

QUANTIFYING TERRORISM RISK FOR INSURED PORTFOLIOS

Gordon Woo
AON Conference, June 2003

ABSTRACT

As with Hurricane Andrew in 1992, and the Northridge Earthquake in 1994, the terrorist attack on the World Trade Center on September 11, 2001, has resulted in major advances in the quantification and management of a class of catastrophe insurance risks. The control of accumulations of exposure in urban areas is a basic principle of insurance portfolio management, but quantitative terrorism risk assessment, which has evolved rapidly as a technical discipline since 9/11, has capabilities extending well beyond this primary beach-head objective. These capabilities are described here within a review of the mathematical concepts and numerical methods used for quantifying terrorism risk. Use of these methods provides insurers and reinsurers with insight into the relative likelihood and loss potential of specific attack scenarios, as well as an overall risk-based appreciation of their terrorism exposures.

TERRORISM MODELLING: CONCEPTS AND CAPABILITIES

It has been said of the Australian army's approach to 21st century modernization (Fry and Forsyth, 2002), that it is 'concept-led and capability-based'. With any complex technical challenge, the response should be led by establishing key concepts, and be based on a realistic understanding of practical capabilities. So it is with quantifying terrorism risk. In addressing this challenge, the starting point is the identification and elaboration of concepts relevant to quantifying terrorism risk. These differ from the standard physical and engineering principles which underlie the quantification of risk from natural hazards. Inevitably, the innovative concepts needed for terrorism risk assessment include game theory, which is the formal mathematical theory of conflict. Concepts from the control theory of stochastic systems are also useful in as much as counter-terrorism suppression of terrorist activity is a control process. Furthermore, the theory of nonlinear complex adaptive systems affords insight into the frustrating and elusive virus-like nature of terrorism, and social network theory is instructive for charting and forecasting the evolution and destabilization of terrorist networks.

The capabilities of a terrorism model reflect the degree to which these various concepts can be formalized and incorporated in an efficient way, and the extent of accessible information about the terrorist threat which governs the scope and reliability of model parametrization. Progress has been made in developing all these concepts, and in compiling information on terrorist activity, thus facilitating probabilistic terrorism risk assessment. The different facets of terrorism risk modelling are outlined below, beginning with the analysis of attack scenario loss modelling.

DETERMINISTIC SCENARIO LOSS MODELING

Coincident with the shock of the 9/11 Al Qaeda attack itself was the shock realization that an enormous aggregate insurance exposure, across many lines of business, was concentrated around a single location. Whatever the terrorist threat may be, insurers need to understand and control their exposure accumulations. Accordingly, Geographic Information System (GIS) software tools are being licensed which allow insurers to map their exposures, and to evaluate loss potential within designated spatial footprints. These losses may vary significantly from one footprint to another, so worthwhile insight is gained by compiling a loss-footprint table. Ranking these losses by severity leads inexorably to the critical actuarial issue of estimating probable maximum loss.

In earthquake insurance, a deterministic approach may be taken to estimate Probable Maximum Loss (PML) using the following three-step procedure: identify the fault posing the greatest threat to the portfolio; assign the maximum credible earthquake to the fault; calculate the portfolio loss assuming this sized event occurs on this fault. For terrorism insurance, this kind of deterministic PML approach may also be attempted, assuming that a maximum credible size of weapon is deployed at the spot where it would cause the worst portfolio loss. This deterministic approach largely removes the human behavioural component from PML estimation, since it assumes pessimistically that the terrorist will have the upper hand in his conflict with counter-terrorism forces, and be allowed the weapon and target of his choosing.

This conservative assumption reduces PML estimation to a series of problems in the domain of the engineering, physical, chemical and biological sciences: evaluating the blast effect of a bomb detonation; the extent of fire from a fuel tanker explosion; the radiation fall-out from a radiological dispersal device; the spread of contagion from a smallpox outbreak etc.. These problems may still be technically complex and challenging, but at least the core mathematical models for blast analysis; conflagration; atmospheric dispersion, pollution transport, epidemiology etc. are well established.

