

Strategic entry deterrence and terrorism: Theory and experimental evidence

John Cadigan · Pamela M. Schmitt

Received: 27 December 2007 / Accepted: 14 July 2009
© Springer Science+Business Media, LLC 2009

Abstract Using a two stage rent-seeking framework, we present a simple model of strategic entry/terrorism deterrence and test the model using laboratory experiments. Our contest success function highlights the potential for strategic spillovers. The theory illustrates that, relative to a cooperative outcome, negative externalities lead to over-spending on deterrence and positive externalities lead to under-spending on deterrence. Our experimental results are broadly consistent; subjects in the negative externality treatment had higher expenditures. In contrast to theoretical predictions, participation decisions, while primarily driven by the probability of winning a contest, were influenced by a subject's ability to participate in multiple contests.

Keywords Terrorism · Rent-seeking · Experiments · Strategic entry deterrence

JEL Classification C70 · C91 · D72 · D74

1 Introduction

Several studies of terrorism utilize game theoretic techniques to model the decision making of governments and terror groups (Enders and Sandler 1995; Arce and Sandler 2005; Ferrero 2006; Rosendorff and Sandler 2004; Sandler and Arce 2003; Bueno de Mesquita 2005a, 2005b, 2007; and Bapat 2006), and a growing empirical literature employs sophisticated time series econometric methods to examine whether the structure, timing, and type of terrorist events has changed over time (Enders and Sandler 1993, 2005; Enders et al. 1992; Sandler and Enders 2004). As befits a game theoretic approach, the models specify objective functions for the players involved that highlight the strategic interdependence of their actions. Empirical evaluation of the models is complicated by the lack of naturally occurring field data, in particular regarding the benefits and costs of various actions as perceived by those involved. While difficult to measure in the field, these parameters can be precisely

J. Cadigan · P.M. Schmitt (✉)
Department of Economics, U.S. Naval Academy, 589 McNair Road, Annapolis, MD 21402, USA
e-mail: pschmitt@usna.edu

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE DEC 2007		2. REPORT TYPE		3. DATES COVERED 00-00-2007 to 00-00-2007	
4. TITLE AND SUBTITLE Strategic entry deterrence and terrorism: Theory and experimental evidence				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) United States Naval Academy (USNA), Department of Economics, Annapolis, MD, 21402				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT see report					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 20	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

controlled in the lab. Motivated by recent research on terrorism, we present a simple conflict model and utilize laboratory experimental methods to evaluate its equilibrium predictions. In our view, laboratory experiments provide an important and underutilized method to examine game theoretic models of terrorism, and can serve as a complement to existing approaches based on field data.¹

In particular, our theoretical model highlights the spillover effects associated with government expenditures to deter terrorism, which can create either positive or negative externalities. Consider, for example, US expenditures on the military conflict in Afghanistan that followed the terrorist attacks on 9–11. Damaging Al Qaeda's infrastructure and personnel decreased the capacity of the terror network to attack the United States but also benefited other countries that may have been potential targets. As another example, expenditures on intelligence gathering provide positive spillovers to other countries who share the information. Alternatively, negative externalities are likely to emerge when terror groups choose methods and locations for their attacks that are responsive to the probability of success, risk of capture, benefits, and costs of the attacks.² In this case, a country's expenditures to 'harden' targets may displace an attack to another location. This negative externality may lead to over-deterrence if countries 'race' with one another, increasing expenditures in an effort to avoid presenting the most vulnerable target.

Borrowing from the rent seeking framework developed by Tullock (1967, 1980), we examine a two stage game in which first movers determine the probabilities associated with various contest outcomes and second movers decide whether to participate in the contests. In a terrorism context, the first movers can be interpreted as governments choosing defense expenditures to influence the chance a terror group's attack is successful, with the second movers being the terror groups deciding whether to attack.³ The two key distinctions in our theoretical model are whether spending on deterrence has positive or negative externalities and whether terrorists can attack all, or only a subset, of the targets. The ideology of some terror groups limits their set of potential targets (such as Tamil Tigers and the Irish Republican Army), while other groups have multiple potential targets (such as Al Qaeda and Hamas).

One variant of the model utilizes a contest success function where the probability associated with each contest outcome is determined by the aggregate spending of all first movers. This induces a collective action problem similar to that associated with the provision of public goods that leads to free riding and under-spending relative to the cooperative outcome. While we keep the model simple to ease implementation in a laboratory environment, our intent is to capture the free riding incentives that characterize the global collective action problem associated with combating terror networks. In a second variant of the model, the probabilities associated with the contest outcomes are determined independently by the expenditures of each first mover. Importantly, the negative externalities associated with the race to avoid being the low spender among first movers only materialize in connection with

¹An interesting example is provided by Bueno de Mesquita (2006) who uses a case study to compare violence, concessions, and counter-terrorism policy with predictions from a theoretical model.

²See, for examples, Landes (1978) and Enders and Sandler (1993), who find that the installation of metal detectors in airports reduced skyjackings but increased other kinds of attacks including barricade missions and kidnappings.

³In an industrial organizational context, the choice to protect market share for one firm in a duopoly market lowers the probability of a successful entry, creating a positive externality for the other firm in the duopoly. Additionally, the choice to protect market share for one firm in a duopoly market increases the probability of successful entry into a different market, creating a negative externality for other markets.

the assumption that the terror group can attack at most a subset of its potential targets.⁴ We view this as an appropriate and important restriction; while the Al Qaeda network has launched attacks against multiple targets simultaneously, their resources are not unlimited. Our experimental results indicate that the presence of multiple targets, in connection with this restriction, has important behavioral effects that are not captured by standard game-theoretic models.

Our modeling approach also has structural similarities to strategic entry deterrence games in industrial organization where the first movers are firms trying to protect their market shares (by using advertising, reputation, patents or product variety) and the second mover is a potential entrant.⁵ The use of laboratory experimental methods connects our research to a limited experimental literature on entry deterrence (Jung et al. 1994; Mason and Nowell 1998; Brandts et al. 2005). Jung et al. (1994) examine the decisions of a monopolist that can adopt a ‘strong’ or ‘weak’ pricing strategy in the presence of entry by a series of subjects. For a large share of subjects playing the role of the monopolist, the weak strategy maximizes round payoffs in the presence of entry but, by signaling the player as a weak type, may generate additional entrants in subsequent rounds that lower aggregate payoffs. They find that a significant fraction of weak monopolists deter entry by playing ‘strong’ in the early rounds of the experiment. While this strategy is effective in the early rounds, they observed an increase in the proportion of subjects choosing to enter against the monopolist in later rounds. Mason and Nowell (1998) examine a two stage, two person non-cooperative game in which the sub-game perfect equilibrium (SPE) has ‘incumbent’ types choose a level of output sufficient to make a positive output choice by a potential entrant yield negative payoffs. They find that while a significant number of incumbent types followed the SPE, a number of entrants choose positive output levels when it was not optimal to do so. This behavior persisted through the final rounds of the experiment, although at lower frequencies than early in the experiment. Brandts et al. (2005) examine a multistage environment in which an incumbent strategically can ‘pre-install’ production capacity and a potential entrant makes an investment decision. In separate treatments, the potential entrant either had or did not have the opportunity to pre-install after the incumbent made its pre-installation decisions. The two treatments yield differing equilibrium predictions, with one having the incumbent deter entry and another (based on a forward induction argument) in which the entrant’s ability to pre-install mitigates the incumbent’s first mover advantage. They find significant levels of deterrence by the incumbent in both treatments, concluding that players perceive a first mover advantage even without the first mover having to pre-commit to capacity.

