Modeling Terrorism:
A Game Theoretical Approach

By Adam Samson[1]
Presented to Professor Tom Parker

[1] Adam Samson is an economics major at New York's Bard College. His contact information can be found at http://www.adamsamson.com/contact.shtml

“*However beautiful the strategy, you should occasionally look at the results…*”

--*Winston Churchill*

Game theory is a branch of mathematics that considers so-called strategic interactions between parties and has a wide range of uses in quantifying fields that are not often quantified. The application that this paper will consider is game theory as applied to terrorism and counterterrorism. Terrorism is a natural concept to analyze using game theory because when reduced to its simplest level, many terrorist events can be summed up to simple or complex strategic interactions.

This paper will examine terrorism on the micro level and the macro level. Micro level terrorism is individual terrorist situations such as hijackings or hostage takings. Macro terrorist events are events that entire countries respond to, such as large scale terrorist attacks, or attempting to destroy terrorist groups. Before we can use game theory to model terrorist situations, it's important to develop the basics of how game theory works. Equally, it's important to define the terms that are used in the game theoretic and terrorism dialogue. Clearly, this paper does not purport to be either a guide to game theory, or quantitative terrorism analysis, rather we hope that it serves as a "jumping off point" for the non-mathematician to contemplate terrorism in a quantitative light.

Game Theory: A Primer

Game theory is a branch of mathematics that seeks to mathematically model strategic situations. New York's New School has produced a substantial amount of work regarding game theory and defines its purpose as follows:

“The main purpose of game theory is to consider situations where instead of agents making decisions as reactions to exogenous prices ("dead variables"), their decisions are strategic reactions to other agents actions ("live variables").” (The New School)

Not surprisingly, that statement is filled with technical jargon, but is important because it provides a basis for why game theory matters for the field of terrorism and counter-terrorism. The outcome of a terrorist situation, say a hostage taking, depends largely on interaction between the hostage-takers and those who are negotiating for the release of the hostages. Largely every “move” that the hostage takers make will force the negotiators to make a make a move, and vice versa.

Each party will take into account the utility that they will derive from the outcome, but the variables are very much alive and constantly changing. This type of system doesn’t tend toward a set equilibrium, but rather the equilibrium depends upon the utility each player derives in every interaction. Another way of thinking about it is that game theory attempts to model the actual outcome of a situation, and not the most efficient one. Hostage takings, and terrorist situation in general tend to be rather complex, multi-step “games,” so in this section we’ll use a very simple (and popular) two player game as an example, and then once we have access to a few game theory tools, we’ll begin analyzing terrorist situations.

The game that we are going to use is referred to as the prisoners’ dilemma. It was developed by two mathematicians in the 1950s and formalized in the late 1990s (Stanford Encyclopedia of Philosophy). We will alter the setup of the game for our own purposes, but keep it similar enough as to make it reflective of the original intent. In game theory it is often useful to begin with a written narrative of the game to be analyzed before attempting to quantify it. The narrative of our own version of the prisoners' dilemma follows:

> Phillip and Christopher walk down a deserted street in Midtown Manhattan and see a lovely BMW car with its doors unlocked and keys inside. The two are by no means experienced car thieves, but they have committed several small scale petty-thefts, and routinely use illegal drugs, however they haven't been caught for in the past. They do know a crime of opportunity when they see it, however, and they know that they will have to act quickly to steal the car and get away. So, they hop into the car and drive about 100 miles and pull into a small motel that they feel as far enough from the city that the car will not be recognized.
>
> They park the car, and book separate rooms. Unbenounced to them, being naive thieves, they made a stupid error. It turns out that the car they stole had a GPS system within it. The owner of the car, obviously, informed the police immediately after it was stolen and it wasn't long before the police found the two culprits in their rooms and arrested for possessing small amounts of illicit substances.
>
> The suspects, who never discussed an exit strategy should they be arrested, were thrown into two separate interrogation rooms and were interrogated. The interrogators (who, as it turns out, only have enough

> evidence to implicate each man on drug possession charges, and not grand theft auto) made each suspect the same offer: implicate your partner, and you'll go free and he'll serve twenty year, unless he also implicates you as well in which case you'll both serve ten years. If they both refuse to talk then they'll each serve one year, but neither suspect knows what the other suspect is going to do.

