Terror Alerts and Beliefs about Terrorism*

J. Atsu Amegashie
Department of Economics
University of Guelph
Guelph, ON. N1G 2W1

Edward Kutsoati†
Department of Economics
Tufts University
Medford, MA 02155

January 14, 2004

Abstract

We provide a simple model to study the U.S. government's nationwide terror alert system. We show that the alert level will depend as much on the public's perception of the risk of an attack as it does on the government's intelligence information. If the public perceives an attack to be very likely or severe, they take adequate private protection (e.g., reduction in social and economic activities) so that alert level is optimally set below the full–information level. When the public is relatively uncertain about the nature of an attack, then the government may put the nation on a higher alert (relative to the full–information case), which may cause the public to become 'complacent' if no attack occurs. Our results demonstrate the challenges faced by the U.S. government and show that the public's response to changes in the terror alert level could be partly explained by information asymmetry.

Keywords: Terror alerts, Asymmetric information, Signalling.

JEL classification: D74, D82

*This paper has benefited from conversations with Sharun Mukand. We are however responsible for all errors.
†Corresponding author: Tel. (617) 627 2688; email: edward.kutsoati@tufts.edu.

1 Introduction

The September 11, 2001 attacks (hereafter 9/11) triggered new ways of dealing with security in the U.S. Perhaps, the most significant transformation is the creation of the Department of Homeland Security (DHS). The DHS merges 22 existing law enforcement agencies to coordinate efforts in preventing attacks and to protect Americans. The DHS also uses a security advisory system, started in March 2002, to warn the public, state and local authorities of potential attacks. The five color–coded alert system consists of *green* for low risk, *blue* for general risk, *yellow* for significant risk, *orange* for high risk, and *red* for severe risk of an attack.[1] The objective of the advisory system is to inform the public in a very simple and easily identifiable way of the necessity to be extra vigilant and take adequate steps to protect oneself against potential attacks.

This paper provides a theoretical analysis of some of the challenges faced by the DHS in its implementation of the terror advisory system. Most Americans by now believe that the country will remain at risk to terrorist attacks for the foreseeable future. But a poll conducted by Opinion Dynamics Corporation one year after its implementation, showed that close to 60% of Americans do not think the color–coded alert system is helpful. Yet, a heightened state of alert has sometimes led to a tremendous increase in safety concerns and a subsequent drop in participation in social activities. For example, the heightened alert on February 7, 2003 caused a lot of panic as households sought to buy items (such as duct tapes) for protection from an attack. How can these two behaviors be reconciled? More generally, do heightened alerts necessarily change the public behavior and attitudes towards risk of attacks? And if so, when are these changes likely to occur?

We analyze a simple model of information asymmetry about potential attacks that can be used to answer these, and other, questions. From its intelligence sources, the DHS learns more about the nature or likelihood of a potential attack (beyond what the public knows) and issues a nationwide alert to warn the public. The government has a dual objective: to secure the nation while promoting a sustainable economic growth. The government knows that if individuals perceive an attack to be likely or severe, then they will be less willing to participate in social or economic activities, as they protect themselves and assets against attacks. Keohane and Zeckhauser (2003) call this "avoidance" behavior, and they also argue

[1] The alert level has been at *yellow* for most of the time, but has been raised 5 times from *yellow* to *orange*: September 9, 2002; February 7, March 17, May 20, and Dec 21, 2003. It was lowered from *orange* to *yellow* on September 24, 2002; February 28, April 16, May 30, 2003 and January 9, 2004. There has been no other terror–related incident since 9/11.

that this could have large negative effects.[2] In addition, the government will also be concerned about how far citizens will be willing to sacrifice their some civil liberties for increased safety.[3] Thus, the competing demands of security and economic freedom can cause the government to issue an alert level below what will be warranted by its information.

As a benchmark case, we first identify the state of alert if the information about a potential threat were publicly available. We then analyze the realistic case whereby, for obvious reasons, information gathered by intelligence sources is not made available to the public. Under asymmetric information, any revision to the alert level, and its impact on the public's behavior, depends as much on underlying beliefs about acts of terrorism as it does on the DHS's intelligence. For example, suppose the public believes that an attack is *very likely.* Then they all take adequate measures to protect themselves, including a reduction in participation in social and economic activities. In this case, the government may have an incentive to *not* confound the public by putting the nation on a very high alert: a further heightened alert level will lead to a further decline in social and economic activities. Hence for this range of prior beliefs, there will be a pooling equilibrium in which the government issues a lower alert level (relative to the full–information case), regardless of its information.

But if the public is uncertain about likelihood or severity of an attack, then any information that can be gleaned from intelligence sources becomes very useful. In this case, the alert level is very effective as the public responds to it as they would in the full–information case.

Finally, we also show that for intermediate priors about the nature or likelihood of an attack, the government is more likely to put the nation on a higher alert level (relative to the full–information case). The intuition behind this result is relatively simple: When the government also has an incentive to promote economic activity, there is an incentive to choose a lower alert level to assure the public of their safety (especially when security agents are well prepared to deal with a possible attack). If the public recognized this, they may be prone to attach lesser importance to the advisory system. This makes it difficult for the government to use the advisory system to effectively warn the public of a potential attack. One solution is to raise the alert level higher than what would be expected in the full–information case. However, if adopted too often, this strategy may have the tendency of causing the public too become

[2]Charles Stein of the `Boston Globe`, puts this succinctly: "A fear tax has been levied on the economy and [Americans] are paying it... The fear tax shows up in consumer spending, business investment, and stock prices. Add them all and the total is significant." (`Boston Globe, February 23, 2003, pp. F1`)

[3]For example, the USA Patriot Act, signed into law on October 26, 2001 to empower law enforcement agencies in their efforts to prevent future attacks is seen by some critics as a blow to the civil liberties of ordinary Americans. See also a study by Viscusi and Zeckhauser (2003) on targeted searches and passenger screening at airports.

complacent of an attack, in particular if no attack occurs during a heightened state of alert. Our results therefore show that the public's response to changes in the terror alert level could be partly explained by information asymmetry.

Our analysis therefore complements the literature on terrorism, which has largely focused on the use of terror attacks (e.g., skyjacking, kidnaps, bombings, threat, etc) to force a government to concede to certain demands. For example, Lapan and Sandler (1988) investigates the time inconsistency of government's policy of never to negotiate with terrorists groups who take hostages, while Lee (1988) examines the issue of optimal retaliation by a government in the case of state–sponsored terrorism. In another paper, Lapan and Sandler (1993) analyzed a signalling game in which the government is uniformed about the terrorist group's capabilities to attack. The government learns about the group from earlier attacks and then decides whether to fight back or concede. They showed that the government may prefer a partial–pooling equilibrium to a never surrender equilibrium. Enders and Sandler (1993) also provide evidence suggesting that terrorists seem to concentrate their efforts on "soft targets" whenever security around a "hard target" is tightened. Our paper departs from this trend: our objective is to provide an economic analysis of how the government, through a terror advisory system, engages the public in the fight against terrorism.

The rest of the paper proceeds as follows. Section 2 outlines a model of the terror alert system. In section 3, we provide the equilibrium analysis, first in a benchmark case when all information is public and later analyze the case of private information. A brief discussion of the main results is provided in section 4, and section 5 concludes the paper. Proofs of all results are in an appendix.

