Hiding in Plain Sight – Using Signals to Detect Terrorists*

Atin Basuchoudhary
Department of Economics and Business
Virginia Military Institute
Lexington, VA 24450.
Email: Basua@vmi.edu
Phone: (540) 464 7450

Laura Razzolini
Department of Economics
Virginia Commonwealth University
Richmond, VA 23284
Email: lrazzolini@vcu.edu
Phone: (804) 828 7187

May 2005

Abstract

In this paper, we study the interaction between a governmental security agency, such as the Transportation Security Agency, or the Immigration and Naturalization Service, and a terrorist organization, such as Al Qaeda. The governmental agency wants to stop the terrorists, but first must infer whether a visa applicant or an airline passenger is a terrorist or not, on the basis of some observable signal. On the other hand, the terrorist organization's objective is to get past security to commit murder and mayhem. We derive the equilibrium strategy under these circumstances. With a signaling model we evaluate specific anti-terrorist policies such as the creation of the new Homeland Security Agency and increased airport security screening.

* We would like to thank Daniel Arce, Dek Terrel and participants at the University of Arkansas seminar for providing helpful and insightful comments on the paper.

Introduction

Terrorism has emerged as a growing problem and a serious threat to civilization in a post-cold war world. The word “terrorism” was first used in France during the Reign of Terror in 1793. Originally, terror was considered the best way to defend liberty, but as the French Revolution progressed, the word took on other meanings of violence and guillotines. Today, the State Department defines terrorism as "premeditated, politically motivated violence perpetrated against noncombatant targets by sub-national groups or clandestine agents, usually intended to influence an audience." (US Department of State, 2002). Academics in the social sciences have studied this phenomenon in an effort to counter it. This has taken a number of different approaches. One approach has tried to understand why terrorists are rational (e.g. Landes, 1978, Pape, 2003). Another approach has been to show that terrorists are indeed rational (Hassan 2001, Margalit, 2003). A third approach has been to use rational choice models to understand the implication of various policy alternatives in an effort to counter terrorism (e.g. Sandler, Enders, and Lapan, 1991, Hartley and Sandler, 1995).

In this paper we focus on the interaction between two rational players – a governmental security agency, such as the Transportation Security Administration (TSA), or other agencies dealing with homeland defense, such as the Coast Guard, the Border Patrol, the Customs Service, or the Immigration and Custom Enforcement (ICE) department of the Department of Homeland Security (DHS), and a terrorist organization such as Al Qaeda. The governmental agency’s objective is to stop the terrorists, while letting the non-terrorists through.[1] The agency, however, must first infer whether a visa applicant or an airline passenger is a terrorist or not, and must base this decision only on some easily observable signal – thus saving on information gathering costs. The terrorist organization's objective, on the other hand, is to get past security to commit murder and mayhem. We study the behavior of these two actors, and derive their equilibrium strategies. We find a number of interesting implications for security policy as a result of this derivation, regarding, for instance, how often we should check passengers or applicants of a given type, or whether it is optimal to publicly announce that certain characteristics may signal terrorist activity.

The use of signals, such as gender or race, and statistical discrimination has also been studied in the context of police searches and criminal behavior. A recent literature pioneered by Knowles, Persico and Todd (2001) has characterized the Nash equilibrium of a simultaneous move game between the police and a specific group of the population, such as motorists, who may commit a crime. The police's objective is to minimize crime or maximize the guilt rate when deciding which vehicle to search, while motorists decide whether to carry contraband or not. In equilibrium unequal investigation rates across different races or groups can be attained even by unbiased police officers if members of one group incur higher costs to carrying contraband than the other group. Our paper, however, even though related to this literature, takes a different approach as it assumes that the criminal group of the population makes a rational decision not only whether or not to commit the crime, but also whether to use strategically a particular signal, such as race or gender.

[1] At the same time, the governmental agency wants to minimize delays and disruption. As some scholars argue, one of the terrorists' goals is simply to interfere and disrupt normal life.

Furthermore, in this paper, we do not try to understand the motivation of terrorists. Indeed, we assume that their motivation is rational. Experts agree that there is almost always a strategy behind terrorists' actions. Whether it takes the form of bombings, shootings, hijackings, or assassinations, terrorism is not random, spontaneous, or blind; it is a deliberate use of violence against civilians for political or religious reasons. Therefore, following the spirit of most game theoretic literature on terrorism (for an excellent survey, see Sandler and Arce, 2003), we model terrorists as rational actors, who follow the classical constrained optimization behavior, familiar to economists. In this sense, the present paper follows the Sandler and Arce's (2003) view that "… game theory is an appropriate tool for examining terrorism…."

