

Outbidding and the Overproduction of Terrorism

Andrew Kydd* Barbara Walter†

October 14, 2009

Abstract

We model a conflict in which the population on one side is uncertain about the preferences of the organizations which strive to represent them. In such a situation, the organizations have an incentive to compete with each other in attacks upon the enemy which may have no prospect of success and may even be counterproductive, but which serve to demonstrate dedication to the cause. This process is sometimes called outbidding and it may serve as a mechanism which explains the overproduction of terrorism in certain long running conflicts, in comparison to the level that would be observed if the side were represented by only one organization, or if the organizations' preferences were known.

*Harvard University, Department of Government, akydd@wcfia.harvard.edu.

†University of California, San Diego, bfwalter@ucsd.edu.

Some terrorist attacks seem to achieve their objectives in relatively short order, such as the bombing of the Marine barracks in Lebanon in 1983 which led to the withdrawal of the U.S. presence there. Other terrorist attacks form parts of long term campaigns that seem to have little prospect of success, and even seem to worsen the situation of the side perpetrating them, such as the attacks by Hamas and Islamic Jihad during the second intifada. Such campaigns could be simply long term wars of attrition in which the parties accept temporary costs in order to inflict costs on the enemy in an effort to get the enemy to make concessions. However, it seems possible that there are dynamics at work which prolong these campaigns longer than would be optimal from the attacking side's perspective, considered as a unitary actor. For instance, the second intifada has resulted in an Israeli withdrawal from Gaza, but this goal had largely been achieved by the Oslo process, and the building of the wall around the West Bank will possibly result in a territorial arrangement worse than what could have been obtained in negotiations.

This suggests that terrorism may be overproduced in comparison to what would happen if the side were represented by a single organization with the preferences of the median member of society at heart. We argue that two conditions are important for such a result to occur. First, there must be multiple groups competing for leadership over the side in question, so that each group feels competitive pressure from the other. If one group fails to demonstrate its commitment to the cause adequately, another group will and as a result will gain support from the society. Second, the society must be uncertain about the preferences of the organizations that claim to represent it. That is, the organizations' preferences must be the subject of uncertainty, such that the society believes that there is some chance the organization is willing to sell out the interests of the group. This

combination, competition plus uncertainty about preferences, generates the incentive for competitive attacks on the enemy, even if these attacks have little chance of achieving any goal. This syndrome is sometimes called *outbidding*, each party attempts to garner support by exceeding the other's professions of loyalty and commitment to the cause. A third condition can make outbidding counterproductive even from the perspective of the attacking side. If there is uncertainty about whether terrorism is productive of concessions or not, outbidding may frustrate the transmission of information to the society, in the sense that any group which refuses to conduct an attack on the grounds that they are counterproductive will be accused of betrayal by groups willing to claim that they will be productive and should be carried out. Thus society will fail to learn or will learn more slowly than it could about the counterproductive nature of terrorist attacks.

In what follows we briefly review the literature on terrorism and outbidding, and then develop a model of this process and analyze its implications. Then we discuss the implications with reference to the historical record of long term terrorist campaigns.

Terrorism and Outbidding

For obvious reasons there has recently been a surge of interest in terrorism and this is reflected in a growing literature on the subject. Many analysts, influenced by the horror of the acts, the suicidal aspect and the influence of religion, focus on psychological motives and try to understand the motivations of the individual bomber (Juergensmeyer 2000, Stern 2003). The political science literature has tended to see a strategic logic at work (Walter & Kydd 2005).¹ Robert Pape exemplifies this trend,

arguing that suicide terrorism is primarily a tool by which weak groups with nationalist goals attempt to pursue a war of attrition designed to coerce withdrawal by foreign occupying forces (Pape 2005).² Game theoretic studies have examined other strategic logics of terrorism; examples include Kydd and Walter and Ethan Bueno de Mesquita (Kydd & Walter 2002, Bueno de Mesquita 2005*a*, Bueno de Mesquita 2005*b*). De Figueiredo and Weingast offer a related model of outbidding in ethnic conflict (de Figueiredo Jr. & Weingast 1999). Siqueira analyzes a public goods model of terrorist activity in which militant and political wings attempt to mobilize support (Siqueira 2005). However he finds ambiguous results when analyzing whether violence is more likely when the groups cooperate or act non-cooperatively.