Let's begin quantifying this rather complex sounding narrative. First, we're going to set up what is referred to as a payout matrix, that is we need think about every possible outcome and figure out the what ends up happening to each suspect (the payoff). We know the following:

a) If each suspect implicates the other then they both serve ten years in prison;
b) If one suspect implicates the other, but the other one stays silent, then the accuser goes free, and the silent one serves twenty years;
c) If both stay silent, then each suspect serves a one year on drug charges.

The first aspect to observe is the efficient outcome, which if we look at the options is both individuals staying silent. The problem is that each suspect has imperfect information about what the other suspect is going to do since they didn't previously make a plan, and they can't talk to each other. Let's now consider the possible interactions between the two suspects. To do this, we're going to actually draw the payoff the payoff matrix.

	Phillip		
Christopher		Implicate (i)	Stay Silent
	Implicate	10,10 (a)	0,20 (b)
	Stay Silent (ii)	20,0 (b)	1,1 (c)

The matrix, which is the basis for mathematically representing games in game theory, looks a lot more intimidating than it actually is. The payoff matrix represents each the payoff that each player receives from each interaction in the game. The boxes on the top and left before the names represent the actions that each player can take. So column (i) represents Phillip implicates Christopher, and row (ii) represents Christopher remaining silent. We read payoff matrices such that the number before the comma represents the payoff of the player to the left, and the payoff after the comma represents the payoff of the player on top (Schelling).

Let's take a look at the upper left box that contains "10,10." This box corresponds to (a) in the list of items that we know above. That is, if both Phillip and Christopher implicate each other, they will both serve ten years. The bottom left box and top right box are the same concept applied to each player. If Phillip implicates Christopher, and Christopher stays silent, then Christopher will serve twenty years, and Phillip will go free (20,0). If Christopher implicates Phillip, and Phillip stays silent, then Phillip will serve twenty years, and Christopher will go free (0,20). Now that we understand how the payoff matrix is constructed, let's look at how the interaction will play out.

First, we need to make an assumption about the mindset of each of the suspects. This proves easy in this type of scenario, but will prove much more difficult when we consider terrorist situations. We can rather safely assume that both suspects want to minimize the time that they serve, and care little about how much time the other one serves. This selfishness might sound mean, but it represents rationality and makes our example work more cleanly. To analyze this game all we have to do is look at each perspective on the payoff matrix. We will provide sub-matrices for each decision and keep a running tally on the full matrix at the end for use later.

If Christopher believes Phillip will implicate him, what should he do?

	Phillip	
Christopher		**Implicate**
	Implicate	**10,10**
	Stay Silent	20,0

Christopher can either implicate Phillip or have to serve ten years in prison, or stay silent and get twenty years in prison. The obvious choice under this assumption is for Christopher to implicate Phillip.

What if Christopher believes Phillip is going to stay silent?

	Phillip	
Christopher		**Stay Silent**
	Implicate	**0,20**
	Stay Silent	1,1

Again, Christopher can either implicate Phillip and go free or Stay Silent and serve one year. Obviously, Christopher is going to choose to implicate Phillip under this assumption. We can perform exactly the same analysis on the other two sub-matrices and arrive at this final matrix in which we kept a running tally of each decision.

	Phillip		
Christopher		Implicate	Stay Silent
	Implicate	**10,10 ****	0,20 *
	Stay Silent	20,0 *	*1,1*

The box with the most tick marks turns out to be the upper left box. This box is referred to as the "Nash equilibrium" for this particular game. Games can have one, several, or no Nash equillibria. A Nash equilibrium, basically speaking, is an equilibrium wherein all sub-matrices are optimally selected (ISCID Encyclopedia of Science and Philosophy). It's important to note that the Nash equilibrium doesn't take into account the entire matrix, as individual players seldom consider the entire payoff matrix when making their decisions; rather they only consider their own sub-matrices. This makes sense on the conceptual level, after all, one can control his own choices, but really doesn't know how another individual or entity will react. That is why in this case it is not surprising that the Nash equilibrium is for both suspects to implicate the other suspect. If each suspect *rationally* thinks through the options, in every case he will choose to implicate because that creates the optimal situation for that suspect. It ends up damning both suspects, in a sense, because they both end up getting ten years in prison whereas they could have easily had one year each (the most efficient outcome) if they just worked together.

Terrorism and Game Theory

This is where the real beauty of game theory in the context of terrorism and counter terrorism begins to show. It is often the case that players on both sides of the terrorism game, the terrorist and the player that is trying to fight the terrorist (be it a police department, a state, *et cetera*) quite often only consider their own sub-matrices when deciding what moves to make. The terrorist generally seeks to achieve his goals, and the other player usually seeks to keep the terrorist at bay (Dietrich).