Our paper is also a direct descendant of those that focus on deterring terrorist behavior (see Lapan and Sandler 1988, or Cauley and Im 1988). These papers, however, tend to present the game in a sequential framework, and deterrence is modeled as an insurance premium, so that the governments' choice of the optimal level of deterrence depends on its tolerance to risk. We, on the other hand, focus on a signaling framework, in which the governmental agency tries to deter terrorists by detecting them before they act. Note that our paper is also different from Lapan and Sandler's (1993), in which a government has to decide whether to resist or to capitulate to terrorist demands. In this framework, the uninformed government uses first period attacks by the informed terrorists as a signal to detect terrorist capability.

We divide our paper into three sections. In Section 1 we describe the game. In Section 2 we solve for, and discuss, the equilibrium strategies. We conclude in Section 3. The game tree is available in Figure 1.

1. The Game

In this game we have two players, an airline passenger or a visa applicant and a government agency. The passenger or applicant may either be a terrorist (T) or not (NT). The government agency or security checks (C) or does not check (NC) the passengers or the applicants. The passenger knows his/her own type, while security knows only the probability that a passenger is a terrorist; that is, P (T) = α. Therefore, the probability a passenger is not a terrorist is P (NT) = 1 – P (T) = 1– α.

The government agency must infer whether a passenger is a terrorist or not, on the basis of some attribute displayed by the passenger or applicant. For instance, attributes such as race, name, and country of origin, travel records, or clenched fists may be a signal of whether a passenger is a terrorist or not. A passenger of either type may have the attribute (A) or not have it (NA). The signal is public knowledge.[2]

The players are rational and try to maximize their objective function or payoff. The elements of each type of player's payoffs are described as follows: A passenger or applicant who is a terrorist derives a positive utility, $W > 0$, if his objective is met. On the other hand, if he is checked by the governmental agency, he incurs a cost $C_c > 0$ and, assuming that all checked terrorists are apprehended, he enjoys $W = 0$.[3] Terrorist passengers who display the signal A (respectively, NA) incur a cost C_A (respectively, C_{NA}). C_A or C_{NA} can be interpreted as the cost incurred by the terrorist organization to

---

[2] One might argue that these signals are not or should not be public. However, we focus on the class of signals that the Homeland Security Department wants us to watch out for (people wearing trench coats in summer), as well as signals such as race. These signals are public.

[3] To keep the model simple, we assume perfect monitoring by security. Alternatively, we could assume that checked terrorists are detected with some probability λ. However, the main results of the paper would not change, but just be scaled by the factor λ.

recruit persons displaying the particular attribute or signal.[4] A passenger who is not a terrorist derives positive utility from having the security in place, $N > 0$[5]. If he is checked he also incurs a cost $C_c > 0$ for being checked.[6] Finally, security derives a positive benefit $B > 0$ from checking and catching a terrorist while incurring a cost $C_s > 0$ for checking passengers. On the other hand, $B = 0$ if a "checked" person is not a terrorist. C is the cost of not checking a terrorist. Figure 1 represents the extensive form game tree associated with this game.

In what follows we assume that

$$B > C_S - C, \tag{1a}$$

$$W - C_i > 0 \quad i=A, NA \tag{1b}$$

$$W + C_C > 0. \tag{1c}$$

Condition (1a) requires the benefit from checking and capturing a terrorist to be greater than the cost of checking net of the cost of not capturing a terrorist. Intuitively, this means that for security the cost of checking cannot be prohibitive compared to the cost of not apprehending a terrorist. If the cost of checking is so high that when compared to the cost of letting terrorists through it exceeds the benefit from catching a terrorist, then rationally we should not be catching terrorists. However, since we do want to catch terrorists, this condition seems plausible.

---

[4] Non-terrorist passengers do not incur the cost C_A or C_{NA}, as they do not consciously recruit people showing the particular attribute, but just happen to possess and display the attribute.

[5] We assume that all terrorist passengers get 0 benefit from security (i.e. $N = 0$).

