Keynote paper

Terrorism & game theory

Todd Sandler
University of Southern California

Daniel G. Arce M.
Rhodes College

This keynote paper examines how game-theoretic analyses of terrorism have provided some policy insights that do not follow from nonstrategic analyses. Some new game-theoretic applications are indicated that concern terrorist targeting of businesses, officials, and the general public, where targets can work at cross-purposes as they attempt to deflect the attack. Other novel applications involve government choice among alternative antiterrorism policies, and government concessionary policy when terrorists are either hardliners or moderates in their viewpoint. Directions for future research are also indicated.

KEYWORDS: *game theory; terrorism; transnational terrorism; deterrence; preemption; concessionary policy; target choice; asymmetric information*

Over the past two decades, a small group of analysts in economics and political science have applied game theory to study terrorism, which involves the premeditated use or threat of use of violence or force on the part of terrorists to achieve a political objective through intimidation or fear.[1] For example, Sandler, Tschirhart, and Cauley (1983) present some rational-actor models that depict the negotiation process between terrorists and government policy makers for incidents where hostages or property are seized and demands are issued. In their model, terrorists' valuation of the likely concession to be granted by a government is based on a probability distribution conditioned on past governmental concessions. Their analysis illustrates that the terrorists' choices and actions are influenced by those of the government and vice versa. Moreover, each adversary acts on its beliefs of the opponent's anticipated actions.

Since 11 September 2001 (hereafter 9/11), there are many articles being written by scholars who apply game theory to the study of terrorism. Game theory is an appropriate tool for examining terrorism for a number of reasons. First, game theory captures the strategic interactions between terrorists and a targeted government, where actions are interdependent and, thus, cannot be analyzed as though one side is passive. Second, strategic interactions among rational actors, who are trying to act according to how they think their counterparts will act and react, characterize the interface among

SIMULATION & GAMING, Vol. 34 No. 3, September 2003 319-337
DOI: 10.1177/1046878103255492

terrorists (e.g., between hardliners and moderates) or among alternative targets (e.g., among targeted governments, each of which is taking protective measures). Third, in terrorist situations, each side issues threats and promises to gain a strategic advantage. Fourth, terrorists and governments abide by the underlying rationality assumption of game theory, where a player maximizes a goal subject to constraints.[2] Empirical support for terrorists' rationality is given credence by their predictable responses to changes in their constraints—for example, the installation of metal detectors in January 1973 led to an immediate substitution away from skyjackings into kidnappings (Enders & Sandler, 1993, 1995; Sandler & Enders, in press). Fifth, game-theoretic notions of bargaining are applicable to hostage negotiations and terrorist campaign-induced negotiations over demands. Sixth, uncertainty and learning in a strategic environment are relevant to all aspects of terrorism, in which the terrorists or government or both are not completely informed.

The purpose of this keynote article is to review how game theory has been applied in the literature on terrorism. Another purpose is to present some new applications that include terrorists' choice of target (i.e., business people, officials, and tourists), governments' choice between preemption and deterrence, and government concessionary policy when terrorists are of two minds—hardliners and moderates. These applications illustrate how game theory can be fruitfully employed to enlighten policy making.

A brief look at the literature

One of the pillars of U.S. antiterrorism policy is never to negotiate or capitulate to the demands of hostage-taking terrorists (U.S. Department of State, 2002, p. xii). This same no-negotiation stance has been taken by other countries, such as Israel. The logic to this policy is that if a target adheres to its stated no-negotiation policy, then would-be hostage takers would have nothing to achieve and so would stop abducting hostages. This outcome implicitly assumes that terrorists only gain from achieving their demands and that both sides are completely informed so that the subgame perfect equilibrium is to pledge not to concede. Obviously, something is incomplete about this logic because terrorists continue to take hostages and even the staunchest advocates of the no-negotiation policy have reneged on their pledge. For example, the Reagan administration bartered arms for the release of Rev. Benjamin Weir, Rev. Lawrence Jenco, and David Jacobsen during the 1985-1986 "Irangate" scandal; Israel was prepared to trade prisoners for the schoolchildren taken in Maalot in May 1974.

Lapan and Sandler (1988) elucidate the policy's incompleteness with a game in extensive form where the government first chooses the level of deterrence, which, in turn, determines the logistical failure or success of terrorists when they engage in a hostage mission. A higher level of deterrence elevates the likelihood of logistical failure. Based on their perceived likelihoods of logistical and negotiation success, the terrorists decide whether to attack. If their expected payoffs from hostage taking are positive, then they attack. The game can end in four ways: no attack, an attack that results in a logistical failure, a successful attack that ends with the terrorists obtaining their demands, and a successful attack that results in no concessions. Information is incomplete because the government does not know the payoffs associated with not capitulating prior to hostage incidents. If a sufficiently important person is secured, then the government may regret its no-negotiation pledge because the expected costs of not capitulating may exceed that of capitulating. That is, the no-negotiation policy is time inconsistent if sufficiently valuable hostages are captured. Even when the government's pledge not to negotiate is believed by the terrorists, a fanatical group may still engage in a hostage mission when a positive payoff is associated with either a logistical or negotiation failure by advertising the cause or achieving martyrdom.

Lapan and Sandler (1988) demonstrate that the effectiveness of the no-negotiation strategy hinges on the credibility of the government's pledge, the absence of incomplete information, the terrorists' gains being solely tied to a negotiation success, and sufficient deterrence spending to eliminate logistical success. In practice, each of these implicit assumptions is suspect. Efforts to restrain a government's discretionary action in hostage scenarios—say, through a constitutional amendment or costly punishment—are required to eliminate a government's ability to renege on its stated policy. Lapan and Sandler examine the importance of reputation costs in a multiperiod model. Their analysis stands in stark contrast to nonstrategic hostage-taking models by Islam and Shahin (1989) and Shahin and Islam (1992), where there are no explicit strategic interactions and feedback between adversaries.

Atkinson, Sandler, and Tschirhart (1987) use an extension to Nash's bargaining game, where time involved in negotiations is included (Cross, 1969, 1977). The duration of the incident and the bargains consummated depend on the costs of bargaining, the impatience of the adversaries, the duration of the incident, and the discovery of bluffs (i.e., threats not carried out). These authors utilize real-world data from hostage-taking incidents to test the underlying model of bargaining.