2 A model of an alert system

We consider a society that is susceptible to a terrorist attack, and concerned about the severity of an attack. We assume that the severity of an attack could be high (H) or low (L).[4] Initially, it is known by all that the probability that the severity of an attack will be high is $p(H) = \omega > \frac{1}{2}$. The assumption that $\omega > \frac{1}{2}$ is consistent with the fact that most people view the society to be at a risk to terrorist attacks.

The government then observes a private signal, s ($s =$ high or low), about the threat level through its intelligence sources. We assume that $P(s = h|H) = P(s = l|L) = \phi$ and $\phi > \omega$, so

[4]Our results are qualitatively the same if instead we assume the severity of an attack is known, but there is uncertainty about the likelihood of an attack. However, this alternative formulation complicates the analysis without adding any new insights to the results.

that the government's information is *more reliable* than the prior information. After assessing its information, the government then issues an alert level $\mathcal{A}$ to the public. For the purpose of tractability, we assume $\mathcal{A}$ is continuous, as opposed to the discrete form of the DHS's advisory system. This assumption does not affect the nature of our results. In particular, the DHS color–coded system could be made continuous if we add a duration component: i.e., an *orange* alert can be considered higher than another if it has a longer duration.

It is assumed that individual members of the society derive utility from economic/social activity, and that the level of utility is impacted by the severity of an attack. In particular, we assume that the utility from a given level of activity, X, is

$$U = \begin{cases} HX^{\rho} & \text{if severity of attack is HIGH;} \\ LX^{\rho}, & \text{if severity of attack is LOW,} \end{cases}$$

where $L > H > 0$ and ρ is a constant, $0 < \rho < 1$. Thus a bad attack lowers the utility derived from a given level of economic activity.[5] There are several ways to interpret the utility function. We can think of X as an aggregate participation in social activities (e.g., attending a professional baseball game, where the fun at the game could be ruined if the stadium had to evacuated due to a potential terrorist attack). Alternatively, X could be interpreted as aggregate level of investment and U as the level of production. A severe attack therefore disrupts productive activity and leads to very low output for any given investment level. More generally, a low aggregate X can be thought of as private protection against potential attacks.

The increased security measures that accompanies a high alert imposes an "inconvenience cost" on individuals. For example, if the country is on a high alert, individuals will have to undergo lengthy security procedures or random searches when boarding a plane, attending a sporting event (e.g., a baseball game), etc. We assume that inconvenience costs increase exponentially in the alert level. Thus, each a member of the society chooses X to:

$$\max_{X} U = E[S|\Omega]X^{\rho} - Xe^{\mathcal{A}}, \tag{1}$$

where $E[S|\Omega]$ is the individual's expectation of S given her information Ω. An individual's information set, Ω, includes whatever is publicly observable and inferred from the government's warning to the public.

Given an alert level $\mathcal{A}$ and individuals' beliefs about the nature of an attack, aggregate economic activity can be determined from the representative individual's maximization problem. The first order condition from utility maximization gives $X = (\rho E[S|\Omega])^{\frac{1}{1-\rho}} e^{-\mathcal{A}}$, and the

[5] Notice the unconventional use of notation here stemming from our assumption that $L > H$. We choose this notation to differentiate between a LOW and HIGH attack.

(natural) logarithm of this can be normalized as

$$\mathcal{X} = \ln E[S|\Omega] - \mathcal{A}. \tag{2}$$

Society's loss from a potential terrorist attack will depend on the severity of the attack, the aggregate level of economic activity at that time, and the preparedness of federal agents to respond to an attack. Define $\overline{\mathcal{X}}_S = \ln \overline{X}_S$, where $\overline{X}_L > \overline{X}_H$ are constants and, for simplicity, let $\overline{\mathcal{X}}_L - \overline{\mathcal{X}}_H = 1$. We then assume that the cost from an attack with severity S is given by $\left(\mathcal{X} - \overline{\mathcal{X}}_S - \gamma\mathcal{A}\right)^2$, where $\gamma > 0$ captures the nation's preparedness for an attack. So, given the severity of an attack, S $(S = H, L)$, it is assumed that the cost of a successful attack will be higher if aggregate economic activity is high, and in particular, exceeds some desirable level, $\overline{\mathcal{X}}_S$. Secondly, security agents play a vital role both in mitigating the loss from an attack. That is, the more prepared security agents are in dealing with an attack, the less costly (or destructive) it is to the society. So for example, suppose security agents at a baseball game suspect a potential attack. Then the speed with which they can evacuate fans to a safe area can greatly reduce the losses if an attack occurs. However, such efforts can be hampered in a sold–out game.

Notice that the cost of attack is decreasing in the alert level. This is because the government commits more resources to detect and prevent terrorist activities if the alert level is raised: the public may attach less credibility to a high alert if anti–terrorism activities are not increased accordingly. For example, we observe many more checks at airports and ballparks when the alert level is raised. We also capture the effectiveness of anti–terrorism investment with the parameter γ. The higher γ is, the better prepared agents are, and hence an attack will be less costly. For example, in May 2003, the cities of Seattle and Chicago held a five–day bioterrorism drill that included hundreds of firefighters and police officers responding to a mock explosion of a radioactive "dirty" bomb. Such exercises can be interpreted as increasing γ in our model.[6]

Our formulation of the cost function can be interpreted as there being some threshold level of activity such that it is equally costly to exceed this threshold just as it is to go below it. If economic activity is high, the cost of an attack will also be high because more is destroyed. On the other hand, if there is very little economic activity, then less will be destroyed. But a relatively low stock of economic activity reduces the economy's capacity to grow. This is

[6]We have abstracted from budgetary considerations in the model, but this is by no means less important. Some estimates suggests that putting the nation on a higher costs about $1 billion a week. With several states facing a budget shortfall, resources can therefore be a constraint on efforts to prevent and/or mitigate a terror attack. In our model, varying amounts of resources can be viewed as changing γ, so that the qualitative nature of our results is preserved without explicitly introducing a budget constraint.

consistent with the argument in Azariadis and Drazen (1990) that economic growth is slower below some threshold level of human capital or stock of knowledge.

Given its signal, s, the government chooses the alert to maximize the trade–off between the benefits from increased economic and social activity and the cost of an attack:

$$\max_{\mathcal{A}} \; E[V|s] = E[\mathcal{X}|s] - \frac{\lambda}{2}(\mathcal{X} - \overline{\mathcal{X}}_S - \gamma\mathcal{A})^2, \tag{3}$$

where $E[\mathcal{X}|s]$ is the expected level of economic activity given the signal, and λ is a positive constant. This objective function captures, in a very simple way, the government's desire to provide security to a worried nation, but at the same time encourage everyone to go about their business as usual.[7] This implies that λ can not be too low nor too high. For simplicity, and in order to study the effects of the public's perception of terrorism risks on the alerts, we set $\lambda = 1$.

