Terrorism and Profiling

Andrew Kydd*

March 6, 2008

Abstract

Terrorist attacks on mass transportation systems pose the problem of how large numbers of passengers can be screened most efficiently to discover terrorists and prevent attacks. Some argue that certain social groups from which terrorists are more likely to come, such as young Muslim men from the Middle East, should be screened more heavily than people not belonging to such groups. Critics deride this as racial profiling, and point to failures in police use of racial profiling to stop crime, as well as the lack of fairness. I examine a model of searching for terrorists among a population divided up into groups that vary in their potential reliability as agents of terrorist attacks. I show that *seemingly unfair profiling*, in which individuals from groups that are *less* likely to be chosen to execute a mission are actually *more* likely to be searched, is a feature of many equilibria of the model.

*University of Wisconsin, kydd@wisc.edu. I would like to thank the participants at the International Relations Colloquium at the University of Wisconsin for helpful comments.

The terrorist attacks of September 11, 2001 were committed by 19 young Muslim men from the Middle East, 15 of whom were from Saudi Arabia. The train bombings in Madrid in 2004 and the subway bombings in London in 2005 were also committed by young Muslim men. In response, transportation security services have increased their screening of passengers dramatically. The plot uncovered in Britain in the summer 2006 to use liquids to blow up airplanes led to worldwide bans on carrying liquids aboard, despite the obvious fact that the likelihood that a grandmother on her way home from Mexico City to Chicago would destroy an aircraft with a bottle of water was effectively zero. This has lead several observers to suggest that transportation security, and anti-terrorism policy more generally, should overcome politically correct scruples about profiling and openly engage in discriminatory practices when it comes to searching for terrorists among the general public. Conservative pundits argued that police should explicitly focus on young Muslim men, particularly from the Middle East or North Africa, and subject people who do not fit into those categories to much less scrutiny.[1] Critics argue that this would be a betrayal of American principles, would increase hostility in Muslim communities in the United States and abroad, and would be easily circumvented by terrorist groups switching to people not likely to get searched under the profiling scheme.

However, little real strategic analysis has been done to assess under what conditions profiling makes sense and what the response would be to it. Philip Carpenter (2006) raises the issue but does not explore it fully. A pathbreaking study of racial profiling and crime by Nicola Persico (2002) is suggestive, but the context of terrorism differs from that of ordinary street crime because it is orchestrated by a strategic actor, rather than a spontaneous result

[1] Charles Krauthammer, "Give Grandma a Pass," *Washington Post* 29 July 2005: A23.

of individual action by self interested criminals. Robert Powell's model of hardening targets against terrorist attacks is also extremely helpful, but the profiling context differs slightly from the target hardening context in ways that make it worth considering the problem carefully in its own right (Powell 2005).

1 The Problem of Profiling with Strategic Actors

The debate between supporters and opponents of profiling may appear unreconcilable. Supporters of profiling point to the fact that most terrorists, in certain contexts at least, are male and Islamic and have some connection with the middle east. Why should people who do not fit these categories be subject to the cost of being searched, delayed, etc, when they are clearly not being recruited by Al Qaida or similar organizations to launch attacks upon the United States, or other western countries? The counter argument also seems compelling, if young muslim men were searched and others let go, Al Qaida would have a strong incentive to recruit someone not fitting this description to carry out attacks for them, even if they were somewhat harder to recruit or less reliable, because they could slip through without detection. Hence we might as well screen everyone, since this also seems more fair.