Another application of game theory involves terrorists' choice of targets for a three-player game involving two targeted nations and a common terrorist threat (Sandler & Lapan, 1988; Sandler & Siqueira, 2003). Each nation independently chooses its deterrence expenditures, which again determines the terrorists' logistical failure probability on that nation's soil. The terrorists pick the venue with the highest expected payoff for their attack. Each nation's choice of deterrence confers benefits and costs on the other target. By transferring the attack abroad, each nation imposes an external cost on its counterpart; however, by limiting attacks and their severity at home, each nation provides an external benefit to foreign residents. Moreover, an external benefit arises whenever the deterrence efforts of the nations sufficiently degrade the terrorists' expected benefits, so that they attack no one. The more fanatical the terrorists, the less likely the no-attack scenario. Sandler and Lapan (1988) show that the Nash equilibrium, where each nation chooses its deterrence in isolation, may result in too much or too little deterrence when compared with a social optimum, depending on the pattern of external costs and benefits.[3] If, for example, attacks in either country lead to no collateral damage on foreign residents or interests, then the countries will engage in a deterrence race as each tries to transfer the potential attack abroad, where it has no residents.[4] In a globalized society, where a country's risks from a terrorist attack are equal everywhere, independent deterrence choices imply too little deterrence as each country fails to account for the protection that its efforts confer on foreign residents (Sandler & Siqueira, 2003).

The game-theoretic approach reveals a couple of paradoxes. Countries may work at cross-purposes when deterring terrorist attacks. Although the United States is the target of approximately 40% of all transnational terrorist attacks, virtually all of these attacks occurred abroad in recent years, with 9/11 being a noticeable exception (Sandler, in press). U.S. overdeterrence means that it experiences attacks where it has little authority to do anything about them. Additionally, efforts to share intelligence on terrorists' preferences and resources may exacerbate this overdeterrence if deterrence decisions are not coordinated as nations use this information to augment efforts to transfer the attacks abroad (Enders & Sandler, 1995; Sandler & Lapan, 1988). This is a standard second-best result in economics in which there are two relevant policy variables—share intelligence and coordinate deterrence—but joint action only involves a single variable. This result highlights the beauty of a strategic approach. Standard intuition suggests that pooling information should enhance welfare, but this is not the case if this sharing worsens the deterrence race that wastes resources without necessarily increasing security against a determined terrorist group.

Another game-theoretic representation analyzes a situation of asymmetric information where the terrorists know their true strength but the targeted government must guess the terrorists' resources based on the level of their attacks. These attacks are intended to apply sufficient pressures, in terms of costs, to a government, so that it concedes to terrorist demands. In a deterministic setting, the outcome of the struggle between the adversaries would be known even before the first play of the game because, in the absence of ties, a finite game of perfect information has a unique subgame perfect equilibrium. If, for example, the government is aware that the terrorists possess sufficient resources to force the government to surrender eventually, then the optimal strategy is for the government to concede at the outset and suffer no attack damage. If, moreover, a well-informed terrorist group understands that it has insufficient resources to obtain its political demands, then it is optimal either to abandon the campaign or expend all of its resources at the outset.

A more interesting and relevant scenario is when the government is incompletely informed about the terrorists' capability.[5] Lapan and Sandler (1993) analyze this scenario in which a signaling equilibrium may allow a government to limit its expected costs from attacks, even though the likelihood of surrender may increase. In this scenario, the extent of terrorist incidents may provide information to the government about the type of terrorist group—strong or weak—that it confronts. Attacks, therefore, serve as a signal that the government can proceed to adjust its posterior beliefs concerning the resources of the terrorists (see also Overgaard, 1994). Such updated beliefs permit the government to decide whether to capitulate or resist. The terrorists face an interesting tradeoff—the use of large amounts of their resources at the outset may correctly or incorrectly convince the government that they are strong, but this outlay results in less future attacks if the government is unconvinced. A perfect Bayesian equilibrium for the two-period signaling game is derived in which the government prefers the associated partial-pooling equilibrium, where the government surrenders to groups whose first-period attacks exceed a certain threshold, over the never-surrender equilibrium. The pooling equilibrium is associated with some regret when the government misjudges the terrorists' true strength based on initial attacks. Intelligence is valued because it can reduce this regret by curtailing the variance of government priors.

Another interesting application of game theory to terrorism involves accommodations reached between terrorists and a host government (Lee, 1988; Lee & Sandler, 1989). In such scenarios, a terrorist organization has an implicit understanding that it can operate with impunity, provided that its attacks do not create collateral damage for the host country. This accommodation can undo efforts of other countries to retaliate against a terrorist group by reducing their cooperative payoffs. Thus, nations now have three options in their reaction to terrorists and their sponsors: do nothing, retaliate against the terrorists and their sponsors, or accommodate the terrorists. The last option helps the terrorists at the expense of the cooperating nations. Lee (1988) shows that this third option dominates the other two, thereby resulting in a PRISONER'S DILEMMA (Tucker, 1950/2001) where some nations seek such accommodations and, in so doing, undo the accomplishments of others to curtail the terrorist threat. Once again, pursuit of self-interest may harm others, owing to strategic considerations.

Proactive versus reactive policies

Governments' antiterrorism policies are either proactive or reactive. Proactive policy involves aggressively going after the terrorists and eliminating their resources, infrastructure, and personnel, whereas reactive policy concerns protective measures either to divert the attack or limit its consequences. A preemptive strike against the terrorists or their state sponsors (e.g., the Taliban in Afghanistan) is an example of a proactive policy. Because a preemptive attack, if successful, eliminates the terrorist threat for all potential targets, there is a tendency to free ride or rely on the efforts of others.

This is illustrated in Matrix *a* in Table 1, in which two players—the United States and the European Union (EU)—must decide whether to preempt a common terrorist threat. Suppose that preemption by each country confers 4 in benefits on both countries at a cost of 6 to the country doing the preemption. If, therefore, the United States preempts and the EU free rides, then the EU receives 4 in benefits, whereas the United States nets –2 (= 4 – 6) as costs of 6 are deducted from derived benefits of 4. The payoffs are reversed when the United States free rides while the EU takes action. If, however, both countries preempt, then each receives 2 in net benefits as preemption costs are deducted from gross benefits of 8 ($= 2 \times 4$). The resulting game is a PRISONER'S DILEMMA where no one takes an aggressive stance against the terrorists.

