QUANTITATIVE TERRORISM RISK ASSESSMENT

By Dr. Gordon Woo
Risk Management Solutions Ltd.

Traditionally, the pricing of terrorism risk has been discovered from the balance of supply and demand in the insurance market, rather than evaluated from actuarial principles. Risk selection through the use of site security surveys has helped reduce the number of inferior risks, and systematic portfolio risk aggregation has limited the Probable Maximum Loss from any attack scenario. With such basic risk management procedures in place, it has been possible for international terrorism to be commercially underwritten in Europe and Asia. Of course, the tragic events of September 11 have irrevocably changed the market place of terrorism insurance. Terrorism is now a US catastrophe risk, and as with natural perils that have the power to cause catastrophic insured losses, the development of computerized tools for portfolio risk management has become a topic of urgent practical interest.

The task of quantifying terrorism risk should not be confused with predicting the next terrorist attack. This important distinction between risk assessment and event prediction exists also with natural perils. No seismologist is capable of predicting the time, place and magnitude of the next major earthquake in California, but it is possible for a seismic risk analyst to evaluate the annual exceedance probability of loss to a California property portfolio. Large earthquakes are impossible to predict because of a haphazard random element in the manner in which the rupture of a geological fault propagates and eventually stops. This randomness contributes to the so-called *aleatory* uncertainty in earthquake occurrence (Woo [1999]), which is readily accounted for within a probabilistic risk assessment, but confounds attempts at deterministic event prediction. This distinction between risk assessment and event prediction matters to civic authorities and insurers. Being responsible primarily for public safety, civic authorities would ideally like to have reliable predictions to warn against imminent hazard events, or prioritize urgent defensive measures (Cordesman [2002]). Insurers, on the other hand, seek to quantify risk not over a brief time window of a few days, but spread over a number of months. For this less ambitious purpose, quantitative risk assessment is achievable.

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 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.

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: the human dimension to conflict has to be incorporated.

A RAND suggestion (Ronfeldt et al. [2001]) 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, 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 practioners 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, through learning from experience and emulating the successful behavior of others, people are able to discover relatively optimal patterns of attitudes, beliefs and behaviors. Al-Qaeda operatives are known to be highly adaptive in learning from past terrorist successes and failures (Gunaratna [2002]). 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.

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. 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.

STRUCTURE OF TERRORIST ORGANIZATIONS

An immediate observation made in the aftermath of September 11 (Hoffman [2001]) 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, scientific 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 adaptability of organization and deftness of manoeuver. Despite being vastly inferior in numbers, weaponry and combat capability, 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. 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. The wave of Hamas suicide bombings in Israel in 2002 illustrate this point. 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.

The names by which terrorist groups are known reflect their organization. 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. To the French security service, the English capital city is 'Londonistan', because of the congregation of Islamic militants who claim refuge across the English Channel. Britain has long been a forward base for Islamic militants; 20% of Osama bin Laden's telecommunications came to Britain. Spanning several continents, an international network cynically exploits national differences in the tolerance of foreign terrorists, in the liberality of laws of asylum and extradition, and in the preservation of civil liberties.

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]). 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.

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. 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.

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. These cells may be geographically spread over a wide area, or even around the world, but would, like submarines, be capable of swarming in for a coordinated terrorist attack. 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 may be mainly manned not by long-term terrorist suspects, whose movements may be tracked via regular surveillance, but by fresher recruits who happen to have drifted towards the terrorist cause.

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 democracy, 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. 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 jihad on the radio, television, internet, in the radical Islamic literature (e.g. Hamza [2002]). A cluster of like-minded individuals, who may never have actually met, could collectively contrive a terrorist act, using global communications such as provided on internet chat-rooms. An emergent network is essentially a virtual one, in respect both of physical presence and web-based communication.

