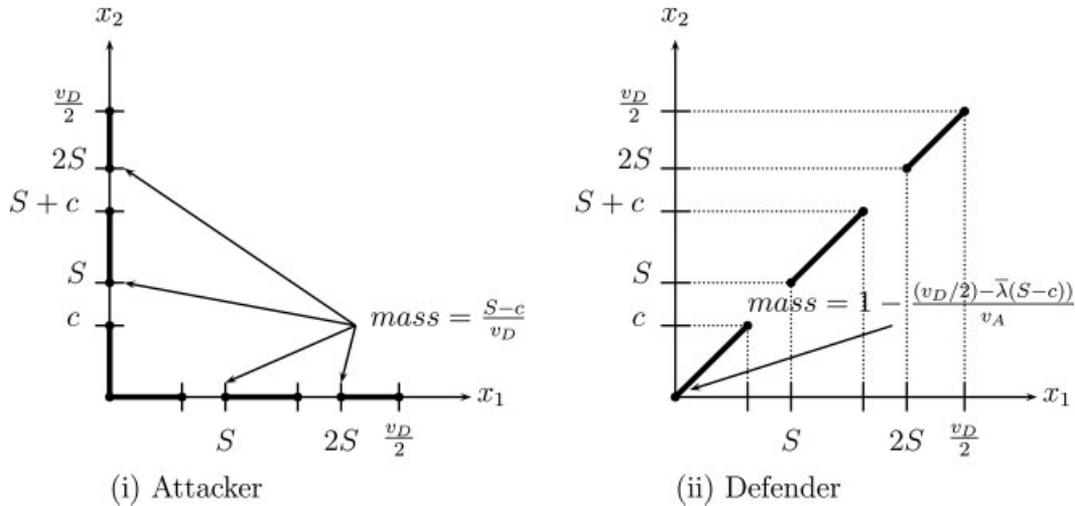


Figure 2. Supports of case (b), (i) equilibrium joint distributions with $\bar{\lambda} = 2$ and $n = 2$.

across the targets within the weakest-link network.²⁰ In contrast, much of the existing literature [e.g., Bier et al. [7] and Rosendorff and Sandler [43] among others] assumes that the number of terrorist attacks (which is usually set to one) is exogenously specified. Additionally, several of the existing models which allow for the attacker to endogenously choose the number of targets to attack²¹ obtain the paradoxical result, that even when (as in a weakest-link network) the attacker’s objective is to destroy a single target, the attacker optimally chooses to attack every target with certainty. Conversely, we find that the attacker optimally chooses to attack at most one target, but each target is chosen with positive probability.

Given the endogenous number of targets that are attacked in the equilibrium given in Theorem 1, we now examine (i) the probability of a suicide attack conditional on an attack being made and (ii) the expected number of suicide operatives that are utilized conditional on a suicide attack being launched. Recall that if $\lambda S \leq \hat{a}_i < (\lambda + 1)S$ then an optimizing attacker has implicitly set $s_i = \lambda$. Let \mathbf{c} and $\mathbf{0}$ denote the n -tuples (c, \dots, c) and $(0, \dots, 0)$, respectively. The conditional probability that the attacker launches at least one suicide attack is given by $(1 - \hat{P}_A(\mathbf{c})) / (1 - \hat{P}_A(\mathbf{0}))$, where $\hat{P}_A(\cdot)$ is player A’s distribution of effective force. Recall that $\hat{P}_A(\mathbf{c}) = \Pr\{\hat{a}_i \leq c \text{ for all } i\}$ is the probability that no attack exceeds level c — and therefore does not require suicide operatives — and $1 - \hat{P}_A(\mathbf{c})$ is the probability of at least one suicide attack. Similarly, $\hat{P}_A(\mathbf{0})$ is the probability of no attack and $1 - \hat{P}_A(\mathbf{0})$ is the probability that at least one attack is made. In case (a), the conditional probability of suicide

attack is $1 - (c / (v_A + \lfloor \frac{v_A}{c} \rfloor (S - c)))$. In case (b), the conditional probability of suicide attack is $1 - (nc / v_D)$. Although the upper bound of the number of equilibrium suicide operatives $\lfloor \frac{v_A}{c} \rfloor$ is not continuously differentiable with respect to v_A and c , it follows that — in cases (a) and (b) and for all marginal changes which hold $\lfloor \frac{v_A}{c} \rfloor$ constant — the conditional probability that the attacker utilizes a suicide attack is decreasing in the cost of a suicide operative, c .