In Matrix *b*, a different scenario is depicted where, unlike the EU, the United States gains a net benefit from its own preemption because it is the favorite target for transnational terrorists. Suppose that U.S. preemption gives it 8 in benefits while conferring

TABLE 1: Three Alternative Game Forms for Preemption

	European Union	
United States	*Preempt*	*Do Not Preempt*
Matrix *a*: PRISONER'S DILEMMA		
Preempt	2, 2	–2, 4
Do not preempt	4, –2	Nash 0, 0
Matrix *b*: asymmetric-dominance equilibrium		
Preempt	6, 2	Nash 2, 4
Do not preempt	4, –2	0, 0
Matrix *c*: COORDINATION		
Preempt	Nash 2, 2	–4, 0
Do not preempt	0, –4	Nash 0, 0

just 4 in benefits to the EU as the United States counters its greater threats. Further suppose that preemption by the EU gives 4 in benefits to both countries. Preemption is again assumed to cost 6. If the United States preempts alone, then it nets 2 (= 8 – 6), whereas the EU still gets a free riding gain of 4. If, instead, the EU preempts alone, then the United States receives 4 as the free rider and the EU receives –2. When both the United States and the EU preempt, the United States nets 6 (= [8 + 4] – 6) and the EU nets 2. Now, the United States has a dominant strategy to preempt and the EU has a dominant strategy to free ride, leading to the asymmetric-dominant Nash equilibrium in the upper right-hand cell of Matrix *b*. This game representation may well characterize the U.S. position after 9/11, where U.S. action was going to yield high payoffs to the U.S. government. With the collapse of the twin towers of the World Trade Center and the damage to the Pentagon, the United States had to take some kind of decisive action to maintain legitimacy.

Up to this point, the status quo of doing nothing resulted in zero payoff. Suppose, however, that no action whatsoever leads to a world under siege by terrorists. In this scenario, the absence of any preemption may imply a negative payoff, which for Matrix *a* would result in a CHICKEN game (Rapaport & Chaman, 1969) whenever this payoff is less than –2 (not shown in Table 1), the payoff from acting alone. For CHICKEN, the Nash equilibria has some country responding. A COORDINATION game (Watson, 2002) may also apply to preemption if both countries must combine forces to achieve the positive payoffs of 2 (= 2 × 4 – 6). This game is displayed in Matrix *c* in Table 1. If only a single country preempts, then there are no benefits, but the preemptor incurs a cost of 4—thus, the off-diagonal payoffs are (–4, 0) and (0, –4). The two pure-strategy Nash equilibria in Matrix *c* correspond either to both countries preempting or neither preempting. Thus, preemption may be consistent with a number of different game forms depending on symmetry, penalties for the status quo, preemption thresholds, or other considerations. Even though preemption confers free rider benefits, a PRISONER'S DILEMMA may not result.

Next, consider other proactive policies and their associated game forms. In the case of group infiltration, the nation conducting the operation often secures benefits over

TABLE 2: Policy Choices and Underlying Games

Policies	*Alternative Game Forms*
Proactive policies	
Preemption	PRISONER'S DILEMMA, CHICKEN, COORDINATION, asymmetric-dominance equilibrium
Group infiltration	Asymmetric-dominance equilibrium, PRISONER'S DILEMMA
Retaliation	PRISONER'S DILEMMA, CHICKEN, COORDINATION, asymmetric-dominance equilibrium
Intelligence	Asymmetric-dominance equilibrium
Reactive policies	
Deterrence	PRISONER'S DILEMMA (deterrence race) with action
Embassy fortification	PRISONER'S DILEMMA with action
United Nations conventions	PRISONER'S DILEMMA in terms of enforcement

and above those conferred on other potential targets. This follows because the infiltrator can exploit the intelligence first and will target those terrorist groups that pose the greatest risks to its interests. As a consequence, an asymmetric dominance is likely to apply, not unlike the preemption scenario in Matrix *b* in Table 1. If, however, the infiltrator receives no special advantages, then a PRISONER'S DILEMMA is anticipated. Retaliation against a state sponsor of terrorism, such as the U.S. raid against Libya in April 1986 or the U.S.-led attack against Afghanistan following 9/11, is an analogous situation to preemption with lots of potential free rider benefits. As such, the same game forms as those associated with preemption are relevant. When intelligence is collected as a proactive policy, the asymmetric-dominance scenario is appropriate, especially if one nation is more often the target of transnational terrorist attacks. The intelligence collector gains relative to free riders.

Reactive responses include deterrence, embassy fortification, and United Nations (UN) conventions. For deterrence, each target takes protective actions to divert the attack. Any nation that spends less on deterrence becomes the more desirable target. This scenario often leads to a deterrence race, best described as a PRISONER'S DILEMMA with too much action (Sandler, in press).[6] An analogous scenario characterizes embassy fortification because the least fortified embassy becomes the target of opportunity for the terrorists. All potential targets increase fortification expenditures but do not necessarily eliminate the attack if the terrorists are bent on attacking someone. For UN conventions, a PRISONER'S DILEMMA results when it comes time to enforce the convention as each nation sits back waiting for others to act for the good of all. UN conventions on outlawing terrorism and its modes of operation (e.g., skyjackings, attacks against diplomatic personnel) have been shown to have no impact whatsoever in accord with the Nash equilibrium of no enforcement, associated with the PRISONER'S DILEMMA representation of enforcing UN conventions (Enders, Sandler, & Cauley, 1990). Table 2 summarizes the different proactive and reactive policies and their associated game forms.