Being spontaneously generated, and with minimal personal outside links to the secular world, such a group would be almost impossible to infiltrate. Referring to the September 11 hijackings, Osama bin laden noted that 'those who were trained to fly didn't know the others. One group of people didn't know the other group'. The sparse links between the nineteen hijackers have been charted by Krebs [2002], who has highlighted the way in which covert networks trade efficiency for secrecy, and the avoidance of detection. This

trade-off is symptomatic of the patient and diligent approach to al-Qaeda operational planning, which tends to prolong the preparation of major attacks.

A STOCHASTIC TERRORISM MODEL

While British Prime Minister, 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. The absolute number of attacks within a year, i.e. the rhythm of terror, might ultimately be determined as much by publicity goals and the political anniversary calendar 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, and concentrate the minds of politicians on the terrorist's agenda. In the modern era of instant global news communication, it only requires a few major successful attacks for a terrorist's message to be retained by the public.

Thwarting Terrorist Attacks

The challenge of thwarting an attack by al-Qaeda has analogies with hunting down a swarm of submarines. In anti-submarine warfare, a key defensive tool is signal processing to extract the submarine signal from the background noise of the sea. In the context of al-Qaeda, the problem is to search for anomalies in the vast global electronic transaction space covering money transfer, credit cards, education, travel, immigration, transportation, housing and medicine. But just as submarines strive to minimize their acoustic signature, so terrorists will try to minimize their transaction space signature. However, through electronic searching for traces of terrorist activity, augmented by human intelligence and covert surveillance, the dots of a planned attack may be joined up; the conspirators identified and tracked; and the attack pre-empted.

Since September 11, a significant number of al-Qaeda attacks around the world, from Europe to the Middle East and Asia, have indeed been thwarted. But some have been successful. As the IRA reminded Margaret Thatcher after its abortive assassination attempt, terrorists have only to be lucky once. The WTC hijackers had their share of good fortune, and those following in the path of Allah would need their share in accomplishing a martyrdom operation. Under pressure from counter-terrorism forces, many things can go wrong for a terrorist group: suspicion may be created, information may leak out and be acted upon; their plans may be uncovered through layers of security checks, or come awry through a diverse litany of technical shortcomings. Analysis of this event-tree of detrimental operational factors allows calculation of the probability that a planned attack is successful.

Markov Model

The term *macroterrorism* has been coined to describe a spectacular act of terrorism, (which may be a multiple strike at several locations), which causes more than $1 billion of loss, or 500 deaths. Minor (micro) terrorist acts, such as house bombing, may occur haphazardly, but not signify a change in the terrorism environment. However, this is not the case with macroterrorism. Following an act of macroterrorism, security and border controls are inevitably strengthened, and emergency government funding made available for improving protective measures. Civil liberties may be temporarily curtailed as suspects are detained without trial, and minority communities potentially supporting sleeper cells are placed under tight surveillance.

Although copycat attacks may be attempted in the aftermath of a successful strike, they are likely to fail due to the heightened security. In the harsher security regime soon after a successful strike, terrorists may rationally decide to lie low, and delay any further action until security is relaxed, border controls are eased, civil liberties lawyers intervene, and public risk awareness fades: circumstances which would give a later attack a higher chance of success. There are other reasons favoring a delay. Logistically, resources may need to be replenished after a macroterror attack. Furthermore, once a terrorist's message has been delivered across the media through a spectacular macroterrorism event, (perhaps after a series of failures), a publicity reminder may not be needed for a while.

The change in system state following a sucessful macroterrorism event implies that, rather like great earthquakes, such events do not satisfy the prerequisites of a Poisson process. Although it would require an elaborate Monte Carlo simulation to realize the temporal pattern of successful al-Qaeda macroterror attacks, the simplest representation is a two-state Markov process. In the first state, security is comparatively relaxed, and conducive to a successful macroterror attack. In the second state, security is comparatively strict, and not conducive to a successful macroterror attack. With the almost infinite payoff of paradise promised to martyrs, patience in waiting for security weaknesses is an optimal strategy. Indeed, it is known that Osama bin Laden has expected very high reliability levels for martyrdom operations.

