LONE WOLF TERRORISM AND SPREES OF VIOLENCE

By
Dr Peter J Phillips[1]
University of Southern Queensland

The purpose of this paper is to apply economic science to the analysis of the opportunities and choices of single individual 'lone wolf' terrorists whose attacks are characterised by 'sprees' of violence that last only for a relatively short period of time. Unlike 'serial' lone wolves who engage in violent terrorism over a prolonged period of time and unlike those lone wolves who engage in a single crudely planned terrorist attack and are apprehended without inflicting significant injuries or fatalities, the attacks of spree lone wolves are concentrated in a very short period of time, perhaps no longer than several days and sometimes as little as a few hours, and may generate significant levels of human tragedy. The spree lone wolf also emerges suddenly. Having previously allocated no resources to violent terrorism, he suddenly and all at once allocates 100 percent of his resources, including time, to violent terrorism. The first step to providing guidance to governments and their security and law enforcement agencies is to encompass some important elements of the spree lone wolf's opportunities and choices within an economic analytical framework. The first steps towards this encompassment are undertaken in this paper by exploring the opportunities and choices of the spree lone wolf from a risk-reward perspective and a treatment of the spree lone wolf as an individual who, while attempting to maximise his expected utility, shuns the risk-reduction benefits of 'time diversification' and suddenly plunges all of his resources into violent terrorism within a single time period.

Key Words: lone wolf, terrorism, violent terrorism, expected utility, spree, serial, plunge.

Word Count: 12,600

Note: At the time of writing (December 2011), this paper (or a revised version) is intended for presentation at the 12th Jan Tinbergen European Peace Science Conference in June 2012, Berlin. This paper will also be included in my new book on lone wolf terrorism (in manuscript form at the time of writing).

[1] Corresponding Author: Peter J Phillips, School of Accounting, Economics and Finance, University of Southern Queensland, Toowoomba, Queensland, 4350. Telephone: 617 46315490. Email: phillipsp@usq.edu.au.

LONE WOLF TERRORISM AND SPREES OF VIOLENCE

A lone wolf terrorist operates independently, alone, outside of a command structure. Some lone wolves have been able to inflict greater amounts of human tragedy than some terrorist organisations. Lone wolf terrorists are dangerous. Because of their independence, the pre-emption of lone wolves by law enforcement is extremely difficult. Lone wolves are not a part of a terrorist organisation or a part of a terrorist network. There is nothing for law enforcement agents to infiltrate. After the lone wolf strikes, he reveals his existence and evidence may be gathered by law enforcement. With each attack, more evidence accumulates and the unique features of the attacks begin to point towards a unique offender. This approach to the law enforcement and security operations pertaining to the pursuit of the lone wolf is well suited to the pursuit of the 'serial' lone wolf who strikes, withdraws and strikes again over an extended period of time. The 'spree' lone wolf is different and represents a different challenge to law enforcement. The terrorist actions of the spree lone wolf are concentrated within a very short period of time. Law enforcement efforts directed towards the pursuit of the 'spree' lone wolf must be condensed and investigative efforts must proceed on the basis of either limited evidence or rapidly evolving and accumulating evidence. In some cases there may be no time for an investigative process at all.

If defence economics is to contribute to the pre-emption or pursuit of the spree lone wolf, the spree lone wolf's opportunities and choices must be delineated. The spree lone wolf's opportunities are characterised by the payoffs that he may expect from engaging in violent terrorism and the risks that attend each of the expected payoffs. The lone wolf's choices are characterised by an attempt to optimise his expected payoffs at the level of risk that he is willing to bear. A lone wolf who engages in violent terrorism must choose an attack method. Each attack method is expected to inflict a particular level of human tragedy: injuries and fatalities. Each attack method is risky. The actual level of human tragedy that is inflicted when a particular attack method is used may be greater or less than that which was expected. For example, a bombing device may detonate prematurely and inflict more injuries and fatalities than the terrorist expected; or the terrorist might be apprehended whilst planting a bombing device; or it may fail to detonate; or fewer people are present at the targeted location than expected. When the spree lone wolf chooses an attack method, he chooses an expected payoff and risk. The lone wolf's opportunities are his feasible combinations of expected payoff and risk. The lone wolf's choices are his attempt to choose the attack method that will yield the highest expected payoff at the level of risk he wishes to bear.

