

QUANTIFYING INSURANCE TERRORISM RISK

*By Dr. Gordon Woo
Risk Management Solutions
Gordon.Woo@rms.com*

*Prepared for the National Bureau of Economic Research meeting,
Cambridge, Massachusetts, 1 February 2002*

ABSTRACT

The World Trade Center disaster was a stark reminder to the insurance industry of the potential dire consequences of accumulating high concentrations of insured value and underestimating a hazard to which they are exposed. By imposing strict coverage limits, or stopping to offer terrorism cover for large commercial policies, initial steps can be taken to address the accumulation problem. Subsequently, exploration of the impact of some hypothetical future terrorist scenarios can guide the control of risk accumulations.

But estimating Probable Maximum Loss, as well as the pricing of terrorism risk, require the hazard issue also to be addressed. This will never be resolved as well as hurricane or earthquake hazard, but some insight into its ranking as a peril is urgently needed. With the insurance industry struggling with the uncertainty of how to deal with terrorism risks, hopes will be placed on a reduction of the terrorist threat, now that there is a concerted global will to combat terrorism. Terrorism hazard is considered here in the wake of this international governmental resolution, and the destruction of the al-Qaeda training camps in Afghanistan. The frequency and severity of attacks depend crucially on organizational structure. To minimize detection by counter-terrorist forces, advantage may be taken by terrorists of alternative forms of network architecture, adopted by drugs syndicates, pirates, and other criminals, within which sporadic pulsing swarm attacks might be effectively launched.

The constraints of this network architecture, with sustained pressure from counter-terrorist forces, will influence the relative likelihood of different scenarios favored by al-Qaeda. A terrorism cost function, involving planning time, technical difficulty, and consumption of resources, may be defined to quantify relative scenario likelihood, and thereby allow a loss severity curve to be derived. As part of the task of normalizing this curve to strike frequency, a logical event-tree is outlined for computing the probability that a planned attack is successful. Any probabilistic framework for quantifying terrorism risk, however logically designed, will ultimately have to involve a measure of expert judgement. Extensive expert consultation exercises have already been commissioned by the Pentagon, and should prove as insightful to the insurance industry as for government agencies.

Introduction

On September 11th 2001, the world-wide insurance community suffered its worst ever loss. The shock of such an enormous loss was compounded by the realization that this was a loss stemming from a risk which was accepted, but not quantified. Underwriters were as surprised as civic authorities and security agents. Rough calculations on risk exposure may have been made on the basis of past claims. But, as with all low frequency, high severity perils, whether natural or man-made, a rudimentary actuarial approach to risk quantification, based upon historical loss experience data, is inadequate for a variety of reasons: in particular, the questionable relevance of much of the historical record, and the disparity between its short length and the return period of extreme events.

Every catastrophic loss teaches a hard lesson in risk management: to underwrite a catastrophe peril ignorant of the scale of the risk is to invite further financial trouble. Where catastrophe risks are underwritten, a diligent attempt should be made to quantify them, challenging though this task may be. The major property insurance losses of recent times, such as Hurricane Andrew and the Northridge earthquake, have not merely been outlying statistics on a loss experience plot, they have propelled forward methodological advances in risk modeling. Terrorism cover will continue to be provided after 2001, and from Ground Zero will have to rise up quantitative methods for modeling terrorism risk.

This paper addresses the challenge of quantifying terrorism risk. The classic definition of risk is that it is a product of hazard and vulnerability. The second factor deals with the loss inflicted if a specific terrorist scenario were to occur. Such a scenario might involve the crash of a plane into an urban area, a city bomb blast, a harbor LNG ship explosion, detonation of a nuclear device, etc.. Modeling a specific scenario is essentially a complex engineering problem, not dissimilar, in principle, to the scenario analysis conducted for natural perils such as windstorms and earthquakes. Given the dynamics of the energy source and the geometry of energy dissipation, the vulnerability of engineering construction of different types may be evaluated. Modeling studies of various hypothetical terrorist scenarios are in progress.

For natural perils, hazard modeling may be technically and computationally demanding, but modelers can take comfort from Einstein's dictum that 'Nature may be subtle, but is not malicious'. Terrorists, on the other hand, are both subtle and malicious. So how can a hazard model for terrorism be developed? Obviously, a different approach is required to the traditional reductionist bottom-up approach used for modeling the inanimate world of engineering physics.

A RAND suggestion (Ronfeldt et al., 2001a) is to focus on the network behavior of a terrorist organization, and its capacity to wage a netwar. A theoretical framework for this modern mode of conflict does exist, based on the principles of complexity (see e.g. Waldrop, 1992), which shows how key features of the organizational structure of complex interacting systems emerge. This theory has been successful over the past decade in pioneering quantitative understanding of many aspects of the social behavior of biological organisms. Doubtless as oblivious of the finer mathematical points of

complexity theory are the seasoned netwar practitioners among the criminal fraternity; drugs, immigration, and smuggling racketeers. It is a basic tenet of complexity theory that network characteristics are not consciously moulded by its components, but rather emerge spontaneously from their individual actions.

In applying the broad ideas of complexity theory to the sociological context, account must be taken of human factors such as intelligence and social interaction. As sociologists have remarked (Kennedy et al., 2001), through learning from experience and emulating the successful behavior of others, people are able to discover relatively optimal patterns of attitudes, beliefs and behaviors. For social groups, in which individuals learn from, emulate, and adapt to other group members, there is thus a collective intelligence, which is geographically distributed. This collective dispersed intelligence is a prime facet of terrorist organizations; a hydra-like feature which makes them inherently more threatening, powerful and evasive than lone terrorists.