Recalling that the normalized relative strength of the attacker is the relevant measure of the symmetry of the conflict, consider two simple symmetry increasing transformations corresponding to the attacker having a normalized relative strength advantage and disadvantage, respectively. The simple transformation for the case in which the attacker has a normalized relative strength disadvantage [case (a) of Theorem 1], which we term a “cost invariant increase in the attacker’s relative strength,” corresponds to an increase in the expression $v_A + \lfloor \frac{v_A}{c} \rfloor S$, where again we focus on marginal changes which hold $\lfloor \frac{v_A}{c} \rfloor$ constant.²² In case (a), any simple transformation of this form results in an increase in the normalized relative strength of the attacker which approaches one from below. The simple transformation for the case in which the attacker has a normalized relative strength advantage [case (b) of Theorem 1], which we term a “relative increase in the defender’s strength,” corresponds to an increase in the expression (v_D / n) . In case (b), a relative increase in the defender’s strength leads to a decrease in the normalized relative strength of the attacker which approaches one from above.

In case (a), the normalized relative strength of the attacker is less than one, and for all cost invariant increases in the

²⁰ See also Kovenock and Roberson [35].

²¹ Most closely related is Clark and Konrad [11] who, utilizing the Tullock contest success function, also examine a weakest-link network.

²² This restriction allows for all marginal changes such that S increases and/or v_A increases, subject to $\lfloor \frac{v_A}{c} \rfloor$ remaining constant.

attacker’s relative strength the conditional probability of suicide attack is increasing. Similarly, in case (b), the normalized relative strength of the attacker is greater than one, and for all relative increases in the defender’s strength the conditional probability of suicide attack is increasing. That is, the more symmetric the conflict the more likely the attacker is to utilize suicide operatives when an attack is launched. These properties of the conditional probability of suicide attack are summarized in corollary 2.

COROLLARY 2: In cases (a) and (b) of Theorem 1, the conditional probability that the attacker utilizes at least one suicide operative is: (i) decreasing with respect to the cost of suicide operatives, and (ii) increasing with respect to our two simple symmetry increasing transformations of the environment.

The characterization above indicates that although terrorist organizations attack at most one target, suicide operations are not the exclusive modus operandi even when such operatives are available; the terrorist leadership randomizes over conventional and suicide tactics. For example, al-Qaeda has been associated with conventional (non-suicide) events such as the Madrid train station bombing, the Al-Khobar massacres in Saudi Arabia, and 2007 attempted car bombings of Piccadilly Circus and Trafalgar Square. As noted by Sandler et al. [45], terrorists broaden their audience beyond the immediate victim by making their attacks and tactics appear to be random, so that everyone feels at risk. Furthermore, the use of suicide operatives is an increasing function of the relative symmetry of terrorists and target governments. This is a novel insight given that the literature on terrorism almost exclusively emphasizes the resource asymmetry between target governments and terrorists but does not incorporate alternative technologies of attack or defense. Our measure of symmetry, α , captures the potential for the weakest-link technology to balance resource disparities. In particular, symmetry within a weakest-link framework leads to an increased likelihood of suicide attack.

Moreover, terrorist organizations not only randomize over the use of suicide and conventional tactics but also the level of effective force. In the case of a suicide attack, this involves randomization over the number of suicide operatives that are utilized. Recalling that the probability that the attacker launches a suicide attack is $1 - \hat{P}_A(\mathbf{c})$, the case (a) expected number of suicide operatives conditional on a suicide attack being launched is

$$\frac{\sum_{i=1}^{\lfloor \frac{v_A}{c} \rfloor - 1} i (\hat{P}_A(i\mathbf{S} + \mathbf{c}) - \hat{P}_A((i - 1)\mathbf{S} + \mathbf{c})) + \lfloor \frac{v_A}{c} \rfloor (1 - \hat{P}_A(\lfloor \frac{v_A}{c} \rfloor - 1)\mathbf{S} + \mathbf{c})}{1 - \hat{P}_A(\mathbf{c})}$$

Table 1. Expected number of suicide operatives conditional on the launch of a suicide attack.