The models are founded on standard scientific principles, and the model results have some validation against observational or surrogate data. For bomb blasts in particular, data are available for past events, and sophisticated computer codes such as AUTODYN account for the reactive dynamic response of buildings. However, as with the vulnerability of individual buildings against earthquake or windstorm loading, lack of detailed information about a building's protection against a terrorist attack limits the resolution of site-specific loss estimation. Another limiting factor for workers' compensation and other casualty risks is the temporal variation in the size of population within the area under terrorist attack. Generic assumptions may be made for this, as for supply chain bottlenecks which may indirectly affect business interruption.

The uncertainty in scenario loss modeling is one reason why, as with natural hazards, a deterministic terrorism PML approach can only be partially satisfactory. Another cogent

reason is that the probability of extreme loss is not addressed. Detonation of a nuclear device, or the sabotage of a nuclear plant, might lead to massive losses, but these hypothetical contingencies should not dominate terrorism PML evaluation, since these are very unlikely, even if conceivable, scenarios. PML is best inferred from the tail of a loss exceedance probability distribution (Woo, 2002a), which can only be constructed through a probabilistic terrorism risk model. Assignment of probabilities to the terrorist attack scenarios is a task which has an intrinsic human behavioral dimension, and so requires a new set of mathematical modeling tools, and makes greater recourse to the elicitation of expert judgement than for natural hazards.

THE USE OF EXPERT JUDGEMENT

Prior to the Earth Science revolution of plate tectonics in the 1960's, earthquake risk assessment was predominantly judgement-based. There were historical catalogues on past earthquakes, but no adequate theory by which seismic phenomena might be properly understood. Even into the 1980's, professional seismologists, who monitored earthquakes, were sceptical of the objectivity of practical engineering seismic hazard analysis. But with the increasing power of computation and theoretical developments, earthquake modeling has become more quantitative, and less subjective. It is hard to avoid a fair measure of expert judgement in terrorism risk assessment, but minimizing subjectivity is key to the scientific evolution of terrorism risk modelling.

Since Al Qaeda operates in almost a hundred countries across the globe, the use of international experts is crucial. Terrorism is governed by intent, capability and opportunity. These key factors must be well researched by experts, who should have active contact with terrorists and a current appreciation of their modus operandi. One such expert, having all the requisite credentials, is Rohan Gunaratna, author of the book, 'Inside Al Qaeda', whose base in Singapore allows him particularly extensive coverage of terrorist operations in Australasia. In the wake of 9/11, he was called to address the United Nations, the US Congress and the Australian parliament.

Because each expert is privy to his own sources of intelligence often gained verbally from debriefings, and has his own security clearances, there is no common database of information upon which all experts can form their judgements. If there were a common literature, as in the sciences, relevant publications could be distributed to all experts, whose opinions might then be elicited on an individual basis. But the world of intelligence is opposite to that of science: the most crucial information is often the most confidential. Accordingly, expert judgement is well elicited through decision conferences, at which intelligence and other confidential information can be pooled and opinions shared. Where some experts are unable to attend, their opinions can be elicited via a Delphi procedure, which provides for feed-back on the opinions of other experts.

Transparency is a virtue in risk assessment, hence the most rigorous approach to any risk model development spurns the excessive use of expert judgement, and terrorism risk is no

exception. The use of expert judgement can be minimized through exploring and developing mathematical models and simulations of the underlying causative processes, which can then be parametrized from observational data. In the following sections, a review is given of mathematical concepts which have already found their way into advanced terrorism risk modeling, and an outline is given of additional ideas which currently are being researched.

STOCHASTIC MODEL FOR TERRORIST ATTACKS

Randomness plays a significant part in any human conflict. This is reflected in Bismarck's perceptive comment that when you draw the sword, you roll the dice. But there are causal factors as well, which shape the conflict landscape, including the temporal pattern of successful attacks. In constructing a stochastic model of terrorist attacks, these non-random factors need to be taken into account through invoking an appropriate methodological paradigm, such as cybernetics. Magnus Ranstorp, director of the Center for the Study of Terrorism and Political Violence at St. Andrews University, has referred to Al Qaeda operatives as parasites on globalization. In common with other prey-predator situations, the conflict between the forces of terrorism and counter-terrorism may be represented using the principles of cybernetics. In particular, the time development of the Al Qaeda conflict is a stochastic process which may be described by a controlled Markov chain model.