As a didactic illustration, consider the binary situation where successful macroterror attacks only take place during the relaxed security state. If the rate of successful macroterror attacks in this first state is U, and the erosion rate of security in the second state is V, then, assuming a successful macroterror attack causes a state transition from 1 to 2, the limiting proportion of time spent in state 1 is V/(U+V), and the limiting frequency of successful macroterror attacks is UV/(U+V). The effect of maintaining security measures is to keep V low, and hence suppress the limiting frequency of successful macroterror attacks.

THE LOSS SEVERITY DISTRIBUTION

Before war was declared on terrorism, al-Qaeda could afford to take time, and devote resources, to plan meticulous attacks against targets which gained legitimacy through being emblematic of US economic, political and military power: the greater the loss, the more attractive to al-Qaeda. 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. 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. The IRA campaign provides illustrations of the effectiveness of heightening security, and cutting off supplies of armaments, in reducing the options for terrorist action.

Relative Likelihood of Attack Modes

'Avoid strength, and attack weakness', a saying of the legendary military strategist Sun Tzu, is a fundamental precept for the terrorist conduct of asymmetric warfare against a much more powerful adversary. For al-Qaeda, this may be expressed in the succinct language of physical science as: *follow the path of least resistance.* In hydrology, the principle of minimum energy expenditure governs the pattern of river drainage networks. In a similar way to the flow of water, the flow of al-Qaeda terrorism activity is towards weapon modes and targets, against which the technical, logistical and security barriers to mission success are least. Since September 11, the counter-terrorism environment for the development of new weapons and planning complex strategic operations has become oppressive for al-Qaeda. Increasingly, like a scavenging bear in the wild, it is becoming obliged to take targets of opportunity. Accordingly, it may look towards off-the-shelf, ready-to-use weapons, (such as SAM or Stinger missiles, hijacked aircraft, and propane tankers), or improvised conventional explosive devices, which do not involve intricate and potentially failure-prone technological development. There is continued financial support from international Muslim communities to buy weapons, but skills to develop new weapons are becoming more scarce, given the crackdown on al-Qaeda. In general, the more complex the weapon, the more terrorists there are with knowledge which may be compromised. For example, dozens of al-Qaeda operatives would need to be involved in the assembly of a nuclear detonation device.

The Selection of Targets

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. The dependence of target location on vulnerability introduces a nonlinear feedback in risk computation, which would tend to escalate the loss potential. This feedback may be recognized explicitly using the mathematical theory of conflict, i.e. game theory. The hijacking of a Singapore Airlines plane in 1991 by Pakistani militants provided an early case study of the use of game theory in the context of terrorist negotiations, but target prioritization is a further area of terrorism application.

Consider the potential targets in the USA, and suppose that they have been ranked in discrete city tiers and type classes (skycrapers, bridges, nuclear plants etc.) by terrorism experts, according to their attractiveness (or utility) to al-Qaeda. Symbolic and publicity value, name recognition in the Middle East, economic and human loss consequence would be factors in gauging target utility. In order to express target prioritization in a quantitative way, the ranking by city and target type has to be converted into mathematical form. This interpolation is simply achieved by invoking Fechner's Law, which states that an arithmetic progression in perceptions requires a geometrical progression in their stimuli. This implies a logarithmic formula for the utility of the C'th city tier, and for the T'th type class. This form of rank interpolation allows the logarithm of the utility to be written parametrically as:

$$Log\{U[C,T]\} = k_0 - k_1 C - k_2 T$$

In order to arrive at a target probability distribution, a mathematical expression needs to be obtained for the functional dependence of target probability on utility. For this, game theory is required. It is known that al-Qaeda is committed to achieving success, is watchful that missions are cost-effective, and is sensitive to target hardening. In order to ensure, as far as possible, that a strike will be successful, irrespective of defensive action by security forces, al-Qaeda will effectively seek to minimize the impact of target hardening. From knowledge of its modus operandi, this goal is attained by al-Qaeda by adopting a mixed strategy of randomizing its target selection, meticulously undertaking

surveillance on targets and avoiding targets where the level of security is very uncertain; and switching targets if the original target has hardened. This might happen if the National Guard were deployed, or the police were working over-time; resources which are sparingly used, because they are expensive for civic authorities to procure.