Case (a)	$\lfloor \frac{v_A}{c} \rfloor - \left(\frac{\lfloor \frac{v_A}{c} \rfloor (\lfloor \frac{v_A}{c} \rfloor - 1) S}{v_A + \lfloor \frac{v_A}{c} \rfloor (S - c) - c} \right)$
Case (b) (i)	$\bar{\lambda} - \left(\frac{\bar{\lambda} (\frac{\bar{\lambda} - 1}{n}) S}{\frac{v_D}{n} - c} \right)$
Case (b) (ii)	$(\bar{\lambda} + 1) - \left(\frac{\bar{\lambda} (\frac{\bar{\lambda} + 1}{n}) S}{\frac{v_D}{n} - c} \right)$

where again the bold notation $i\mathbf{S} + \mathbf{c}$ denotes the n -tuple $(iS + c, iS + c, \dots, iS + c)$ and the term $\hat{P}_A(i\mathbf{S} + \mathbf{c}) - \hat{P}_A((i - 1)\mathbf{S} + \mathbf{c})$ is the probability that the attacker allocates exactly i suicide operatives. In case (b) (i) [case (b) (ii)], the expected number of suicide operatives conditional on a suicide attack being launched is similarly calculated by replacing each $\lfloor \frac{v_A}{c} \rfloor$ in the above expression with $\bar{\lambda}$ [$(\bar{\lambda} + 1)$]. Table 1 provides the expected number of suicide operatives conditional on a suicide attack being launched in each of the three cases of Theorem 1.

As was the case with the conditional probability that the attacker launches a suicide attack, the expected number of suicide operatives conditional on a suicide attack being launched is decreasing with respect to the cost of each suicide operative. Furthermore, in case (a), the expected number of suicide operatives conditional on a suicide attack being launched is increasing for all cost invariant increases in the attacker’s relative strength, and in case (b) the expected number is increasing for all relative increases in the defender’s strength. That is, the expected number of suicide operatives conditional on a suicide attack being launched increases as the conflict becomes more symmetric, according to the normalized relative strength of the attacker. The properties of the expected number of suicide operatives conditional on a suicide attack being launched are summarized in Corollary 3.

COROLLARY 3: In cases (a) and (b) of Theorem 1, the expected number of suicide operatives conditional on the attacker launching a suicide attack is: (i) decreasing with respect to the cost of each suicide operative, and (ii) increasing with respect to the two simple symmetry increasing transformations of the environment.

As highlighted above, the level of symmetry in the conflict, which depends on both the characteristics of the players and those of the weakest-link network, is a pivotal determinant of the optimal attack and defense strategies. In particular, note that in case (a) the attacker launches at most one attack and launches no attacks with positive probability. However, the probability that the attacker launches an attack is weakly increasing as the normalized relative strength of the attacker

approaches unity (i.e., as the conflict becomes more symmetric). Thus, for both of the simple symmetry increasing transformations of the environment: (i) the probability of a terrorist event weakly increases, (ii) the conditional probability that such an event involves a suicide attack increases and (iii) the expected number of suicide operatives conditional on a suicide attack increases. Although the logic of this result is straightforward, this does complicate the conventional wisdom that an increase in the frequency and magnitude of terrorist attacks (of either the conventional or suicide variety) signals desperation on the part of a weakened terrorist organization. In particular, this popular characterization applies only in the case that the attacker has a normalized relative strength advantage. If the attacker is disadvantaged with respect to his normalized relative strength, then an increase in the frequency and magnitude of terrorist attacks signals that the terrorist has actually become relatively stronger and the conflict has become more symmetric.

4. CONCLUSION

In this article, we examine a model of terrorism which focuses not on the rationality of suicide operatives, but on the tradeoffs facing a terrorist organization that has the ability to utilize either or both suicide terrorism tactics and conventional tactics. A second feature of our focus is weakest-link networks of targets and the structural asymmetries between attack and defense. In this context, we find that the attacker endogenously launches at most one attack. The attacker randomizes over exclusively using a conventional attack and exclusively using a suicide attack each with positive probability. Conditional on an attack being launched, the probability of a suicide attack depends on both the structural asymmetry arising in the weakest-link network and asymmetry between the characteristics of both the attacker and the defender. Indeed, we show that the strategic implications of asymmetry between terrorists and target governments cannot be fully captured by differences in available resources but must also take into account the technologies of attack and defense. The availability of suicide operatives acting against a weakest-link defense can lead to a previously unrecognized symmetrization of conflict. As the conflict becomes more symmetric, suicide attacks are more likely to occur, and, conditional on a suicide attack being launched, the expected number of suicide operatives is increasing.