At any moment in time, the predator (e.g. Al Qaeda) is in some specific state of attack preparedness, whilst the prey (e.g. USA) is in some corresponding state of defense preparedness. In a democracy, there are rigorous checks and balances imposed on the law enforcement and security services. Accordingly, the counter-terrorism response has to be commensurate with the terrorism threat: draconian measures (e.g. detention without trial) are only tolerable when the threat level is high. Democracies are prevented constitutionally from mounting an unlimited war on terrorism.

A Markov chain is defined by the series of states that Al Qaeda occupies, and makes transitions to and from. This is a controlled Markov chain because, whatever state Al Qaeda occupies, the police and security forces counter the prevailing threat with actions which aim to control terrorism. These actions are commensurate with the threat, and hence are a function of the Al Qaeda state. Because of these controlling counter-actions, the process of attack occurrence is not Poissonian, as is generally assumed for natural hazards. In mathematical terms, these counter-actions are termed the *Markov feedback policy*. The Markovian concept of a system state is well suited to the fluctuating dynamics of the terrorism conflict, with the need for periodic updating of the threat situation. System states are distinguished from one another in respect of significant differences in the terrorists' organization, attack capability, and modus operandi. In some substantive degree, the threat parameters vary from one state to another. In increasing threat order, the alternative states of the terrorist network range from destabilization, to facility at launching conventional attacks, to capability of attempting attacks using weapons of mass destruction.

The term *macroterrorism* has been coined to describe a spectacular act of terrorism, (which may be a multiple strike at several locations), which causes large economic and/or human loss. Minor ‘potboiler’ terrorist acts, such as house bombing, may occur haphazardly, but the occurrence of spectacular macroterrorism events, which are deliberately intended to cause massive loss, does not satisfy the prerequisites of a Poisson process. Once a terrorist’s message has been delivered successfully across the media through a spectacular event, perhaps after a series of failures, a publicity reminder may not be needed for a while. On the counter-terrorism side, following an act of macroterrorism, security and border controls are inevitably strengthened, and extra government funding made available for improving protection. Civil liberties may be temporarily curtailed as suspects are detained without trial, and the human rights pleas of asylum seekers and refugees may be denied. Such heightened security is a deterrent to another spectacular attack, but as time elapses uneventfully after a spectacular attack, security tends to be relaxed, making a further attack then more likely. The result is a cycle of terrorism, which is a notable feature of terrorist campaigns, one that can be modelled using mathematical methods from economics (Faria, 2003).

ADAPTIVE LEARNING 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*. The notion that this principle may guide the probability distribution of certain human actions was originally developed by Zipf in his quantitative sociological studies, and may be considered in the context of attack mode preferences. One of the main signposts on the path of least resistance is adaptive learning. Al Qaeda is eager to learn from past terrorist experience – the successes and failures of attacks perpetrated by its own network, and by other terrorists around the world. Al Qaeda would tend to ‘copycat’ methods which either have proven to be successful, or are perceived to have the potential to be successful. If an attack mode has demonstrated effectiveness, or has the promise of being effective, it is likely to be an attack option. Statistical learning models may be more relevant than pure frequency models in quantifying attack mode likelihood.

The basic arsenal for terrorists contains a range of conventional weapons: improvised explosive and incendiary devices, and standard military weapons such as automatic rifles, grenades, mortars, and surface-to-air missiles. Sticking with off-the-shelf or tried-and-tested weapons might seem to be the easiest strategy, but substitution of alternative weapons may be forced if such weapons become harder to procure. In any case, further variety in attack modes is necessary from time to time as it keeps counter-terrorism forces guessing.

This necessity leads to the invention of unconventional attack modes: industrial, infrastructure and agricultural sabotage, hijacked jets, helicopters and ships, bomb-laden boats and planes, chemical-biological-radiological-nuclear (CBRN) weapons, cyberspace hacking, food and drink contamination etc..