[6] All types of passengers bear this cost if they are "checked." C_c can be thought of as the cost of being thought a terrorist when you are not – this may be pure embarrassment or could be a "real" cost of being caught for some other "crime," like immigration or customs violation, discovered because of the additional scrutiny.

Condition (1b) requires that for terrorists the expected benefit, net of recruiting costs, when non-checked by security to be positive; and Condition (1c) requires that either the benefit from a successful mission or the cost of being checked by security to be positive.

2. Equilibrium Strategies

In solving the game, we will analyze first separating equilibria, then pooling equilibria and finally we will consider mixed strategies equilibria.

2.1. Separating Equilibrium

There are two possible pure-strategy separating equilibria in this game.

1. All passengers with the suspicious attribute (A) are terrorists, while all passengers without the suspicious attribute (NA) are not terrorists.
2. All passengers with the suspicious attribute (A) are not terrorists, while all passengers without the suspicious attribute (NA) are terrorists.

Consider Case 1 first: all terrorists choose to send applicants/passengers with the signal A, while all non terrorist passengers do not display the signal. Under these circumstances, security should check only people displaying the A attribute, since all are terrorists. But if only the people sending the A signal get checked, then it will be optimal for the terrorists to deviate from this equilibrium, unless the utility from choosing A and being checked is greater than the utility from choosing NA:

$$-C_C - C_A > W - C_{NA}\ . \qquad (2)$$

That is, if the cost of sending the NA type is extremely high, then the attribute could serve to separate the terrorists from the non terrorists because it would be optimal for all the terrorists to choose A. However, given Condition (1b) (i.e., $W - C_{NA} > 0$)

condition (2) will not hold for positive values of C_C and C_A , and a separating equilibrium does not exist.

Similarly, in Case 2 while all non terrorist passengers happen to display the signal A, all terrorists choose NA. Again security will check only people displaying the NA attribute, since all are terrorists. But if only the individuals sending the NA signal get checked, then it will be optimal for the terrorists to deviate from this equilibrium, unless the utility from choosing NA and being checked is greater than the utility from choosing A:

$$– C_{NA} – C_C > W – C_A , \qquad (3)$$

which, given condition (1b), is not possible for reasonable value of C_C and C_{AN}.

Result 1. *There is no pure strategy separating equilibrium in this game for positive value of C_C, C_{NA} and C_A.*

Neither of the two possible separating equilibria is optimal because in both cases the passenger or applicant, given the security agency's strategy, can improve his/her payoff by deviating. That is, if security's strategy is to play C if A and NC if NA, then the best response for a terrorist organization would be to send terrorists passengers without the suspicious attribute.

This result suggests that a policy of complete racial profiling may be counterproductive. If security only checks passengers of a certain racial type and not others, then terrorist groups would choose passengers of the non-checked racial type as hijackers or suicide bombers.

This result may also suggest that public announcements that certain types of attributes may indicate or signal that a person is a terrorist may be counterproductive. After all, announcing that people should be on the lookout for individuals with clenched fists as a possible indicator of suicide bombers, might make the potential terrorists unclench their fists. In these circumstances, it may, instead, be a better policy for the security agency to keep such information private. However, in this game we assume that the identity of the signal is not, or cannot be, private information.

2.2. Pooling Equilibrium

In this game there are two possible pooling equilibria.

1. Only passengers with the A attribute travel; that is, the terrorists choose A and all non terrorists display A.
2. Only passengers with the NA attribute travel; that is, the terrorists choose NA and all non terrorists display NA.

For both of these equilibria, security's decision to check or not depends on their expected payoffs from checking, and their expected payoff, in turn, depends on their belief over the type of passenger they are checking. In this case, security's belief (based on Baye's rule) that a particular type of passenger is a terrorist or not is the prior distribution of types – the probability of being a terrorist is α and the probability of being a non-terrorist is 1 – α. That is,

$$P(T \mid A) = \alpha.$$

Security is indifferent between checking and not checking a person displaying the A attribute if

$$P(T \mid A)(B - C_S) + [1 - P(T \mid A)](-C_S) = P(T \mid A)(-C),$$

which reduces to:

$$P(T \mid A) = \frac{C_S}{(B + C)}. \qquad (4)$$

There are three possible cases in this situation as we compare the right hand side of (4) with the expected probability of being a terrorist given the attribute, which is equal to the prior probability of being a terrorist, α.