The concept of swarm intelligence has been developed to describe populations which exhibit certain basic collective patterns of behavior, arising not so much from leadership direction, but rather emerging from the actions of individual members. The social insect metaphor has been a powerful tool in exploring some crucial characteristics of social organizations, including flexibility, robustness, distributed functioning and autonomy.

Although originally developed in the context of cellular robotic systems, the foremost paradigm for swarm intelligence is that of the ant colony, which has the remarkable capability of collectively arriving at solutions to almost intractable mathematical problems (Bonabeau et al., 1999). Edward Wilson's deep fascination with the way that the mass organized behavior of ant populations may be understood from single ant behavior drew him towards a theory of sociobiology (Wilson, 1975). Of course, ants are not human beings. Ants are genetically programmed to dedicate themselves to their pre-assigned tasks; to kill and be killed rather like automatons; to be as prepared to die for the colony as to live.

If the ideas of swarm intelligence are applicable to any group of human beings, it would be to zealous and fanatical terrorists, bound together as one by the Islamic bond of brotherhood; as absolute as that shared by blood relatives. Such a terrorist group could not be adequately represented simply as a set of single-minded individuals, espousing a common cause. Much of their special strength resides in the unusual qualities of their collective swarm intelligence, which govern the group's fearsome capability of executing extreme acts of terror and escaping prior detection. Accordingly, in order to understand the nature of terrorism hazard: its spatial and temporal characteristics, the frequency and severity of attacks; it is necessary to comprehend the structure of terrorist organizations.

Structure of Terrorist Organizations

An immediate observation made in the aftermath of September 11 was the meticulous planning and precise execution of the surprise assault on the United States. The inference was that this well-coordinated assault had to have been masterminded by a very highly organized terrorist network. However well resourced and armed, terrorist groups can never match the economic and technological capability of nation states. As in all conflicts involving an imbalance of military assets, the lesser party can only hope to achieve its objectives through efficiency and smartness of organization and deftness of manoeuver. Despite being vastly inferior in overall numbers and weaponry, at the moment of attack, terrorist forces may coalesce to form an over-powering danger.

The effectiveness of the attacks which a terrorist group might be capable of launching depends much on the structure of its organization. The less centralized and hierarchical, the more resilient the organization will be to counter-terrorist action (Ronfeldt et al., 2001). Hamas, for example, is much less centralized than the Palestine Liberation Organization (PLO), so the detention or death of its members causes little disruption to its capability of launching numerous attacks, most of which are comparatively modest in scale. Although a hierarchical army-style organization is more vulnerable to counter-terrorist action, for as long as its command and control center is functional, it may have the potential to launch highly destructive raids of military proportions. Accordingly, the frequency and severity of attacks by terrorists depend on their organizational structure.

The names by which terrorist groups are known reflect the manner in which they are collectively organized. Some may be self-styled as liberation or freedom-fighting organizations, armies, brigades or fronts, but no appellation is as frustrating to national security services as that of the network, most notably as of the late 1990's, the al-Qaeda network. The French security service ironically called the English capital city, "Londonistan", because of the congregation of Islamic militants who claimed refuge and human rights across the English Channel. Spanning several continents, an international network prospers from national differences in the tolerance of foreign terrorists, and in the liberality of laws of asylum and extradition.

Dispersed over a multitude of host countries, al-Qaeda is in fact a hybrid of hierarchical and network forms of organization; a terrorism conglomerate with both vertical and horizontal command elements (Ronfeldt et al., 2001). Notorious for the leadership of Osama bin Laden, within al-Qaeda there are semi-autonomous groups, which have horizontal coordination. If al-Qaeda had a standard hierarchical army structure, then the capitulation or removal of its leadership might signal its demise as a terrorist force. If this were the case, then the hazard stemming from al-Qaeda would be greatly reduced. This may be wishful thinking. There are a variety of alternative network architectures that al-Qaeda, or one of the other dozen major terrorist organizations, might adopt. Each architecture poses a different challenge to the security services, and to life and property. To avoid the targeting of leaders, a network design may encourage the spatial diffusion of leadership, minimizing the dependence on central control.

Multi-hub Networks

One possible architecture for a terrorist network involves multiple independent hubs, each serving as a control center for a number of satellite cells. To maximize the chance of surviving concerted counter-terrorist action, these hubs may be dispersed over different countries, if not continents. This multi-hub structure is illustrated in Fig.1. The cells attached to a given hub would, for information security reasons, be isolated from one another, with instructions restricted to a 'need-to-know' basis. But the cells might be linked up for major operations. Traditional terrorist organizations, such as the Irish republican army IRA and the Basque separatist group ETA developed complicated cell structures to combat infiltration and monitoring by the security services.




Fig.1 Illustration of a simple two hub network, with the hubs being independent

This hub architecture is partly hierarchical, in that financial, logistical and training support for the peripheral cells are centrally sourced, and strategic planning would be directed from the hubs. However much of the tactical strike planning, cell recruitment and management would be handled locally. This kind of network architecture would, through the ambition of the hub leadership and continuity of the senior high command, enable centralized planning to be conducted on a long-term basis, and facilitate the coordination of complex operations. Ultimately, this would tend to cut the failure rate of attempts at major spectacular attacks.