Our experiments differ from these in several respects. Because we focus our motivation on models of terrorism, we allow potential entrants to choose to enter a contest *or* contests based on their probability of success. Depending on the form of the contest success function, this either induces competition between the first movers to avoid being chosen as a target or induces a collective action problem among the first movers (when the contest success function focuses on aggregate effort choices). Our results indicate that negative externalities associated with the competition to avoid being a low spender are mitigated by the presence of multiple contests in which a potential entrant can participate. Furthermore, while spending on deterrence in the positive externality treatments was lower than the baseline or negative

⁴We discuss this further in the modeling section. If the terror group can attack all potential targets, each contest is essentially independent in a strategic sense. Strategic interdependence among the first movers is achieved when the terror group is forced to pick among a subset of potential targets.

⁵See, for examples, Dixit (1980), Kreps and Wilson (1982), and Milgrom and Roberts (1982).

externality treatments this was not a result of complete free riding by many of the subjects in our experiments. Also, while the probability of winning was a primary determinant for entry decisions, subject participation decisions were also heavily influenced by the number of contests available.

While our approach is linked to the study of strategic entry deterrence games, we focus our motivation on terrorism prevention. Although our game theoretic predictions and laboratory results are stylized, we believe that isolating specific incentives, spillovers, or the number of targets is a beneficial starting point for merging behavioral and experimental economics to the terrorism literature.⁶ Our results illustrate behavioral implications only when strategic spillovers are present, ignoring any impact on how government and terrorist interactions may change either player's future behavior.⁷ Therefore we do not provide policy implications because we acknowledge that our one-shot laboratory results may not be robust to the dynamic aspects of the real world.

Finally, we note that in all of our treatments, subjects 'over-invested' relative to the SPE prediction. Overinvestment is a well established result in the literature covering experiments on rent seeking (see for examples, Millner and Pratt 1989, 1991; Shogren and Baik 1991; Davis and Reilly 1998; Potters et al. 1998, and Öncüler and Croson 1998), and, as we utilize a modified rent seeking approach, this result was not surprising. In the experimental literature on rent seeking, our approach is most closely connected to the experiments reported in Cadigan (2007). Similar to the results presented below, using a two stage structure with a contest success function based on aggregate effort choices, he finds little evidence of free riding, concluding that team oriented considerations or concerns for others' payoffs may play an important role in subject decision making. The remainder of the paper is organized as follows: Sect. 2 presents the model and theoretical results, Sect. 3 details the experimental design and procedures, Sect. 4 examines the experimental results, and Sect. 5 concludes.

2 The model

2.1 Baseline model

We begin with the simple case in which two players, X and Y , compete for a contest prize, B , in the following fashion. In the first stage, player X chooses an effort level (x) that determines the probability associated with the contest outcome. The marginal cost of effort for player X is a constant, C . In the second stage, player Y chooses whether to pay a non-refundable participation fee (F) to enter the contest. If player Y does not enter the contest, the prize is awarded to player X . If player Y enters, a probabilistic draw is held to determine which player is awarded the contest prize. Specifically, let contest success function be given by:

$$P_y = \frac{1}{1+x},$$

⁶In terms of our modeling approach, the work most relevant to ours is Sandler (2005).

⁷For example, other studies have examined the impact of negotiating with terrorists, as in Bapat (2006), how promised concessions or replacement of terrorist leadership may impact terrorists' behavior, as in Bueno de Mesquita (2005b), or how election processes impact the behavior of governments as in Bueno de Mesquita (2007).

where P_y is the probability player Y wins the contest. Entering the contest is a best response for player Y if:

$$\left(\frac{1}{1+x}\right)B - F > 0,$$

which occurs when

$$x < \frac{B}{F} - 1.$$

Thus, player X can deter player Y from participating by exerting effort $x \geq \frac{B}{F} - 1$ in the first stage of the contest. If player X anticipates participation by player Y , x is chosen to maximize:

$$\left(\frac{x}{1+x}\right)B - Cx.$$

This leads to optimal expenditures $x^* = \sqrt{\frac{B}{C}} - 1$, which is less than the amount required for deterrence if $B > \frac{F^2}{C}$. As such, when $B > \frac{F^2}{C}$, the sub-game perfect Nash equilibrium has player X exert effort $x^* = \sqrt{\frac{B}{C}} - 1$, and player Y participate in the contest. Otherwise, the equilibrium has player X choose effort $x = \frac{B}{F} - 1$, and player Y choose not to participate.

2.2 Negative externalities in effort choice: Adding a second contest

Next, consider the effects of adding a second X -type player who chooses an effort level to determine the probability associated with a second contest outcome. Equilibrium predictions in this case depend on whether the Y player is constrained to participate in at most one contest. As Enders and Sandler (1995) argue in a terrorism context, if a terror group chooses to attack one target based on the likelihood an attack is successful, potential targets may increase their expenditures beyond the socially optimal amount. In this case deterrence expenditures, in addition to lowering the probability that a particular target is chosen, generate negative externalities by raising the threat to other potential targets. From a game theoretic perspective, this argument depends critically on the restriction that the terror group attacks a single target.

Case 1: No participation constraint

Assume, for example, that there are two type X players (X_1 and X_2) who choose effort levels x_1 and x_2 , and let the contest success functions (one for each contest) be:

$$P_{y_i} = \frac{1}{1+x_i}; \quad i = (1, 2),$$

where P_{y_i} is the probability player Y wins contest i . After observing the effort choices of both type X players, and thus the probabilities associated with each contest outcome, player Y decides whether to take part in the contests. Specifically, player Y chooses (for each of the contests) whether to pay a fee, F , to participate. This means that player Y can choose to participate in a contest with X_1 , with X_2 , both, or neither. Participation in each contest is associated with a separate fee. Note that player Y 's ability to participate in either of the contests, and the lack of any spillover effects between the contests, leads each of the X -type players to act as if they were facing player Y alone. In other words, because each contest is essentially independent, there are no changes to the equilibrium predictions described

in the baseline model. When $B > \frac{F^2}{C}$, the sub-game perfect Nash equilibrium has both X players exert effort $x^* = \sqrt{\frac{B}{C}} - 1$, and player Y participates in both contests. Otherwise, the equilibrium has both X players choose effort $x = \frac{B}{F} - 1$, and player Y chooses not to participate in either contest. Importantly, the equilibrium prediction does not depend on whether the type X players act cooperatively or non-cooperatively because the independence of the contests implies that the outcome that maximizes the sum of player X earnings (the cooperative outcome) also maximizes each player X 's earnings.