This article contributes to the analysis of the logic of suicide terrorism in finding that suicide operatives represent a discrete increase in terrorists' effective force that can symmetrize their conflict with target governments. This is particularly the case when governments are subject to a weakest-link defense technology (or definition of successful counterterrorism policy), as investigated here. Governments would do well by

deemphasizing the importance of an individual attack and continuing with everyday life, as is often the case in Europe, whereas US policy continues to be cast in terms of publicly emphasizing terrorists' success. Under such a policy change extensions to our model that recognize alternative technologies and/or multiple terror attacks may come into play.

APPENDIX

This appendix contains the proof of Theorem 1 and the statement of Theorem A.1, which provides an equilibrium in the remaining parameter configurations [i.e., the suicide attack technology is prohibitively costly for the attacker ($v_A \leq c$) or the defender is so weak that suicide tactics are suboptimal for the attacker ($\frac{v_D}{n} < c$)].

THEOREM A. 1: For the remaining parameter configurations a Nash equilibrium of the model of terrorism is for each player to allocate his forces as follows:

- (c) If $\alpha < 1$ and $v_A \leq c$, then for player D and $\mathbf{x} \in [0, v_A]^n$,

$$P_D(\mathbf{x}) = \frac{\min_i \{x_i\}}{v_A} \quad (7)$$

Similarly for player A and $\hat{\mathbf{x}} \in [0, v_A]^n$,

$$\hat{P}_A(\hat{\mathbf{x}}) = 1 - \frac{nv_A}{v_D} + \frac{\sum_i \hat{x}_i}{v_D} \quad (8)$$

The expected payoff for player A is 0, and the expected payoff for player D is $v_D - nv_A$.

- (d) If (i) $\alpha \geq 1$ and $(v_D/n) < c \leq S$, then for player D and $\mathbf{x} \in [0, (v_D/n)]^n$,

$$P_D(\mathbf{x}) = 1 - \frac{v_D}{nv_A} + \frac{\min_i \{x_i\}}{v_A} \quad (9)$$

Similarly for player A and $\hat{\mathbf{x}} \in [0, (v_D/n)]^n$,

$$\hat{P}_A(\hat{\mathbf{x}}) = \frac{\sum_i \hat{x}_i}{v_D} \quad (10)$$

The expected payoff for player A is $v_A - (v_D/n)$, and the expected payoff for player D is 0.

PROOF OF THEOREM 1: This proof, which is for case (a), shows that the pair of joint distribution functions P_D and \hat{P}_A form a Nash equilibrium in mixed strategies. In particular, we show that for each player each point in the support of their equilibrium n -variate distribution functions (stated in Theorem 1) results in the same expected payoff, and there are no profitable deviations from this support. The proofs of cases (b)–(d) follow along similar lines. \square

We begin with the support of each player's case (a) equilibrium distribution of force. For $y_k \leq z_k$ for all $k = 1, 2, \dots, n$, let $[\mathbf{y}, \mathbf{z}]$ denote the n -box $B = [y_1, z_1] \times [y_2, z_2] \times \dots \times [y_n, z_n]$, the Cartesian product of n closed intervals. The vertices of an n -box B are the points (v_1, v_2, \dots, v_n) where v_k is equal to y_k or z_k . Recall the following two definitions.

DEFINITION 4: Given an n -variate distribution P , the P -volume of the n -box $[\mathbf{y}, \mathbf{z}]$ is given by

$$V_P([\mathbf{y}, \mathbf{z}]) = \Delta_{y_n}^{z_n} \Delta_{y_{n-1}}^{z_{n-1}} \dots \Delta_{y_2}^{z_2} \Delta_{y_1}^{z_1} P(\mathbf{t})$$

where

$$\Delta_{y_k}^z P(\mathbf{t}) = P(t_1, \dots, t_{k-1}, z_k, t_{k+1}, \dots, t_n) - P(t_1, \dots, t_{k-1}, y_k, t_{k+1}, \dots, t_n)$$

DEFINITION 5: The support of an n -variate distribution function, P , is the complement of the union of all open sets of \mathbb{R}^n with P -volume zero.