3 Equilibrium analysis

We are interested in analyzing the perfect Bayesian equilibria to the signaling game. In such an equilibrium, the government has a welfare maximizing behavioral strategy that describes the alert level $\mathcal{A}(\omega, s)$, taking into account the level of economic activity such a security advice will elicit. The public's decision to undertake social or economic activity is a function of their beliefs about the terror–related information received by the government and the alert level, with the aim of maximizing utility.

We first analyze a benchmark case in which all information is publicly observable, and relate the results to the more realistic case of private information.[8] To ensure that the optimal alert levels in the full–information case has the desirable properties, we make the following assumption:

Assumption 1: $\frac{L}{H} < 2$.

Assumption 1 guarantees that in the full–information case, the nation will be put on a higher alert level if a *high* signal of an attack is observed. Secondly, given the signal observed, the full–information alert level is higher if the public's prior about an attack is high.

[7] For example, citing concerns of a terrorist attack, a member of the U.S. House Select Committee on Homeland Security, Rep. Christopher Shays of Connecticut, advised the public to consider not attending new year's eve celebrations in New York's Times Square on Dec. 31, 2003 because "if there was panic, a lot of injuries would take place." He was sharply criticized by the Mayor of the city of New York for instilling fear into the public.

[8] In practice, the government does not disclose this information, since doing so will compromise its intelligence sources.

We also assume that the information gathered through intelligence sources is very *informative*: with almost certainty, the government sees a *high* signal if the threat of an attack is high.

Assumption 2: The private signal is *very* informative; i.e., $\phi \to 1$.

3.1 Equilibrium when signals are public information

Suppose that information about terrorist attacks is known by all; i.e., the government has no private information. Then individuals' information set will be given by $\Omega = \{\omega, s\}$ and the government's payoff, if a signal $s = h, l$ is received, will be

$$\begin{aligned} E[V|s] &= \ln E[S|s] - \mathcal{A}_s - \frac{1}{2}P(H|s)\left[\ln E[S|s] - \overline{\mathcal{X}}_H - (1+\gamma)\mathcal{A}_s\right]^2 \\ &- \frac{1}{2}P(L|s)\left[\ln E[S|s] - \overline{\mathcal{X}}_L - (1+\gamma)\mathcal{A}_s\right]^2 \end{aligned} \tag{4}$$

where $E[S|s] = P(H|s)H + P(L|s)L$, and $s = h, l$.

Differentiating $E[V|s]$ w.r.t. $\mathcal{A}_s$, and solving the first order condition for $\mathcal{A}_s$, gives the first–best alert levels to be

$$\mathcal{A}_h^o(\omega) = \frac{1}{1+\gamma}\left(\ln E[S|h] - P(H|h)\overline{\mathcal{X}}_H - P(L|h)\overline{\mathcal{X}}_L - \frac{1}{1+\gamma}\right) \tag{5}$$

when the signal is high; and

$$\mathcal{A}_l^o(\omega) = \frac{1}{1+\gamma}\left(\ln E[S|l] - P(H|l)\overline{\mathcal{X}}_H - P(L|l)\overline{\mathcal{X}}_L - \frac{1}{1+\gamma}\right) \tag{6}$$

when $s = l$. Subtracting (6) from (5) gives

$$\mathcal{A}_h^o(\omega) - \mathcal{A}_l^o(\omega) = \frac{1}{1+\gamma}\left(P(H|h) - P(H|l) - \ln\frac{E[S|l]}{E[S|h]}\right) \tag{7}$$

since it is assumed that $\overline{\mathcal{X}}_L - \overline{\mathcal{X}}_H = 1$.

By Bayes rule, $P(H|h) - P(H|l) > 0$, so that $\mathcal{A}_h^o$ can either be greater than or less than $\mathcal{A}_l^o$, depending on the value of L relative to H. In particular, a very large L encourages a higher level of economic activity, especially if the publicly available information suggests a lower threat of an attack. This however makes the society susceptible to a high degree of destruction if an attack occurs. To mitigate this, the government issues a higher alert level if it sees a lower threat of an attack. However throughout the analysis, we restrict parameter values in our model so that the nation will be on a higher alert if a *high* threat is publicly observed, for every value of ω. We show in the following lemma that our assumption that $L < 2H$ ensures that $\mathcal{A}_h^o > \mathcal{A}_l^o$, and therefore helps us to capture this realistic feature into the model.

Lemma 1 *If assumptions 1 and 2 hold, then (i) $\mathcal{A}_h^o(\omega) > \mathcal{A}_l^o(\omega)$ for all values of ω; and (ii) $\mathcal{A}_s^o(\omega)$ increases in ω, for $s = h, l$.*

Put differently, Lemma 1(i) says that if information gathered from intelligence sources is *sufficiently reliable*; i.e, $\phi \to 1$, and H is not too low (relative to L), then in the full–information case, the nation will be on a higher alert level if the there is sufficient information to believe that the threat of an attack is high. Part (ii) of Lemma 1 says that, given the threat signal observed by the public, the full–information alert level increases in the society's initial perception of a terrorist attack.

Finally, by substituting $\mathcal{A}_h^o$ and $\mathcal{A}_l^o$ into equation (2), it is easy to see that a low threat level leads to a higher economic activity. Hence, when information about an attack is known to the government only, this may provide an incentive for the government to lower the alert level (below the full–information level). As we shall see, such incentives exist when society members perceive an attack to be imminent. We explore this fully in the next section.

3.2 Equilibrium when information is private

Now suppose, as it is the case, that the government does not fully disclose all of its information and/or the source of its intelligence. The question we ask is, will the government always find it optimal to choose the alert level that corresponds to the full–information case?

Let the government be denoted by type–h if its private signal suggests a severe attack, and by type–l if the signal points to a lower threat of an attack. We focus our attention on the incentives faced by a type–h for a given prior, ω, of a severe attack. We first look for the conditions under which a separating equilibrium exist at $\mathcal{A} \in \{\mathcal{A}_h^o(\omega), \mathcal{A}_l^o(\omega)\}$. Consider a deviation from $\mathcal{A}_h^o(\omega)$ by the type–h government. If the public incorrectly interprets this to mean a *low* level of risk, then the government's net gain from deviation will be given by $\Delta \equiv E[V|h, \mathcal{A}_l^o] - E[V|h, \mathcal{A}_h^o]$, the difference in payoffs between deviating and following the signal. This is given by

$$\begin{aligned}
\Delta(\omega) &= \frac{1}{1+\gamma}\left[P(H|h) - P(H|l) + \gamma \ln\left(\frac{E[S|l]}{E[S|h]}\right)\right] - \frac{1}{1+\gamma}(P(H|h) - P(H|l)) \\
&- \frac{1}{2}[P(H|h) - P(H|l)]^2 \\
&= \frac{\gamma}{1+\gamma}\ln\left(\frac{E[S|l]}{E[S|h]}\right) - \frac{1}{2}[P(H|h) - P(H|l)]^2 \qquad (8)
\end{aligned}$$