The process of terrorist attack mode selection can be simulated as follows. At any given time, there is a small probability that a new terrorist attack mode will be chosen, and a complementary probability that one of the existing attack modes will be chosen. In keeping with the principle of adaptive (copycat) learning, the relative likelihood that a specific existing attack mode will be preferred may be assumed to be an increasing function of the amount of its previous usage. The more often an attack mode has been used, the more likely it is to be re-used in another terrorist operation. This usage growth pattern is common to a number of sociological contexts, where this type of stochastic growth modeling has proved instructive. An outcome of this simulation is insight into the empirical probability distribution of attack mode preferences: some of the key modes dominate the distribution, with a long tail of other ancillary attack modes.

DEVELOPMENTS IN TERRORISM SIMULATION

The simplest approach to modeling a conflict is by considering the interplay between two opposing force blocks. This is the approach described above, with the terrorist organization opposing the counter-terrorism organization. But extra insight into the dynamics of a terrorist network can be gained by looking inside: analysing the social network of inter-connections between network nodes, which correspond to individual terrorists. The French magistrate, Jean-Louis Bruguière, has aptly likened Al Qaeda to a virus. In order to survive, a virus must mutate faster than its environment changes. Similarly, the Al Qaeda network has shown flexibility in adapting to survive counter-terrorism action. This adaptation process can be simulated by evolving the social network according to a set of basic rules.

Nodes communicate with one another to exchange information, financial and logistical resources, subject to the risk that any communication might be detected by security services. Local cells are autonomous to a substantial degree, and recruit attack team members and carry out target reconnaissance. Spectacular attacks are planned, but the larger and more ambitious that an attack becomes, the higher the chance of it being compromised by one of the attack team. If any node is removed from the network, there is a chance that any node connected to it might also be named and removed. Thus, the more hierarchichal the network, the greater the chance of destabilization through the arrest of senior leaders.

Through repeated computational network simulation following these rules, an ensemble of different network evolutions can be generated, the contrasting patterns of network structure can be studied, and the relative likelihood of different network configurations can be gauged. From analysis of these simulations, cell statistics and the network capability to launch major multiple attacks can be assessed probabilistically.

Such network analysis has to cope with the problem of missing data. As in the war on terrorism, massive amounts of uncertainty and dearth of data plague decision-makers on the military battlefield. Where should we attack? When should we attack? What weapons should we use? These are some of the critical questions facing military leaders. In this parallel warfare context, battle decisions might just be left to the judgement of generals, rather as terrorism insurance decisions might be left to the judgement of underwriters. Creditably, instead of an air of technical resignation pervading the Pentagon, massive investments of resources are being made to provide quantitative decision-support tools for the military. Sophisticated methods of combat modeling are being developed which incorporate all manner of extraneous factors that impact upon the decision-making processes of battlefield commanders. This wargaming has been substantially advanced through a variety of quantitative means: mathematical equations, also large-scale simulations, and distillations (i.e. relatively simple simulations that capture the salient features of the situation, without trying to model all the details). Similar bottom-up combat modeling initiatives are being explored for the war on terrorism. One of the purposes of the analysis is to identify emergent dynamic behavior, such as might arise if there were a concentration of terrorist resources in a particular threat mode, as defined by choice of weaponry, target and mode of attack delivery.

TERRORIST TARGET SELECTION

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. As with burglar alarms, self-protection has the externality of shifting risk to one's neighbours. This effect may be recognized explicitly using the mathematical theory of conflict, i.e. game theory, which is a collection of tools designed to help understand the interaction of decision-makers.

The two fundamental precepts underlying game theory are that the protagonists are rational and intelligent in strategic reasoning. These are justifiable for macroterrorism. As a weaker force confronting a nation state with far greater military and economic resources, a terrorist organization needs to have a smart strategy to survive and launch spectacular attacks: terrorists poor in strategic reasoning fade rapidly into oblivion. Indeed, Dr. George Habash, co-founder of the Popular Front for the Liberation of Palestine, referred to terrorism as a thinking man's game. Like Dr. Habash, Dr. Ayman Al-Zawahiri, the Al Qaeda chief strategist, was an eminent doctor before turning to terrorism, and noted for his brilliance.