The most extensive discussion of terrorism and outbidding is by Bloom (2005). She compares the Palestinian case and the Sri Lankan case and notes that suicide bombing flourished in the Palestinian case because of the competition between groups, whereas once the Liberation Tigers of Tamil Elam (LTTE) consolidated power they reduced their use of suicide bombing. However, as she points out, other factors help account for the LTTE's decreased use of violence, including increased war weariness of the Tamil population and decreased funding from abroad. Bloom's analysis while very suggestive, however, leaves several questions unanswered. For instance, why would the public support groups that engage in suicide bombing when by doing so they signal that they may have more extreme preferences than the majority of the public? Placing extremists in charge of the negotiations could lead to prolonging the conflict beyond what the public would prefer. More generally, the literature on outbidding ignores the fact that there could be a diametrically opposed process of "underbidding," in which competition between groups

forces more extremist groups to moderate their stances because the population is believed to favor peace. The median voter theorem, after all, predicts that party platforms will converge on the preference of the median voter, so if the median voter prefers peace then political parties should as well (Black 1948, Downs 1957). The question then becomes, when do central tendencies prevail and when are parties, candidates or other groups pushed towards the extremes.

One context that may favor extreme candidates is bargaining. It is sometimes argued that employing agents with different preferences from oneself can be advantageous in bargaining. Such agents could enable one to commit to rejecting deals that one would otherwise accept, forcing the other side to offer a better deal (Muthoo 1999, 230-32). Similarly, the literature on voting and coalition formation suggests that voters may vote for parties that are more extreme than themselves, if they know that in the coalition formation process they will have to make compromises with other parties (Kedar 2005, Austen-Smith & Banks 1988, ?).³ Similarly, Alesina and Rosenthal find that moderate voters have an incentive to split their ticket or vote for candidates of the opposite party from the President in off year elections to make moderate outcomes more likely (Alesina & Rosenthal 1996).

Other scholars have argued that uncertainty over what the best policy choice is can also frustrate incentives for moderation. For instance, Canes-Wrone and Shotts analyze a model in which the policy maker has some information about what the best policy is and the voters are uncertain about the policy maker's preferences (Canes-Wrone & Shotts 2005).⁴ They show that if the public cannot verify whether the policy chosen was correct or not, there is no electoral incentive for moderation at all, incumbents will simply choose whichever policy the voters ex-ante think is best, regardless of the policy maker's

information. Even if the voters are able to verify whether the policy chosen was correct, there are only rarely incentives for moderation, usually, incumbents pick one policy or the other and ignore their information. Thus information that would be valuable to the public is lost because they do not trust the incumbent. The incumbent might choose a policy out of bias, rather than because of its information about the policy.

The literature therefore suggests the following points for investigation. When representatives are elected to bargain with opponents, then extremist candidates are favored, who will be firm in the bargaining game. If there is uncertainty about the candidates' preferences, then terrorist attacks can signal low costs for fighting, or greater resolution, which in turn can lead to support by the electorate. If there is uncertainty about the wisdom of terrorist attacks in gaining concessions, then the previous consideration may frustrate information transmission about this question, leading to terrorism that is counterproductive even from the attacking side's perspective. In the model we develop below we will investigate these points.

The Model

Consider a model in which there are two sides, A and B , engaged in a dispute about something of value to both. On side A there are three actors, two organizations, labeled $A1$ and $A2$, and "society" AS or the median voter on that side. The B side is unitary.

Notation is summarized in Table 1.

The issue in dispute is represented by the unit interval and a possible issue resolution by $x \in [0, 1]$. Each player has single peaked preferences over the possible issue

resolutions, $u_{Ai}(x)$ and $u_B(x)$ and ideal points I_{Ai} and I_B . Given that the B side is unitary, we normalize $I_B = 0$ and assume that $\min\{I_{AS}, I_{A1}, I_{A2}\} > 0$, so that in the conflict of interest between the two sides, the A side prefers higher outcomes and B prefers lower ones. However, the A side may not be united over how high an issue resolution it prefers. We normalize the maximum value for all utility functions to 1 and the minimum to zero, so any payoff in the game is bounded above by 1. The issue space and utility functions for AS and B are illustrated in Figure 1.