i) $\alpha < \frac{C_S}{B+C}$. In this case, the probability that a passenger is a terrorist is very small and security never checks. In either possible pooling equilibrium there is no incentive for any type of passenger to deviate: For instance, when only passengers with the A attribute travel, terrorists have no incentive in sending a passenger without the attribute. These are stable equilibria. This suggests that whenever the probability that a passenger is a terrorist is very small it is hard to justify having a security apparatus. Moreover, since all passengers send the same signal, the signal provides no information to the security agency.

ii) $\alpha > \frac{C_S}{B+C}$. In this case, the probability that a passenger is a terrorist is higher than the threshold, and security always checks, while passengers of both types have an incentive to deviate. In the pooling equilibrium where only passengers with the A attribute travel, a passenger who is also a terrorist has an incentive to deviate from the pooling equilibrium, since $0 > W - C_{NA} > -C_C - C_A$. Thus, pooling on A is not a stable equilibrium. Similarly, in the pooling equilibrium where only passengers with the NA attribute travel, the terrorist type passengers have an incentive to deviate. Thus pooling on NA is not possible either. Consider, for example, the program National Security Entry Exit Registration System (NSEERS), which was enacted for some time after 9/11. One component of this program required that Arab immigrants already in the US be subject to special registration rules simply because of their nation of origin.[7] These immigrants could be either terrorist or non-terrorist. Similarly, white (or black) US citizens could also be of both types. However, the program did not require for them to register or to submit themselves to a background check. The lack of a pooling equilibrium may suggest that Al Qaeda, when faced with increased surveillance toward people of Arab origin, might try and find recruits among white (or black) US citizens, even if they are more costly to recruit.

iii) $\alpha = \frac{C_S}{B+C}$. In this case, the probability that a passenger is a terrorist is equal to the threshold, and security is indifferent between checking and not checking. This implies that security could randomize between checking and not checking. It would be foolish for terrorists to pool with other passengers with either A or NA attributes.

This discussion leads to:

Result 2. *A pure strategy pooling equilibrium exists for* $\alpha < \frac{C_S}{B+C}$. *For* $\alpha \geq \frac{C_S}{B+C}$ *there is no pure strategy pooling equilibrium.*

2.3. <u>A Mixed Strategy Equilibrium.</u>

A mixed strategy equilibrium exists for values of $\alpha \geq \frac{C_S}{B+C}$. Our game has two information sets in which a mixed strategy may be optimal and is possible. In either set, the terrorist's strategy is P(A|T); that is, the terrorists must optimally determine how often

[7] This component of NSEERS was remarkable for the fear and loathing it generated among immigrants – the very people who could be useful in generating human intelligence about terrorist acts. Between September 11, 2002 and January 17, 2003 there were 23,400 registrations out of which 164 resulted in detention (Porter, 2003). No terrorists were charged or detained under this program (Romero, 2004).

they want to send a person displaying the A attribute to complete their mission, given that non terrorist individuals travel with and display both attributes. Security's strategy is P(C|A) or P(C|NA), depending on which information set it finds itself, which indicates how often security wants to check somebody who is displaying the A or the NA attribute. We do not calculate the non-terrorist type passengers' strategy because it would be hard to imagine that they can coordinate their strategy as a group. Thus, P(A|NT) and P(NA|NT) describe the proportion of the population with the respective attribute rather than a strategy. Also, since security is mainly concerned with catching terrorists, we consider only its strategy with respect to terrorists, rather than the optimal strategy toward non-terrorists.[8]

At the right-hand information set, security's beliefs about the probability whether somebody with the A attribute is a terrorist or not, P(T|A) and [1 - P(T|A)], is given by Bayes' rule. So,

$$P(T \mid A) = \frac{\alpha P(A \mid T)}{\alpha P(A \mid T) + (1-\alpha) P(A \mid NT)}. \qquad (5)$$

As we have seen before, security is indifferent between checking and not checking if it sees the attribute A if

$$P(T \mid A) = \frac{C_S}{(B+C)}. \qquad (6)$$

Substituting (6) into (5) and rearranging, we get

[8] It turns out that security's optimal strategy in this case would be to check passengers with the A attribute less than they would check passengers with the NA attribute, if they wished to avoid embarrassing the passengers displaying the A attribute. One would imagine, however, that catching terrorists would drive the security agency more than avoiding embarrassment to passengers.