What if we could model the terrorist's sub-matrix? Then we could, theoretically, think of ways to interact that leave us with the outcome that we desire. Indeed, this is usually the goal of hostage negotiation situations. Game theory allows us to both extend these negotiation tactics to other parts of terrorism and quantify and formalize the field. It is quite possible that this is already done by state players unconsciously as they draft plans for responding to terrorists, but this formalization offers an interesting and possibly quite useful perspective. Let's consider a fictional situational case to how game theoretical analysis would apply to a terrorist.

The Micro Level

As we did in the simpler prisoner's dilemma scenario, we will begin with a narrative about the situation involving a non-suicidal bomber. The following scenario was suggested by Weaver *et al.* in their paper "Modeling and Simulating Terrorist Decision-making." Although we retain the general form of the game, we have amended (and simplified) it to suit the needs of this paper[2].

> A terrorist group decides that it would be worthwhile to attack an oil derrick in a village in Northern Iraq. The group realizes that harming a single oil derrick wouldn't actually cause a major disruption in the oil supply, but they deem that it is of strategic importance because an attack on an oil derrick would most likely cause speculation in worldwide oil markets that could easily lead to a spike in oil prices on an international level. Since the terrorist group is attempting to propel itself onto the world stage, this, they feel, is an excellent move.
>
> The United States has an army checkpoint in front of the oil derrick for that exact reason. The oil derrick is a relatively small one, so the checkpoint consists of a 50-caliber gun and a couple of guards. The plan is to drive by the guards, shoot them as they pass, drive by, plant the bomb, detonate the bomb, and run away.

[2] It is interesting to note that NATO actually uses a similar scenario derived from the same paper in a terrorism brief that they produced. This serves as a "proof of concept" for this paper.

The group has instilled in the bombers that although it is preferable to come back from bombing alive, it is better to die rather than get caught (which is an embarrassment). In particular, it would be especially embarrassing to get caught be a guard that is caught be surprise.

This situation is one that many counter-insurgency forces and states face each day. Protecting strategic assets using checkpoints is very common, and certainly something that militaries across the globe have become quite familiar with. Let's use a game theoretical analysis to quantify the situation. The possible outcomes are as follows:

a. Terrorist Drives Through Checkpoint Unnoticed;
b. Terrorist drives through the checkpoint and shoots the guard, killing him;
c. Terrorist drives through the checkpoint, shoots the guard, misses, and gets killed;
d. Terrorist gets caught by the guard.

With each outcome, we can assign a corresponding utility for the terrorist and the guard (who can either be ready for the attack or unready). In so-doing we're creating an informal qualitative payoff matrix that will be used to construct a more formal quantitative payoff matrix.

	Terrorist			
	Drive through unnoticed	Shoot the guard and escape	Shoot the guard and die	Get caught
Ready Guard	Embarrassing, Excellent	Very Embarrassing, Good	Very Embarrassing, OK (remember, death > capture) Good	Excellent, Very Embarrassing
Unready Guard	Very Embarrassing, Very Good	Most Embarrassing, Good	OK (at least the terrorist is dead), Good	Very Good, Most Embarrassing

Next, we turn the qualitative payoff matrix into the proper quantitative one:

	Terrorist			
	Drive through unnoticed	Shoot the guard and escape	Shoot the guard and die	Get caught
Ready Guard	-1, **3**	-2, 3	0, 2	3, -2
Unready Guard	-2, 3	-3, 2	0, 1	2, -3

Although this is a fictional situation, there are several factors that we can glean from this. The first, albeit obvious, observation is that a ready guard is of far more use than an unready guard. The utility of a ready guard is always higher than that of an unready guard. This makes sense, but is further formalized in the game theoretical sense. There are also less obvious observations that one can draw. One is that getting caught is the only situation that provides the terrorist with deep embarrassment. Planting the bomb and escaping as well as shooting the guard and escaping provides the greatest level of utility (slightly less so if they are noticed by an unready guard). This is because both of those situations allow the terrorist to complete the objective and return home and continue helping out the group. The next level down is shooting the guard and dying. Although seemingly nonsensical, it is preferable for the terrorist to die in a shoot-out with guards and die then to be caught and imprisoned.

The utility values are determined based on cardinality that implies that the individual numbers don’t matter, but rather the relations of each number to other numbers does matter. It would be possible, the Weaver paper suggests, for those protecting against terrorist attacks to conduct data-mining and come up with actual utility values. This data would be invaluable to people on the ground because it would enable them to make the most intelligent strategic interactions in their particular situation. Knowing how much more the terrorist prefers dying than being caught is important in coming up with strategic plans to fight terrorists in this situation and similar situations.