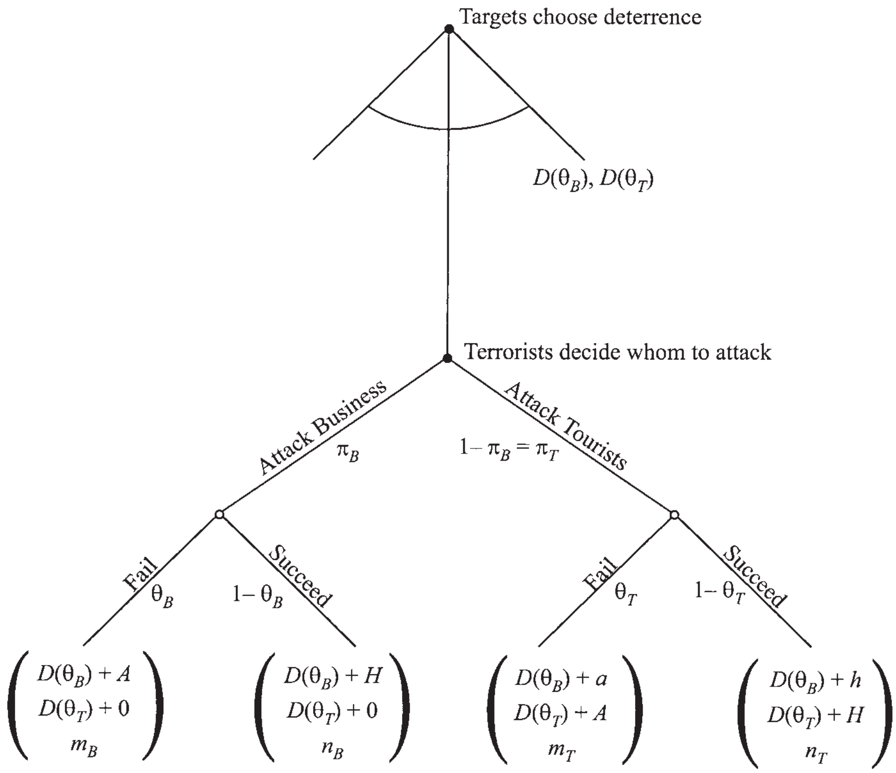

FIGURE 1: Deterrence Game Tree for Targets

Alternative targets

To illustrate the deterrence decision, we apply the model from the literature (Sandler & Lapan, 1988; Sandler & Siqueira, 2003) to a novel choice where a terrorist group can target a business (*B*) or tourist (*T*) venue. As such, this terrorist choice of targets may involve domestic or transnational terrorism. The terrorists are assumed to stage their attack at a single venue in each period.[7] Additionally, the terrorist group is assumed to be fanatical (i.e., gaining a net benefit even if the mission fails), so that the group will attack one of the two venues.[8] Figure 1 depicts the associated deterrence game tree where the targets go first and choose their level of deterrence or deterrence costs, that is, $D(\theta_B)$ for the business target and $D(\theta_T)$ for the tourist target. By choosing their deterrence expenditure to thwart attacks, the two targets affect the terrorists' perceived probability of success or failure, where θ_B and θ_T are the probabilities of logistical failure when the business or tourist venue is attacked, respectively. Thus, $1 - \theta_B$ and $1 - \theta_T$ are probabilities of logistical success for attacks at these two venues. Deterrence costs increase with θ_B at an increasing rate, so that $D'(\theta_B) > 0$ and $D''(\theta_B) > 0$. The same is true for deterrence costs spent to protect tourists.

The terrorists move second where they decide whom to attack, in which their attack probability π_i against target $i(= B, T)$ depends on their perceived probabilities of failure, that is, $\pi_i(\theta_i, \theta_j)$ for $i, j = B, T$ and $i \neq j$. This probability function is assumed continuous with $\partial\pi_i/\partial\theta_i < 0$ and $\partial\pi_i/\partial\theta_j > 0$, so that information is not complete with respect to terrorists' beliefs and values. The assumed partial derivatives indicate that efforts to decrease success in venue i through greater deterrence lowers the likelihood of attack there and transfers the attack to the other venue, so that independent deterrence decisions may work at cross-purposes.

Because terrorists are assumed fanatical, the game can end in four outcomes: terrorist failure or success at the business target, or terrorist failure or success at the tourist target. Within the four bold parentheses in Figure 1, the payoffs to business, tourists, and the terrorists are listed in descending order. Both targets confront costs that they want to minimize, whereas the terrorists receive benefits that they want to maximize. Terrorists' payoffs m_i and n_i ($i = B, T$) are not necessarily known with certainty, so that the terrorists must make an educated choice of target based on their anticipated gain from success and failure at each of the two potential targets.

We shall focus on the actions of the targets. The business (tourist) target must pay deterrence costs regardless of the game's outcomes; hence, this expense is like an insurance premium, paid in good and bad states. In Figure 1, the payoffs are depicted so that there is no collateral damage on tourists at a business venue. There is, however, collateral damage of a or h on business interests when a tourist venue is attacked. That is, a terrorist attack on an airport is sure to affect tourists and business people, whereas a terrorist attack on a specific business is unlikely to harm tourists. Some symmetry of costs is assumed in which a direct attack on a business or tourist location causes damage of A for a failure and H for a success, where $H > A$. For collateral costs to business interests, terrorist success is more costly that failure, that is, $h > a$.

Based on Figure 1, the expected costs to business from a business attack is

$$l(\theta_B) = \theta_B A + (1 - \theta_B)H, \quad (1)$$

whereas the analogous costs to tourists from a tourist attack is

$$l(\theta_T) = \theta_T A + (1 - \theta_T)H. \quad (2)$$

The collateral damage from a terrorist attack on business interests is

$$\nu(\theta_T) = \theta_T a + (1 - \theta_T)h, \quad (3)$$

whereas the collateral damage from a business attack on tourist interests is

$$\nu(\theta_B) = 0. \quad (4)$$

Given the assumption on A, H, a, and h, $l(\theta_i)$ decreases as θ_i increases for $i = B, T$, and $\nu(\theta_T)$ decreases as θ_T increases because expected damage falls as terrorist failure becomes more likely. When acting independently, the expected costs of terrorism to business—denoted by C_B—is

$$C_B = D(\theta_B) + \pi_B l(\theta_B) + \pi_T \nu(\theta_T). \tag{5}$$

Owing to the absence of collateral damage, the cost to tourists is

$$C_T = D(\theta_T) + \pi_T l(\theta_T) + 0. \tag{6}$$

A Nash equilibrium corresponds to each target choosing its respective deterrence level to minimize these expressions while taking the other target's deterrence as given.