$$P(A|T) = \frac{(\alpha - 1)C_S P(A|NT)}{\alpha(C_S - B - C)}. \quad (7)$$

Terrorists, on the other hand, are indifferent between sending passengers with the A attribute or with the NA attribute if

$$(-C_C - C_A)\,P(C|A) + (W - C_A)\,P(NC|A) =$$

$$(-C_{NA} - C_C)\,P(C|NA) + (W - C_{NA})\,P(NC|NA).$$

Since, $P(NC|A) = 1 - P(C|A)$ and $P(NC|NA) = 1 - P(C|NA)$, we get:

$$P(C \mid A) - P(C \mid NA) = \frac{C_{NA} - C_A}{(C_C + W)}. \quad (8)$$

We are, therefore, able to state the next result:

Result 3. *The optimum mixed strategy for the security agency is to randomize such that* $P(C \mid A) - P(C \mid NA) = \frac{C_{NA} - C_A}{C_C + W}$*, and the optimum mixed strategy for the terrorists is to randomize such that* $P(A|T) = \frac{(\alpha - 1)C_S P(A|NT)}{\alpha(C_S - B - C)}$.[9]

Figure 2 shows the two players' reaction correspondences given $\alpha \geq \frac{C_S}{B + C}$. From Equation (7) we see that security should always check if $P(A|T) > \frac{(\alpha - 1)C_S P(A|NT)}{\alpha(C_S - B - C)}$, while from Equation (8) we derive that the terrorist group should send terrorists displaying the particular attribute if the probability that security check a passenger with the attribute is less that the critical value, that is $P(C \mid A) < P(C \mid NA) + \frac{C_{NA} - C_A}{C_C + W}$. The

[9] Without condition (1a) holding, we would get a negative probability.

intersection between the two reaction functions represents the mixed strategy equilibrium.

The security agency's equilibrium strategy suggests that the frequency of security checks of individuals with the A attribute ultimately depends on the difference in the recruitment cost between the two types. That is, the greater the cost to recruit NA types versus the A types, the more often security should check the latter. Thus, it is optimal to check passengers of a certain racial type, as long as the cost to recruit terrorist of different racial origins is relatively higher.

Considering the terrorists' optimal mixed strategy, it is interesting to observe a direct linear relationship between P(A|T) and P(A|NT), the proportion of the general population displaying the particular attribute. If a certain attribute is more common, P(A|NT) raises, terrorists are more likely to choose that attribute even when the security agency is checking passengers with that attribute. Hiding in open sight is easier if they can blend in. Figure 3 shows this direct relationship for specific values of the parameters α, B, C_S, and C. If P(A|NT) equals zero then P(A|T) become zero as well. In a world in which nobody displays a certain attribute, terrorists as well will be forced not to display it.

Moreover, P(A|T) falls as α increases, as shown in Figure 4. This may appear counter-intuitive. However, this result is true for a given value of P(A|NT) (in Figure 3, P(A|NT) is assumed to be 0.1). This suggests that as the number of terrorists rise relative to the general population, but the proportion of the population with the suspicious attribute does not, the terrorists are less likely to send people with that suspicious attribute.

Finally, considering the left-hand information set, observe that the security's beliefs P(T|NA) and 1 - P(T|NA) are, as previously, determined by Bayes' rule. In this case, again the optimum mixed strategy for the security agency is to randomize so that $P(C \mid A) - P(C \mid NA) = \frac{C_{NA} - C_A}{C_C + W}$ and for the terrorists to randomize such that $P(A \mid T) = \frac{(\alpha - 1) C_S P(A \mid NT)}{\alpha (C_S - B - C)}$. The comments we made for the right-hand information set are also quite valid here.

Conclusion

We have derived the optimal strategy for a security agency that has to infer whether a passenger or an applicant is a terrorist or not based on certain visible attributes, as well as the optimal strategy for a terrorist who has to decide what attributes to hide from a security agency. We find that there is no separating equilibrium for reasonable positive values of the players' payoffs and costs. It is, therefore, sub-optimal to only check people with a certain type of attribute and not those with other attributes. In other words, always checking people of a certain ethnic origin or racial type, while not checking people of other ethnic origins or racial types could be counterproductive toward national security. Similarly, public announcements that certain characteristics may be a signal that a person is a terrorist may be not as effective. After all, it is likely that potential terrorists will unclench their fists or dispose of their bulky jackets, if the general public is on the lookout for people with clenched fists and bulky jackets.