While the two strategic logics seem contradictory, they can be dealt with together in very simple game theoretic frameworks. Consider the simplest possible profiling problem in which there are two groups, 1 and 2. The terrorist organization is composed of group 1 members, and fights on their behalf. Group 2 represents the rest of society. The terrorist group must select an operative who either belongs to group 1 or to group 2. The government must search either group 1 or group 2. If the government searches the group that the terrorists send, I

Table 1: Profiling as Matching Pennies

		The Terrorist Group	
		Send Group 1	Send Group 2
The Government	Search Group 1	1, 0	1-r, r
	Search Group 2	0,1	1, 0

assume that the government blocks the attack achieving a payoff of 1 while the terrorists get zero. If the government searches the wrong group, the terrorists operative eludes detection. If it is a group 1 member, we assume that it succeeds in carrying out its attack, earning a payoff of one for the terrorist group and zero for the government. If it is a group 2 member, I assume it carries out its attack with probability $r < 1$, for reliability, so the terrorist's payoff is r and the government's payoff is $1 - r$. This captures the idea that people in group 2 are likely to be less reliable agents of group 1 based terrorists attacks. Group 2 members may be more likely to get cold feet, talk to non group 1 members about the plan and get persuaded out of it or exposed to the authorities, etc. The resulting game is illustrated in in Table 1.

The game is a version of the canonical "matching pennies" game, and it is quickly apparent that there is no pure strategy equilibrium. The government wants to match the terrorists choice by searching the group the terrorist sends, and the terrorists want to have a mismatch, with the government searching the wrong group. The equilibrium in the game is therefore in mixed strategies. The government will search each group with a certain probability, say p_1 for group 1 and p_2 for group 2. The terrorists will send someone from group 1 with probability q_1 and from group 2 with probability q_2. These probabilities are calculated to make the other side indifferent between its two options, thereby making them willing to mix. The likelihood that the terrorist group sends a group 1 member is obtained by equating the government's payoff for searching group 1, $q_1 \times 1 + (1 - q_1) \times (1 - r)$ with the payoff for searching group 2, $q_1 \times 0 + (1 - q_1) \times 1$, and solving for q_1 which gives

$$q_1 = \frac{r}{1 + r}. \tag{1}$$

A similar calculation on the terrorist's payoffs shows that the likelihood that the government searches group 1 must be

$$p_1 = \frac{1}{1 + r}. \tag{2}$$

These searching and sending probabilities have several implications for the debate between advocates and critics of profiling outlined above. First, the government searches both groups in equilibrium, but it searches group 1 members more frequently, that is, $p_1 > p_2$. If group 2 were perfectly reliable, $r = 1$, then $p_1 = \frac{1}{2}$ and both groups would be searched equally. But as long as group 2 is less reliable from the terrorist's perspective, group 1 will be searched more intensively. If this were not the case, then the terrorists would have an incentive to send group 1 members exclusively, since they would be equally likely to be searched and would be more reliable if undetected. More counterintuitively, the terrorists are more likely to send group 2 members, $q_2 > q_1$. If the groups were equally reliable, q_1 would also equal one half, but as r shrinks, q_1 declines, indicating that the terrorists are more likely to send group 2 members.

Thus even in this very simple game we can capture the logic of the profiling debate and show why both sides are incomplete. It is true that the government cannot search only some fraction of society, because the terrorists could easily evade this by sending someone who would elude detection. However, it is also true the the government must search more reliable groups, from the terrorist's perspective, more intensely in order to deter them from sending only their own group members. The terrorist group's strategic response is to send less reliable groups more frequently, in order to keep the government willing to search both groups with some level of effort. This pattern in which the government searches more reliable groups more intensely but the terrorists send less reliable groups more frequently I call *seemingly unfair profiling* because of the glaring contradiction between the fact that the government searches the terrorists core group more intensely despite the fact that the terrorists send this group less frequently. While this sort of profiling appears normatively objectionable, it can be a product of rational strategic interaction.

Definition 1 Seemingly unfair profiling *is a pattern in which a group more closely associated with a terrorist organization is searched more intensely than some reference group, even though the terrorist group is less likely to send a member from the associated group than from the reference group.*

While the matching pennies game captures some essential elements of the profiling dilemma, it fails to address many other issues. Most crucially, the government's search intensities across the two groups are not necessarily exclusive, that is, the government could choose to search both groups more intensely or both less intensely, and the degree of intensity would be reflected in increased costs. Furthermore, there are many possible groups, not just two. A more complicate game drawing from the literature on "Colonel Blotto" games and Robert Powell's analysis of target hardening enables us to address these issues.