Case 2: Participation constraint

We now compare Case 1 with the outcome under the restriction that player Y can participate in only one contest. The best response function for player Y can be summarized as:

Participate in Contest 1 if: $x_1 < x_2$ and $x_1 < \frac{B}{F} - 1$,

Participate in Contest 2 if: $x_2 < x_1$ and $x_2 < \frac{B}{F} - 1$,

Randomize participation in Contest 1 and Contest 2 if $x_1 = x_2 < \frac{B}{F} - 1$, and

Do not participate in either contest if $\frac{B}{F} - 1 < \text{Min}(x_1, x_2)$.

This means that player Y is deterred from entering either contest if both X -type players choose effort $\frac{B}{F} - 1$. If stage 1 effort levels do not deter participation, player Y either enters against the type X player exerting the lowest effort or (in the case of matching effort levels) participates in each contest with probability $1/2$. If the type X players anticipate participation by player Y , a cooperative solution in which the type X players match effort levels exists.⁸ Aggregate payoffs for the X -type players can be written as:

$$B + \frac{1}{2} \frac{x_1}{1+x_1} B + \frac{1}{2} \frac{x_2}{1+x_2} B - Cx_1 - Cx_2.$$

Because player Y can participate in at most one contest, the player X s are guaranteed at least one prize, which accounts for the initial B in the expression. The next two terms reflect the randomization associated with player Y 's participation decision, multiplying the probability of participation ($1/2$) by the expected prize associated with the contest. The final terms represent the aggregate effort costs. For the symmetric outcome ($x_1 = x_2 = x$) this can be rewritten as:

$$B + \frac{1}{2} \frac{2x}{1+x} B - 2Cx.$$

The optimal effort choice is given by:

$$x_1 = x_2 = x^* = \sqrt{\frac{B}{2C}} - 1.$$

Therefore, relative to Case 1, less effort is needed by each player X to deter entry by player Y when $B > \frac{F^2}{2C}$. In Case 2, the cooperative outcome (which yields the highest group payoff for the type X players) has $x_1 = x_2 = x^* = \sqrt{\frac{B}{2C}} - 1$, and player Y participates in one of the contests (chosen at random). Otherwise, the equilibrium has both X players choose effort

⁸ Given the symmetry of the model, it is straightforward to verify that the cooperative outcome maximizing the sum of type X player payoffs has identical effort choices.

$x = \frac{B}{F} - 1$, and player Y chooses not to participate in either contest. Relative to Case 1, in which player Y can enter both contests, the cooperative equilibrium has lower expenditures by both of the type X players, because the participation constraint (by lowering the probability player Y enters a contest) lowers the marginal benefit associated with exerting effort.

Importantly, while this outcome maximizes the joint payoff of the type X players, it is not a non-cooperative Nash equilibrium, and this gets at the heart of the externality argument. For any effort choice by X_2 that is less than the amount required for deterrence ($x_2 < \frac{B}{F} - 1$), X_1 could earn a higher payoff by setting $x_1 = x_2 + \varepsilon$. Specifically, X_1 prefers to spend $x_2 + \varepsilon$ to matching if:

$$B - C(x_2 + \varepsilon) > \frac{1}{2}B + \frac{1}{2} \frac{x_2}{1 + x_2} B - Cx_2.$$

This can be rewritten as:

$$B \left(\frac{1}{2} - \frac{1}{2} \frac{x_2}{1 + x_2} \right) > C\varepsilon.$$

The term on the left-hand side reflects the increased probability associated with winning the prize by outspending X_2 (which leads Y to enter the contest with X_2 rather than randomize), while the term on the right-hand side reflects the increased effort cost. The incentive to outspend the other X player might lead both players to spend the amount required for deterrence (which would occur, for example, if $B < \frac{F^2}{C}$). Alternatively, if player X effort was capped at a maximum of $\sqrt{\frac{B}{C}} - 1$, spending would converge to this level. In any event, competition between the type X players (reminiscent of the typical Bertrand competition model) leads to overspending on deterrence relative to the social optimum.

2.3 Positive externalities: Changing the contest success function

While competition may lead to negative externalities in some environments, there is also the possibility that effort choices may generate spillover benefits. In the current framework, assume that X_1 and X_2 make effort choices to influence the following contest success function:

$$P_y = \frac{1}{1 + x_1 + x_2},$$

where P_y denotes the (common) probability that player Y wins a contest against X_1 or X_2 . In this case, exertion of effort generates a positive externality in the sense that it lowers the chance that player Y is successful against both type X players, and not just the player exerting effort. Given the common probability of success in the contests, player Y either chooses to participate in both contests (when possible) or in neither contest.

Case 1: No participation constraint

Recall that no participation constraint means that player Y can choose to participate in a contest with X_1 , with X_2 , both, or neither; participation in each contest is associated with a separate fee. Here, player Y chooses to pay the participation fees (one for each contest) if:

$$\sum_i x_i < \frac{B}{F} - 1.$$

Consider the cooperative outcome in which X players match expenditures. The optimal effort choice maximizes:

$$\frac{2x}{1+2x}2B - 2Cx$$

which leads to optimal effort choice:

$$x_1 = x_2 = x^* = \frac{\sqrt{\frac{2B}{C}} - 1}{2}.$$

Aggregate expenditures are less than those required for deterrence if $B > \frac{2F^2}{C}$. Otherwise the cooperative outcome with symmetric expenditures has each player split the deterrence cost, leading to expenditures of $\frac{B-1}{2}$. For this case, expenditures in the cooperative solution are higher than in the non-cooperative Nash equilibrium. Consider the problem of player X_1 , who chooses x_1 to maximize:

$$\frac{x_1 + x_2}{1 + x_1 + x_2}B - Cx_1.$$

The objective function for X_1 , because it includes the effort choice of X_2 , illustrates that the optimal choice of x_1 depends on the effort exerted by X_2 . Specifically, the best response functions are given by:

$$x_1^* = \sqrt{B/C} - 1 - x_2, \quad \text{and} \\ x_2^* = \sqrt{B/C} - 1 - x_1.$$

Essentially, player X_1 's best response is to reduce effort (from the level associated with the one person contest) by one unit for each unit of effort exerted by player X_2 . This illustrates the potential for free-riding. In this case, the symmetric equilibrium has both X players choose effort:

$$x_1^* = x_2^* = \frac{\sqrt{B/C} - 1}{2}.$$

Note, however, that any pair of effort choices such that $\sum_i x_i = \sqrt{B/C} - 1$ can be supported in equilibrium (for example, one equilibrium would have complete free-riding by one of the X players, while the other exerted effort $\sqrt{B/C} - 1$). When $B > \frac{F^2}{C}$, the symmetric equilibrium has both X players exert effort $x^* = \frac{\sqrt{\frac{B}{C}} - 1}{2}$, and player Y participates in both contests. Otherwise, the equilibrium has both X players choose effort $x = \frac{B-1}{2}$, and player Y chooses not to participate in either contest. Note that in the case of participation by player Y , aggregate spending in the Nash equilibrium is lower than that associated with the cooperative outcome. This is due to the positive externality generated by the contest success function.