The Macro Level

Now that we have considered game theoretical approaches to analyze micro, situational terrorist events, let’s move on to assess broader macro-level terrorism and counter-terrorism using game theory. Mathematicians have considered many macro-level situations using game theory. Unfortunately, many such situations are highly cumbersome and advanced mathematical approaches that are out of the scope of this paper. However, Lee and his associates performed a particularly novel and reasonably simple game theoretical analysis of whether countries should be proactive or reactive regarding terrorism.

Countries are often faced with a situation wherein the country and one of its allies has a common terror threat. Each country has two choices: proactively address the threat, or wait until the threatened attack actually materializes and react to it then. There are (qualitative) advantages and disadvantages to each strategy, but game theory can be used to quantitatively address the question. Let's look at three examples of countries that find themselves in such situations. Note that although the particular situations are fictionally based, these could easily translate into real-world scenarios.

Mutual Benefit

Suppose the United States and Great Britain share a threat from a terrorist group that is located in Afghanistan. The group opposes both nations because of their stance regarding the war in Iraq and religious views. The two countries are aware that terror cells located in Afghanistan are training suicide bombers to attack both locations. The terror group isn't biased – they hate the United States and Great Britain equally. Therefore, both the British and the United States will benefit equally from eliminating, or attempting to eliminate the terrorist threat. Now let's think about the possible outcomes that can occur from taking pre-emptive action or not taking pre-emptive action.

	Great Britain		
United States		Preempt	Do not preempt
	Preempt	PA Taken	PA Taken
	Do not preempt	PA Taken	No PA

PA Denotes Pre-Emptive Action

In all boxes except the lower right box, pre-emptive action is taken. The only difference in this situation is who takes the action. If we assume that both the United States and Great Britain have the same terror-fighting abilities, then either both can fight the terrorists, or only one can fight the terrorists and the same outcome will occur. However, we need to consider the costs to figure out which outcome will actually occur. We can arbitrarily set the benefit for fighting terrorism at four, and the price for fighting terrorism at six. That would mean that if the United States chooses to fight terrorism itself, it receives negative two in benefit (4-6=-2) and Great Britain receives a benefit of four (4-0=4). We can continue this logic, and create a payoff matrix:

	Great Britain		
United States		Preempt	Do not preempt
	Preempt	2,2	-2,4 *
	Do not preempt	4,-2 *	**0,0 ****

We can then use the game theoretical approach described in the primer to figure out the Nash equilibrium for this payoff matrix. Below is a simplified breakdown of the sub-matrices:

- If Great Britain preempts, then the United States benefits the most from not preempting;
- If Great Britain does no preempt, then the United States benefits the most from not preempting;
- If the United States preempts, then Great Britain benefits the most from not preempting;
- If the United States doesn't preempt, then Great Britain benefits the most from not preempting.

The equilibrium turns out to be the lower right box, that is, neither country ends up taking pre-emptive action. Clearly, this is not the efficient outcome, the efficient outcome (that with the greatest net benefit, distributed the most equally) would have been in the upper left box. Why would the countries not preempt instead of preempting even when the benefit that is gleaned from both preempting is greater? Game theory suggests that rational players tend to consider their own sub-matrices rather than the full matrix. Game theory is useful here because it illustrates that state actors, who might think that they are acting in the interest of their country, might actually be acting in a counter-productive manner. Furthermore, if game theory is used to proactively think about the moves players can make before they make them, then it is possible that they will consider the entire matrix, which would produce the most efficient outcome.

Unequal Benefit

Let's now flip to a situation wherein the United States benefits more from a preemptive action against terrorists than Great Britain. This time, the United States receives a utility of eight from preemption, when Great Britain only receives a utility of four for preemption. The cost of preemption remains static at six. This changes the structure of the payoff matrix:

	Great Britain		
United States		Preempt	Do not preempt
	Preempt	6,2 *	**2,4 ****
	Do not preempt	4,-2	0,0 *

This situation is referred to as one of Asymmetric Dominance (Sandler). This implies that one player desires a particular outcome more than the other player. The situations that we have previously considered have been perfect symmetric, which is why they have tended toward the passive equilibrium. To determine the Nash equilibrium we will again consider all of the sub-matrices.