2 Profiling with N groups

Assume that there are $n \geq 2$ social groups, which are determined by how people score on certain characteristics, such as religion, skin color, gender, etc. Assume the terrorist group is composed of members of group 1, and is fighting for their cause. Group 1 members are therefore the most likely to support the terrorist group. Assume that the other groups are ordered in decreasing likelihood to support the terrorists, so that group 5 is less likely than group 4 to support the group, etc. Formally, each group has a reliability score, r_i, which indicates how likely a member of this group would be, if it were recruited and trained by the terrorist group, to reliably carry out a terrorist attack if not stopped by the government. I assume that $r_i > r_j$, for all $i < j$.[2] Furthermore, I assume that the terrorist group finds the less reliable groups more costly to recruit from, either in additional training and vetting or simply search costs to find a suitable candidate from a large unsympathetic group. Thus, each group has an associated cost of recruiting, k_i, where $k_i < k_j$, $\forall i < j$. This reflects that it is harder for the terrorist group to recruit a suitable operative from the less reliable groups, regardless of whether the attacker succeeds or not. Finally, each group also has a size, s_i, which indicates what fraction of the relevant population it represents, for instance, what proportion of airline travelers in the airline screening case. The sizes sum to 1, $\sum_i s_i = 1$. The likelihood that the terrorist group sends a member of group i is q_i. I assume that the terrorist group sends someone, so the q_i also sum to 1.

The government side has to decide upon a screening policy. For each social group, the

[2] I assume strict inequality, if two groups have the same reliability score they can be amalgamated for the purpose of the analysis.

government selects an intensity of screening which yields a probability, p_i, that a terrorist from this group would be discovered and thereby thwarted. The likelihood that a terrorist in group i goes undetected is therefore $1 - p_i$. An undetected terrorist attacks with probability r_i, so the likelihood that a terrorist from group i attacks is $(1 - p_i)r_i$.

I assume that the terrorist's payoff is an increasing function of the likelihood of a successful attack, less the costs of recruitment,

$$\sum_{i=1}^{n} q_i\{(1 - p_i)r_i - k_i\}. \tag{3}$$

The terrorists choose a vector of likelihoods of sending an agent from each group, $\mathbf{q}$.

Searches are costly and the more intensely and widely the government searches, the more expensive it is. This is reflected in a cost for searching, $c(p_i, s_i)$, which is an increasing function of the likelihood of successful detection, p_i, and the size of the group subject to that level of search, s_i. I assume that costs increase more steeply the larger a group is, $\frac{\partial^2 c(p_i,s_i)}{\partial p_i \partial s_i} > 0$, and that setting $p_i = 1$ is prohibitively costly.

I assume the government's payoff is equal to the likelihood that there is not a successful attack, minus the search costs, or

$$1 - \sum_{i=1}^{n} q_i(1 - p_i)r_i - \sum_{i=1}^{n} c(p_i s_i) \tag{4}$$

The government selects a vector of search probabilities, $\mathbf{p}$, to maximize this function.

2.1 Equilibria in the Model

The game can be solved for a Nash equilibrium in the searching and sending probabilities. The equilibrium will have some of the characteristics of mixed strategy equilibria. In particular, the government search levels will have to make the terrorists indifferent among the groups which it sends with positive probability. If the payoff to sending one of these groups were higher than another, then the terrorist group would send only the higher payoff group and not the lower.

In any equilibrium, the terrorists will place positive probability in sending agents from the more reliable groups and there may exist some less reliable groups that are not sent or searched in equilibrium.