To ascertain the relative efficiency of the Nash solution, we must find the social ideal and compare it with the Nash equilibrium. This ideal is obtained when the deterrence levels are chosen for the two targets to minimize the aggregate cost C:

$$C = D(\theta_B) + D(\theta_T) + \pi_B(\theta_B, \theta_T)(l[\theta_B]) + \pi_T(\theta_B, \theta_T)(l[\theta_T] + \nu[\theta_T]). \tag{7}$$

For the business target, the comparison is accomplished in a couple of steps. First, we derive the first-order condition for minimizing costs to the business target by taking the derivative of C_B with respect to θ_B. This condition includes marginal deterrence costs, the potential harm to business interests as attacks are diverted to the tourist site, and the marginal benefits of diverting attacks and limiting damage of a business attack. Second, we minimize social cost, C, with respect to θ_B. Third, we evaluate these first-order conditions for social cost at the θ_B that satisfies the Nash equilibrium where $\partial C_B/\partial\theta_B = 0$, denoted by θ_B^N.[9] This evaluation leads to just

$$l\left(\theta_T^N\right)\frac{\partial \pi_T}{\partial \theta_B} > 0$$

because there is no collateral damage on tourists at the business site. This term represents the external costs that the independent deterrence decision of the business target imposes on tourists by transferring more attacks to them. This inequality accounts for the potential deterrence race and implies that business interests independently spend too much on deterrence. Because there are no tourists to protect at the business venue, there is no opposing external benefits coming from this deterrence spending.

A different situation characterizes the deterrence decision of the tourist target because diversion of a potential attack protects business visitors to the tourist venue in two ways: (a) It limits the damage to business interests by increasing the likelihood of terrorist failure, and (b) it makes business interests safer at the tourist site by diverting the attack. To these external benefits, there is also the external cost of a greater likelihood of attack at the business venue as the attack is diverted. Thus, whether the tourists' Nash equilibrium implies underdeterrence or overdeterrence hinges on which of these opposing influences dominates.[10] Compared with the deterrence choice of the business target where there are no opposing external benefits, the tourists will either underdeter or overdeter to a smaller extent.

There is also a collective action rationale why tourists may actually underdeter and are less adept at deflecting attacks than a business target. Business firms or their employees at risk need only act unilaterally to increase their protection; in contrast, tourists at risk must mount a collective response (which is highly unlikely) or lobby the government for help. This lobbying effort is anticipated to take a long time before there is any government response and, in the meantime, tourists or the public at large remain vulnerable. Moreover, tourist targets are varied and diffused, whereas business targets are generally more specific and easier to guard. Tourist attacks have resulted in better protection at airports, monuments, bridges, and some public places, but not every location can be equally protected, which results in targets of opportunity for terrorists.

Next, consider what would happen if the two targets were officials and businesses (or the general public). Now, officials are better equipped to solve the collective action problem because they can allocate public funds to protect themselves as has been done at U.S. embassies and other government buildings. Thus, it is no wonder that the smallest number of transnational terrorist attacks are now against the military and government targets. Both the general public and businesses face the largest number of attacks (U.S. Department of State, 2002, p. 174). As alternative targets divert attacks, those least able to do so become the victims.

Deterrence or preemption

As shown in the last section, many terrorism-related games involve at least three players. Another instance concerns two targeted governments—the United States and the United Kingdom—that must choose whether to focus their antiterrorism policy on deterrence or preemption. Deterrence diverts the attack by making such acts more difficult, whereas preemption seeks out the terrorists by eliminating their base of operations and resources. Each player has two choices: Each government can either concentrate on deterrence or preemption, and the terrorists can either execute a spectacular terrorist event (e.g., 9/11 or the 1998 simultaneous bombings of the U.S. Embassies in Nairobi, Kenya, and Dar es Salaam, Tanzania) or a normal terrorist event.

For a spectacular terrorist event, preemption costs exceed deterrence costs, whereas for a regular terrorist event, deterrence costs exceed preemption costs. The U.S.-led October 2001 attack on the Taliban and al-Qaida in Afghanistan was an effort to preempt future spectacular events and was extremely expensive. When a government protects against a spectacular event, intelligence can limit deterrence costs so that only key sites are afforded increased security. Deterrence can, however, be quite expensive for normal terrorism because potential targets everywhere must be guarded. Because regular events are planned by terrorists with less security precautions than spectaculars, preemption costs are anticipated to be less costly than protecting such a target-

TABLE 3: Two-Country Choice of Deterrence Versus Preemption

	United Kingdom	
United States	*Deterrence*	*Preemption*
Matrix *a*: Spectacular event		
Deterrence	$-S$, 0, S	0, $-pS - C$, pS
Preemption	Nash $-pS - C$, 0, pS	$-C$, $-C$, $-L$
Matrix *b*: Normal event		
Deterrence	$-T - c$, $-c$, T	$-c$, 0, $-l$
Preemption	0, $-c$, $-l$	Nash 0, 0, $-l$

rich environment. To simplify the mathematics without changing the strategic aspects of the underlying game, we normalize deterrence costs for spectaculars to zero and let C be preemption costs above and beyond deterrence costs for such events. Similarly, we normalize preemption costs for regular terrorism to zero and let c be deterrence costs above and beyond preemption costs for such events.

The three-player deterrence-preemption game is displayed in normal form in Table 3, where the United States chooses the row, the United Kingdom the column, and the terrorist group the matrix. Matrix *a* corresponds to a spectacular, whereas Matrix *b* corresponds to a normal terrorist event. In each cell, the first payoff is that of the United States, the second is that of the United Kingdom, and the third is that of the terrorist group. The payoffs depend on the following technologies of deterrence and preemption. To thwart a spectacular event with certainty, the United States and the United Kingdom must preempt. If only one country preempts and the other deters, then the event occurs with probability p in the country doing the preempting. If, however, neither country preempts and a spectacular event occurs, then it is logistically successful. A best-shot technology of preemption applies to a normal terrorist event, that is, preemption by either country is sufficient to make the event fail with certainty.

To compute the payoffs in Table 3, we must define a couple more terms: S is the terrorists' payoff for a successful spectacular, and T is their payoff for a successful normal event, where $S > T$. Deterrence by a single nation deflects the attack to the other target, whereas deterrence by both countries makes the United States the target of choice. With the United States being the target of 40% of all transnational terrorist attacks, this is a reasonable assumption. In Matrix *a*, if both countries deter, then the spectacular is successful against U.S. interests, so that the United States loses S (hence, the $-S$ payoff), and the terrorist group gains S. Owing to our cost normalization, the United Kingdom's deterrence cost is zero. If only one country deters and the other preempts, then the former nets zero and the latter endures an expected cost of $-pS - C$, which equals the expected loss from the attack and the spent preemption costs. The terrorists receive an expected payoff of pS over and above the costs of the operation.[11] In the case of mutual preemption, the spectacular is averted, so that the countries only cover their preemption costs of C, whereas the terrorist group loses the costs of their operation, L, which may include fallen comrades and wasted resources.