The difference in payoffs has two opposing terms: the first term represents the gain from a higher economic activity, and the second term is the cost incurred by the society if an attack occurs. It is easy to see how different values of the model's parameters change the incentive to deviate or not. For example, since $E[S|l] > E[S|h]$ for all $\omega > 1/2$, an increase in L (relative to H) increases the first term on the RHS of equation (8) and thus increases the incentive to deviate. Also, $\frac{\partial \Delta}{\partial \gamma} > 0$, so that the incentive to deviate from the full–information alert level increases in γ: as γ increases, $\mathcal{A}_l^o \to \mathcal{A}_h^o$ (see equation (7)) and the public will not be able to make the distinction between a high threat from a low threat of an attack. In this case, the government sets a lower alert level one. Intuitively, if the government is confident in the steps taken to secure the nation, then the alert level will be lowered to assure the public.[9]

Proposition 1 *Suppose* $2e^{1/4} - 1 < \frac{L}{H} < e^{1/2}$ *and* $\gamma > 1$. *Then there exists some* $\frac{1}{2} < \underline{\omega} < 1$, *such that the government has an incentive to:*
(i) choose the full–information alert level if the public is uncertain about the threat of an attack; i.e. $\Delta < 0$ *if* $\frac{1}{2} < \omega < \underline{\omega}$; *and*
(ii) deviate from the full–information alert level if the public's prior belief of an attack is high; i.e. $\Delta > 0$ *if* $\underline{\omega} < \omega < 1$.

The intuition behind Proposition 1 is very simple. The government responds to both the public's prior perception of attack and information gathered from intelligence sources. Proposition 1(i) says that the government adopts the full–information alert level when, *ex–ante*, there is sufficient uncertainty about the threat of an attack (i.e., $\omega < \underline{\omega}$). The reason is that since the public and the government are very uncertain about the severity of a potential attack, the government puts a significant weight on its intelligence. So if a *high* signal is observed, the government becomes 'confident' that an attack is likely; and vice versa. Hence, it is optimal to adopt an alert level accordingly to the signal received, and this impacts on the public's perception of an attack and their behavior, as in the full–information case. Note that this separating equilibrium does not depend crucially on terrorists behavior. In particular, if terrorist decide to attack during a period of low alert level, then this will be reflected in the government's intelligence and the alert level will be revised accordingly.

[9]Note that there are no incentives to deviate if the government's intelligence sources shows no threat of a severe attack. To see this, define $E[V|l, \mathcal{A}_l^o]$ to be the type–l government payoff from following its signal and setting $\mathcal{A} = \mathcal{A}_l^o$, and let let $E[V|l, \mathcal{A}_h^o]$ be her payoff from deviating to $\mathcal{A}_h$. Then the difference in payoffs between deviating and choosing $\mathcal{A}_h^o$ will be given by

$$E[V|l, \mathcal{A}_h^o] - E[V|l, \mathcal{A}_l^o] = -\frac{\gamma}{1+\gamma} \ln\left(\frac{E[S|l]}{E[S|h]}\right) - \frac{1}{2}[P(H|h) - P(H|l)]^2 < 0.$$

Hence, the type–l government never finds it optimal to set an alert level different from what its information suggests. Doing so hurts the government both in lower economic activity and higher level of uncertainty.

Part (ii) of Proposition 1 says that the government has an incentive not to implement the full–information security advisory if $\omega > \underline{\omega}$ and γ is sufficiently high. To see this, suppose the public perceives an attack to be severe. Then a further heightened state of alert will be economically too costly as it confirms the public's fear of an attack. Hence, if the nation's law enforcement and security agencies preparedness for an attack is sufficiently high (i.e., a high γ), then the government may have an incentive to set the alert level below the full–information case in order to promote more economic activities. So the alerts levels $\{\mathcal{A}_h^o, \mathcal{A}_l^o\}$, together, can not form an equilibrium. The question then is: How does the government convey different threats of an attack to the public?

In turns out that any attempt to resolve this dilemma also involves raising the alert level beyond the full–information level: that is, for some range of ω, the government puts the nation on a higher state of a alert even if the signal received is *low*. Since such an action entails some economic costs, it must be weighed against the benefit of reducing or eliminating complacency in the public. Hence, a higher alert level when the signal is *low*, must satisfy two conditions:

(i) any alert level $\mathcal{A} > \mathcal{A}_l^o$ must not be 'too high,' such that a type-l government finds it too costly to implement;

(ii) at the same time, such alert level $\mathcal{A}$ must be high enough such that the type–h government finds it sub–optimal to mimic this strategy.

We show in Proposition 2 (below) this strategy will be feasible if the public's prior of an alert is not too high. For sufficiently high prior, the economic losses do not justify the use of this strategy. Hence, there is a discontinuity in the warning system: for intermediate values of ω, the government over–reacts to its information and issues a higher alert (relative to the full–information case); whereas for sufficiently skewed priors, there will be pooling equilibrium in which the government keeps the alert level lower than the full–information case, regardless of its information. However, this does not alter the public's perception of an attack and its behavior as everyone chooses an adequate level of private protection.

We first show that such pooling equilibrium exist, and then turn our attention to the case of over-reaction.

3.2.1 Pooling equilibria: under-reaction to private information

Suppose $\omega > \underline{\omega}$. In a pooling equilibrium, the government's alert level provides no information to the public, hence the public continue to hold their prior beliefs about the severity of an attack. To solve for a pooling equilibrium, we first find the optimal alert level that will be adopted if the information from intelligence sources reveal a low level of threat.

The type–l problem will be to choose $\mathcal{A}$ to solve

$$\begin{aligned}\max_{\mathcal{A}} E[V|l,\omega] &= \ln E[S_\omega] - \mathcal{A} - \frac{1}{2}P(H|l)\left[\ln E[S_\omega] - \mathcal{A} - \overline{\mathcal{X}}_H - \gamma\mathcal{A}\right]^2 \\ &- \frac{1}{2}P(L|l)\left[\ln E[S_\omega] - \mathcal{A} - \overline{\mathcal{X}}_L - \gamma\mathcal{A}\right]^2 \end{aligned} \tag{9}$$

where $E[S_\omega] = \omega H + (1-\omega)L$. The optimal alert level is

$$\mathcal{A}^*(\omega) = \frac{1}{1+\gamma}\left(\ln E[S_\omega] - P(H|l)\overline{\mathcal{X}}_H - P(L|l)\overline{\mathcal{X}}_L - \frac{1}{1+\gamma}\right)$$

which is less than $\mathcal{A}_l^o(\omega)$.

Now consider a deviation from $\mathcal{A}^*(\omega)$ by the type–h government. Let the out–of–equilibrium beliefs be given by ρ; i.e., $\text{Prob}(S = H|\mathcal{A} \neq \mathcal{A}^*) = \rho(\omega)$. Then the optimal deviation will be given by

$$\mathcal{A}_\rho = \frac{1}{1+\gamma}\left(\ln E[S_\rho] - P(H|h)\overline{\mathcal{X}}_H - P(L|h)\overline{\mathcal{X}}_L - \frac{1}{1+\gamma}\right), \tag{10}$$

where $E[S_\rho] = \rho H + (1-\rho)L$.

Then as before, the type–h government's net gain from choosing $\mathcal{A}^*(\omega)$ rather than $\mathcal{A}_\rho$ will be given by

$$\Delta_h^* = E[V_\omega|h, \mathcal{A}^*(\omega)] - E[V_\rho|h, \mathcal{A}_\rho] = \frac{\gamma}{1+\gamma}\ln\left(\frac{E[S_w]}{E[S_\rho]}\right) - \frac{1}{2}[P(H|h) - P(H|l)]^2, \tag{11}$$

and will choose the alert level $\mathcal{A}^*(\omega)$ if $\Delta_h^* > 0$.