However, rather like giant battleships in a naval war, hubs are high-profile targets, the elimination of which can spell rapid military defeat. Of course, the more hubs there are, the more redundancy there is, and the harder the network is to defeat. In the global campaign against al-Qaeda, the objective of smashing the organizational hubs has already met with a degree of success. But the al-Qaeda network architecture prior to September 11 may well adapt to survive the subsequent declaration of war on terrorism. The international safe havens where al-Qaeda could operate its hubs without hindrance may disappear, forcing al-Qaeda to assume another organizational guise.

Swarm Intelligence Networks

A more elusive and resilient type of network architecture has no hub, but consists simply of a set of terrorist cells, which may comprise one or more individuals. As perceived by RAND strategists, these cells may be geographically spread over a wide area, or even around the world, but would be capable of swarming in for a coordinated terrorist attack. The kind of architecture is illustrated in Fig.2. The possibility for complex movements of large numbers to be coordinated without any central command and control is familiar from the flocking of birds, and the swarming of insects. Some aerobatic flocking displays seem remarkable without the presence of a lead bird guiding the others. Astonishing feats of spatial intelligence are also achievable by colonies of ants, following an equally elementary set of individual behavioral rules.




Fig.2 Schematic diagram of an 8 cell network with no central control

In the realm of human conflict, an analogy may be made with the predatory tactics of the German U-boat fleet during the second world war. The submarines were dispersed over thousands of square kilometres of the North Atlantic. When a merchant ship was spotted by one of the submarines, information was communicated by radio to the rest of the fleet, and those within range swarmed in for the kill.

With no specific hub to aim at, counter-terrorist forces face a greater challenge in trying to root out all the individual cells making up a swarm intelligence network. The cells would each tend to have a lower security profile than a hub. Already, Iraqi reports claim that a score of trained terrorists are dispersed around the world, prepared to launch a concerted terrorist strike against western interests. Such terrorists constitute at least a tangible search target for counter-terrorist forces; even if special intelligence resources would be needed to locate them. Tough as this challenge may be, more problematic would be a swarm cluster which emerged more or less spontaneously. Such a cluster would be very difficult to identify, and security services would have little warning of its operations.

Emergent Swarm Clusters

Where cells exist with a definite geographical locus, they may become progressively vulnerable to surveillance operations, and infiltration, by counter-terrorist forces. For protection and survival, the dynamics of cell formation may have to be adapted. Harder for security services to thwart would be an attack from an alternative network: one which emerged almost spontaneously from the complex behavior of peripheral sympathisers of the terrorist cause. A swarm attack would be mainly manned not by long-term terrorist suspects, whose movements may be tracked via regular surveillance, but by newer recruits to the terrorist cause, unknown to the security forces. London mosques, with large itinerant congregations, serve as an active recruitment ground for young Moslems, whose British passports afford them freedom of travel without any national allegiance.

The most difficult type of network architecture to deal with is one which has the superficial appearance of random disorganization, with components moving in an apparently chaotic manner, but which actually display a high degree of spatially distributed intelligence. What would be especially puzzling to security forces is the apparently haphazard variation in the commitment of a specific individual to the terrorist cause. Such individuals would not be classified as hard-liners, and would soon disappear from the terrorist radar screen. The charter of the jihad movement, as expressed in the Islamic radicalism of Sayyid Qutb (Mousalli, 1993), is freely available on the internet; attendance at mosques to hear radical imams might be rather sporadic. Grounds for prior arrest or detention as a potential terrorist suspect would accordingly be very thin.

Swarming is an image borrowed from the natural world of three space dimensions. A swarm of bees, for example, is defined by spatial clustering. However, swarming may be defined in any number of dimensions, including non-physical dimensions such as support for jihad; disdain for western culture; etc.. For simplicity, these other dimensions may be collapsed to a single dimension defined by commitment to participate in a terrorist act. The greatest challenge to security forces would arise from swarming in this virtual terrorism dimension, by individuals who might physically be geographical dispersed all over the world. This is illustrated schematically in Fig.3.

These individuals may not themselves have any prolonged history of links with radical groups, so they would be hard to identify in advance as potential suspects. They may be motivated through public exhortations to violence on the radio, television, or the internet. A cluster of like-minded individuals, who may never have actually met, could collectively contrive a terrorist act, using global communications such as the internet. Being spontaneously generated, such a group would be almost impossible to infiltrate. An emergent network is essentially a virtual one, in respect both of physical presence and web-based communication. The capability of militant anarchists and anti-capitalists to mass together and cause mayhem at the economic summits in Seattle and Genoa shows the potency of an emergent network. The ranks of the hard-core anarchists were swelled by middle-class students and young professionals. An alarming future prospect would be the rapid recruitment to the militant Islamic cause of well educated but

disaffected moslems, especially those born and raised in the West, whose loyalty to al-Qaeda may be all but invisible to the security forces.