We showed that an optimal mixed strategy is possible. It turns out that it is optimal to check some passengers with a certain attribute, if the cost to recruit potential terrorists not displaying that attribute is relatively higher. This suggests that using racial or ethnic origin information as a component of a passenger profile and checking those passengers who fit that profile with a certain probability is optimal, as long as it is less costly than to recruit terrorists of different racial and ethnic origin. At the same time passengers of other ethnic origins should be checked with a certain probability as well. This is one way of using openly available, and therefore cheap, information to stymie terrorist organizations.

One obvious restriction in our model is that both the security agency and the terrorists know the identity of the signal used to form inferences about the underlying nature of the passenger: that is, whether the passenger is a terrorist or not. If the terrorist does not know the signal that will form the basis of a security check then they cannot strategize with respect to the signal. So the visibility of the signal as public knowledge drives our game. However, this restriction is not completely unrealistic, since we often hear public announcements about different attributes and characteristics we should be on the lookout for people wearing bulky jackets on warm days, or people with clenched fists. Moreover, a security agency faces an information trade-off – by making the signal public knowledge they may induce strategizing by the terrorists, but at the same time, to the extent the signal is valid, the possibility of catching the terrorist may increase.

Last, it should be clear that in equilibrium, checking by security cannot be a plausible technique to catch all terrorists and eliminate terrorism after all. Thus, public policies with regard to terrorism cannot be uniquely focused on security through checks and deterrence. Clearly the most effective way of reducing and eliminating the threat of terrorism is by affecting the underlying state of nature – in our model this corresponds to reducing the value of α. In this sense, it is important to recognize that our model is embedded in a larger game, in which α would be an endogenous variable. Public policy must not devalue the role of rationally checking and verifying all sources of information. At the same time, public policy must act to change the underlying causes of terrorism. That is the challenge to civilization.

Bibliography

Cauley and Im. 1988. "Intervention Policy Analysis of Skyjackings and Other Terrorist Incidents." *American Economic Review,* 78, pp. 27-31.

Gibbons, R., 1992. *Game Theory for Economists*. Princeton University Press. Princeton, New Jersey.

Hartley, K. and Sandler, T., 1995. *Handbook of Defense Economics*. Amsterdam: North Holland.

Hassan, N. 2001. "An Arsenal of Believers." *The New Yorker*. November 19. pp. 36 – 41.

Knowles, J., Persico, N., and Todd P., 2001. "Racial Bias in Motor-Vehicle Searches: Theory and Evidence." *Journal of Political Economy*, 109, pp.203-29.

Landes, W. M., 1978. "An Economic Study of US Aircraft Hijackings, 1961 – 1976." *Journal of Law and Economics.* 21, pp. 1 – 31.

Lapan, H. E., and Sandler T., 1988. "To Bargain or Not To Bargain: That Is the Question." *American Economic Review*. 78, pp. 16-20.

Lapan, H. E., and Sandler T. 1993. "Terrorism and Signaling." *European Journal of Political Economy*. 93, pp. 383-97.

Margalit, A. 2003. "The Suicide Bombers." *The New York Review of Books*. Jan 16.

Pape, R.A. 2003. "The Strategic Logic of Suicide Terrorism." *American Political Science Review*. 97.

Porter, Charlene, 2003. "Suspected Terrorists, Convicted Criminals Identified in New U.S. Border Checks." *Washington File*. U.S. Department of State International Information Programs.

Romero, Victor, 2004. "Race, Immigration, and the Department of Homeland Security." *St. Johns Journal of Legal Commentary*. 19 (1).

Sandler, T., and Daniel G. Arce. 2003 "Terrorism and Game Theory." *Simulation and Gaming.* 34(3).

Sandler, T.; Enders, W., and Lapan, H.. 1991. "Economic Analysis Can Help Fight Terrorism." *Challenge*. 34, pp. 10-17.

U.S. Department of State, 2002. *Patterns of Global Terrorism.* Washington, DC: US Department of State.

U.S. Department of Homeland Security, 2003. *Changes to National Security Entry/Exit Registration System*. Washington DC: Press Office, DHS.

Figure 1- The Game Tree.