Next, consider the payoffs in Matrix *b*, associated with a normal terrorist attack. Mutual deterrence will lead to a terrorist attack on the United States, whose net payoff is the damage, $-T$, plus the costs of deterrence, $-c$. With no preemption, the attack is not stopped. The United Kingdom loses just its costs of deterrence, which has shifted the attack to the terrorists' preferred target—the United States. The terrorists obtain a gain of T, equal for the simplicity to the U.S. losses. When one country preempts and the other deters, the attack is foiled, so that the former nets zero and the latter pays its deterrence costs. The terrorists lose their logistical costs of l, where l is less than L because more planning and effort goes into a spectacular. Finally, if both countries preempt, then each nets zero and the terrorists lose their logistical costs.

There are two potential pure-strategy Nash equilibria for this game. For normal events, preemption is a dominant strategy for both countries because preemption costs are less than deterrence costs and preemption can remove the threat. Even though only one country needs to preempt, both have incentive to do so despite the redundancy of effort. When the payoffs in the lower right-hand cell in Matrix *b* and Matrix *a* (in Table 3) are compared for the terrorists, they prefer planning the normal event. Thus, the lower right-hand cell is a Nash equilibrium. For a spectacular, the two countries do not necessarily have a dominant strategy; because S is so large compared with C, we anticipate that $(1 - p)S > C$ unless the probability of a terrorist success is near certainty.[12] If, however, this inequality does not hold, then the pure-strategy Nash equilibrium in matrix *a* is in the upper left-hand cell because $S > T$ for the terrorists. Given a Nash equilibrium in Matrix *a* and Matrix *b*, there is also a mixed-strategy equilibrium involving a randomization of strategies. The existence of an equilibrium in both matrices means that the failure of the targeted governments to coordinate their preemption policy will mean that a spectacular will succeed on occasions, despite the huge costs that such events imply. This outcome accords with the facts, where spectaculars with hundreds of deaths occur about once every 2 years (Quillen, 2002a, 2002b; U.S. Department of State, 2002).

The outcome of this game is also descriptive of real-world policy where nations rarely coordinate their deterrence or preemption decisions. In Matrix *a*, this coordination failure leads to insufficient preemption so that spectacular events occur, whereas in Matrix *b*, this coordination failure results in too much preemption. If the Nash equilibrium is as indicated in Matrix *a* of Table 3, then observation of deterrence by the terrorists encourages a spectacular insofar as $S > T$ and $pS > -l$, which holds whether the terrorists observe U.S. or U.K. deterrence. If, however, the terrorists only observe preemption by either country, then they require more information to make the best choice (i.e., $pS > -l$ and $-L < -l$).

Granting concessions and terrorist types

An unresolved issue in the literature on negotiating with terrorists is the unintended consequences of increasing violence by conceding to the demands of the more moderate elements within a terrorist organization. The appeasement of the moderates

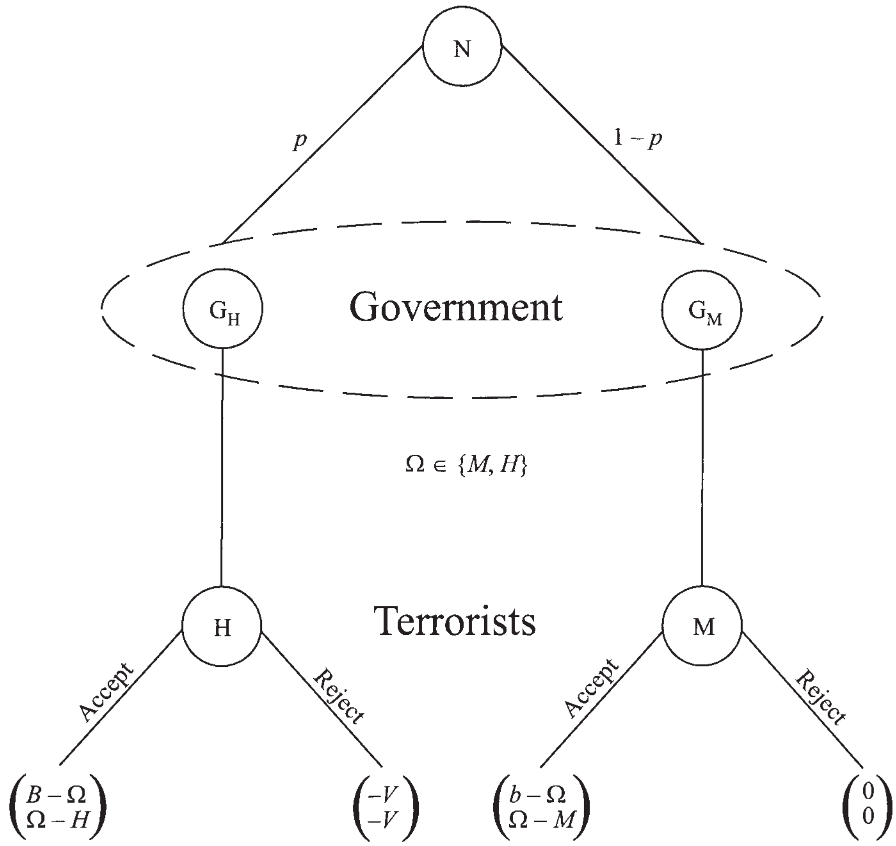

FIGURE 2: Granting Concessions and Terrorist Types

isolates hardliners, thereby leaving an adverse selection of terrorists more inclined to violence. Adverse selection involves one side of a transaction being more informed than another, which implies asymmetric information. In the classic example of a used car market, the uninformed buyers reduce their offers on cars to an average price where many reliable cars are taken off the market, so that a preponderance of "lemons" remain (Akerlof, 1970).

Figure 2 displays a model of bargaining between a government (G) and a terrorist group with moderate (M) and hardline (H) members in proportion p and $1 - p$, respectively. Hardliners make demands of H, whereas moderates make demands of M, where $H > M$. The government moves first and makes an offer, Ω, of either M or H to the terrorist group. The government is uncertain about the actual distribution between hardliners and moderates; hence, nature (N) moves first and selects the terrorist group's composition so that the government faces hardliners (node G_H) in proportion p and moderates (node G_M) in proportion $1 - p$.

Due to the government's lack of information about terrorists, nodes G_H and G_M are contained within the same information set, thereby implying that the government cannot tailor one offer to moderates and another to hardliners. Given $H > M$, there is an incentive for moderates to posture as hardliners because they are better off receiving an offer of H from the government, the acceptance of which yields a payoff of $H - M > 0$, as indicated by the third-from-the-left payoff in Figure 2. In contrast, the acceptance on an offer of M gives the moderates the baseline payoff of zero as their demands are met.