Commitment To Terrorism

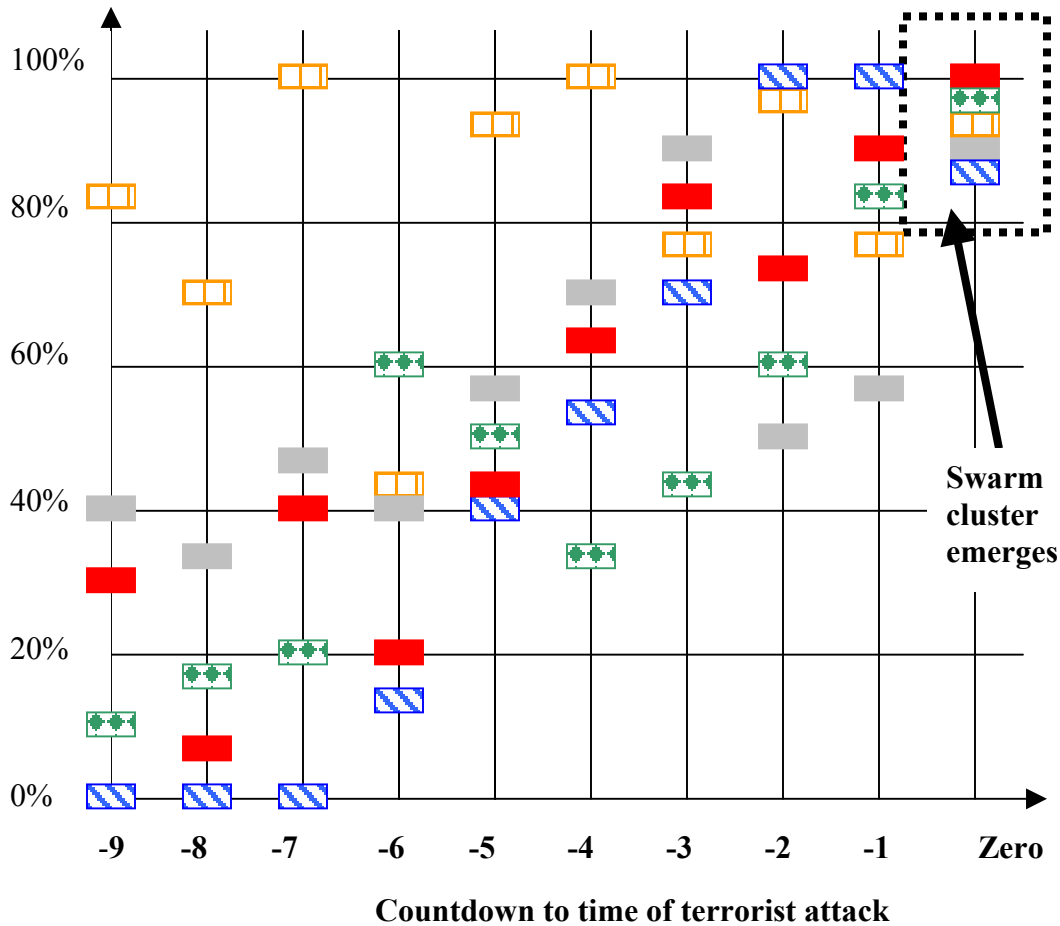


Fig.3 Schematic chart showing how a 5-man cluster can suddenly emerge, without central control, even if the conspirators are geographically widely separated. The commitment of an individual to participating in terrorism is gauged in percentage terms, and graphed as the ordinate on the plot. Clustering is in this dimension, rather than in actual space dimensions, so the detection of a cluster by security services, in advance of the zero hour for attack, would be extremely difficult.

The formation of attack swarms can be computer simulated, using algorithms drawn from the modeling of other emergent swarm intelligence phenomena (Kennedy et al., 2001).

Since cell membership may fluctuate widely from time to time, terrorists would be hard to track down by security forces. The attacks perpetrated by an emergent network would be difficult to prevent, even if intensive counter-swarm measures were implemented to catch nests of terrorists. Given the comparatively low reliability of attack detection, the number of attacks therefore might proliferate. However, given the shorter training available for the attacks, and the involvement of less experienced personnel, the damage inflicted in any action is unlikely to be as extensive as in the attacks planned carefully over many years. Nevertheless, the remote possibility exists that an emergent network, involving apparently respectable professional technologists, educated in the West, might be capable of delivering a strike involving weapons of mass destruction: WMD. The danger is that, in the distorted perception of al-Qaeda, this might be viewed as the perfect network for such a mission.

The Frequency of Planned Attacks

For as long as jihad is used to support or excuse fanatical suicide missions, others will follow in the footsteps of the British-educated Ramzi Yousef, master-bomber of al-Qaeda, who imagined he was advancing the Palestinian cause by plying his terrorist trade around the world. As with the Irish Republican Army's terror campaign in Britain, the frequency of Islamic terrorist attacks may be expected to fluctuate with the swings of political fortune and progress in peace negotiations. Exasperation with the political process may encourage the proliferation of attacks.

In a reference to IRA media coverage, Margaret Thatcher observed that terrorists thrive on the 'oxygen of publicity'. In the days before television, terrorists might signal their presence via an intensive bombing campaign: the IRA exploded 127 devices in Britain during the late 1930's. In contrast, IRA political frustration with the Ulster peace process was vented in 1996 by a showpiece bomb blast at Canary Wharf, London's WTC. As appreciated by the IRA, publicity and fear go together, and these are maximized if terrorists choose not just political and financial targets, but a wide range of targets, including infrastructure, shops, bars, government offices, as well as commercial buildings. This diversity in targeting is some consolation to property insurers, who would not be heavily exposed to loss from all attacks.