Theorem 1 *In any equilibrium, terrorists will place positive probability on sending groups* $1 \ldots t$ *where* $t \leq n$*. All groups less reliable will not be sent or screened,* $p_i = q_i = 0$*,* $\forall i > t$*.*

Proof: Assume $q_{t+1} = 0$. Then $p_{t+1} = 0$ as well because increasing p_{t+1} above zero would produce costs for the government with no reduction in the likelihood of an attack. This implies that for all $j > t+1$, $q_j = 0$. Otherwise, if $q_j = x > 0$, the terrorists could reduce q_j to zero and increase q_{t+1} to x, which would increase the terrorists' payoff since group $t+1$ is not being searched at all and is more reliable than group j. The government will therefore also set $p_j = 0$ $\forall j > t+1$. ■

Next, the terrorist group must be indifferent among the groups from which it selects in choosing whom to send. This implies that the government must search more reliable groups more intrusively, in order to make them less attractive to send.

Theorem 2 *The intensity of being searched and likelihood of being detected,* p_i*, is an increasing function of* r_i*, that is,* $p_i > p_j$*,* $\forall i < j \leq t$*.*

Proof: The terrorist group must be indifferent among the groups it sends with positive probability, so from Equation 3 and $\forall i < j \leq t$,

$$(1 - p_i)r_i - k_i = (1 - p_j)r_j - k_j.$$

which implies that

$$(1-p_j)r_j - (1-p_i)r_i = k_j - k_i$$

Since $k_j > k_i$, the right hand side is positive, so the left hand side must be as well. Since $r_i > r_j$, it must be that $1-p_j > 1-p_i$ which implies that $p_i > p_j$. ■

Thus the government will more heavily search groups that are known to be more reliable for the terrorists, in order to discourage them from sending them. This departure from a common search probability for all groups is a feature of any equilibrium of the game given that setting $p_i = 1$ is prohibitively costly. Note this result holds regardless of the relative size of the groups. Even if a more reliable group is larger, and hence more costly to search, it will be searched more intensely than a less reliable group in equilibrium.

Next, the terrorists are more likely to send a member of a group the *less* reliable it is, and the larger it is, holding the other factor constant.

Theorem 3 *The likelihood of being sent for a group that is sent with positive probability, q_i, $i \leq t$, is decreasing in r_i and increasing in s_i, holding the other factor constant.*

Proof: In equilibrium, the government must equalize the marginal benefit of searching each group. Taking derivatives of Equation 4 with respect to p_i and p_j and equating we have

$$q_i r_i - \frac{\partial c(p_i, s_i)}{\partial p_i} = q_j r_j - \frac{\partial c(p_j, s_j)}{\partial p_j}.$$

Increasing a group's reliability, r_i, must therefore be compensated by a decrease in the likelihood of sending it, q_i. Similarly, an increase in the size of the group, s_i, must be compensated by an increase in the likelihood of sending it. ■

This implies that if the groups were of equal size, the terrorist group's likelihood of sending them would be a decreasing function of their reliability. Coupled with the previous theorem, this would imply that the government is more likely to search a group the more reliable it is, but the terrorists are actually less likely to send someone the more reliable the group is. The government's search probabilities and the terrorists likelihood of sending are therefore inversely related as in the previous matching pennies analysis. To the outside observer, this will appear the rankest form of profiling, yet it will be the product of purely rational behavior by both sides, hence the label *seemingly unfair profiling.*

How can we determine the specific equilibrium values for q_i and p_i? The algorithm developed by Powell (2005) can be employed. Start the government out with no searching, $p_i = 0$, $\forall i$. The terrorist group will naturally set $q_1 = 1$ and $q_i = 0$, $\forall i > 1$ since group 1 is most reliable and no one is being searched. The government then increases p_1 until the terrorists are willing to mix between groups 1 and 2, namely when $(1 - p_1)r_1 - k_1 = r_2 - k_2$. Then the government increases both p_1 and p_2, maintaining $(1-p_1)r_1 - k_1 = (1-p_2)r_2 - k_2$, until $(1 - p_2)r_2 - k_2 = r_3 - k_3$ and the terrorists are indifferent between groups 1, 2 and 3. The government continues until searching more thoroughly becomes too expensive.