Case 2: Participation constraint

Similar analysis can be applied to the case when player Y is constrained to enter at most one contest. Player Y participates in one contest, chosen at random, if:

$$\sum_i x_i < \frac{B}{F} - 1.$$

In this case the symmetric cooperative outcome has each of the X players exert effort:

$$x_1 = x_2 = x^* = \frac{\sqrt{\frac{B}{C}} - 1}{2}.$$

This is an equilibrium if aggregate effort is not sufficient to deter participation, which occurs when $B > \frac{F^2}{C}$. Otherwise the type X players split the cost of deterrence, exerting effort $\frac{\frac{B}{F} - 1}{2}$. In the non-cooperative outcome, the objective function for X_1 changes to:

$$\frac{1}{2}B + \frac{1}{2} \frac{x_1 + x_2}{1 + x_1 + x_2} B - Cx_1.$$

The best response functions can be written as:

$$\begin{aligned} x_1^* &= \sqrt{B/2C} - 1 - x_2, \quad \text{and} \\ x_2^* &= \sqrt{B/2C} - 1 - x_1. \end{aligned}$$

The symmetric equilibrium has both X players choose effort:

$$x_1^* = x_2^* = \frac{\sqrt{B/2C} - 1}{2}.$$

Again, however, any pair of effort choices such that $\sum_i x_i = \sqrt{B/2C} - 1$ can be supported in equilibrium, including an outcome in which one of the players free rides. When $B > \frac{F^2}{C}$, the symmetric equilibrium has both X players exert effort $x^* = \frac{\sqrt{\frac{B}{2C}} - 1}{2}$, and player Y participates in both contests. Otherwise, the equilibrium has both X players choose effort $x = \frac{\frac{B}{F} - 1}{2}$, and player Y chooses not to participate in either contest.

To summarize, equilibrium predictions indicate that, relative to the cooperative solutions, negative strategic spillover effects should increase spending and positive spillovers should reduce spending. In the next section, these predictions are put to an experimental test. In all, five treatments were conducted: a ‘baseline’, two separate negative externality treatments (one with and one without a participation constraint) and two separate positive externality treatments (again, with and without the participation constraint).

3 Experimental design and procedures

All subjects were paid volunteers recruited from the undergraduate population at American University. Prior to volunteering, subjects received an e-mail invitation to participate in a decision making exercise. The invitation indicated that participants would be paid a \$5 ‘show up’ fee in addition to an amount that would depend on their decisions and the decisions of others in the experiment. All payments were made in cash, privately, at the end of the experimental session. Upon arrival at the experiment site, subjects were seated in front of a computer terminal and given the experiment instructions (available by request from the authors), which are summarized below. The experiment was programmed and conducted with the software z-Tree (Fischbacher 2007).

3.1 Baseline treatment

Subjects were informed that there would be eight decision-making rounds in the experiment. In each round, subjects were randomly and anonymously paired, and each subject was endowed with \$5 that could be used for participation in a contest. Rounds were independent in the sense that earnings did not carry over across rounds; the \$5 endowment for a round could be used only to finance expenditures in that round. Subjects were classified as either a player 1 type or a player 2 type, maintained this classification for the duration of the experiment, and each pair of subjects consisted of one player 1 and one player 2. The player 1 subject made a decision first, choosing how much of the endowment to spend on the contest. Spending influenced the probability that player 2 would win the contest prize of \$4.50, if player 2 chose to participate. Specifically, the probability player 2 would win was given by:

$$P_2 = \frac{1}{1 + x_1},$$

where P_2 is the probability player 2 wins the contest, and x_1 denotes the expenditures of player 1. After player 1 made a decision, information regarding player 1's spending and the probability associated with the contest outcome was transmitted to player 2. Player 2 then chose whether to spend \$1 of their \$5 endowment to participate in the contest. If player 2 chose to participate, the prize was awarded to the winner of a draw conducted with a computerized random number generator and according to the probability determined by the player 1's expenditures. If player 2 chose not to participate, the contest prize was awarded to player 1. For player 1 types, earnings for the round consisted of the \$5 endowment, less the amount spent on the contest (which was subtracted regardless of player 2's participation decision), plus the \$4.50 prize if applicable. For player 2 types, earnings from a round consisted of the \$5 endowment, less \$1 if the subject chose to participate, plus the \$4.50 prize if applicable. At the conclusion of each round, a 'round summary' was displayed on each subject's computer screen, with information regarding the decisions made by both subjects in the pair (and only to those subjects; no aggregate information was provided), the contest outcome, and the subject's earnings for the round. The experiment then proceeded to the next round, in which subjects were anonymously and randomly repaired. At the end of the eighth round, subjects were called out of the lab individually, and paid in cash a \$5 show up fee in addition to the sum of their earnings from 2 randomly selected rounds.⁹ Given the parameters used for this treatment, the sub-game perfect Nash equilibrium is for player 1 to spend \$1.12 and for player 2 to participate.

3.2 Negative externality treatments:

The following modifications were made to the baseline for the externality treatments. A second type 1 player was added and this player made expenditures to determine the probability associated with contest for a second \$4.50 prize. The player 1 types were referred to as player 1A and player 1B, and the contests as contest 1A and contest 1B. Player 1 types made their expenditure decisions simultaneously, without any communication or information regarding each other's expenditure decisions. After both player 1 types made their expenditure decisions, player 2 chose whether to participate or not. In the Externality 1 (E1) treatment,

⁹Prior to making a decision in the first round, subjects were informed that while they would make decisions in each of the eight rounds, only two rounds selected at random would be used for payment purposes.

player 2 was constrained to participate in (at most) one of the contests. This means that player 2 could either participate in contest 1A, 1B, or neither. In the Externality 2 (E2) treatment, player 2 could choose among the following options: pay \$1 to participate only in contest 1A, pay \$1 to participate only in contest 1B, pay \$2 to participate in both contest 1A and contest 1B, or not participate in either contest.