Given Definitions 4 and 5, it is straightforward to show that in all feasible case (a) parameter configurations the support of player D's equilibrium distribution of force is uniformly distributed along the following set of line segments.²³ One line segment connects the origin with the point $\mathbf{c} \equiv (c, c, \dots, c)$. For $\mu = 1, \dots, \lfloor \frac{v_A}{c} \rfloor - 1$ there are also line segments connecting the points $\mu\mathbf{S}$ to the point $\mu\mathbf{S} + \mathbf{c}$. If $v_A - c\lfloor \frac{v_A}{c} \rfloor > 0$, then there is also a line segment that connects the point $\lfloor \frac{v_A}{c} \rfloor \mathbf{S}$ to the point $v_A + \lfloor \frac{v_A}{c} \rfloor (\mathbf{S} - \mathbf{c})$. Similarly, the support of player A's effective distribution of force consists of the combination of a set of mass points and mass uniformly distributed along a set of line segments both of which are located on the axes. One mass point of size $1 - \alpha$ is located at the origin. On each axis i there are $\lfloor \frac{v_A}{c} \rfloor$ mass points of size $\frac{S-c}{v_D}$ located at the points $\hat{x}_i = \mu S$ for $\mu = 1, \dots, \lfloor \frac{v_A}{c} \rfloor$. There is one line segment on each axis from the origin to the point $\hat{x}_i = c$. On each axis i and for $\mu = 1, \dots, \lfloor \frac{v_A}{c} \rfloor - 1$, there are also line segments from $\hat{x}_i = \mu S$ to $\hat{x}_i = \mu S + c$. If $v_A - c\lfloor \frac{v_A}{c} \rfloor > 0$, then there is also a line segment on each axis i from $\hat{x}_i = \lfloor \frac{v_A}{c} \rfloor S$ to $\hat{x}_i = v_A + \lfloor \frac{v_A}{c} \rfloor (S - c)$.

For each point in the support of player D's strategy, player D must have the same expected payoff. Let \mathcal{D} denote the set of n -tuples \mathbf{x} such that $\tilde{\mu}S \leq x_i \leq \min\{\tilde{\mu}S + c, v_A + \lfloor \frac{v_A}{c} \rfloor (S - c)\}$ for $\tilde{\mu} = 0, \dots, \lfloor \frac{v_A}{c} \rfloor$ and $i = 1, \dots, n$. Note that the support of player D's equilibrium strategy is a strict subset of \mathcal{D} .

If player A is using the equilibrium strategy \hat{P}_A given in (2), then the expected payoff to player D for any allocation of force $\mathbf{d} \in \mathbb{R}_+^n$ is

$$\pi_D(\mathbf{x}, \hat{P}_A) = v_D \hat{P}_A(\mathbf{d}) - \sum_i d_i. \tag{11}$$

From Eq. (2), the probability that with an allocation of \mathbf{d} player D wins every target i is

$$\hat{P}_A(\mathbf{d}) = 1 - \frac{n(v_A + \lfloor \frac{v_A}{c} \rfloor (S - c))}{v_D} + \frac{\sum_i h(d_i)}{v_D}. \tag{12}$$

Inserting Eqs. (12) into (11) and simplifying, the expected payoff to player D from any allocation $\mathbf{d} \in \mathcal{D}$, is $v_D - n(v_A + \lfloor \frac{v_A}{c} \rfloor (S - c))$. Thus, as the support of player D's equilibrium strategy is contained in \mathcal{D} , each point in the support of the n -variate distribution function P_D results in the same expected payoff.

To show that there are no profitable deviations from this support, note that for $\tilde{\mu} = 0, \dots, \lfloor \frac{v_A}{c} \rfloor$ it is clearly suboptimal for player D to allocate a level of force d_i to any target $i = 1, \dots, n$ such that $\min\{\tilde{\mu}S + c, v_A + \lfloor \frac{v_A}{c} \rfloor (S - c)\} < d_i < (\tilde{\mu} + 1)S$. In any such allocation, player D could decrease his cost without changing his probability of winning all of the targets by setting $d_i = \min\{\tilde{\mu}S + c, v_A + \lfloor \frac{v_A}{c} \rfloor (S - c)\}$. That is, it is suboptimal for player D to allocate a level of force above $v_A + \lfloor \frac{v_A}{c} \rfloor (S - c)$ or to allocate a level of force between $\tilde{\mu}S + c$ and $(\tilde{\mu} + 1)S$ to any target. However, this rules out all n -tuples in $\mathbb{R}_+^n - \mathcal{D}$ from being profitable deviations. As established above all n -tuples in \mathcal{D} yield the same expected payoff. Thus, for player D there are no profitable deviations from the distribution of force P_D given in (1).