For a hierarchical terrorist organization, the rate of planned attacks should be proportional to the number of mature active cells. This number would thus be drastically curtailed by sustained global counter-terrorist action to seek and destroy these cells. However, the spontaneous formation of less detectable emergent terrorist cells would permit quite frequent attacks to be launched, in guerilla style, on a multiplicity of targets around the world. Attack frequency would not be restricted by the number of cells, since these would rapidly form and disperse. The numbers of planned attacks, i.e. the rhythm of terror, might ultimately be determined as much by publicity goals as by the size of the

terrorist ranks. As exemplified by the IRA campaign, well-publicized occasional moderate bomb damage suffices to perpetuate a reign of fear.

The Spatial Distribution of Hazard

The distributed spatial intelligence of trans-continental terrorist networks allows attacks to be made across the globe, by operatives of many nationalities, at locations which may be far distant from any cell. From the bombing of the Khobar Towers in Saudi Arabia in 1996, to the Nairobi and Dar-es-Salaam US embassy bombings in 1998, to the bombing of the U.S.S. Cole in 2000, and the WTC disaster of 2001, al-Qaeda have developed a swarm-like campaign of pulsing attacks from different nodes of its global network.

Unlike natural perils, the hazard from which is spatially concentrated around geological faults, coastlines, flood-plains etc., the free mobility of terrorists means that there is no fixed geography for terrorism hazard. There is an earthquake engineering adage that an earthquake will expose the weakest link in a building. But if a number of structures are randomly distributed in a region, the pattern of seismicity does not alter so that the weakest structure is most likely to be shaken. Yet, with a terrorist threat to a number of prize targets, the most vulnerable may have the highest probability of being attacked. Other things being equal, less planning time and fewer resources should be needed to launch a successful raid. The dependence of hazard location on vulnerability introduces a nonlinear feedback in risk computation, which would tend to escalate the loss potential.

A classic illustration of this principle is Hammersmith Bridge in West London, which has been repeatedly bombed with impunity by the IRA. Of all the major London bridges, this has especially lax night-time security. Closure of this vital transport artery for months at a time has caused significant business interruption and consequent economic loss. Gridlock is just one objective; a terrorist campaign may, following the IRA example, switch intermittently between political, commercial and economic targets.

Should a terrorist organization aim to cause maximal economic disruption, it might preferentially target spatially distributed infrastructure networks of telecommunications, power, water supply and transportation. Not only is there little redundancy in these systems, but they are also highly vulnerable. US critical infrastructure suffers historically from serious neglect by regulatory and law-enforcement agencies. The potential consequences of a breach of security are enormous. If the U.S.S. Cole attack were replicated against a tanker docked at the main terminal in Long Beach, California, the economy of southern California would be crippled in days through lack of oil supplies (Flynn, 2001). Similarly, if Port Everglades in Florida were shut down through terrorist action, millions of southern Florida residents would suffer an almost immediate fuel shortage.

Event-tree for the Detection and Thwarting of Planned Attacks

In common with other long-tail risk phenomena, such as aviation or marine accidents, the task of characterizing the tail of the terrorism loss distribution is assisted by exploration and analysis of the statistics of the more common 'near-miss' events, which narrowly failed to be significant disasters. In the context of terrorist attacks, near-misses arise because they were thwarted by counter-terrorist forces, security checks, or because of some technical or logistical misadventure. Nobody familiar with the prolific professional bombing career of Ramzi Yousef would be content to assess risk on the basis of actual losses alone. In 1993, he came close to toppling one of the towers of the World Trade Center; he bombed a Philippines airliner in 1995, and subsequently plotted to bomb as many as eleven jets over the Pacific. Even though Yousef was jailed, he was a member of a terrorist network, and prudent underwriters might well have double-checked their exposure to multiple WTC and civil aviation disasters.

Apart from the 1993 attack itself, the possibility of the WTC being targeted for terrorist action might have been apparent to underwriters, if not also the security forces, from other developments. The concept of using a fuel-laden commercial airliner as a missile had occurred to Algerian terrorists in 1994, who hatched a plan to crash a plane into the Eiffel Tower in Paris. Not long afterwards, unbeknown to MI5 or CIA, Algerian and Afghan café workers in London were soliciting American signatories on applications for flying training in USA. The long-term planning for September 11, 2001 had already begun.

Not all planned attacks succeed in causing a notable loss. Terrorists may be defeated through good intelligence; good security and policing, and some good fortune. Yousef's audacious multiple airliner bomb plot was foiled through an accidental fire breaking out during the bomb-making process. The Eiffel Tower still stands because the plane hijacked by Algerian terrorists was stormed by French commandos, while being refuelled.

In the probabilistic risk assessment (PRA) of nuclear installations, which provides the methodological basis underlying insurance natural catastrophe modeling (Woo, 1999), and more recently civil aviation risk modeling, the damage consequences of an initiating hazardous event are logically charted via a multi-branch event-tree. The process of systematically dis-aggregating risk into component elements, through an event-tree, is an important aspect of structuring a risk analysis.

The event-tree can be used in the present context to estimate the probability that a planned terrorist attack results in a notable loss. The success of a planned attack is contingent on certain events either occurring or not occurring. These events are the procurement of intelligence on the attack; resolution to act based on any intelligence; interdiction by police or other security officials; and technical or logistical operational failure. A basic event-tree, constructed around these events, is shown in Fig.4.