3.2.2 Separating equilibria: over-reaction to private information

If $\omega > \underline{\omega}$, then any separating equilibrium must consist of the government adopting a much higher alert level when she observes a low threat. But for this to be achievable, the two conditions (stated in the previous section) must be satisfied: (i) it must be that the high $\mathcal{A}$ can effectively separate the *low* from the *high* signal government; and (ii) it must be optimal for the type–l government to implement the extra security measures.

Suppose there is some alert level $\mathcal{A} > \mathcal{A}_l^o$ that satisfies both conditions. That is, the government chooses this $\mathcal{A}$ if it sees a low signal. Now consider a deviation by a type–h from $\mathcal{A}_h^o$ to this level of $\mathcal{A}$, and let the net gain from doing so be given by $\hat{\Delta}_h = E[V|h, \mathcal{A}] - E[V|h, \mathcal{A}_h^o]$. Such deviation will not be optimal if $\hat{\Delta}_h < 0$.

Similarly, the second condition requires that the type–l government finds it optimal to implement $\mathcal{A} > \mathcal{A}_l^o$ rather than pool with *high–signal* type at $\mathcal{A}^*(\omega)$: that is, $\Delta_l^* = E[V|l,\mathcal{A}] - E[V|l,\mathcal{A}^*] \geq 0$. Our second proposition shows that, if ω is not too large, then there is some $\mathcal{A}$ that satisfies both conditions, so that a 'hybrid' separating equilibria exists. For very extreme priors, there will only be a pooling equilibrium with the alert level lower than the full–information case.

Proposition 2 *Suppose $\omega > \underline{\omega}$, then there exists some $\overline{\omega} > \underline{\omega}$, such that*
(i) if $\omega \in (\underline{\omega}, \overline{\omega})$, then a separating equilibrium exists in which the government chooses an alert level, $\hat{\mathcal{A}}(\omega)$ ($\mathcal{A}_l^o < \hat{\mathcal{A}}(\omega) < \mathcal{A}_h^o$), if a low signal is observed; and $\mathcal{A} = \mathcal{A}_h^o(\omega)$ if the threat level is high.
(ii) if $\omega > \overline{\omega}$, the model features a pooling equilibrium in which the government chooses an alert level $\mathcal{A} = \mathcal{A}^*(\omega)$ for all signals.

Proposition 2 establishes a very interesting result about the government's terror advisory system and how it impacts on the public's perception of an attack. We have shown that when $\omega > \underline{\omega}$, the full–information terror advisory policy is no longer an equilibrium. Hence, in order to make it more credible, it is necessary that the government raises the alert level when it sees a low threat of attack, and $\omega \in (\underline{\omega}, \overline{\omega})$. That is, there is a 'hybrid' separating equilibrium in which the alert level is set at the full–information alert level when the signal is *high*, but raised above $\mathcal{A}_l^o$ when the signal is *low*. For these ranges of ω values, the extra vigilance and lower potential cost if an attack occurs (due to the extra security and heightened alert level) more than offsets the economic costs.

However, for extreme priors, issuing a higher alert level is costly. When everyone perceives the threat of an attack to be high, they take adequate private protection against the perceived risks. Given this, the government finds it optimal to issue a lower alert (relative to the full–information level) to calm the panic–stricken public: here, the concern is that individuals will over–protect themselves and not participate in social or economic activities. Hence, we have pooling equilibrium in which the government issues a low alert level, $\mathcal{A}^*(\omega)$, regardless of its information. Nevertheless, this does not change the public's perception of an attack, and the level of private protection.

We illustrate this in Figure 1 below. The bold *solid* curve represents $\frac{1}{2}[P(H|h) - P(H|l)]^2$, the second term of the RHS of equation (8), and the bold *dash* curve is the first term. For $\omega < \underline{\omega}$, the second term dominates the first, and hence it is optimal to adopt the full–information. The *dash-dot* curve represents $\frac{\gamma}{1+\gamma} \ln\left(\frac{E[S_\omega]}{E[S|h]}\right)$, and it lies below the curve $\frac{1}{2}[P(H|h) - P(H|l)]^2$ when

$\omega \in (\underline{\omega}, \overline{\omega})$. For this range of ω values, there are separating equilibria that consist of a higher alert level when a *low* signal is observed. Finally, for sufficiently skewed priors, we have a pooling equilibrium at $\mathcal{A}^*(\omega)$.

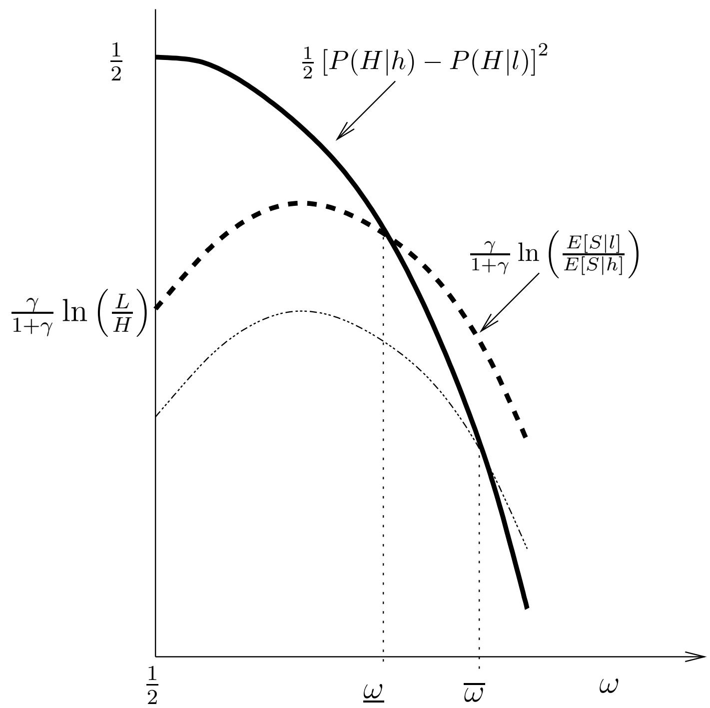

Figure 1: Illustration of Propositions 1 and 2

This diagram illustrates the impact of asymmetric information on the terror alert system. The full–information alert level is adopted when $\omega < \underline{\omega}$. When $\omega > \overline{\omega}$, the government chooses an alert level $\mathcal{A}^*(\omega)$, regardless of its information. For intermediate priors, the alert level is more often higher, relative to the full–information case.