The case of player A is similar. For each point in the support of player A's strategy, player A must have the same expected payoff. Note that the support

of \hat{P}_A consists of all effective force allocations $\hat{\mathbf{a}} \in \mathbb{R}_+^n$ such that there exists exactly one target i in which $\tilde{\mu}S \leq \hat{a}_i \leq \min\{\tilde{\mu}S + c, v_A + \lfloor \frac{v_A}{c} \rfloor (S - c)\}$ for $\tilde{\mu} = 0, \dots, \lfloor \frac{v_A}{c} \rfloor$ and $\hat{a}_{i'} = 0$ for all $i' \neq i$. Clearly, it is cost minimizing for the attacker to set $s_i = \tilde{\mu}$. Thus, $\hat{a}_i = a_i + \tilde{\mu}S$ and, it follows that $0 \leq a_i \leq c$ for the one target that receives a positive level of effective force.

Given that player D is using the equilibrium strategy P_D given in (1) the expected payoff to player A from an effective force allocation $\hat{\mathbf{a}}$ from the support of \hat{P}_A in which $\tilde{\mu}S \leq \hat{a}_i \leq \min\{\tilde{\mu}S + c, v_A + \lfloor \frac{v_A}{c} \rfloor (S - c)\}$ is

$$\pi_A(\hat{\mathbf{a}}, P_D) = v_A P_D\left(\hat{a}_i, \left\{v_A + \left\lfloor \frac{v_A}{c} \right\rfloor (S - c)\right\}_{i' \neq i}\right) - \tilde{\mu}c - a_i \tag{13}$$

where $P_D(\hat{a}_i, \{v_A + \lfloor \frac{v_A}{c} \rfloor (S - c)\}_{i' \neq i})$ is the probability that player A wins target i . Note that $P_D(\hat{a}_i, \{v_A + \lfloor \frac{v_A}{c} \rfloor (S - c)\}_{i' \neq i})$ is the univariate marginal distribution of P_D for the i th target, which we will henceforth denote as F_D^i . From Eq. (1), it follows that for any effective force allocation in the support of \hat{P}_A player A's expected payoff is

$$\pi_A(\hat{\mathbf{a}}, P_D) = v_A \left(\frac{g(\hat{a}_i)}{v_A} \right) - \tilde{\mu}c - a_i = 0$$

as $g(\hat{a}_i) = a_i + \tilde{\mu}c$ for all such points.

We now show that there are no profitable deviations from the support of player A's equilibrium joint distribution. Note that if player A attacks only one target i , then it is clearly suboptimal for player A to allocate a level of effective force \hat{a}_i such that $\min\{\tilde{\mu}S + c, v_A + \lfloor \frac{v_A}{c} \rfloor (S - c)\} < \hat{a}_i < (\tilde{\mu} + 1)S$. That is, it is clearly strictly dominated for player A to allocate an effective level of force above $v_A + \lfloor \frac{v_A}{c} \rfloor (S - c)$ or to allocate an effective level of force strictly between $\tilde{\mu}S + c$ and $(\tilde{\mu} + 1)S$ to target i . The only remaining possible deviation from the support is for player A to allocate a strictly positive level of effective force to two or more targets.

The probability that player A wins both targets i and i' is given by the bivariate marginal distribution $P_D(\hat{a}_i, \hat{a}_{i'}, \{v_A + \lfloor \frac{v_A}{c} \rfloor (S - c)\}_{i'' \neq i, i'})$, which we will denote as $P_D^{i, i'}(\hat{a}_i, \hat{a}_{i'})$. The expected payoff to player A for any allocation of force $\hat{\mathbf{a}} \in \mathbb{R}_+^n$ which allocates a strictly positive level of force to two targets i, i' is

$$\pi_A(\hat{\mathbf{a}}, P_D) = v_A F_D^i(\hat{a}_i) + v_A F_D^{i'}(\hat{a}_{i'}) - v_A P_D^{i, i'}(\hat{a}_i, \hat{a}_{i'}) - (a_i + cs_i) - (a_{i'} + cs_{i'}).$$

Simplifying,

$$\pi_A(\hat{\mathbf{a}}, P_D) \leq -v_A P_D^{i, i'}(\hat{a}_i, \hat{a}_{i'}) < 0$$

where the left-hand weak inequality holds with equality if for $k = i, i'$ there exist $\tilde{\mu}_k \in [0, \dots, \lfloor \frac{v_A}{c} \rfloor]$ such that $\hat{a}_k \in [\tilde{\mu}_k S, \min\{\tilde{\mu}_k S + c, v_A + \lfloor \frac{v_A}{c} \rfloor (S - c)\}]$ and $\hat{a}_k = a_k + \tilde{\mu}_k S$. Furthermore, $P_D^{i, i'}(\hat{a}_i, \hat{a}_{i'}) > 0$ as $\hat{a}_i, \hat{a}_{i'} > 0$, and thus it is unprofitable for player A to allocate a strictly positive level of effective force to two targets.