To parametrize this event-tree, four basic conditional probabilities need to be quantified.

- [1] Given that an attack is planned, what is the probability that there is some prior intelligence about it?
- [2] Given that an attack is planned, and there is some prior intelligence about it, what is the probability that the intelligence is acted upon?
- [3] Given that an attack is planned, and either no intelligence exists or else is not acted upon, what is the probability that the attack is nevertheless detected by border guards, police or other security personnel?
- [4] Given that an attack is planned, but remains completely undetected, what is the probability that it fails to cause significant loss due to technical or logistical shortcomings?

These key questions of probability are elaborated below.

[1] Availability of intelligence

Even with a hub network structure, which is comparatively accessible to surveillance, intelligence on al-Qaeda has been poor. Whatever the enhancement of counter-terrorist intelligence, Pentagon officials (Reeve, 1999) admit that future attacks planned by emergent local networks will be “very, very difficult to stop”. Terrorists may never actually meet, except to carry out a specific operation, with all the planning, the vetting of individuals all taking place through the Internet. Thus the chance that intelligence about such an attack would be forthcoming is very slight.

[2] Action on intelligence

False alarms are the bane of intelligence services. The more disinformation and bogus threats there are, the smaller the probability that intelligence about a specific threat will be acted upon. In the six months following the US embassy bombings in East Africa, US diplomatic facilities around the world were swamped with 650 credible threats from al-Qaeda (Bergen, 2001). These threats were garnered from a range of sources, including informants, telephone surveillance, sightings of a terrorist in a city, and embassy videotaping. Finite resources for checking the reliability of threats greatly restricts the number which can be adequately acted upon. Furthermore, as with other natural peril hazard alarms, the high cost of the disruption resulting from a false alarm has to be weighed against the expected benefit from sounding a correct alarm.

In common with the international transport of illegal immigrants, intelligence is vital. But unless this intelligence is sound, the time and expense of stopping and searching can

be inordinately high. In the case of the cargo ship MV Nisha, purportedly carrying sugar from Mauritius to London, via Djibouti, the detention of the vessel for several days at the end of 2001 turned out to be an unwarranted and costly mistake by British intelligence. No evidence of explosives was found. Given practical tolerance limits to the economic disruption induced by rigorous security checks at borders, searching for terrorist needles in a transportation haystack may well prove to be a largely futile exercise.

[3] Security barriers

Even if no prior intelligence exists about a planned attack, or if any such intelligence is ignored, there is still a chance that it would be detected at some stage by border guards, police, customs officers, or another security barrier. However, recognizing the coarse mesh of the global security net, this chance may be comparatively low. Given the vast traffic flows across international borders, tracking terrorists effectively without unacceptable delays to others may be nigh impossible.

At the world's busiest commercial land-border crossing, which is the Ambassador bridge between Michigan and Ontario, US customs officers have no more than two minutes for each truck. Despite security deficiencies, sometimes terrorists do manage to get themselves caught at border crossings, as happened when the sweating Ahmed Ressam was arrested on a ferry arriving from Canada. In his car were 130 pounds of explosives intended for a millennium eve bombing.

[4] Technical or logistical failure

The track record of terrorist miscalculations, bungles, and other miscellaneous mishaps provides some relevant experience data on mission failure. One of the longest data series, covering a number of decades, comes from the IRA campaign in Britain. In attacks on political targets, the IRA succeeded in several political assassinations, and came close on several occasions, in 1984 and 1991, to assassinating a group of senior government ministers. In attacks targeted at the public, car and truck bombs have taken their toll of bars, shopping centers and offices. But apart from those bombs which have caused notable damage, others have either only partially exploded, failed to detonate, or detonated prematurely. Failures have also occurred for logistical reasons such as in the theft of cars or trucks for bomb transport.

Public sources of information provide some guide as to the proportion of undetected bombing attempts which failed to make a significant impact for technical or logistical reasons. However, allowance needs to be made for those technical or logistical failures which never attracted public attention. In principle, the failure ratio might be assigned on a scenario basis, but this complication may be avoided by using a generic value.