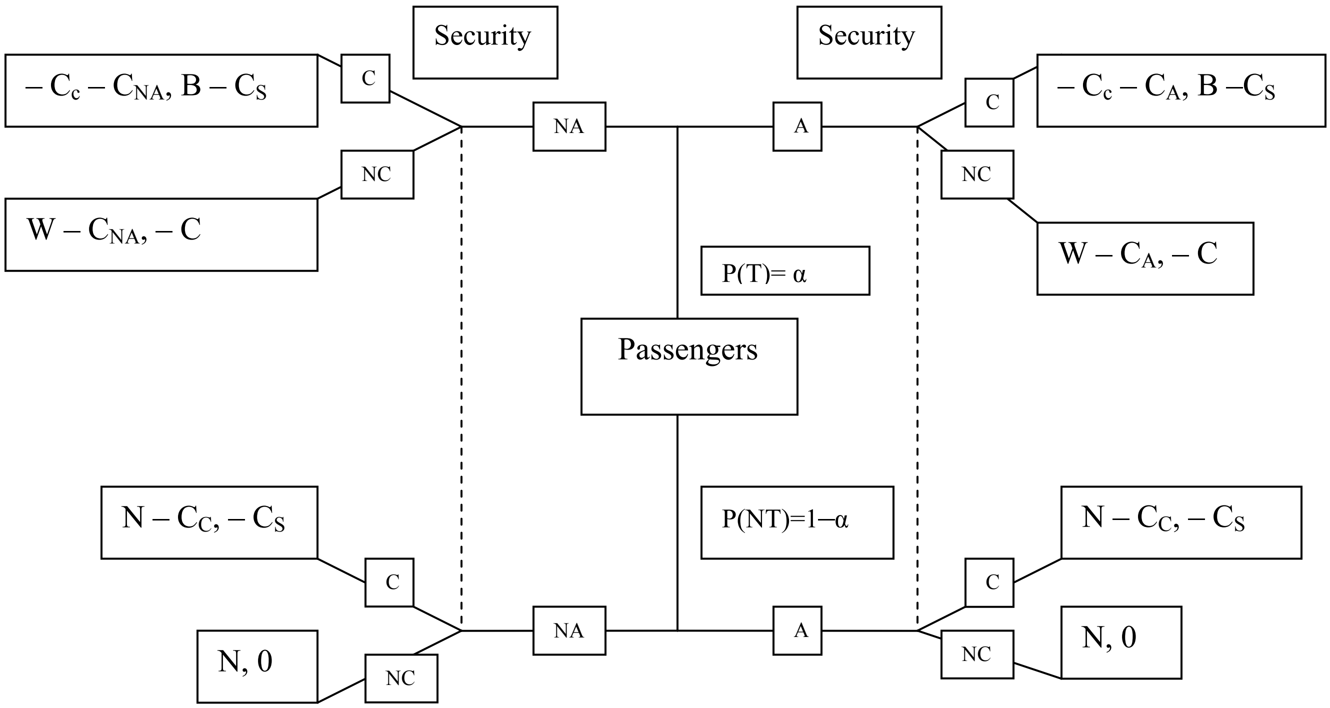

Figure 2 – The two players' reaction correspondences when $\alpha \geq \frac{C_S}{B+C}$.