In applying game theory to terrorism, it is important to leave behind popular notions of rationality, and to return to the formal mathematical definition of rational behavior, namely that actions are taken in accordance with a specific preference relation. There is no requirement that a terrorist's preference relation should involve economic advantage

or financial gain. Much of the purpose of terrorism is psychological: inspiring the global Jihad; whipping up malicious joy at seeing the USA suffering loss; and terrorising the general public. Nor is it necessary that a terrorist's preference relation conform with those of society at large. Game theory is not restricted to any one cultural or religious perspective.

The test of any mathematical risk model is its explanatory and predictive capability. Game theory predicts that, as prime targets are hardened, rational terrorists will tend to substitute lesser softer targets. This prediction is essentially equivalent to the statement made by the CIA director, George Tenet, in his prophetic unclassified but unheeded testimony of February 7, 2001, (prior to 9/11): 'as security is increased around government and military facilities, terrorists are seeking out softer targets that provide opportunities for mass casualties.' Indeed, by the time of this statement, security had been increased so as to thwart attacks against the US embassies in Albania, Azerbaijan, Ivory Coast, Tajikistan, Uganda and Uruguay, which, like the US East African embassies bombed in 1998, lacked modern security. It is to be expected that, in response to a perceived threat to public property, a government will sense a public duty to take protective measures, even if the implications for private property and ordinary citizens are not necessarily thought through. Being symbolic of American global domination, US embassies and consulates around the world are prime terrorist targets. Hardening these targets has resulted in attacks being deflected elsewhere.

In the language of terrorism experts, this phenomenon is called target substitution. A year after the destruction of the World Trade Center, the Bali bombing on 12th October 2002 provided another tragic confirmation of this game theory prediction. A bomb left at the US consulate perimeter fence was too distant at 100 meters to cause any damage, but a bomb-laden truck could park immediately outside a nightclub and kill hundreds of tourists. A few days earlier, on 6th October, the French tanker, Limburg, was holed off the Yemen coast, an attack reminiscent of that on the USS Cole, except that in 2002 US warships had become too hard targets. In May 2003, lightly protected residential compounds housing expatriates were attacked in the Saudi capital Riyadh. Explicit admission of this soft target strategy has come from Khalid Sheikh Mohammed, the Al Qaeda chief of military operations, who was arrested in March 2003. Further validation of the terrorism target prioritization model is provided by analysis of the IRA campaign in Ulster and England, and the GIA campaign in France.

GLOBAL MODEL OF TERRORISM

The success of this game theory model of target selection illustrates the future potential for quantitative terrorism model development. One such development is enlargement of the geographical region of modelling, specifically construction of a global model of terrorism covering the threat from the global Al Qaeda network. Rather as has happened with international drugs trafficking, the tightening of border controls and internal security within one country may divert terrorism to another country, which is less protected against terrorist attack.

Another possibility is that terrorists may exploit lax security in one country, (or, as in Canada, the vigorous upholding of human rights) , and use it as a haven from which to attack other countries. In this situation, a form of prisoner's dilemma arises. If a country unilaterally takes action against terrorism, it may be open to reprisals, which may be costly. However, joint action by neighbouring countries may benefit all. Mistrust between countries may result in inadequate action being taken against terrorism.

The suicide attacks in Casablanca, Morocco, in May 2003 highlight the geographical diversification of Islamic militancy. For a muslim living in a poor urban environment, the promise of eternal paradise as a martyred suicide bomber may seem an attractive rational option to a life of perpetual urban deprivation. As in Casablanca, so also around the world, mosques may serve as recruitment centres for Osama bin Laden's Jihad against Jews and Crusaders. In London's unfashionable Little Algeria district, the Finsbury Park mosque has been associated with such activity. Richard Reid, the airplane shoe-bomber, attended this mosque, where he found fellowship otherwise denied a mixed-race dropout from British society. This mosque was raided by the police shortly after the neighbourhood discovery of ricin in January 2003. Algerian asylum-seekers were making this deadly toxin, but the mastermind behind the operation was based not in Europe, but in the Pankisi Gorge between Chechnya and Georgia, to which lawless area a number of Al Qaeda operatives have fled from Afghanistan.