The case of player A allocating a strictly positive level of force to more than two targets follows directly. Clearly, in any optimal strategy player A never allocates a strictly positive level of force to more than one target. This concludes the proof that in case (1) the pair of joint distribution functions P_D and \hat{P}_A constitute a Nash equilibrium of the model of terrorism with suicide attack. The proofs of cases (b)–(d) follow a similar line of argument.

REFERENCES

- [1] J. Alcalde and M. Dahm, Rent seeking and rent dissipation: A neutrality result, *J Public Econ* 94 (2010), 1–7.
- [2] D. Arce and T. Sandler, Counterterrorism: A game theoretic approach, *J Conflict Resol* 49 (2005), 183–200.

²³ Figure 1 shows that for $\lfloor \frac{v_A}{c} \rfloor = 2$ and $n = 2$ the support of player D's distribution of force P_D is uniformly distributed along the three shaded line segments.

- [3] S. Atran, Genesis of suicide terrorism, *Science* 299 (2003), 1534–1539.
- [4] M.R. Baye, D. Kovenock, and C.G. de Vries, The solution to the tullock rent-seeking game when $R > 2$: Mixed-strategy equilibria and mean dissipation rates, *Public Choice* 81 (1994), 363–380.
- [5] M.R. Baye, D. Kovenock, and C.G. de Vries, The all-pay auction with complete information, *Econ Theory* 8 (1996), 291–305.
- [6] D. Bernhardt and M.K. Polborn, Non-convexities and the gains from concealing defense from committed terrorists, *Econ Lett* 107 (2010), 52–54.
- [7] V.M. Bier, S. Oliveros, and L. Samuelson, Choosing what to protect: Strategic defensive allocation against an unknown attacker, *J Public Economic Theory* 9 (2007), 563–587.
- [8] D.W. Blackett, Blotto-type games, in: *Proceedings Fourth Annual Logistics Conference (Part II-Restricted Session)*, General Services Administration, Washington D.C., 1953, pp. 31–41.
- [9] E. Borel, La théorie du jeu les équations intégrales à noyau symétrique, *Comptes Rendus de l'Académie* 173 (1921), 1304–1308. (English translation by L.J. Savage, 1953. The theory of play and integral equations with skew symmetric kernels. *Econometrica* 21, 97–100).
- [10] G. Brown, M. Carlyle, J. Salmerón, and K. Wood, Defending Critical Infrastructure, *Interfaces* 36 (2006), 530–544.
- [11] D.J. Clark and K.A. Konrad, Asymmetric conflict: Weakest link against Best Shot, *J Confl Resol* 51 (2007), 457–469.
- [12] R. Cornes, Dyke maintenance and other stories: Some neglected types of public goods, *Q J Econ* 108 (1993), 259–271.
- [13] M. Crenshaw, Explaining suicide terrorism: A review essay, *Secur Stud* 16 (2007), 133–162.
- [14] Economist, Special report: Terrorism, *The Economist*, 19 July 2008, pp. 8595.
- [15] J.S. Feinstein and E.H. Kaplan, Analysis of a strategic terror organization, *J Confl Resol* 54 (2010), 281–302.
- [16] R.E. Franck and F. Melese, Exploring the structure of terrorists' WMD decisions: A game theory approach, *Defense Secur Anal* 20 (2004), 355–372.
- [17] M. Gassebner, R. Jong-A-Pin, and J.O. Mierau, Terrorism and electoral accountability: One strike, you're out!, *Econ Lett* 100 (2008), 125–129.
- [18] M. Gassebner, R. Jong-A-Pin, and J.O. Mierau, Terrorism and cabinet duration, *Int Econ Rev* 52 (2011), 1253–1270.
- [19] R. Golman and S.E. Page, General blotto: Games of strategic allocative mismatch, *Publ Choice* 138 (2009), 279–299.
- [20] O. Gross, The symmetric blotto game, RAND Corporation RM-424, 1951.
- [21] O. Gross and R. Wagner, A continuous colonel blotto game, RAND Corporation RM-408, 1950.
- [22] S. Hart, Discrete Colonel Blotto and General Lotto Games, *International Journal of Game Theory* 36 (2008), 441–460.
- [23] K. Hausken, Probabilistic risk analysis and game theory, *Risk Anal* 22 (2002), 17–27.