The difference between a serial lone wolf and a spree lone wolf is the time-concentrated nature of the deployment of the chosen attack methods. If the choices of the spree lone wolf are to be understood, this aspect of his behaviour must be encompassed within an economic analytical framework. As

previous work has shown[2], the opportunities and choices of serial lone wolves can be encompassed within an economic analytical framework that depicts the lone wolf as an individual who attempts to optimise expected utility on the basis of the expected payoffs and risks that characterise each of the available attack methods or attack method combinations. To ensnare the spree lone wolf within this framework requires an exposition of the circumstances or conditions under which a lone wolf terrorist will 'plunge' all of his resources, including time, into violent terrorism. What is also required is an exposition, drawn from economic theory, of the behavioural aspects that will characterise this type of lone wolf terrorist. Once this is accomplished, further conclusions may be reached and inferences may be drawn that will help to shed some light on the structure of the opportunities and choices that underlie the phenomenon of 'spree' lone wolf terrorism.

THE EXPECTED PAYOFFS AND RISKS OF DIFFERENT TERRORIST ATTACK METHODS

RAND identifies ten attack methods that have been deployed by transnational terrorists over the past four decades: armed attacks, arson, assassination, hijacking, hostage-taking, kidnapping, bombing, unconventional attacks, 'other' and 'unknown'. Each known form of terrorist attack method may be placed within one of these categories. If particular types of terrorist attacks emerge as distinct categories over time, they may be designated to a new category. For example, the 9/11 terrorist attacks are currently categorised under the heading of 'unconventional attacks'. However, if such a type of attack were to recur with greater frequency and become both distinct and conventional a new attack method category may be formed. The RAND categories have emerged with the historical evolution of transnational terrorism. They are complete with reference to known forms of attack. Innovations in terrorist attack methods, to the extent that they occur, will shape the RAND categories over time.

Among transnational terrorist incidences, the most commonly used attack method is bombing. Between 1968 and 2007, there were 17,856 transnational terrorist incidences involving the utilisation of bombing. The second most commonly utilised attack method is armed attacks. There were 7,114 incidences of armed attacks during the period. Unconventional attacks such as the 9/11 terrorist attacks were the least commonly used attack method. RAND records only 56 unconventional attacks for the entire period. This indicates that for the most part terrorists deploy tried, tested and more common attack methods. These are also the attack methods for which terrorists may have a greater knowledge of the distribution of injuries and fatalities. These more common types of attack methods have proven to be effective at causing a level of human tragedy. Bombing attacks generated the highest number of fatalities and injuries. During the period 1968 to 2007, there were 120,550 fatalities

[2] Phillips (2011).

and injuries from transnational terroristic bombing attacks. Armed attacks resulted in 26,310 fatalities and injuries between 1968 and 2007. Unconventional attacks resulted in 5,444 fatalities and injuries. Unconventional attack methods have generated the highest level of human tragedy proportionally with the number of incidences of unconventional attacks that have been observed. Just 56 attacks resulted in more than 5000 injuries and fatalities. On the same calculation, hostage-taking ranks second. Just 210 incidences of transnational hostage-taking have generated more than 3000 injuries and fatalities.

Table 1 Attack Methods and Human Tragedy

Attack Type	Total Number of Incidences 1968 to 2007	Total Fatalities and Injuries 1968 to 2007
Armed Attacks	7114	26310
Arson	1019	709
Assassination	2357	4282
Hostage	210	3123
Bombing	17856	120550
Hijacking	229	858
Kidnapping	2164	1641
Other	168	578
Unconventional	56	5444
Unknown	404	836

A terrorist who chooses to engage in violent terrorism and deploy one or more of these attack methods must expect some level of human tragedy to result. Even the non-overtly violent attack method of hostage-taking must be expected to inflict injuries and fatalities. The number of injuries and fatalities that can be expected to result from the deployment of one of the attack methods may be approximated by the arithmetic mean of the injuries and fatalities that have been generated per attack per year for each attack method as reflected in the historical time-series data reported by RAND. Each time a terrorist chooses to deploy a particular attack method he faces the prospect that the actual injuries and fatalities that he inflicts will diverge more or less from the expected (mean or average) amount of human tragedy. This is the risk that he must bear. This risk can be measured by the variance or standard deviation of the injuries and fatalities historically generated per attack per year for each attack method. The greater the variance or standard deviation, the greater is the chance that the deployment of a particular attack method will result in a magnitude of injuries and fatalities that diverges from that which the terrorist had expected.