The game proceeds as follows. First, $A1$ and $A2$ simultaneously choose whether to attack the enemy or not, where attacking imposes costs on everyone, c_{AS} , c_{A1} , c_{A2} and c_B . Denote Ai 's strategy $s_{Ai} \in \{a, w\}$, where a signifies attack and w signifies wait, or not attack and $a \prec w$. Next, AS decides which organization to support, $A1$ or $A2$. Whichever group is supported is said to be in office, and has the power to negotiate with B . Third, B makes a concession x^0 , which establishes a new status quo. An example would be Israel turning over the Gaza strip to the Palestinians. Finally, the group in power either accepts the concession, leading to peace with payoffs $u_{AS}(x^0)$, $u_{A1}(x^0)$, $u_{A2}(x^0)$ and $u_B(x^0)$, or rejects it, leading to renewed conflict. Let player Ai 's response to the offer be denoted $r_{Ai} \in \{n, y\}$, where n means no and y means yes and $n \prec y$. Each player suffers a cost of war, k_{AS} , k_{A1} , k_{A2} , and k_B .

We assume that if the concession is rejected and war occurs the end result will be some outcome, x , and the likelihood of each outcome is given by a distribution $p(x|x^0)$. The likelihood of the various outcomes depends on the concession made. The issue in dispute is assumed to be correlated with the relative power between the two sides, so that the greater a concession B grants to A , the greater A 's chances in the subsequent war if A

decides to not accept the offer but continue the war. In the Israeli-Palestinian situation this reflects the idea that having control over Gaza increases the power of the Palestinians in any subsequent struggle over the West Bank. Formally we assume that the cumulative distribution over outcomes, $P(x|x^0)$, is decreasing in x^0 , which makes higher outcomes more likely, which is less favorable to B and more favorable to the A side. This implies, by stochastic dominance, that B 's war payoff is declining in the concession made to A (Levy 1992). If an actor on the A side has an ideal point equal to 1 (as AS does in Figure 1), then its war payoff is increasing in the concession. Even if the ideal point is not quite equal to 1, the war payoff is likely to be increasing in the concession, because the utility function will be increasing over most of its range. We can therefore write player i 's war payoff as $w_i(x^0, k_i) = \int_0^1 u_i(x)p(x|x^0)dx - k_i$, where the payoff is a function of the concession made and the player's costs of war.

Uncertainty

We assume that AS is uncertain about two parameters affecting the decisionmaking of the organizations. First is the cost of fighting. AS would prefer an organization with low costs of fighting, because that leads to higher resolve and a greater concession from B . If an organization is afraid to fight, then B can afford to make a small concession to it. In order to simplify the analysis, consider two ideal types, a weak type with $c = k = 1$ and a strong type with $c = k = 0$. Types with costs equal to 1 cannot possibly attack in equilibrium to begin with and will never reject any offer from B , no matter how small. These types are sellouts. Types with costs equal to zero can easily attack in the first round and will reject

any offer that falls below their expected value for war, the maximum offer the A side can hope for. These are the tough types. Let the likelihood that A_i is strong be σ_i .

Second, we assume that AS is uncertain about the utility function of the organizations. Utility functions differ in how highly they value moderate outcomes in comparison with the more extreme outcomes. More moderate utility functions have higher valuations for moderate outcomes, and increase only gradually if at all towards the end of the interval. More extreme utility functions place a low value on intermediate outcomes, and only place a high value on extreme outcomes representing total victory. Players with such utility functions will be predisposed to go for broke, reject any compromise solution and go to war. Once again considering ideal types, let there be two types of utility function. Moderate types have a utility function identical to the median voter in the society, $u_{A_i}(x) = u_{AS}(x)$. Extremist types have a utility function that is equal to zero for all $x \in [0, 1)$ and equal to 1 for $x = 1$. Thus strong extremists have no interest in accepting compromise solutions, but will fight at the end of the day. The likelihood that A_i is extremist is ϵ_i .

Combining these two axes of uncertainty, each organization can be one of four possible types, represented in Table 2. The probability of each type is given in the cell, and the types are assumed to be uncorrelated.