The training camps in Afghanistan have gone, but their terrorist legacy remains in the guise of thousands of Al Qaeda operatives around the world, who have attended these camps, and dispersed to sleeper cells in many host countries. In the Middle East, as well as USA, Britain, France, Australia, and other named target countries, there are thousands of potential targets for Al Qaeda. But for government self-protection of national landmarks, the range of plausible targets would be far narrower. As it is, through increased government security, risk is transferred from the public to private sectors. As an example, in London, in the aftermath of the Iraq war, concrete barriers were erected around parliament and other prominent public buildings to deter suicide bomb attacks. As a consequence, a truck bomb attack in London is more likely to be directed at a commercial or residential target, resulting in property or casualty insurance loss.

In astronomy, the darkness of the night sky once seemed paradoxical, given the illumination from myriad stars. In terrorism, a similar paradox arises: why is the surface

of the Earth not constantly lit up by spectacular terrorist attacks, given the presence of thousands of terrorist targets? Minor potboiler attacks, which are not intended to cause significant loss, occur randomly, are inexpensive to launch, and are common. However, the overall global number of spectacular terrorist attacks per year is limited to single figures by logistical resource constraints, and also the infrequent need to issue public reminders that a terrorist organization is still in business: the media reverberation time for a spectacular Al Qaeda attack is measured in months, rather than weeks.

If security levels in target countries were adjusted approximately so as to equalize risk around the world, the annual probability that any one target would be attacked would be extremely small. This is reminiscent of the savannah where a herd of antelope may be hunted by a lion, which requires one kill per day. The chance that any individual prey will fall victim to the predator is very small – there is safety in numbers. Similarly, global security may equilibrate to a situation where terrorism risk is spread very broadly, but thinly, across the world. From Singapore in the east to Los Angeles in the west; from London in the north to Melbourne in the south, cities are exposed to terrorism risk.

CONCLUSION

With the Cold War ended, the terrorism threat from Islamic militants has been called World War IV by James Woolsey, the former CIA director. Already, wars in Afghanistan and Iraq have been waged by the USA in retaliatory response to this perceived threat. Robert Kagan (2003) has elucidated the different roles that America and Europe play in the new world order which is unfolding: Europe plays Venus to America's Mars. Only America has the stature and power to deal with threats militarily. But if US action is taken against Syria or Iran, Hezbollah may be incited to retaliate via terrorist attacks in the US homeland. Already, reconnaissance missions are being undertaken on potential targets. Unlike Al Qaeda, Hezbollah will seek to launch attacks on a proportionate scale, so the use of weapons of mass destruction is unlikely.

The price of Middle Eastern intervention by the USA is likely to include increasing hostility amongst muslim populations around the world. Islamic extremists form the apex of a pyramid of global muslim discontent at US military, economic, technological, and cultural domination (Revel, 2002). As security within the USA is increased, so countries sharing the same democratic tradition and religious heritage as the 'Great Satan' are liable to be substituted as terrorist targets. Consequently, the need for prudent insurance management of Islamic militant terrorism seems assured across the western world. In servicing this need, the methods for quantifying terrorism risk for insured portfolios will continue to evolve with increasing mathematical and computational sophistication, in the same way that the methods for quantifying windstorm and earthquake risk have systematically advanced over the past decade since the devastation caused by Hurricane Andrew and the Northridge earthquake.

REFERENCES

- Faria J.R. 2003. Terror Cycles. Studies in Nonlinear Dynamics and Econometrics, 7(1).
- Fry A., Forsythe A. 2002. The Australian Army and Project Albert: Pursuing the Leading Edge of Military Thinking and Technological Development. In, Horne G., Johnson S. (Editors) Maneuver Warfare Science, USMC: 1-16.
- Gunaratna R. 2002. Inside Al Qaeda. C. Hurst & Co., London.
- Harris J.W. 2002. Building Leverage in the Long War. Policy Analysis 439: 1-14.
- Kagan R. 2003 Paradise and Power. Atlantic books, London.
- Revel J.-F. 2002. L'Obsession Anti-Américaine. Plon, France.
- Woo G. 2002a. Natural Catastrophe Probable Maximum Loss. British Actuarial Journal, Vol.8, Part V: 943-959.
- Woo G. 2002b Quantitative Terrorism Risk Assessment. Journal of Risk Finance, Vol.4, No.1: 7-14.