Table 2 Expected Injuries and Fatalities (Payoffs) and Risks

Attack Type	Average Fatalities and Injuries Per Attack Per Year	Standard Deviation
Armed Attacks	5.32	15.00
Arson	0.72	1.85
Assassination	1.54	0.71
Hostage	11.46	38.18
Bombing	5.88	6.09
Hijacking	3.91	10.82
Kidnapping	0.46	0.35
Other	1.14	3.02
Unconventional	7.48	42.94
Unknown	1.11	4.04

Each of the attack methods is characterised by an expected payoff and risk. The availability of the various categories of attack methods with their attendant expected payoffs and risks presents the terrorist with a set of opportunities. His opportunities are, fundamentally, opportunities to obtain a particular expected payoff in return for bearing the associated risk. His opportunities are expected payoff-risk 'pairs'. There are no risk-free attack methods. In order to obtain expected payoffs, the terrorist must bear risk. The more risk that he is willing to bear, the more payoff he can expect. Of course, with higher risk comes the increased chance that the actual outcome of his attacks will be different from that which he expected. Hostage-taking and unconventional attacks are characterised by very high risk but have very high expected payoffs. Assassination, arson and kidnapping are characterised by relatively low risk but have relatively low expected payoffs. The lone wolf who contemplates the deployment of any of these attack methods confronts an inherent risk-reward trade-off. It is a trade-off that he cannot escape. If he wishes to generate higher expected payoffs from his engagement in violent terrorism he must bear greater risk.

THE ATTACK METHOD CHOICES OF THE LONE WOLVES

The attack method choices of the lone wolves have not been distinctly different from the choices of the transnational terrorists[3]. Armed attacks and bombings, the attack methods most commonly chosen by transnational terrorists, are also the most common choices of attack method among lone wolves. In the United States, at least 32 lone wolf terrorists were actively engaged in violent terrorism in the period between 1940 and 2007. The majority of these lone wolves—28 of the 32—inflicted injuries and fatalities. In Europe, 39 lone wolf terrorists were active between 1968 and 2007. Of the 71 American and European lone wolf terrorists engaged in violent acts of terrorism, armed attacks and

[3] Instituut voor Veiligheids en Crisismanagement (2007).

bombings constituted the attack method choices for 49 of the terrorists. Other, less popular, attack methods include hijacking or hostage-taking (chosen by fourteen of the European lone wolves), assassination (chosen by three of the lone wolves) and arson (chosen by four of the lone wolves). In deploying these attacks, the lone wolf can expect to inflict injuries and fatalities. The actions of the transnational terrorists have produced a historical record of injuries and fatalities that have been associated with each of the attack methods. Over a long period of time an average or expected number of injuries and fatalities has come to characterise each attack method. The extent to which the outcomes of a particular attack may diverge from the outcomes that were expected is reflected by the variance or standard deviation which characterises the historical time series of injuries and fatalities. The lone wolf cannot escape from this inherent trade-off between expected payoffs and risk.

The trade-off between the injuries and fatalities that can be expected and the risk that the actual deployment of a particular attack method will result in more or less human tragedy has proven to be adequate to entice individuals to engage in violent lone wolf terrorism. The result has been a long list of fatalities and an even longer list of injuries. Between 1940 and 2007, the 32 lone wolf terrorists actively operating within the United States were responsible for 87 fatalities and 322 injuries. To place this in perspective, the 2nd of June Movement, predominantly deploying the same attack methods of armed attack and bombing, inflicted 5 fatalities and 3 injuries in a decade-long engagement in violent terrorism. The Army of God, predominantly deploying bombing and assassination, inflicted 3 fatalities and 124 injuries in a fourteen-year engagement in violent terrorism. Black September, predominantly deploying bombing, inflicted 77 fatalities and 121 injuries in a six year engagement in terrorism that involved striking targets in cities across the world, including New York, London, Cairo, Geneva, West Berlin and Rome. The Red Army Faction inflicted 16 fatalities and 62 injuries during a thirty-year-long campaign. There are, of course, many other terrorist groups who engaged in violent terrorism but who inflicted no injuries or fatalities[4]. Well-organised and financed terrorist organisations have sometimes been no more dangerous than a relatively small number of lone wolves operating alone and outside of a command structure.