For an attack using a specific weapon against a target in category [C,T] with defense D, let P_D be the probability that the defense is unable to prevent or stop the attack. As increasing defensive resources are applied to protect targets of high utility, the marginal improvement in security diminishes. This is reflected by a defense saturation condition, the power-law form of which is motivated by the fractal nature of defense-in-depth hierarchy (Paparone et al. [2002]), as realized, for example, in multiple defence barrier models (Major [2002]):

$$\frac{\partial P_D}{\partial D} \propto -U[C,T]^{-\lambda} \qquad (\lambda > 0)$$

For a mixed attack strategy, designed so as not to be impaired by changes in defense strategy, game theory optimization analysis suggests that the probability of selecting a target $P(U[C,T])$ may be expressed as follows:

$$1 / P(U[C,T]) \propto -U[C,T]\frac{\partial P_D}{\partial D} \propto U[C,T]^{1-\lambda}$$

Combining formulae, the following result is obtained:

$$Log\{P(U[C,T])\} = a - (\lambda - 1).k_1 C - (\lambda - 1).k_2 T$$

Substituting b_1 for $(\lambda - 1).k_1$, and b_2 for $(\lambda - 1).k_2$,

$$Log\{P(U[C,T])\} = a - b_1 C - b_2 T$$

From this equation, one can see that the relative likelihood of targets being selected depends simply on the two parameters b_1 and b_2. The form of this equation is reminiscent of the elegantly simple Gutenberg-Richter relation in seismology, which is the cornerstone of seismic hazard analysis. As with the Gutenberg-Richter relation, the parsimony of the equation, which has just two parameters to determine, compensates for the approximate nature of the model. The notation, b_1 and b_2, is chosen to echo the Gutenberg-Richter b-value. As with the seismological b-value, the parameters b_1 and b_2 may be estimated from empirical data, supplemented by expert judgement, where data are sparse. Note that the game theory impact is to shift the target probability distribution away from the targets of highest utility, which are likely to be well defended.

PROBABILISTIC RISK ANALYSIS

The preceding sections outline a framework for a probabilistic risk analysis. This basic framework for a probabilistic model of terrorist activity is the same for any terrorist organization, but in as much as their motives, strategies and objectives differ, so also will the model. An organization intent on attracting publicity to a political cause, but not to kill thousands of civilians, will not be motivated to search for weapons of very high lethality; this may alienate their popular support, and so be counter-productive. Their purposes may be served by sporadic conventional bomb attacks, preceded possibly by bomb warnings. However, Islamic militants such as al-Qaeda, who have made clear a determined intent to inflict large scale damage and human suffering on the USA, will be concerned with lethality issues of weapon effectiveness.

Hausken [2002] has remarked on the applicability of game theory concepts to natural hazard Probabilistic Risk Assessment (PRA), in circumstances, such as the joint maintenance of flood defenses, where individuals incur costs to increase system reliability, which is interpreted as a public good. In so far as terrorists will choose the moment and place of an attack to exploit defensive weaknesses, explicit use of game theory is a feature of terrorism PRA.

Whatever the underlying theoretical foundation, given the dependence on the modus operandi of a terrorist organization, any risk calculation inevitably involves a number of subjective probability assignments, variability of which amplifies the epistemic uncertainty. These assignments may be made informally by in-house risk analysts, but, in respect of a global network such as al-Qaeda, these are best made through eliciting the expert judgement of international terrorism experts, familiar with terrorist operations on all continents. The nature of the threat is better understand, and the uncertainty reduced, if analyzed on a global rather than national US scale.