- [24] G. Heal and H. Kunreuther, IDS models of airline security, *J Confl Resol* 49 (2005), 201–217.
- [25] J. Hirshleifer, From weakest-link to best-shot: The voluntary provision of public goods, *Publ Choice* 41 (1983), 371–386.
- [26] B. Hoffman and G.H. McCormick, Terrorism, signaling, and suicide attack, *Stud Confl Terrorism* 27 (2004), 243–281.
- [27] R. Hortala-Vallve and A. Llorente-Saguer, A simple mechanism for resolving conflict, *Games Econ Behav* 70 (2010), 375–391.
- [28] R. Hortala-Vallve and A. Llorente-Saguer, Pure-strategy nash equilibria in non-zero sum colonel blotto games, *Int J Game Theory* 40 (2012), 331–343.
- [29] L.R. Iannaccone, The market for martyrs, *Interdiscipl J Res Religion* 2 (2006), Article 4.
- [30] D. Jacobson and E.H. Kaplan, Suicide bombings and targeted killings in (counter-) terror games, *J Confl Res* 51 (2007), 772–792.
- [31] E.H. Kaplan, M. Kress, and R. Szechtman, Confronting entrenched insurgents, *Oper Res* 58 (2010), 329–341.
- [32] P. King, Editorial: The mechanics of terrorism, *NATO Review*, Volume 1, 2008.
- [33] T. Klumpp and M.K. Polborn, Primaries and the New Hampshire effect, *J Publ Econ* 90 (2006), 1073–1114.
- [34] K.A. Konrad and D. Kovenock, Multi-battle contests, *Games Econ Behav* 66 (2009), 256–274.
- [35] D. Kovenock and B. Roberson, Terrorism and the optimal defense of networks of targets, 2010a.
- [36] D. Kovenock and B. Roberson, "Conflicts with multiple battlefields," *Oxford handbook of the economics of peace and conflict*, M.R. Garfinkel and S. Skaperdas (Editors), Oxford University Press, Oxford, 2010b.
- [37] M. Kress, The effect of crowd density on the expected number of casualties in a suicide attack, *Nav Res Logisit* 52 (2005), 22–29.
- [38] D. Kvasov, Contests with limited resources, *J Econ Theory* 136 (2007), 738–748.
- [39] S. Macdonell and N. Mastronardi, Colonel blotto equilibria: A complete characterization in the two battlefield case, University of Texas, Mimeo, Austin, Texas, 2010.
- [40] B. Roberson, The Colonel Blotto game, *Econ Theory* 29 (2006), 1–24.
- [41] B. Roberson and K. Dmitriy, The non-constant-sum Colonel Blotto game, *Econ Theory*, 10.1007/s00199-011-0673-z, 2011.
- [42] T. Rosenbaum, The president has kept us safe, *Wall Street J* 251 (2008), A15.
- [43] B.P. Rosendorff and T. Sandler, Too much of a good thing? The proactive response dilemma, *J Confl Resol* 48 (2004), 657–671.
- [44] M. Sageman, *Understanding terror networks*, University of Pennsylvania Press, Philadelphia, 2004.
- [45] T. Sandler, D. Arce, and W. Enders, Transnational terrorism, copenhagen consensus, 2008 challenge paper. Available at: www.copenhagenconsensus.com/Default.aspx?ID=1152, 2008. Accessed on July 17, 2012.
- [46] M. Shubik and R.J. Weber, System defense games: Colonel Blotto, command and control, *Nav Res Logistics Q* 28 (1981), 281–287.
- [47] J.M. Snyder, Election goals and the allocation of campaign resources, *Econometrica* 57 (1989), 637–660.
- [48] B. Szentes and R.W. Rosenthal, Three-object two-bidder simultaneous auctions: Chopsticks and tetrahedra, *Games Econ Behav* 44 (2003a), 114–133.
- [49] B. Szentes and R.W. Rosenthal, Beyond chopsticks: Symmetric equilibria in majority auction games, *Games Econ Behav* 45 (2003b), 278–295.
- [50] P. Tang, Y. Shoham, and F. Lin, Designing competitions between teams of individuals, *Artif Intell* 174 (2010), 749–766.
- [51] US Congress 2002, Joint Inquiry into Intelligence Community Activities Before and After the Terrorist Attacks of September 11, 2001, Washington, DC: Senate Report No. 107-351, House Report No. 107-792.
- [52] J. Weinstein, Two notes on the Blotto game, Northwestern University, Mimeo, Evanston Illinois, 2005.