In the E1 treatments the group optimum is for player 1s to match expenditures at \$0.50 each and for player 2 to participate in one contest chosen at random. However, this outcome is not a Nash equilibrium; both player 1 types could win their contest for sure by outspending the other. The independence of the contests in the E2 treatments leaves the equilibrium predictions unchanged from the baseline treatment. Specifically, in each round the equilibrium is for player 1 to spend \$1.12 and for player 2 to participate in both contests.

3.3 Positive externality or ‘public goods’ treatments

In the public goods treatments, the expenditures of two player 1 types were used to determine the probability associated with the contest outcome. As in the negative externality treatments, each of the player 1 types was associated with a separate contest having a \$4.50 prize; the player 1 types were referred to as Player 1A and Player 1B, and the contests as contest 1A and contest 1B. Specifically, the contest success function was:



$$P_2 = \frac{1}{1 + x1a + x1b},$$

where P_2 denotes the (common) probability that player Y wins contest 1A or 1B.

The player 1 types made their expenditure decisions simultaneously, without any communication or information regarding each other's expenditure decisions. After both player 1 types made their expenditure decisions, the player 2 type choose whether or not to participate. In the Public Goods 1 (PG1) treatment, player 2 was restricted to participate with at most one of the player 1 types. In the Public Goods 2 (PG2) treatment, the Player 2 type could choose to participate with either Player 1A or Player 1B, both or neither. For the PG1 treatment, the symmetric equilibrium is for both player 1s to spend \$0.25, and for player 2 to participate in one of the contests, chosen at random. For the PG2 treatment, the symmetric equilibrium is for both player 1s to spend \$0.56, and for player 2 to participate in both contests.

The baseline, E1, and PG1 treatments utilized a total of 180 subjects (60 per treatment), while the PG2 and E2 treatment had 126 subjects (63 per treatment). While the baseline treatment had 30 player 1 types and 30 player 2 types, each of the other treatments had about 40 player 1 types and 20 player 2 types. Sessions lasted approximately 45 minutes with average earnings of approximately \$18.25 (inclusive of the \$5 show up fee). Given the parameters used in the treatments, a summary of the equilibrium predictions is as follows:

Baseline	Player 1 spending \$1.12	Player 2 decision Participate
E1	No pure strategy equilibrium, want to outspend other Player 1 type by ε	Enter against low spender if either Player 1 spends less than \$3.50
E2	\$1.12	Participate in both contests
PG1	\$0.25	Participate, choice of contest is random
PG2	\$0.56	Participate in both contests

Fig. 1 Average spending per period**Fig. 2** Average spending per period

4 Experimental results

Our modeling approach and experimental design generate two central predictions regarding player 1 spending decisions: relative to the baseline and E2 treatments, spending should be *higher* in the E1 treatment and *lower* in the PG treatments. To assess whether the differences in spending choices across treatments were statistically significant, we use a subject's average level of spending from all eight rounds of the experiment and compare across treatments using a Mann-Whitney rank sum test (all p -values are two-tailed unless otherwise specified). We also examine whether, as predicted, spending in PG1 was less than PG2, and whether spending in E2 was different than the baseline case.

4.1 Player 1 spending: Analysis of aggregate data

Figures 1 and 2 display the average player 1 spending by round for the treatments.

Average spending across all rounds was \$2.06, \$1.57, and \$1.85 for the baseline, E1 and E2 treatments, respectively.¹⁰ In contrast to our expectations, spending in the E1 treatment was lower than either the baseline or E2 treatment, and the differences are statistically significant at the 10% level (Mann Whitney p -values of 0.051 for E1 vs. Baseline and 0.1 for E1 vs. E2). Consistent with the equilibrium prediction, we are not able to reject the null hypothesis of no difference in spending between the baseline and E2 treatments (Mann Whitney p -value of 0.21). For the PG1 and PG2 treatments, average expenditures across all rounds were \$1.36 and \$1.38, respectively. As shown in Fig. 2, subjects in the baseline treatment averaged higher expenditures than those in PG1 or PG2 in each round, and the

¹⁰Consistent with the vast majority of rent seeking papers utilizing experiments, player 1 subjects 'over-invested' relative to the equilibrium point predictions in all of our treatments.

differences in subject expenditures were statistically significant (Mann Whitney p -values of 0.006 for Baseline vs PG1 and 0.003 for Baseline vs PG2).¹¹ However, in contrast to the equilibrium prediction, differences in spending between the PG1 and PG2 treatments were not significant (Mann Whitney p -value 0.889).

Based on these results, we make two conclusions with respect to our main hypotheses. First, changing the contest success function to emphasize aggregate effort choices in the PG treatments resulted in significantly lower subject expenditures. The intuition behind this result is clear—subjects correctly identified the positive externality associated with individual spending, and the resultant free riding incentive led to lower expenditures. It is important to note, however, that the subject decisions do not reflect *complete* free riding. Although expenditures were lower, player 1 subjects spent 0 in only 22 cases (approximately 7%) for the PG1 treatment and six cases (approximately 2%) for the PG2 treatment. Moreover, as discussed in the next section, player 1 subjects in the PG treatments do not appear to have responded to the spending choices of the other type 1 player in their group from the prior period. Although our analysis takes place in a different context, we view our results as consistent with the well documented experimental evidence regarding positive (but small) contributions in the more typical VCM-type public goods games.

Second, while the equilibrium predictions from our theoretical model indicate that the negative externalities in the E1 treatment should *raise* expenditures, we find that spending in the E1 treatment was significantly *lower*. This is an important result that demonstrates, in our view, one of the benefits of utilizing experimental methods to study game theoretic models of terrorism. While the strategic incentives are clear in theory, important behavioral considerations (isolated in the empirical analysis presented in the next section) are introduced when the type 2 players make choices to participate in multiple contests. The analysis presented in the next section indicates that subjects in the E1 and E2 treatments *did* respond to higher expenditures from the other type 1 player in their group by raising expenditures in the following period. The type 2 players, however, were less likely to participate in a given contest when they had multiple contests to choose from (even with controls for the probability of winning in place). This illustrates two competing influences associated with adding additional contests. On one hand, increasing expenditures makes sense because (as captured by the theory) outspending other player 1 types makes a subject's contest the hardest to win (from the perspective of the player 2), and may secure the contest prize. Alternatively, lower player 2 participation rates in multiple contest cases (a behavioral result not captured by the theory) decrease the probability that a particular player 1's prize is contested, leading to a reduction in subject expenditures. While the existing theoretical literature in the study of terrorism is focused almost exclusively on the first effect, our experimental results suggest the second (behavioral) effect is important—in fact, given the parameters used in our design, the second effect was stronger. As a potential extension, consider the likely behavioral consequences for player 1 spending associated with adding additional potential targets in the E1 treatment. While we did not conduct such a treatment, we believe that although the incentive to overspend remains in theory, it is likely that the perceived vulnerability to attack would decline and this would further reduce Type 1 expenditures.