The government's benefits from an agreement (prior to deducting the offer) are B when hardliners accept and b when moderates accept, where $B > b \geq M$ and $H > b$. Although either terrorist type accepts an offer of H, the government experiences a cost, $b - H < 0$, from appearing weak when it concedes H to moderate terrorists. If the government offers H to hardliners, then the terrorists accept, giving the hardliners a payoff of zero and the government a net payoff of $B - H$, which may be positive. When, alternatively, hardliners reject an offer of M, we assume that they resort to terrorism and violence. For simplicity, we assume that this terrorism costs hardliners V to commit while it inflicts damages of V on the government—the $(-V, -V)$ payoffs. We assume that $M - H < -V$, so that hardliners prefer a terrorist attack to accepting a moderate offer. Rearranging this inequality makes the definition of a hardliner more apparent: $M < H - V$ implies that hardliners are willing to engage in violence at cost V to obtain demand H in lieu of the moderate concessions. Because moderates are satisfied with M, there is no net benefit to them or to the government if they reject such an offer, leading to the far-right (0, 0) payoffs.

What kind of offer should a government make? If it concedes a moderate offer, then hardliners will reject it, whereas moderates will accept it. Given the uncertainty over the composition of the terrorist organization, the government's expected payoff for a moderate offer is

$$-pV + (1 - p)(b - M). \tag{8}$$

In comparison, both hardliners and moderates will accept an offer of H. The expected payoff for the government is

$$p(B - H) + (1 - p)(b - H). \tag{9}$$

We are concerned when a moderate offer leaves an adverse selection of hardliner terrorists for the government to contend with. The government makes a moderate offer only if the expected payoff in Equation 8 is as least as great as that in Equation 9, which holds for

$$(H - M)/(V + B - M) \geq p. \tag{10}$$

If the government's belief that it is facing hardline terrorists, p, is less than the left-side of Equation 10, then an adverse selection results, in which moderates are placated and hardliners resort to violence. This outcome is contrary to the ultimate goal of the government to end violence.[13] A notable feature of Equation 10 is that the benefits to making a moderate offer, b, do not even figure into the equilibrium condition.

The possibility of adverse selection increases with hardliners' demands, H. Moreover, this likelihood falls with higher costs of violence, V. As B approaches H in value, so that government's and hardliners' preferences converge, V becomes the sole determinant of the adverse-selection equilibrium. This may explain why the Diaspora or state sponsors are essential in perpetuating violence, even though the government and hardliners agree on the need for a solution. Subsidies from outside interests reduce V and lead to an adverse selection that perpetuates hostilities.

Unanswered questions and future directions

There are many unanwered questions with respect to terrorism, for which game theory can be fruitfully applied. In terms of noncooperative game theory, there is no true multiperiod analysis of terrorist campaigns, where the terrorist resource allocation is studied over time. The closest analysis is that of Lapan and Sandler (1993) and Overgaard (1994), where terrorists signal their alleged strength in the initial period, but only a two-period analysis is presented. A many period investigation is required where the conflict between terrorists and the government has an unknown (probabilistic) endpoint that is influenced by actions of the adversaries. Ideally, information must be treated as incomplete or imperfect when investigating the temporal and strategic aspects of terrorist campaigns.

Another area for study involves the use of differential game theory to examine how terrorist organizations—their personnel and resources—are influenced by successful and failed operations. By applying a differential game framework, the analyst can display the dynamics of the strategic choices of the terrorists and the government in which the underlying constraints capture the rate of change over time of resource supplies based on terrorists' operations and the government's policy choices. The genesis and demise of terrorist groups can be analyzed based on strategic considerations. If, for example, this demise is understood, then governments may better plan their antiterrorist policies. Do some policies (i.e., harsh reprisals) actually encourage recruitment, opposite to the government's intent? This question and others can then be addressed.

To date, cooperative game theory has not been applied to the study of terrorism. In contrast to governments that have not cooperated effectively, terrorists have formed elaborate cooperative networks (Arquilla & Ronfeldt, 2001). Terrorists share training facilities, intelligence, operatives, innovations, and logistical methods. This cooperation is motivated by the weakness of terrorists who must pool resources and knowledge if they are to threaten much stronger governments. Terrorists often harbor common resentment of governments that unite their interests. Moreover, terrorists are in repeated interactions with one another, where leaders have no known endpoints unlike officials in democratic governments whose time horizons are quite finite. These repeated interactions allow terrorist groups to maintain agreements through trigger tit-for-tat mechanisms that punish defections. A careful cooperative game-theoretic analysis of terrorist groups can provide some useful insights in assessing the true threat of terrorism. Terrorists' cooperation allows them to prod government defensive measures to uncover the weakest link, to which they dispatch their best-shot team to create maximum damage.

Concluding remarks

Strategic interaction between terrorists and governments, among targeted governments, and among terrorists make game theory an appropriate tool to enlighten policy makers on the effectiveness of antiterrorist policies. This article takes stock of past game-theoretic applications to the study of terrorism. The article also presents novel applications and suggests some future applications. Since 9/11, academic interest in modeling and studying terrorism is increasing greatly.

Notes

1. The relevant literature includes Lapan and Sandler (1988, 1993), Lee (1988), Overgaard (1994), Sandler (in press), Sandler and Enders (in press), Sandler and Lapan (1988), Scott (1991), and Selten (1988).

2. Terrorists' high success rate and their ranking of tactics based on risk, time, and potential confrontation with authorities also support the rationality assumption (Sandler, Tschirhart, & Cauley, 1983).

3. At the social optimum, the sum of the two countries deterrence spending is minimized to identify the cooperative outcome.

4. As for an arms race, a PRISONER'S DILEMMA applies in a 2×2 representation (Sandler, in press).

5. Scott (1991) examines terrorism when the terrorists are uninformed about the type of government that they confront, but Scott does not consider a signaling equilibrium.

6. This deterrence race is analogous to the problem of the open-access commons, where mutual action in the form of exploitation of the shared resource leads to a PRISONER'S DILEMMA and a suboptimal Nash equilibrium (Sandler & Arce M., in press).