The pooling equilibrium alert level has some interesting implications. First, it is worth noting that the pooling equilibrium alert level also depends on ω, and it is less than $\mathcal{A}_l^o(\omega)$. It is also easy to see that $\mathcal{A}^*(\omega)$ decreases in ω: differentiating $\mathcal{A}^*(\omega)$ w.r.t. ω gives

$$\frac{\partial \mathcal{A}^*(\omega)}{\partial \omega} = \frac{1}{1+\gamma}\left[\frac{(H-L)}{\omega H + (1-\omega)L} + \frac{\phi(1-\phi)}{(\omega(1-\phi) + (1-\omega)\phi)^2}\right], \tag{12}$$

which is negative for $\phi \to 1$. So the alert level is lowered further (relative to the full–information level) as the public becomes increasingly concerned about a potential terrorist attack. But such assurances from the government do not change individuals' resolve to protect themselves: aggregate economic activity is lower when the society perceive an attack as extremely likely. That is, for $\omega > \overline{\omega}$, $\frac{\partial \mathcal{X}^*(\omega)}{\partial \omega} < 0$. Recall from Lemma 1 that, in the full–information case, the nation is always put on a higher alert level whenever the public perceives a severe attack to be more likely. The result in (12) therefore suggests the opposite: Intuitively, the government finds it comforting when people choose to be more vigilant and take adequate steps to protect themselves. In such cases, the economic costs from a further heightened alert level more than offsets any gains.

Finally, differentiating $\mathcal{A}^*(\omega)$ w.r.t. ϕ gives

$$\frac{\partial \mathcal{A}^*(\omega)}{\partial \phi} = -\frac{1}{1+\gamma}\left[\frac{\omega(1-\omega)(\overline{\mathcal{X}}_L - \overline{\mathcal{X}}_H)}{(\omega(1-\phi)+(1-\omega)\phi)^2}\right] < 0.$$

Recall that $\omega < \phi$. Hence, for a fixed and sufficiently high ω, the alert level is higher if the intelligent information is made a little less informative. This is an interesting result: it shows that if the public perceives an attack to be imminent, then the government would raise the alert level, when it receives a more general intelligence information about a potential terrorist attack.

4 Discussion of main results

Designing an effective terror alert system has several challenges. On one hand, there are economic agents that do not want their lives to be changed completely due to risk of an attack, as this would be seen as "giving in" to terrorists. On the other hand, there is a government that would want the public to heed to a terror warning, so that they can be extra vigilant and know how to protect themselves against an attack. Compounding this is the enormous task of sieving through intelligence and knowing when to act on different pieces of information.[10] What is clear though is that the government wants Americans to be watchful but not panicky.

But the five–color alert system has been the object of many jokes and drawn several criticisms from lawmakers and the public. Our paper shows that the information asymmetry that lies between the public and the government, in and of itself, presents some of the challenges faced in implementation of the terror alert system. Given that information received by the government cannot be shared with the public (and for good reasons), the public is more likely to attach less importance to the alerts. This means that the government may sometimes raise the alert level beyond what it would be were intelligence information made public. This was shown in our result in Proposition 2(i). But such a strategy can become a source of complacency: even though static, our model can be used to explain why if adopted too often in a

[10] Some DHS worry that operatives of the terrorist networks were deliberately trying to spoof U.S. and international intelligence networks aimed at uncovering terrorist threats by planting misinformation on lines of communications they believed were monitored.

dynamic context, there is a risk that the public will become blasé about repeated heightened alerts (especially, if an attack does not occur during a heightened state of alert).[11]

Secondly, our results that the alert level may be set at a lower level (relative to the full information case) even when the public's perceives an attack to be very likely can be related to the fact that the alert level has not been raised above orange in its short history. The main reason is that, when the public thinks an attack is imminent, they will take adequate protection against such an attack, and this behavior is less likely to be influenced by the state of alert. They are also more likely to trade–off a "loss of freedom" for extra security (e.g., being receptive to extra security at airports and other public places). Hence, the alert level will be set below the full–information level for this range of prior beliefs of an attack. Intuitively, the government may have an incentive *not* to further confound the public's anxiety by explicitly raising the alert level. We also show that given a high prior of an attack, the alert level is set high when the government receives a more general intelligence (i.e., $\phi \rightarrow \omega$). One could use risk aversion to explain this behavior. But even without risk aversion, our model shows that the government might choose to raise the alert level in this case for strategic reasons.

These results, taken together, seem consistent with the fact that the alert level has never been at *blue* or *green* since the system was instituted. So far, the lowest alert level has been *yellow*. We therefore agree with Mr. Tom Ridge, the Secretary of the DHS, for suggesting that it could be decades before the government would be able to lower the threat alert to *green*.

5 Conclusion

The events on 9/11 has significantly changed the way many governments now view terrorism, and has triggered new strategies to deal with terrorism. As a way of engaging the public in its strategies, the government devised a security advisory system to inform the public of when to be extra vigilant. It can probably be argued that the terror alert system will continue to be an integral part of the U.S. government's efforts to crack down on terrorism. But these innovations have also raised concerns about the line between security and civil liberties. With this background, this paper attempts to provide an economic analysis of how the government, through a terror advisory system, engages the public in the fight against terrorism.

We show that a full–information alert policy, which prescribes a higher alert level *both* when the intelligence sources indicate a higher threat and when the public is fearful of such attacks,

[11]Indeed, strong criticisms by counter–terrorism experts, and much of the public, led the DHS to revamp the system in September 2003. Tougher internal guidelines must now be met for raising the alert level.

may not be implemented all the time. When the public perceives an attack to be very likely, then a further heightened alert level can cause an enormous panic in the society and substantial reduction in participation in both social and economic activities. That is, individuals over–protect themselves from any possible attack. So a relatively lower alert level will be adopted to encourage more economic participation, especially if extra vigilance is also seen as “giving in” to terrorists. Lakdawalla and Zanjani (2002) also argued that the the idea of not “giving in” to terrorists can be used as a justification for government subsidies in terror insurance. They argue that without such subsidies, individuals will over–invest in self–protective behavior and this may be bad for “national prestige.”[12]

Finally, we also find that in order to make the advisory system more credible, the government may also decide to raise the alert level (above the full–information level) to avoid the public being complacent of an attack. Our results can therefore be used to shed some light on alerting the public of potential attacks, and how such alerts changes, or not, the public’s perception of terrorism and participation in both economic and social activities.

[12]See also Brown, Kroszner and Jenn (2002) for a more thorough analysis of the case for federal intervention in the market for insurance against terrorism.

Appendix

Proof of Lemma 1:(i) Using Bayes' rule, $P(H|h) = \frac{\omega\phi}{\omega\phi+(1-\omega)(1-\phi)}$ and $P(H|l) = \frac{\omega(1-\phi)}{\omega(1-\phi)+(1-\omega)\phi}$. Hence in the limit as $\omega \to 1$,

$$\lim_{\phi\to 1;\omega\to\phi}[P(H|h) - P(H|l)] = 1/2 \quad \text{and} \quad \lim_{\phi\to 1;\omega\to\phi} \ln \frac{E[S|l]}{E[S|h]} = \ln \frac{L+H}{2H}.$$