Note, if screening is very expensive, or group 1 is very large (high s_1), the government may be unwilling to search group 1 thoroughly enough to make the terrorist group indifferent between groups 1 and 2. In this case the equilibrium would exhibit *fair profiling* because the government would search only group 1 but this would be justified by the fact that the terrorists only send group 1 members. Also, if the government were willing to search groups 1 and 2, but group 2 was much smaller than group 1, it could be that the government searches group 1 more intensely, and the terrorists also send this group more frequently, also resulting in fair profiling. Thus, at airports where group 1 members are numerous and group 2 members scarce, say, in Saudi Arabia, fair profiling may be enacted. At airports in target countries, however, where group 1 members are likely to be scarce in comparison to group 2 and higher members, profiling will be seemingly unfair.

2.2 Implications of not Profiling

If the government were constrained to search all groups at the same rate ($p_i = p_j, \forall i, j$), then the terrorist's rational response is to send exclusively group 1 members, for they are the most reliable and are not any more likely to be searched than anyone else. The government then faces a slightly different optimization problem, namely to maximize

$$1 - (1 - p)r_1 - c(p, 1)$$

where p is the common screening intensity.

The first order condition is

$$r_1 = \frac{\partial c(p, 1)}{\partial p}$$

which equates the marginal benefit from increased screening, r_1, with the marginal cost, assuming everyone is being searched at the common level. Assuming that the marginal cost is increasing, there will be an equilibrium level of searching intensity in the common search case, p^*.

The key result is that in the non-profiling case, the government does worse and the terrorists do better, so not profiling involves an unambiguous welfare loss for the government side and a gain for the terrorists.

Theorem 4 *If the government is constrained to implement a common search level, p^*, the government's payoff will decrease while the terrorist's payoff will increase.*

Proof: (In progress)

■

3 Extensions

1. Terrorists know more about reliability than government. Could have terrorists know whether agent is reliable or not, and only sends reliable agents. Similar model, then k takes over role of r and is interpreted as the cost of finding a reliable agent in group i. Could have more complicated information structure in which the terrorists know r and the government only knows a distribution from which r is drawn.

2. There is some normative cost for deviating from a common search procedure. Say the government faces an additional cost equal to the variance of the p_i. This variance is presumably an increasing function of progress in the Powell algorithm, so maximized at the end. This implies that in the new equilibrium all groups would be searched with lower probability? If cost were high enough could make common p rational. Two group version would be simple case.

3. Behavioral profiling. Search intensity a function of group membership plus behavioral clues. Not sure this affects things, since terrorists will try to send operatives trained to avoid this behavior from whatever group, and group membership doesn't seem to be correlated with ability to avoid such clues.

4 Conclusion

Profiling in the context of strategic terrorist groups is rational, when considered solely on the basis of what will stop more terrorist attacks. Other factors may constrain the United States from profiling, but the logic of interaction with terrorist groups suggests a powerful dynamic in its favor.

Table 2: Notation in the Game

n	Number of social groups
r_i	Reliability of group i (to terrorists)
k_i	Cost of recruiting an operative from group i
s_i	Size of group i
p_i	Likelihood government searches a member of group i
q_i	Likelihood terrorists select member of group i
c	Cost of searching
t	Least reliable group that is sent and searched

References

Carpenter, Phillip E. 2006. Suicide Bombers: Can Racial Profiling be Used to Find the Needle in the Haystack? PhD thesis Harvard University School of Extension Studies.

Kydd, Andrew H. & Barbara F. Walter. 2006. "The Strategies of Terrorism." *International Security* 31(Summer):49–80.

Persico, Nicola. 2002. "Racial Profiling, Fairness, and Effectiveness of Policing." *The American Economic Review* 92(December):1472–1497.

Powell, Robert. 2005. "Defending Against Terrorist Attacks with Limited Resources.".