7. Of course, multiple attacks can be easily addressed, but for simplicity, we assume a single attack per period.

8. Other scenarios are addressed in Sandler and Lapan (1988). As groups have become more religious-based over the past two decades, the act itself gives a positive payoff to the terrorists.

9. For more complex and different scenarios, these steps are displayed in detail in Sandler and Siqueira (2003).

10. Overdeterrence or underdeterrence depends on the sign of the following composite term:

$$\pi_T v'\left(\theta_T^N\right) + v\left(\theta_T^N\right)\frac{\partial \pi_T}{\partial \theta_T} + l\left(\theta_T^N\right)\frac{\partial \pi_B}{\partial \theta_T},$$

where the first two terms are negative and the third term is positive. If the overall sum is negative (positive), then there is underdeterrence (overdeterrence).

11. This payoff could be easily changed to $pS - L$ with no change in the analysis. Moreover, a cost less than L could be deducted. Our simple model in Table 3 only allows the government to choose between preemption and deterrence. In reality, governments preempt and deter at the same time. To allow for this

possibility, one should view the government strategies in the model as a greater reliance on deterrence or preemption. Different payoff arrays and alternative underlying "technologies" of deterrence and preemption can be incorporated into the model. A rich number of interesting models and outcomes can be displayed, as noted by Hirofumi Shimizu.

12. This inequality implies that $-pS - C > -S$.

13. There is, indeed, nothing in the model that restricts us from setting B equal to V.

References

Akerlof, G. (1970). The market for lemons: Qualitative uncertainty and the market mechanism. *Quarterly Journal of Economics*, *84*(3), 488-500.

Arquilla, J., & Ronfeldt, D. (2001). *Networks and netwars*. Santa Monica, CA: RAND.

Atkinson, S. E., Sandler, T., & Tschirhart, J. T. (1987). Terrorism in a bargaining framework. *Journal of Law and Economics*, *30*(1), 1-21.

Cross, J. G. (1969). *The economics of bargaining*. New York: Basic Books.

Cross, J. G. (1977). Negotiations as a learning process. *Journal of Conflict Resolution*, *21*(4), 581-606.

Enders, W., & Sandler, T. (1993). The effectiveness of anti-terrorism policies: Vector-autoregression-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., & Cauley, J. (1990). UN conventions, technology and retaliation in the fight against terrorism: An economic evaluation. *Terrorism and Political Violence*, *2*(1), 83-105.

Islam, M. Q., & Shahin, W. N. (1989). Economic methodology applied to hostage-taking in light of the Iran-Contra affair. *Southern Economic Journal*, *55*(4), 1019-1024.

Lapan, H. E., & Sandler, T. (1988). To bargain or not to bargain: That is the question. *American Economic Review*, *78*(2), 16-20.

Lapan, H. E., & Sandler, T. (1993). Terrorism and signaling. *European Journal of Political Economy*, *9*(3), 383-397.

Lee, D. R. (1988). Free riding and paid riding in the fight against terrorism. *American Economic Review*, *78*(2), 22-26.

Lee, D. R., & Sandler, T. (1989). On the optimal retaliation against terrorists: The paid-rider option. *Public Choice*, *61*(2), 141-152.

Overgaard, P. B. (1994). Terrorist attacks as a signal of resources. *Journal of Conflict Resolution*, *38*(3), 452-478.

Quillen, C. (2002a). A historical analysis of mass casualty bombers. *Studies in Conflict & Terrorism*, *25*(5), 279-292.

Quillen, C. (2002b). Mass casualty bombings chronology. *Studies in Conflict & Terrorism*, *25*(5), 293-302.

Rapaport, A., & Chaman, A. M. (1969). The game of CHICKEN. In I. R. Buchler & H. G. Nutini (Eds.), *Game theory in the behavioral sciences* (pp. 151-175). Pittsburg, PA: University of Pittsburg Press.

Sandler, T. (in press). Collective action and transnational terrorism. *World Economy*.

Sandler, T., & Arce M., D. G. (in press). Pure public goods versus commons: Benefit-cost duality. *Land Economics*.

Sandler, T., & Enders, W. (in press). An economic perspective on transnational terrorism. *European Journal of Political Economy*.

Sandler, T., & Lapan, H. E. (1988). The calculus of dissent: An analysis of terrorists' choice of targets. *Synthèse*, *76*(2), 245-261.

Sandler, T., & Siqueira, K. (2003). *Global terrorism: Deterrence versus preemption*. Unpublished manuscript, University of Southern California.

Sandler, T., Tschirhart, J. T., & Cauley, J. (1983). A theoretical analysis of transnational terrorism. *American Political Science Review*, *77*(4), 36-54.

Scott, J. L. (1991). Reputation building in hostage incidents. *Defense Economics*, *2*(2), 209-218.

Selten, R. (1988). A simple game model of kidnappings. In R. Selten (Ed.), *Models of strategic rationality* (pp. 77-93). Boston: Kluwer Academic Press.

Shahin, W. N., & Islam, M. Q. (1992). Combating political hostage-taking: An alternative approach. *Defense Economics*, *3*(4), 321-327.

Tucker, A. (2001). PRISONER'S DILEMMA: A two person dilemma. In E. Rasmussen (Ed.), *Readings in games and economics* (pp. 7-8). Malden, MA: Blackwell. (Original work published 1950)

U.S. Department of State. (2002). *Patterns of global terrorism*. Washington, DC: Author.

Watson, J. (2002). *COORDINATION—Strategy: An introduction to game theory*. New York: W. W. Norton.

Todd Sandler is the Robert R. and Katheryn A. Dockson Professor of International Relations & Economics, University of Southern California. He has used game theory in the study of terrorism since 1983. In 2003, he was a corecipient of the National Academy of Sciences' Estes Award for excellence in science for his research on terrorism.

Daniel G. Arce M. is the Robert D. McCallum Distinguished Professor of Economics & Business, Rhodes College. He has applied game theory in economics and international relations to a wide variety of problems. His teaching and research interests involve many aspects of game theory.

ADDRESSES: TS: School of International Relations, University of Southern California, Von Kleinsmid Center 330, Los Angeles, CA 90089-0043, USA; telephone: +1 213-740-9695; fax: +1 213-742-0281; e-mail: tsandler@usc.edu. DGAM: Department of Economics, Rhodes College, 2000 North Parkway, Memphis, TN 38112-1690, USA; telephone: +1 901-843-3121; fax: +1 901-843-3736; e-mail: arce@rhodes.edu.