4.2 Player 1 spending: An econometric analysis of individual decision making

Additional insight can be gained from analyzing individual subject decision making. At the end of each period, subjects in our experiment were shown a 'round summary' screen

¹¹ Differences between PG1 and E1 as well as differences between PG2 and E2 are also significant at the 1% level.

Table 1 Random-effects GLS regression results. Dependent variable: Player 1 expenditure

Treatment	E1	E2	PG1	PG2
Number of observations	280	294	280	294
Overall R^2	0.53	0.31	0.68	0.73
$Spending_{t-1}$	0.763 (17.01)***	0.583 (10.43)***	0.826 (23.62)***	0.853 (27.96)***
Win_{t-1}	-0.431 (-3.45)***	-0.514 (-3.79)***	-0.018 (-0.15)	-0.281 (-3.04)***
$Othersp_{t-1}$	0.086 (2.10)**	0.164 (3.27)***	0.016 (0.47)	0.039 (1.26)
Period	-0.028 (-1.27)	-0.01 (-0.38)	-0.027 (-1.28)	0.012 (0.67)
Constant	0.75 (4.43)	0.949 (4.89)***	0.342 (2.04)**	0.283 (2.01)**

Coefficients (z-stats in parentheses)

** Significant at the 5% level

*** Significant at the 1% level

that displayed the spending of the type 1 player (or players) in their group, the participation decision(s) of the type 2 player in their group, the contest outcome(s), and round earnings for the subject. Although this information does not affect the theoretical predictions associated with spending or participation in subsequent rounds, we believe (and our econometric results suggest) that it has important behavioral effects. Our data consist of multiple observations for each subject (i.e., individuals played eight one-shot games, randomly being repaired for each one-shot game) and therefore previous outcomes and learning may impact behavior in future periods. While reputation effects do not exist, learning effects may.

As it is panel data, each dimension has the potential for heterogeneity: time (period) and individual. To allow for time effects we include a period trend in our specification. To capture possible heterogeneity across individuals, which is expected in such experiments, we use GLS random effects. The regression results are presented in Table 1. The dependent variable in each model is player 1 subject's expenditure, with independent variables including period, spending in the previous period ($Spending_{t-1}$), whether the subject won a contest in the previous period (Win_{t-1}), and spending of the other type 1 player in a subjects group form the previous period ($Othersp_{t-1}$).¹²

We include $Spending_{t-1}$ because we view prior spending as a natural reference point, with subjects making adjustments relative to this level of spending depending on the results from the previous period. With respect to Win_{t-1} , two behavioral hypotheses seem reasonable: a subject may be positively reinforced and increase expenditures or, alternatively, players winning in the previous period might lower their expenditures in the hope that they would still win the contest but have spent less to do so. Our main variable of interest is $Othersp_{t-1}$. The strategic logic associated with best responding in the E1 treatment

¹²Including whether player 2 participated in the subject's contest in the previous period ($Attack_{t-1}$) does not qualitatively affect the other coefficient estimates or their significance, and the $Attack_{t-1}$ coefficient, while positive, does not approach conventional levels of significance.

involves forming an expectation about what the other type 1 subject in a group will spend. Lagging the other type 1 player's spending may serve as a proxy for what a player 1 subject expects in the current period. As such we expect higher lagged other's spending to generate higher own spending in the current period for the E1 and E2 treatments, because the type two player in the group may reasonably choose to attack the low spender in the group. Note that the incentive to increase one's own expenditures in the current period, in response to higher spending by the other type 1 player, is not present in the PG treatments. Here, higher lagged other's spending may decrease current period spending by sharpening the free riding incentives, or alternatively, it may increase current period spending if subjects display a variant of 'conditional cooperation.' The differential impact of the lagged other's spending variable by treatment, as well as its indeterminate value in the baseline treatment, raises several complications with respect to pooling the data across treatments. As such, we include in Table 1 the results of separate regressions for each treatment, rather than pooled results.

The coefficient on the $Spending_{t-1}$ variable is positive and statistically significant in each treatment, which supports our view that it serves as an important reference point and control variable. The coefficient on the Win_{t-1} , is negative in each treatment, and statistically significant at the 1% level in three of the four treatments (E1, E2, and PG2). Thus winning in the prior period led subjects to reduce their expenditures in the current period, a strong and in our view sensible behavioral effect. The coefficient on the lagged other's spending variable is positive and statistically significant in the E1 and E2 treatments, but is not statistically significant for the PG treatments. As such, subjects in the E1 treatment did increase their spending in response to higher spending levels from the other type 1 player in their group in the prior period. This is consistent with the underlying strategic logic that leads to the overspending hypothesis for the E1 treatment. Although the theoretical model does not suggest that the overspending incentive is present in the E2 treatment, we note that the coefficient on the lagged other's expenditure variable approximately doubles in the E2 treatment (relative to E1), demonstrating another important behavioral effect—although in theory the two contests in the E2 treatment are independent, type 1 subjects strongly responded to the spending decisions of other players.

The positive and significant coefficient on lagged other's spending in the E treatments, taken together with the result (based on the analysis of aggregate data) that spending in the E1 treatment was lower, suggest another consideration absent from the theoretical model. Based on our analysis of the aggregate data, spending was highest in the baseline, followed (in order) by the E2, E1, PG2 and PG1 treatments. While the relative ranking of the PG treatments is consistent with theoretical expectations, the rankings of the E1 and E2 treatments are not, and this does not appear to be driven by a lack of 'best responding' in the negative externality treatments. We suggest an alternative explanation based on the behavior of the type 2 player that is related to perceived vulnerability of the type 1 players across our treatments. The key distinctions between the baseline and negative externality treatments involve the number of contests available to the type 2 player. In the baseline, there was a single potential target. In the E2 treatment, there were two potential targets, and the type 2 player had the resources to compete in both. In the E1 treatment, there were two potential targets, but type 2 players could participate in at most one of the contests. Our hypothesis is that type 1 players perceived the likelihood of attack to be highest in the baseline case (where they presented the only target), followed by the E2 treatment (in which they were one of two potential targets) and the E1 treatment (in which they were one of two potential targets, but for which the type 2 player could participate in at most one contest). This hypothesis is supported by the analysis of type 2 player decisions presented in the next section.