Terrorism Alternative Risk Transfer

In the post-September 11 security environment, the possibility of the issuance of a terrorism catastrophe bond has been raised, but is it more than of academic interest? Given the notional risk ambiguity, a coupon of somewhere between 12% and 20% has been suggested as necessary to entice investors (Kunreuther [2002]). Such a high coupon range would make terrorism cover extremely expensive, perhaps prohibitively so. But this range does not reflect any underlying quantitative risk analysis, nor the potential for dicing the risk up into more affordable, more diversified, components.

Quantitative terrorism risk models may provide impetus for securitizing terrorism risk, or at least some parts of it. For example, one might conceive of a workers compensation terrorism catastrophe bond triggering well above the macroterrorism level, at 1000 employee fatalities. Depending on the territorial region and trigger threshold, this might

be competitively priced. Another prospect would be a catastrophe bond to cover life insurers against massive losses following an attack using a weapon of mass destruction.

A further opportunity which may develop, once terrorism risk is sufficiently quantified, is in the field of contingent finance. The economic fallout from another terrorism showpiece may leave some vulnerable corporations disappointed by their bankers. A prior guarantee of finance is thus worth having, at a price. It is not just the direct damage caused by another major al-Qaeda attack that is of concern, the consequent business interruption may ripple through the US economy. In setting alight his explosive-laden shoes on board a transatlantic jet, Richard Reid has admitted that he intended to bring down not just the plane, but the civil aviation industry as well.

The economic malice of al-Qaeda operatives, such as Richard Reid, engenders a very strong correlation between macroterrorism and declines of the financial markets. This was demonstrated by the sharp falls on Wall Street following the September 11 attacks, and the market depression preceding July 4th, when fears abounded of an Independence Day atrocity. So what asset manager, extending the efficient frontier between risk and reward, might be tempted to hold terrorism bonds? Possibly the sale of terrorism catastrophe bonds might be targeted at bearish investors or hedge funds adopting a 'black swan' strategy of selling the stock market short. As evidenced even from threat announcements, the impact of a macroterror attack should turn a profit for hedge funds who have purchased a large number of out-of-the-money puts. During lulls between macroterror attacks, the bond coupons could defray the cost of continuously maintaining their bearish positions.

REFERENCES

Abu Hamza. *Allah's governance over the Earth*. Deluxe printers, England, 2002.

Arquilla J., Ronfeldt D., Zanini M. "Information-age terrorism". *Current History*, 99 (2002), pp.179-185.

Bergen P.L. *Holy war, Inc.* Weidenfeld and Nicholson, London, 2002.

Cordesman A.H. *Terrorism, asymmetric warfare, and weapons of mass destruction.* Praeger Publishers, Westport, 2002.

Gunaratna R. *Inside Al-Qaeda*, C. Hurst & Co., London, 2002.

Hausken K. "Probabilistic risk analysis and game theory". *Risk Analysis*, 22, No.1, (2002), pp.17-27.

Hoffman B. “Re-Thinking terrorism in light of a war on terrorism”. Testimony to the House of Representatives, RAND Report, 2001.

Krebs. V.E. “Uncloaking terrorist networks”. *First Monday*, issue7_4., 2002.

Kunreuther H. “The role of insurance in managing extreme events: implications for terrorism coverage”. *Risk Analysis*, 22, No.3, (2002) , pp.427-438.

Major J. “Advanced techniques for modeling terrorism risk”. *Journal of Risk Finance*, 2002.

Paparone C.R., Crupi J.A. “Janusian thinking and acting”. *Military Review*, Jan.-Feb., 2002.

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

Woo G. *The mathematics of natural catastrophes*. Imperial College Press, London, 1999.