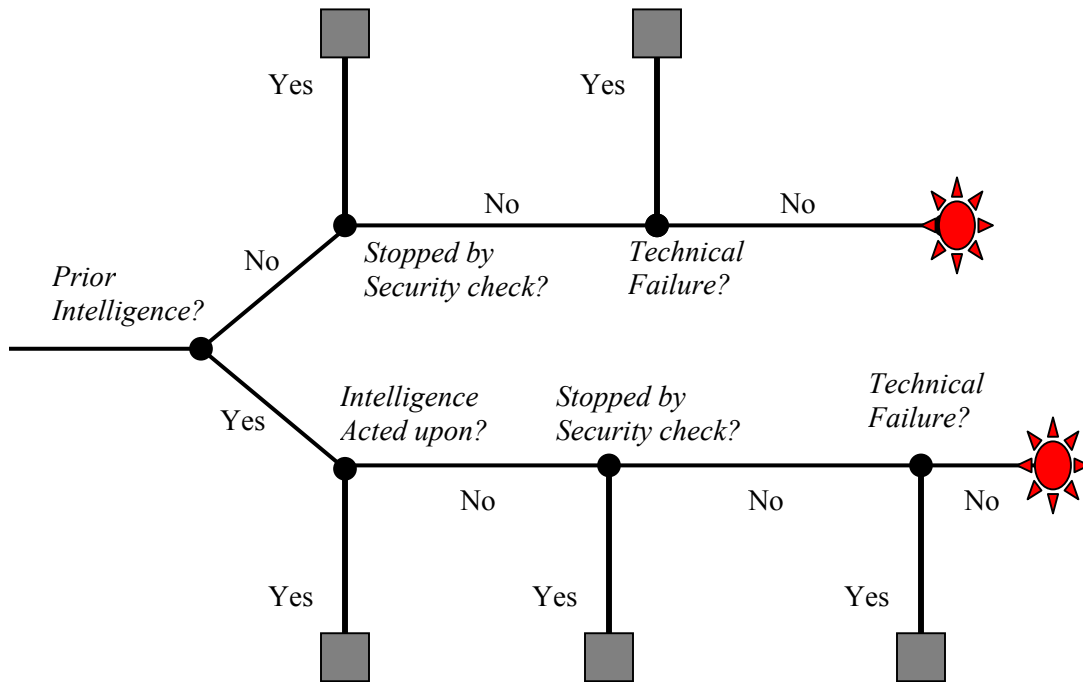


Fig.4 Event-tree depicting the alternative outcomes for a planned attack. Each branch is assigned a probability weight. The five square termination points signify that the attack is thwarted; the two starred termination points signify that some notable loss results.

The systematic assignment of probability weights remains to be undertaken, but preliminary estimates suggest that the overall chance of a planned attack being successful would be high for an emergent network. Weaknesses in the defensive shield against such terrorist attacks will be hard to rectify. Whatever the international will to fight terrorism, a substantial residual risk is likely to remain.

Loss Severity Distribution

Before war was declared on terrorism, al-Qaeda could afford to take time, and devote resources, to plan attacks which would perpetuate their reign of terror, and also inflict substantial loss: the larger the loss, the more likely the scenario would be chosen. With this positive feedback, the loss severity distribution prior to September 11 would have been skewed towards heavy losses; a risk characteristic consistent of course with the WTC attack.

Adapting from a hub structure to an emergent network architecture, al-Qaeda may become less visible to the spreading force of counter-terrorists, but the organization would pay a penalty: it would be hampered with more coordination and supply problems. The impairment of coordination and restriction of resources should make it more difficult for spectacular massive single-site strikes to be successfully delivered, and more tricky to synchronize contemporaneous strikes at different locations. Furthermore, the nonlinear feedback dependence of scenario likelihood on loss will be much diminished; high loss scenarios may be attractive to al-Qaeda, but they may also be especially hard to execute under pressure. Counting synchronized strikes as a single event, the relative likelihood of one event generating enormous losses is correspondingly reduced with the hubless swarm network architecture. In earthquake parlance, one would say that the emergent network architecture yields a higher b-value of the Gutenberg-Richter magnitude-frequency relation; the b-value being a measure of the relative frequency of small to large events.

This argument can be made computationally explicit when expressed in terms of an ensemble of realistic terrorism scenarios, which are the basic building blocks of a terrorism stochastic model. GIS tools for mapping the geography of exposure, developed originally for natural hazard applications, are being adapted to map the spatial footprint of loss across a spectrum of high-publicity showpiece terrorist scenarios, considered as realistic in terms of their potential attraction to al-Qaeda. This mapping procedure is quite arduous and lengthy, because these scenarios can be quite diverse in their initiating event, and in their urban impact.

For an emergent network, under constant pressure from international counter-terrorist forces, the types of attacks which can be attempted will be constrained by available resources, and limited by shortness of planning time. The IRA campaign provides illustrations of the effectiveness of heightening security, and cutting off supplies of armaments, in reducing the options for terrorist action. In future, al-Qaeda will have more of an operational research allocation problem to solve: to achieve its high-profile publicity goals, subject to fairly stringent time and resource constraints. Thus whatever the personal predilections of the individual terrorists, the relative likelihood of a showpiece scenario, (missile strike, LNG tanker explosion, nuclear device detonation or whatever), will be governed by a cost function, which is a measure of the overall difficulty in execution, allowing for a variety of factors, such as planning time, personnel effort, technical problems, and consumption of financial and material resources. (For

martrydom missions, there is of course no cost consideration for danger to the terrorists themselves.)

By quantifying relative likelihood on the basis of the cost function, a value of which is assigned to each realistic scenario, a loss severity distribution can be constructed by ranking the terrorism scenarios according to loss. Clearly, the form of this loss severity distribution will depend on the choice of the cost function. But, since the cost function is used for assigning relative rather than absolute frequencies, the final risk results should be reasonably robust against changes, the effect of which, in any case, can be explored in sensitivity studies.

Loss Exceedance Probability Estimation

Once the loss severity distribution is parametrized, the annual loss exceedance probability may be calculated using a best estimate, (or hypothetical time-dependent distribution), for the overall frequency of planned attacks, and the event-tree estimate for the probability of an attack being thwarted. This risk calculation inevitably involves a number of subjective probability assignments. These may be made informally by in-house risk analysts, but, in order to maintain the principles of fairness, neutrality and accountability (Cooke, 1991), these are best made through eliciting the expert judgement of terrorism experts.