Table 2 Random-effects Probit regression results. Dependent variable: Player 2 participation decision

	E1	E2	PG1	PG2	Base
Number of observations	280	294	280	294	210
Wald Chi ²	46.52	45.62	6.62	20.82	18.68
(<i>p</i> value)	(0.00)	(0.00)	(0.08)	(0.00)	(0.00)
Probwin	6.67 (6.75)***	6.93 (6.72)***	1.14 (2.22)**	3.48 (4.32)***	5.91 (3.05)***
Win_{t-1}	0.21 (0.84)	0.01 (0.04)	0.023 (0.10)	0.15 (0.70)	−0.33 (−1.27)
Period	0.00 (0.07)	0.03 (0.72)	0.04 (0.97)	0.03 (0.83)	0.02 (0.35)
Constant	−3.28 (−5.97)***	−2.73 (−5.44)***	−0.98 (−2.78)***	−1.35 (−3.61)***	−0.95 (−1.62)

Coefficients (z-stats in parentheses)

** Significant at the 5% level

*** Significant at the 1% level

4.3 Player 2 participation decisions

In theory, player 2 participation decisions are driven *solely* by the observed probability of winning a contest. Type 2 players should participate in *each* contest with a positive expected value (given sufficient resources) or pick the contest with the highest expected value (if resources are limited). Our experimental results, however, demonstrate that the number of available contests played an important role. In particular, type 2 subjects were less likely to participate in a given contest when they could choose from a set of contests. For example, given the observed player 1 expenditures in the baseline case, participation in the contest was a best response (i.e., had positive expected value) in 210 of the 240 contests. Type 2 players chose to participate in a total of 177 (approximately 84%) of these contests. For the PG2 and E2 treatments, participation was a best response in 265 and 304 cases, respectively. Yet type 2 subjects in these cases chose to participate in 119 (about 55%) of the PG2 cases and 124 (about 59%) of the E2 treatments. Below we analyze participation decisions at the individual level using a random effects Probit regression model. The dependent variable for our regressions is a binary variable taking the value 1 if the subject chose to participate in a contest, with independent variables including period, the probability of winning the contest (*Probwin*), and a dummy variable for whether the subject won a contest in the previous period (Win_{t-1}).¹³

Table 2 presents the results of separate regressions for each treatment and Table 3 presents the pooled results (across the baseline PG2 and E2 treatments for Model 1 and across all treatments for model 2).

As seen in Table 2, the coefficient on the probability of winning variable is positive and statistically significant in each treatment, indicating that subjects were more likely to enter

¹³ Adding a variable for the subjects lagged participation decision does not substantively change the other coefficient estimates or the overall significance of the results, and generates an estimated coefficient that is not statistically significant.

Table 3 Pooled Random-effects Probit regression results.
Dependent variable: Player 2 participation decision

	Model 1	Model 2
Number of observations	798	1358
Wald Chi ²	116.81	200.77
(<i>p</i> value)	(0.00)	(0.00)
Probwin	5.33 (9.58)***	4.02 (12.04)***
LagWin	−0.05 (−0.37)	0.02 (0.20)
E1		−1.46 (−8.98)***
PG1		−1.26 (−8.06)***
E2	−1.06 (−6.84)***	−.98 (−6.29)***
PG2	−0.70 (−4.82)***	−0.73 (−4.79)***
Period	0.02 (0.71)	0.02 (0.86)
Constant	−0.974 (−3.80)	−.596 (−3.02)***

Coefficients (*z*-stats in parentheses)

** Significant at the 5% level

*** Significant at the 1% level

a contest the higher was its expected value. Yet Models 1 and 2 in Table 3 illustrate the importance of the number of available contests. The coefficients on the treatment dummies are negative and statistically significant, indicating that relative to the baseline, subjects were less likely to participate when there were multiple contests to choose from. The estimates in Model 1 are particularly noteworthy—in this case subjects had the resources to participate in each contest, but controlling for the probability of winning, we find that they were less likely to participate in the treatments with two contests. As noted previously, participation was a best response in the vast majority of cases for each of our treatments. When type 2 players had a single contest with positive expected value to consider, as in the baseline treatment, they were likely to enter. Player 1 subjects, in recognition, were more likely to have higher expenditure levels. In the E2 and PG2 treatments however, type 2 subjects faced with two contests (each having positive expected value) were more likely at the margin to choose to participate in a single contest. This behavioral effect, absent from the theoretical models, had potential feedback effects on type 1 expenditures. Facing a lower probability of attack, the type 1 players in the E2 and PG2 treatments were able to further decrease their expenditure levels. Taken as a whole, our experimental results offer qualified but strong support for the central hypothesis derived from our theoretical model, and also illustrate important behavioral considerations.

5 Conclusions

We present a simple model of strategic entry deterrence with the potential for strategic spillovers and test the model using laboratory experimental methods. While exploiting similarities with the entry deterrence models typically associated with industrial organization,

our primary motivation is the study of game theoretic approaches to terrorism. In particular we examine whether the game theoretic predictions associated with strategic spillovers are consistent with laboratory behavior. In our view, lab experiments provide an important and underutilized complement to field studies of terrorism because they allow for precise control over several important parameters that are difficult to assess using naturally occurring data, and we believe that this paper is a beneficial starting point for merging behavioral economics to the terrorism literature. Experimental methods are also important because they can be used to identify and evaluate behavioral effects not captured with existing theory and may result in changes to theoretical assumptions.

We do not draw policy implications because we acknowledge that our one-shot laboratory results may not be robust to the dynamic aspects of the real world (voting in new governmental leaders or reputation effects in negotiating when repeated interaction with the same terrorists). We acknowledge that laboratory experiments show the comparative statics of behavior when the impact of strategic spillovers changes (from positive externalities to negative externalities) and do not encompass all behavioral implications (such as the fact that more terrorist cells emerge when countries go on the offense for terrorism protection). To our knowledge, our paper is the first attempt to examine mathematical models of terrorism in the lab; future work could include the dynamic element of repeated interactions or add a third player, voters, who decide to change the government.

Our theoretical results reinforce those in the existing literature regarding terrorism prevention. In particular, negative externalities associated with competition to avoid becoming the most vulnerable ‘target’ lead to overspending relative to the cooperative solution. This has led some to argue that countries are spending too much on terror prevention. Alternatively, when there are positive spillovers, the theoretical results highlight the potential for free riding, leading to lower equilibrium expenditures that raise concerns regarding under-spending on terror prevention.

Our experimental results are broadly consistent with theoretical expectations. Relative to the baseline, type 1 subjects significantly lowered expenditures in the PG treatments. Our data also support the hypothesis that for the negative externality treatments, type 1 subjects increase expenditures in response to higher expenditures from the other type 1 player in their group. Yet, in contrast to theoretical expectations, spending in the E1 treatment was lower than either the baseline or E2 treatment. This result is related to behavioral considerations regarding type 2 participation decisions not captured by existing theory. Controlling for the probability of winning, type 2 subjects were less likely to participate in a given contest when they had multiple contests to choose from. This reduced vulnerability to attack appears to have had a feedback effect in the sense that type 1 players reduced their expenditures in the E2 and E1 treatments. Although our treatments utilized only one or two player 1 types, we hypothesize that the inclusion of additional player 1 types would further dampen the competitive pressures, and as such concerns regarding overspending in the presence of negative externalities may be misplaced.