We assume that AS is the only actor who is uncertain about the types of the organizations striving to represent it. That is, the organizations know each other's types, and so does the opposition, B . It might seem strange that the target side should know the organization's preferences, and indeed in previous work we have assumed it does not (Kydd & Walter 2002). In this context, however, it provides a simple and interesting illustration of

the perverse incentives generated by this strategic situation. Given that the organizations know each other's types and they know that the opposition does as well, strong moderate types have every incentive to self-identify as such and identify any competing group that is extremist as such. Strong moderates know that the opposition's reaction to an extremist in power will be to not make any concessions but rather to fight. However, the strong moderate may be unable to credibly communicate this to AS , even though AS would benefit greatly from the knowledge. One can think of this as a simplification of a more complex game in which B is not perfectly informed, but the organizations are better informed about B 's information than AS is, which is entirely plausible.

Equilibria

We solve for the unique perfect Bayesian equilibrium.⁵

Starting from the end of the game, we know that weak types, both moderate and extreme, will accept any offer from B and strong moderates will accept any offer such that $u_{AS}(x^0) > w_{AS}(x^0, 0)$. A strong extremist organization will not accept any offer short of $x^0 = 1$.

B , therefore, will offer $x^0 = 0$ to weak types and strong extremists. The weak types will be content with nothing and the strong extremists will simply use the increase in strength from any concession to renew the attack. If the organization is a strong moderate, B will offer $x^0 = x^*$ such that $u_{AS}(x^*) = w_{AS}(x^*, 0)$ if B in turn prefers that to war, $u_B(x^*) > w_B(x^*, k_B)$. This will generally be the case if both B and AS have risk neutral or risk averse preferences over the issue, and if the concession does not have too great an

impact on relative power, and we assume these conditions are met.⁶ Both B and AS get a surplus over their war payoffs in this case, B because it gets to pick the offer and AS because it is using an agent with the same utility over deals but lower war costs, which compels B to make a better offer than it would to AS directly.

Combining these considerations, we can derive a preference ordering for AS over the different types of organization: strong moderates \succ strong extremists \succ weak types. Strong moderates will negotiate a good deal with B , strong extremists will continue the war, while weak types will capitulate to the enemy. This ordering will influence AS 's choice of which organization to support based on their behavior.

The decisions about which organization AS will support and whether $A1$ and $A2$ will attack or not are linked. Weak organizations will not attack. Strong organizations therefore have an incentive to attack to distinguish themselves from the weak. In the unique equilibrium in the game, strong types, both moderate and extremist, attack. Given this fact, AS knows that if an organization attacks, it is strong, and if it does not, it is weak. With two organizations, there are four possible states of the world. If both organizations wait, both are identified as weak and AS will be indifferent between the two, and can randomize. If one waits and the other attacks, then AS will support the attacking organization. If player i waits and j attacks, the payoff for supporting i will be $u_{AS}(0)$ and the payoff for supporting j will be $\epsilon_j w_{AS}(0, k_{AS}) + (1 - \epsilon_j) u_{AS}(x^*)$. Since the offer to a strong moderate is preferred to war with no concession, which is in turn preferred to capitulation, the second always beats the first, regardless of the probability that player j is extremist. Therefore if one organization attacks and one does not, AS supports the attacking organization. Finally, if both organizations attack, AS will support whichever

organization is believed to be less likely to be an extremist, since both are identified as strong and AS would prefer to get a strong moderate.

AS 's equilibrium strategy in response to the organization's behavior is summarized in Table 3. We assume, without loss of generality, that organization $A1$ is the one most likely to be moderate, that is $\epsilon_1 < \epsilon_2$.⁷ Therefore, $A1$ gets the nod if both organizations attack. If both players wait, AS can flip a coin over whom to support, which will provide the proper incentives to the organizations.

Results

The results of the model are fairly stark. The median voter on the A side would prefer to be represented by a strong moderate. If there were complete information and a strong moderate organization were available, AS would support it and the result would be a successful negotiation with B , leading to a peace deal that both sides prefer to war.

Terrorism is unproductive in terms of the relations with B , because since B knows the preferences of the organizations, it will not alter its offer in response to a terrorist attack.