The procedure for eliciting expert judgement indeed has its origins in assessments made of Cold War threats in the 1960's, by the RAND corporation, and its use in the present terrorist context closes a circle of application. This reference to RAND's expertise in strategic thinking is particularly apposite, because the seminal ideas of netwar which underpin this paper originate from RAND experts (Arquilla et al., 2000). Since the 1960's, the expert elicitation procedure has been used widely in natural hazard risk studies, including the most rigorous site-specific analyses of engineering seismic hazard. It has also been used to assess the risk of malicious intrusion into a radioactive waste disposal site (Woo, 1989). The theft of radioactive material, and its dispersal in an urban population center, has a terrorist precedent in the Chechen campaign against Russia, and is very much a major concern for the future.

The value of consulting a group of leading terrorism experts was clearly demonstrated in a secret study called 'Terror 2000', conducted by the Pentagon, intended to help the intelligence world meet the terrorism threat. This study was facilitated by the president of Forecasting International, who involved forty leading experts, including a former KGB head; a senior official from Mossad; and Prof. Paul Wilkinson, the noted British terrorism analyst. One of the prescient conclusions of the study was that terrorists would soon try to conduct simultaneous bombings, perhaps in different countries, to maximize the devastation and publicity. This was just one of a number of findings which recipients of the report, both in government and industry, found to be unrealistic as well as unpalatable.

Knowledgeable as these experts would have been of the foiled Algerian attempt to crash a plane into the Eiffel Tower, had they been asked in 2000 to assign a probability to an aviation attack on the World Trade Center, the answer would surely have prepared insurers better, both commercially and psychologically, for the disaster of the following year. Probable Maximum Loss values maintained by some insurers for the WTC would have been tenable only if the probability were minuscule.

In order to provide insurers with a quantitative guide as to the extent of terrorism exposure, the probabilistic framework outlined here for quantifying terrorism risk could be parametrized systematically with the assistance of a convocation of international terrorism experts. Simple calculations suggest that, despite international counter-terrorist action, the risk is currently substantial, as indeed it was before September 11, 2001.

Conclusions

Although insurance is provided against many types of criminal behavior, quantifying terrorism risk is recognized to be a formidable task for the insurance industry. Yet it should be remembered that the challenge of treating human factors, such as fraud, vandalism, sabotage, and human error is met in the course of many industrial probabilistic risk assessments.

An approach to tackling the terrorism risk quantification task has been outlined, from which the following principal observations may be drawn:

- The frequency and severity of planned attacks will depend critically on the network architecture of the terrorist organization.
- Pressurized increasingly by counter-terrorist forces, terrorist organizations may adapt to form emergent swarm clusters. These rapidly forming virtual cells, communicating via internet, will be very hard to detect and stop.
- Emergent networks will facilitate the execution of more frequent, but less ambitious and generally less damaging, planned attacks.
- An event-tree may be constructed to estimate the probability that a planned attack will succeed, depending on the availability and usage of intelligence; the effectiveness of security barriers; and technical and logistical mishaps.
- The loss severity distribution may be derived by mapping losses from realistic showpiece terrorism scenarios, and assigning a cost function to each. The cost function reflects practical logistical factors such as planning time, technical difficulty, and consumption of scarce resources.
- The overall computation of a terrorism loss exceedance curve can be achieved, provided that the assignment of subjective input probabilities is made using the formal elicitation of expert judgement, such as has been invoked already by government security agencies.

As with any risk analysis, the derivation of probabilistic loss exceedance curves for terrorism is not an end in itself, but an aid for insurers to make better decisions under uncertainty. Even if these loss curves themselves are uncertain, general guidance on the steepness and scale of these loss curves, and their sensitivity to model assumptions, would be welcome technical support for insurers in setting PML's and in pricing terrorism cover.

References

Arquilla J., Ronfeldt D., Zanini M. (2000) Information-age terrorism. *Current History*, Vol.99, pp.179-185.

Bergen P.L. (2001) *Holy war, Inc.* Weidenfeld and Nicholson, London.

Cooke R.M. (1991) *Experts in uncertainty.* Oxford University Press, Oxford.

Flynn S.E. (2001) *The unguarded homeland.* In: *How did this happen?* (J.F. Hoge and G. Rose, Eds.), PublicAffairs Ltd., Oxford.

Johnson S. (2001) *Emergence.* Allen Lane, the Penguin Press.

Kennedy J., Eberhart R.C. (2001) *Swarm intelligence.* Morgan Kaufmann Publishers.

Moussalli A.S. (1993) *Radical Islamic fundamentalism: the ideological and political discourse of Sayyid Qutb.* Syracuse University Press.

Reeve S. (1999) *The new jackals.* Andre Deutsch, London.

Ronfeldt D., Arquilla J. (2001) *Networks, Netwars, and the fight of the future.* *First Monday*, issue6_10.

Ronfeldt D., Arquilla J. (2001) *Networks and Netwars: The Future of Terror, Crime and Militancy.* RAND Corporation.

Waldrop M.M. (1992) *Complexity.* Viking Press, London.

Wilson E.O. (1975) *Sociobiology.* Harvard University Press, Cambridge MA.

Woo G. (1989) *Is the risk of human intrusion exaggerated?* In *Proceedings of NEA workshop on: risks with human intrusion at radioactive waste disposal sites*, OECD, Paris.

Woo G. (1999) *The mathematics of natural catastrophes.* Imperial College Press.