In our view, game theoretic approaches to terrorism, and empirical studies relying on naturally occurring field data, offer important insights that can inform policy makers. However, based on the underlying mathematical structure and the difficulty of replicating it in a field environment, we view experimental studies as an important complement to this research. Our results suggest that concerns over free riding incentives are well placed. Future research designed to evaluate alternative mechanisms to overcome free riding incentives in a terrorism context is warranted. Our evaluation of the overspending incentives associated with the potential for target displacement, however, suggests the limited resources of terror

groups and the vast number of potential targets available may reduce perceived vulnerability to attack, and the resultant incentive to reduce defensive expenditures may act as an important counterbalance to the strategic incentives.

Acknowledgements We are grateful to Peter Barkley for initial z-tree programming assistance. Two anonymous referees, the associate editor, and the editor in chief provided very valuable comments and suggestions and have significantly improved the paper. The paper has also benefited from presentation comments at the 2007 SEA meeting and the 2007 SABE meetings. Financial support from the United States Naval Academy Faculty Development Fund, the United States Naval Academy Research Council (NARC, 2006 and 2007), and the United States Naval Academy Department of Economics is gratefully acknowledged.

References

- Arce, D. G., & Sandler, T. (2005). Counterterrorism: A game-theoretic analysis. *Journal of Conflict Resolution*, 49(2), 183–200.
- Bapat, N. A. (2006). State bargaining with transnational terrorist groups. *International Studies Quarterly*, 50, 213–229.
- Brandts, J., Cabrales, A., & Charness, G. (2005). *Forward induction and entry deterrence: An experiment*. Available at SSRN: <http://ssrn.com/abstract=486182>.
- Bueno de Mesquita, E. (2005a). Conciliation, counterterrorism, and patterns of terrorist violence. *International Organization*, 59, 145–176.
- Bueno de Mesquita, E. (2005b). The terrorist endgame: a model with moral hazard and learning. *Journal of Conflict Resolution*, 49(2), 237–258.
- Bueno de Mesquita, E. (2006). *Conciliation, counter-terrorism, and patterns of terrorist violence: A comparative study of five cases*. Available at <http://home.uchicago.edu/~bdm/PDF/adverseterrorcases.pdf>.
- Bueno de Mesquita, E. (2007). Politics and the suboptimal provision of counterterrorism. *International Organization*, 61(1), 9–36.
- Cadigan, J. (2007). Two stage team rent-seeking: Experimental analysis. *Southern Economic Journal*, 74(1), 85–103.
- Davis, D., & Reilly, R. (1998). Do too many cooks always spoil the stew? An experimental analysis of rent-seeking and the role of a strategic buyer. *Public Choice*, 95, 89–115.
- Dixit, A. (1980). The role of investment in entry deterrence. *Economic Journal*, 90, 95–106.
- Enders, W., & Sandler, T. (1993). Effectiveness of antiterrorism policies: Vector-auto regression-intervention analysis. *American Political Science Review*, 87(4), 829–844.
- Enders, W., & Sandler, T. (1995). Terrorism: Theory and applications. In K. Hartley & T. Sandler (Eds.), *Handbook of defense economics* (Vol. 1, pp. 213–249). Amsterdam: North-Holland.
- Enders, W., & Sandler, T. (2005). After 9/11: Is it all different now? *Journal of Conflict Resolution*, 49(2), 259–277.
- Enders, W., Parise, G. F., & Sandler, T. (1992). A time-series analysis of transnational terrorism: Trends and cycles. *Defense Economics*, 3(4), 305–320.
- Ferrero, M. (2006). Martyrdom contracts. *Journal of Conflict Resolution*, 50(6), 855–877.
- Fischbacher, U. (2007). z-Tree—Zurich toolbox for readymade economic experiments. *Experimental Economics*, 10(2), 171–178.
- Jung, Y., Kagel, J., & Levin, D. (1994). On the existence of predatory pricing: An experimental study of reputation and entry deterrence in the chain-store game. *The RAND Journal of Economics*, 25(1), 72–93.
- Kreps, D., & Wilson, R. (1982). Reputation and imperfect information. *Journal of Economic Theory*, 27, 253–279.
- Landes, W. M. (1978). An economic study of U.S. aircraft hijackings, 1961–1976. *Journal of Law and Economics*, 21(1), 1–31.
- Mason, C., & Nowell, C. (1998). An experimental analysis of subgame perfect play: the entry deterrence game. *Journal of Economic Behavior and Organization*, 37, 443–462.
- Milgrom, P., & Roberts, J. (1982). Predation, reputation, and entry deterrence. *Journal of Economic Theory*, 27, 280–312.
- Millner, E., & Pratt, M. (1989). An experimental investigation of efficient rent-seeking. *Public Choice*, 62(2), 139–151.
- Millner, E., & Pratt, M. (1991). Risk aversion and rent-seeking: An extension and some experimental evidence. *Public Choice*, 69(1), 81–92.

- Önçüler, A., & Croson, R. (1998). Rent-seeking for a risky rent: A model and experimental investigation. *Journal of Theoretical Politics*, 17(4), 403–429.
- Potters, J., de Vries, C., & van Winden, F. (1998). An experimental examination of rational rent-seeking. *European Journal of Political Economy*, 14, 783–800.
- Rosendorff, B. P., & Sandler, T. (2004). Too much of a good thing? The proactive response dilemma. *Journal of Conflict Resolution*, 48(5), 657–671.
- Sandler, T. (2005). Collective versus unilateral responses to terrorism. *Public Choice*, 124(1–2), 75–93.
- Sandler, T., & Arce, D. (2003). Terrorism and game theory. *Simulation and Gaming*, 34(3), 319–337.
- Sandler, T., & Enders, W. (2004). An economic perspective on transnational terrorism. *European Journal of Political Economy*, 20(2), 301–316.
- Shogren, J., & Baik, K. (1991). Reexamining efficient rent-seeking in laboratory markets. *Public Choice*, 69(1), 69–79.
- Tullock, G. (1967). The welfare costs of tariffs, monopolies and theft. *Western Economic Journal*, 5, 224–232.
- Tullock, G. (1980). Efficient rent-seeking. In J. M. Buchanan, R. Tollison, & G. Tullock (Eds.), *Toward a theory of the rent seeking society* (pp. 267–292). College Station: Texas A&M Press.