Because of competition between groups and uncertainty about the organization's preferences, however, strong moderate groups engage in terrorist activity along with strong extremist groups. This imposes a cost on society, c_{AS} , and on the other side, c_B . Despite this cost, information revelation is incomplete and AS may still end up supporting an extremist organization that will catapult it back to war, despite AS 's desire for a negotiated settlement. The desire to avoid supporting sellouts leads to a preference for strong organizations, even if they may be extremists.

Discussion

Empirical discussion of long running campaigns.

Conclusion

Outbidding among terrorist groups is driven by the fact that they are at least tacitly bargaining with their target for greater concessions, their supporters are uncertain about their preferences, and they face competition from other groups that strive to demonstrate their fidelity to the cause. This can generate incentives for terrorist attacks even on the part of organizations that know that the attack will simply impose costs and have no strategic impact in terms of getting the target to grant greater concessions. This logic suggests that consolidation of terrorist groups is to be encouraged. The downside of consolidation, of course, is that united groups may be stronger than the sum of their parts. Which of these considerations dominates depends on the relative strength of the two effects.

Appendix

We know weak types wait, so with two players, two strong types, and two choices, there are 16 possible combinations of behavior. All but one of these can be ruled out as equilibrium candidates.

If only strong moderates attacked, then *AS* would support an organization on seeing an attack, since strong moderates are preferred. This would give an incentive to strong extremists to attack, negating the equilibrium. (Note, the strong types at least weakly prefer to be in office rather than have a different type in office, even without assuming a direct benefit to officeholding. Weak moderates prefer that a strong moderate be in office, weak extremists prefer either themselves or a strong extremist.) This eliminates six possibilities.

Next, if the strong types of one organization pool on waiting, then an attack by that organization takes the game off the equilibrium path. The equilibrium refinement of divinity argues that in such cases, beliefs should converge on the type with the least cost of making such an error, in this case the strong types (Banks & Sobel 1987). This means that in the pooling equilibrium attacking would signal strength, and so the strong types would deviate to attacking, eliminating the equilibrium. This eliminates another six possibilities.

Finally, if a strong extremist attacks and a strong moderate does not, then if *AS* would support an attacker, the strong moderates would have an incentive to switch to attacking, while if *AS* supported waiting, then the strong extremist would have an incentive to switch to waiting, either one of which would violate the equilibrium. This eliminates three possibilities. The one remaining candidate is where strong types of both

players pool on attacking.

Notes

¹See *The Political Economy of Transnational Terrorism*, a special issue of *The Journal of Conflict Resolution*, Volume 49, Number 2, and Sandler & Arce M. (2003) for a review of game theoretic applications to terrorism. Early analyses of terrorism from a rationalist perspective include Crenshaw (1981) and Thornton (1964).

²For pioneering signaling models along these lines see Lapan & Sandler (1993) and Overgaard (1994).

³There is a large related literature on whether voters vote for parties that are near their preferred policy (called “proximity” voting) or parties that are on their “side,” and the more extreme the better (called “directional” voting) e.g. Lewis & King (1999). Directional voting simply assumes a preference for extreme candidates; far more interesting is the problem of why voters who prefer moderate outcomes would nonetheless vote for extreme candidates.

⁴See Morris (2001) for an interesting related model.

⁵Subject to the equilibrium refinement of divinity, see Appendix.

⁶If they are not then there is no type in the model to which B is willing to make a concession. Interestingly, strong types would still attack at the outset, because AS would prefer to continue the war rather than capitulate and so still wants to support a strong organization.

⁷For the case where $\epsilon_1 = \epsilon_2$, AS can flip a coin when deciding which to support if both

attack or both wait.

References

- Alesina, Alberto & Howard Rosenthal. 1996. "A Theory of Divided Government." *Econometrica* 64(6):1311–1341.
- Austen-Smith, David & Jeffrey Banks. 1988. "Elections, Coalitions and Legislative Outcomes." *American Political Science Review* 82(2):405–422.
- Banks, Jeffrey S. & Joel Sobel. 1987. "Equilibrium Selection in Signaling Games." *Econometrica* 55(3):647–661.
- Black, Duncan. 1948. "On the Rationale of Group Decisionmaking." *Journal of Political Economy* 56(1):23–34.
- Bloom, Mia. 2005. *Dying to Kill: The Allure of Suicide Terrorism*. New York: Columbia University Press.
- Bueno de Mesquita, Ethan. 2005a. "The Quality of Terror." *American Journal of Political Science* 49(3):515–530.
- Bueno de Mesquita, Ethan. 2005b. "The Terrorist Endgame: A Model with Moral Hazard and Learning." *Journal of Conflict Resolution* 49(2):237–258.
- Canes-Wrone, Brandice & Kenneth W. Shotts. 2005. "When Do Elections Promote Policy Moderation by Privately-Informed Officials?".
- Crenshaw, Martha. 1981. "The Causes of Terrorism." *Comparative Politics* 13(4):379–399.