Secondly, $\frac{\partial}{\partial w}(P(H|h) - P(H|l)) = \frac{\phi(1-\phi)(2\phi-1)(1-2\omega)}{(P(h)P(l))^2} < 0$, where $P(s)$ is the unconditional probability of observing a signal s $(s = h, l)$. Hence $P(H|h) - P(H|l)$ falls in ω and maximized at $\omega = \frac{1}{2}$ with $[P(H|h) - P(H|l)]|_{\omega=1/2} = 1$. Again, differentiating $\ln \frac{E[S|l]}{E[S|h]}$ w.r.t. ω, we get

$$\begin{aligned} \frac{\partial}{\partial w} \ln\left(\frac{E[S|l]}{E[S|h]}\right) &= \frac{(L-H)}{E[S|h]}\left[\frac{\partial}{\partial w}(P(H|h) - P(H|l)) - \frac{(E[S|l] - E[S|h])}{E[S|l]}\frac{\partial P(L|l)}{\partial w}\right] \\ &> \frac{L-H}{E[S|h]}\left[\frac{\partial}{\partial w}(P(H|h) - P(H|l))\right] \end{aligned}$$

since $\frac{\partial P(L|l)}{\partial w} < 0$ and $E[S|l] - E[S|h] > 0$ for all ω. So, $\ln \frac{E[S|l]}{E[S|h]}$ attains a maximum at some $\omega > 1/2$ (when $P(H|h) - P(H|l)$ is decreasing). Combining this with the limiting points, it is easy to see that if $\ln \frac{L+H}{2H} < \frac{1}{2}$, then $P(H|h) - P(H|l) > \ln\left(\frac{E[S|h]}{E[S|l]}\right)$, for all values of ω, and $\mathcal{A}_h^o(\omega) > \mathcal{A}_l^o(\omega)$. But this holds if $\frac{L}{H} < 2 < 2e^{1/2} - 1$.

(ii) Differentiating $\mathcal{A}_s^o(\omega)$ with respect to ω gives

$$\begin{aligned} \frac{\partial \mathcal{A}_s^o}{\partial w} &= \left(\frac{1}{1+\gamma}\right)\frac{\phi(1-\phi)}{(P(s))^2 E[S|s]}[H - (L-H)P(H|s)] \\ &> \left(\frac{1}{1+\gamma}\right)\frac{\phi(1-\phi)}{(P(s))^2 E[S|s]}[2H - L] \end{aligned}$$

By assumption 1, $\frac{\partial \mathcal{A}_s^o}{\partial w} > 0$. Hence, the full–information alert levels will be increasing in ω. QED

Proof of Proposition 1: (i) Since $\lim_{\phi\to 1;\omega\to\phi} \ln \frac{E[S|l]}{E[S|h]} = \ln \frac{L+H}{2H}$ and $\lim_{\phi\to 1;\omega\to\phi}[P(H|h) - P(H|l)]^2 = \frac{1}{4}$, the limit as $\omega \to 1$ of equation (8) will be given by

$$\lim_{\phi\to 1;\omega\to 1-\phi} \Delta(\omega) = \frac{\gamma}{1+\gamma} \ln \frac{L+H}{2H} - \frac{1}{2}\left(\frac{1}{4}\right).$$

Therefore, if $\ln \frac{L+H}{2H} > \frac{1}{4}$, which implies that $\frac{L}{H} > 2e^{1/4} - 1$, then $\lim_{\phi\to 1;\omega\to\phi} \Delta(\omega) \geq \frac{1}{4}\left(\frac{\gamma}{1+\gamma} - \frac{1}{2}\right)$ which is positive if $\gamma > 1$. Secondly, for ω close to $\frac{1}{2}$, we have $\lim_{\phi\to 1;\omega\to\frac{1}{2}} \ln \frac{E[S|l]}{E[S|h]} = \ln \frac{L}{H}$ and $\lim_{\phi\to 1;\omega\to\frac{1}{2}}[P(H|h) - P(H|l)]^2 = 1$. Hence, if $\ln \frac{L}{H} < \frac{1}{2}$, which implies that $\frac{L}{H} < e^{1/2}$, then $\lim_{\phi\to 1;\omega\to\frac{1}{2}} \Delta(\omega) < 0$ for all γ. This implies that there must be some $\omega = \underline{\omega}$ such that $\Delta < 0$ if $\omega < \underline{\omega}$.

(ii) The proof of part (i) shows that for all other values of ω, $\Delta < 0$ if $2e^{1/4} - 1 < \frac{L}{H} < e^{1/2}$ and $\gamma > 1$. QED

Proof of Proposition 2:(i) Let $\hat{\Delta}_h = E[V|h,\mathcal{A}] - E[V|h,\mathcal{A}_h^o]$ be a type-h government's gain from deviating from the high alert level by announcing the alert level $\mathcal{A}$, where $\mathcal{A}_l^o < \mathcal{A} < \mathcal{A}_h^o$. Then a deviation will not be optimal if $\hat{\Delta}_h < 0$. This condition becomes

$$\begin{aligned}\ln\left(\frac{E[S|l]}{E[S|h]}\right) - (\mathcal{A} - \mathcal{A}_h^o) \quad \leq \quad & \frac{1}{2}P(H|h)\left[\ln E[S|l] - \mathcal{A} - \overline{\mathcal{X}}_H - \gamma\mathcal{A}\right]^2 + \\ & \frac{1}{2}P(L|h)\left[\ln E[S|l] - \mathcal{A} - \overline{\mathcal{X}}_L - \gamma\mathcal{A}\right]^2 - \frac{1}{2}\left(\frac{1}{1+\gamma}\right)^2 \\ - \quad & \frac{1}{2}P(H|h)P(L|h). \end{aligned} \tag{13}$$

However, for any the alert level $\mathcal{A}_l^o < \mathcal{A} < \mathcal{A}_h^o$ to be feasible, it must be that the type–l government finds it optimal to implement it rather than pool with *high–signal* type at $\mathcal{A}^*(\omega)$: that is, $\Delta_l^* = E[V|l,\mathcal{A}] - E[V|l,\mathcal{A}^*] \geq 0$, or

$$\begin{aligned}\ln\left(\frac{E[S|l]}{E[S_\omega]}\right) - (\mathcal{A} - \mathcal{A}^*) \quad \geq \quad & \frac{1}{2}P(H|l)\left[\ln E[S|l] - \mathcal{A} - \overline{\mathcal{X}}_H - \gamma\mathcal{A}\right]^2 + \\ & \frac{1}{2}P(L|l)\left[\ln E[S|l] - \mathcal{A} - \overline{\mathcal{X}}_L - \gamma\mathcal{A}\right]^2 - \frac{1}{2}\left(\frac{1}{1+\gamma}\right)^2 \\ - \quad & \frac{1}{2}P(H|l)P(L|l). \end{aligned} \tag{14}$$

Define $\sigma_H = \ln E[S|l] - \mathcal{A} - \overline{\mathcal{X}}_H - \gamma\mathcal{A}$ and $\sigma_L = \ln E[S|l] - \mathcal{A} - \overline{\mathcal{X}}_L - \gamma\mathcal{A}$. Then subtracting the inequality (14) from (13) gives

$$\begin{aligned}\ln\left(\frac{E[S_w]}{E[S|h]}\right) - (\mathcal{A}^* - \mathcal{A}_h) \quad \leq \quad & \frac{1}{2}\left[P(H|h) - P(H|l)\right]\left(\sigma_H^2 - \sigma_L^2\right) - \\ & \frac{1}{2}\left[P(H|h)P(L|h) - P(H|l)P(L|l)\right] \\ = \quad & \frac{1}{2}(P(H|h) - P(H|l))\left[2\ln E[S|l] - (\overline{\mathcal{X}}_H + \overline{\mathcal{X}}_L) - \right. \\ & \left. 2(1+\gamma)\mathcal{A}\right] - \frac{1}{2}\left[P(H|h)P(L|h) - P(H|l)P(L|l)\right] \end{aligned}$$