- If Great Britain preempts, then the United States benefits the most from preempting;
- If Great Britain does not preempt, then the United States benefits the most from preempting;
- If the United States preempts, then Great Britain benefits the most from preempting;
- If the United States does not preempt, then Great Britain benefits the most from not preempting.

In this case the Nash equilibrium can be founding the upper right cell. That is the Nash equilibrium is that the United States will work to preempt the terrorist attack, and Great Britain will not preempt the terrorist attack. In the case of asymmetric dominance we find that if one country benefits more from preempting a terrorist attack, then another country that also benefits from the preemption might receive a free ride. This has important international relations implications because it is often the case that one country will benefit the most from preempting terrorist attacks. In this case the aggressive stance ends up being taken, but once country has to foot the bill. What if two countries choose to cooperate in fighting terrorism instead of maintaining their own separate counter-terrorism strategies?

Cooperation

We can change the game once again to a situation where the United States and Great Britain have to work together to successfully thwart the terrorist attack. In this case, if a single country preempts the attack, the utility is zero (Rapaport). So, if a country tries to preempt the attack without the other country, they stand to receive a utility of negative four (and the other country receives no utility). Let's take a look at the payoff matrix for this situation:

	Great Britain		
United States		Preempt	Do not preempt
	Preempt	**2,2 ****	-4,0
	Do not preempt	0,-4	**0,0 ****

The corresponding sub-matrix analysis follows:

- If Great Britain preempts, then the United States benefits the most from preempting;
- If Great Britain does not preempt, then the United States benefits the most from not preempting;
- If the United States preempts, then Great Britain benefits the most from preempting;

- If the Untied States does not preempt, then Great Britain benefits the most from not preempting.

In this case, often referred to as one of coordination in game theory, we find that there are two separate Nash equilibria. Either the United States and Great Britain work together to prevent the terrorist attack or the attack occurs. Again, this case has significant political implications because it shows that there are times when a nation, on its own, simply can't preempt a terrorist attack.

Conclusion

Whether we seek to quantitatively analyze terrorist situations on the micro level, or on the macro level, game theory can prove an invaluable resource. It is easy to create payoff matrices for a wide variety of events, that can afford policy makers, or even those that are involved in standoffs with terrorists with important theoretical data. To make the game theoretical analysis significantly more robust, it would be useful to collect actual data about utilities and payoffs.

In the case of the micro-situation, researchers, or perhaps more appropriately intelligence professionals could work on finding actual utilities instead of theorized ones. When looking at macro-events, political scientists could analyze history and current events and derive actual utilities. Overall, game theory, as a concept, certainly can enhance our understanding of terrorism and in the future can actually serve a functional purpose.

Works Cited

Dietrich, Franz. "Terrorism prevention: A General Model." Unpublished essay, Feb. 2004. University of Konstanz.

Enders. "The effectiveness of anti-terrorism policies: Vector-autoregression-intervention analysis." American Political Science Review 87 (1993): 829-844.

"Game Theory." New School Economics Department. The New School. 13 May 2008 <http://cepa.newschool.edu/het/schools/game.htm>.

Kuhn, Steve. "Prisoner's Dilemma." Stanford Encyclopedia of Philosophy . 4 Sept. 1997. Stanford University. 4 May 2008 <http://plato.stanford.edu/entries/prisoner-dilemma/>.

Lee, Darin. "A note on the individualistic foundations of the core in economies with asymmetric information." Elsevier 61.2 (1998).

Myerson, Roger B. Game Theory: Analysis of Conflict . N.p.: Harvard University Pres, 1997.

"Nash Equilibrium." ISCID Encyclopedia of Science and Philosophy. The International Society for Complexity, Information, and Design. 7 May 2008 <http://www.iscid.org/encyclopedia/Nash_Equilibrium>.

NATO & OTAN. Game Theory as a Tool for Analyzing Terrorism. By Vito Fragnelli. Allesandria: n.p., 2006.

Rapaport. "The game of Chicken." Pittsburg University Press (1969).

Sandler, Todd, and Daniel M Arcre. "Terrorism and Game Theory." Simulation & Gaming 34 (Sept. 2003): 1-30.

Schelling, Thomas C. The Strategy of Conflict. N.p.: Harvard University Press, 2007.

Watson. Strategy: An introduction to game theory. New York: W. W. Norton, 2002.

Weaver, Ransom, and Barry Silverman. "Modeling and Simlulating Terrorist Decision-making: A 'Performance Moderator Function' Approach to Generating Virtual Opponents." Unpublished essay. University of Pennsylvania.