- de Figueiredo Jr., Rui J. P. & Barry R. Weingast. 1999. The Rationality of Fear: Political Opportunism and Ethnic Conflict. In *Civil Wars, Insecurity and Intervention*, ed. Barbara Walter & Jack Snyder. New York: Columbia University Press pp. 261–302.
- Downs, Anthony. 1957. *An Economic Theory of Democracy*. New York: Harper and Row.
- Juergensmeyer, Mark. 2000. *Terror in the Mind of God: The Global Rise of Religious Violence*. Berkeley: University of California Press.
- Kedar, Orit. 2005. “When Moderate Voters Prefer Extreme Parties: Policy Balancing in Parliamentary Elections.” *American Political Science Review* 99(2):185–199.
- Kydd, Andrew & Barbara F. Walter. 2002. “Sabotaging the Peace: The Politics of Extremist Violence.” *International Organization* 56(2):263–296.
- Lapan, Harvey E. & Todd Sandler. 1993. “Terrorism and Signaling.” *European Journal of Political Economy* 9(3):383–398.
- Levy, Haim. 1992. “Stochastic Dominance and Expected Utility: Survey and Analysis.” *Management Science* 38(4):555–593.
- Lewis, Jeffrey B. & Gary King. 1999. “No Evidence on Directional vs. Proximity Voting.” *Political Analysis* 8(1):21–33.
- Morris, Stephen. 2001. “Political Correctness.” *Journal of Political Economy* 109(2):231–265.
- Muthoo, Abhinay. 1999. *Bargaining Theory with Applications*. Cambridge: Cambridge University Press.

Overgaard, Per Baltzer. 1994. "The Scale of Terrorist Attacks as a Signal of Resources."

Journal of Conflict Resolution 38(3):452–478.

Pape, Robert A. 2005. *Dying to Win: The Strategic Logic of Suicide Terrorism*. New York:

Random House.

Sandler, Todd & Daniel G. Arce M. 2003. "Terrorism and Game Theory." *Simulation and*

Gaming 34(3):319–37.

Siqueira, Kevin. 2005. "Political and Militant Wings within Dissident Movements and

Organizations." *Journal of Conflict Resolution* 49(2):218–236.

Stern, Jessica. 2003. *Terror in the Name of God: Why Religious Militants Kill*. New York:

Ecco-Harper Collins.

Thornton, Thomas Perry. 1964. Terror as a Weapon of Political Agitation. In *Internal*

War: Problems and Approaches, ed. Harry Eckstein. London: Free Press of Glencoe

pp. 71–99.

Walter, Barbara F. & Andrew H. Kydd. 2005. "Strategies of Terrorist Violence."

Table 1: Notation in the Game

x	A possible issue resolution
$u(x)$	Utility functions over x
I	Ideal points
c	Cost of terrorist attack
$s \in \{a, w\}$	Organization's terrorism strategy, attack or wait
$r \in \{y, n\}$	Organization's decision on the offer, yes or no
p	Distribution over outcomes if there is a war
k	Cost of war
w	War payoff
σ	Likelihood that organization is strong
ϵ	Likelihood that organization is extremist

Table 2: The Four Types of Organization

	Moderate	Extremist
Strong	$\sigma_i(1 - \epsilon_i)$	$\sigma_i\epsilon_i$
Weak	$(1 - \sigma_i)(1 - \epsilon_i)$	$(1 - \sigma_i)\epsilon_i$

Table 3: *AS*'s Equilibrium Strategy

	<i>A2</i> Attacks	<i>A2</i> Waits
<i>A1</i> Attacks	Support <i>A1</i>	Support <i>A1</i>
<i>A1</i> Waits	Support <i>A2</i>	Randomize

Figure 1: The Bargaining Space