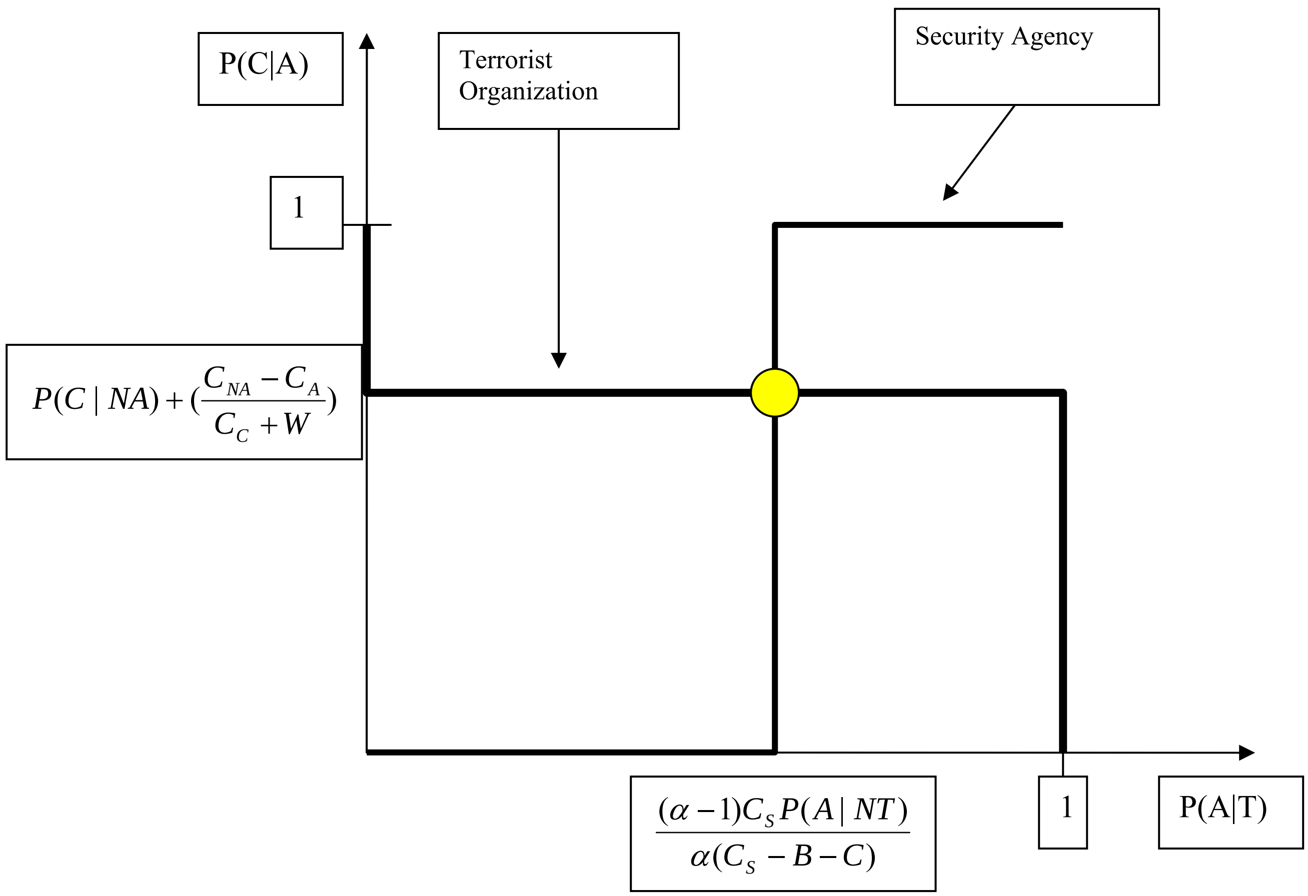

Figure 3 - Terrorist strategy as a function of the probability of finding an attribute among non-terrorists.

[$\alpha = 0.45$, $B = 100$, $C_S = 50$, and $C = 25$]

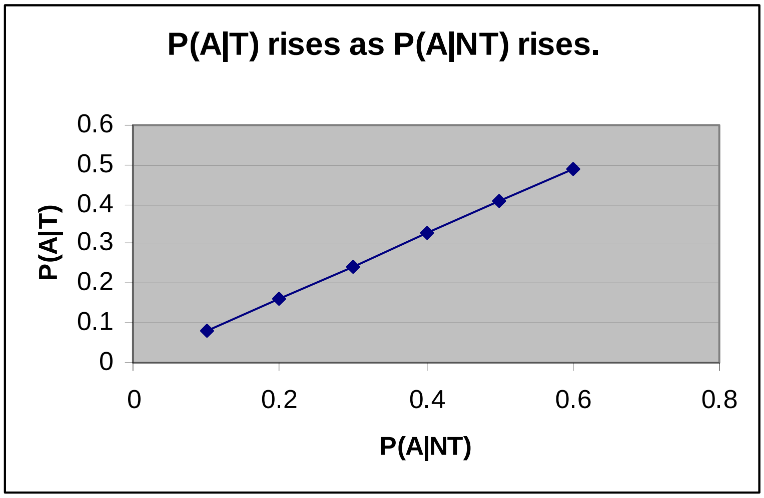

Figure 4 - Terrorist strategy as a function of the underlying probability of being a terrorist.

[P(A|NT) = 0.1]