This implies that

$$\begin{aligned}\mathcal{A} \leq \overline{\mathcal{A}} \quad \equiv \quad & \frac{1}{2(1+\gamma)}\left[2\ln E[S|l] - (\overline{\mathcal{X}}_H + \overline{\mathcal{X}}_L) - \frac{P(H|h)P(L|h) - P(H|l)P(L|l)}{P(H|h) - P(H|l)}\right. \\ - \quad & \left.\frac{2}{1+\gamma}\left(\frac{\gamma\ln\left(\frac{E[S_\omega]}{E[S|h]}\right)}{P(H|h) - P(H|l)} + 1\right)\right] \end{aligned}$$

since $\ln\left(\frac{E[S_w]}{E[S|h]}\right) - (\mathcal{A}^* - \mathcal{A}_h) = \frac{1}{1+\gamma}\left(\gamma \ln \frac{E[S_w]}{E[S|h]} + (P(H|h) - P(H|l)\right)$. But for $\mathcal{A}$ to be part of a separating equilibrium, it must be that $\mathcal{A}_l^o < \overline{\mathcal{A}}$. That is,

$$\begin{aligned}
\frac{2\gamma}{1+\gamma}\left(\frac{\ln\left(\frac{E[S_\omega]}{E[S|h]}\right)}{P(H|h)-P(H|l)}\right) &< \left[P(H|h) - P(L|h) - \frac{P(H|l)P(L|l) - P(H|h)P(L|h)}{P(H|h)-P(H|l)}\right] \\
&= P(H|h) - P(H|l)
\end{aligned}$$

Hence, there exists an $\hat{\mathcal{A}} > \mathcal{A}_l^o$ that satisfies both constraints iff

$$\frac{\gamma}{1+\gamma}\ln\left(\frac{E[S_\omega]}{E[S|h]}\right) < \frac{1}{2}\left[P(H|h) - P(H|l)\right]^2. \tag{15}$$

We know from Proposition 1 that $\frac{\gamma}{1+\gamma}\ln\left(\frac{E[S|l]}{E[S|h]}\right) > \frac{1}{2}\left[P(H|h) - P(H|l)\right]^2$ if $\omega > \underline{\omega}$. But since $E[S_\omega] < E[S|l]$, $\forall\omega$, it implies that there exists some $\overline{\omega}$, such that the inequality in (15) holds when $\underline{\omega} < \omega < \overline{\omega}$. That is, there exists a *hybrid* separating equilibrium, involving a higher alert lever $\hat{\mathcal{A}}(\omega) > \mathcal{A}_l^o$ if the signal is *low*, and $\underline{\omega} < \omega < \overline{\omega}$.

(ii) Let $\overline{\omega} < \omega < \frac{1}{2}$, and suppose that the public's beliefs about the risk of an attack is given by $\text{Prob}(H|\mathcal{A} \neq \mathcal{A}^*) = \rho(\omega)$. Then optimal deviation ,$\mathcal{A}_\rho$, from $\mathcal{A}^*$ is as given in equation (10). We first show that $\mathcal{A}_\rho \not< \mathcal{A}^*$. Notice that

$$\mathcal{A}_\rho - \mathcal{A}^*(\omega) = \frac{1}{1+\gamma}\left(P(H|h) - P(H|l) - \ln\frac{E[S_w]}{E[S_\rho]}\right)$$

For sufficiently skewed priors (in particular $\omega \to 1$), we have $P(H|h) - P(H|l) \to 1/2$. Hence $\mathcal{A}_\rho - \mathcal{A}^*(\omega) \approx \frac{1}{1+\gamma}\left(\frac{1}{2} - \ln\frac{E[S_w]}{E[S_\rho]}\right)$. Suppose $\rho = 0$ — the most favorable case for any deviation below $\mathcal{A}^*$. Then for $\omega \to 1$, we get $\mathcal{A}_\rho < \mathcal{A}^*(\omega)$ iff $\frac{1}{2} < \ln(H/L)$ or $\frac{L}{H} < e^{-1/2}$. But this violates the condition that $\frac{L}{H} \geq 2e^{1/4} - 1$, stated in Proposition 1. Hence $\mathcal{A}_\rho \not< \mathcal{A}^*$.

Next, we know (from equation (11)) that a pooling equilibrium exists if

$$\Delta_h^* = E[V_\omega|h, \mathcal{A}^*] - E[V_\rho|h, \mathcal{A}_\rho] = \frac{\gamma}{1+\gamma}\ln\left(\frac{E[S_w]}{E[S_\rho]}\right) - \frac{1}{2}[P(H|h) - P(H|l)]^2 > 0.$$

We also know that $E[S_\omega] > E[S|h]$, $\forall\omega$. So for $\omega > \overline{\omega}$, we can choose $\rho \geq P(H|h)$ such that the above inequality holds. Hence, a pooling equilibrium at $\mathcal{A}^*(\omega)$ when $\omega > \overline{\omega}$ can be supported with out–of–equilibrium beliefs given by $\text{Prob}(H|\mathcal{A} \neq \mathcal{A}^*) = \rho(\omega) \geq P(H|h)$. That is, for these values of ω, the public perceives the likelihood of a severe attack to be at least $P(H|h)$ when an alert level $\mathcal{A} > \mathcal{A}^*$ is announced. QED.

References

[1] Azariadis, C., and Drazen, A. (1990), "Threshold Externalities and Economic Development", *Quarterly Journal of Economics*, 105, 501-26.

[2] Brown J.R, Randall S. Kroszner and Brian Jenn (2002), "Federal Terrorism Risk Insurance," *National Tax Journal*, 55 (3), 647–57

[3] Enders, Walter and Todd Sandler (1993), "The Effectiveness of Anti–Terrorism Policies: A Vector–Autoregression Intervention Analysis," *American Political Science Review,* 87 (4), 829–844.

[4] Keohane, Nathaniel O. and Richard J. Zeckhauser (2003), " The Ecology of Terror Defense," *Journal of Risk and Uncertainty,* 26 (2-3), 201-229

[5] Lakdawalla Darius and George Zanjani (2002), "Insurance, Self– Protection and the Economics of Terrorism," Working Paper, RAND.

[6] Lapan, Harvey and Todd Sandler (1988), "To bargain or not to bargain: That is the question," *American Economic Review,* 78, 16–20.

[7] — (1993), "Terrorism and signalling," *European Journal of Political Economy,* 9, 383–97.

[8] Lee, Dwight R. (1988), "Free Riding and paid riding in the fight against terrorism," *American Economic Review,* 78, 22–26

[9] Viscusi, W. Kip and Richard J. Zeckhauser (2003), "Sacrificing Civil Liberties to Reduce Terrorism Risks" *Journal of Risk and Uncertainty,* 26 (2-3), 99–120