Lurking behind the macro-statistical data are important differences between the lone wolves. Although the attack method choices are relatively uniform or, at the very least, can be placed within one of the RAND attack method categories, the lone wolves remain idiosyncratic in many of the ways in which they deploy these attack methods. A lot of these idiosyncrasies are absorbed by the fact that all attack methods are fundamentally characterised by an expected payoff and risk. Regardless of the uniqueness of the lone wolf's approach, his opportunities and choices are characterised by 'pairs' of expected payoffs and risks. An important idiosyncrasy that is not completely transcended by this

[4] All data are from the Global Terrorism Database (GTD) or RAND.

characterisation is the difference between the 'spree' lone wolf and the 'serial' or 'standard' lone wolf terrorist. An archetypal serial lone wolf is Theodore Kaczynski, a.k.a. 'The Unabomber'. Between 1978 and 1995 Kaczynski mailed sixteen parcel and letter-bombs to various targets around the United States. The attacks occurred intermittently over this seventeen year period with large periods of time passing without an attack. Despite the intermittent nature of the attacks, Kaczynski persistently pursued his campaign for almost two decades. The stable structure of the evidence produced by his engagement in terrorist violence permitted the construction by the Federal Bureau of Investigation of an offender profile. Evidence gradually accumulated and pointed towards a single offender. This stands in direct contradistinction to the actions of those lone wolf terrorists who have condensed their engagement in violent terrorism into a very short period of time. These are the lone wolf terrorists who have engaged in 'sprees'.

A shooting spree is the type of attack that immediately springs to mind. This is, no doubt, due to the popular usage of the word 'spree' in describing various acts of concentrated violence, especially those involving the use of firearms. However, a spree need not only involve shooting or armed attacks and it need not take place all on one day. A spree might involve bombing and the acts of violence may be continuous over a period of several days. If this definition is accepted, both Joseph Christopher and David Copeland are examples of the 'spree' lone wolf. Copeland engaged in a 13-day bombing spree in London in 1999. He planted bombs in crowded areas on Saturday April 17, Saturday April 24 and Friday April 30. The bombs, laden with nails to maximise the number of injuries and fatalities, resulted in three fatalities and 129 injuries. Although Christopher has characteristics that could define him as a serial lone wolf[5], four of his fatalities were inflicted during a three-day shooting spree in New York State in September 1980. The challenges presented by spree lone wolves to investigators are clearly evident in both of these cases. Neither Christopher nor Copeland was apprehended at the scene of their attacks. Christopher might not have been apprehended at all but for the fact that he later bragged about his crimes. Copeland was identified by a colleague who recognised him from the security-camera vision police had released to the public. This information was provided to police just before the final attack—which killed three and injured 79. Even though police now knew who the offender was, his last deadly bombing could not be prevented[6].

[5] Christopher is known as both the '0.22 Calibre Killer' and 'The Midtown Slasher'. After his initial spree, Christopher changed his preferred murder weapon to a knife, which he used to attack his victims and stab them to death.

[6] These facts were obtained from the relevant Wikipedia entries and Instituut voor Veiligheids en Crisismanagement (2007) and checked against the original news sources.

The archetypal modern example of the spree lone wolf is almost certainly Anders Behring Breivik. In Norway in July 2011 Breivik engaged in a bombing and shooting spree in that resulted in the deaths of 77 people. The attacks were undertaken in two parts. First, a car bomb was left outside the government building that housed the offices of the Prime Minister. The bomb was detonated and killed eight people. This is exactly what would be expected from an 'average' bombing attack. Breivik then travelled approximately 40 kilometres to a youth camp where he murdered another 69 people in an hour-long shooting spree. For a terrorist attack that might be categorised as an 'armed assault', 69 fatalities is many standard deviations from the mean number of fatalities historically generated in such attacks[7]. Although many details remain to be confirmed, on the basis of currently available information, Breivik's attacks represent an archetypal example of lone wolf terrorism and a rare example of combination of attack methods by a spree lone wolf. By combining attack methods, the lone wolf terrorist can obtain increases in the expected numbers of injuries and fatalities without bearing additional risk (Phillips 2009; 2011). This is a property of the correlation structure that characterises the expected payoffs to different attack methods. In this particular case, Breivik's combination of attack methods appears also to have contributed to the increase in expected injuries and fatalities by diverting law enforcement attention to the location of the bombing whilst the shooting spree took place some distance away. Whether this was planned or, as has been reported[8], the shooting spree was undertaken after it became clear that the car bomb did not collapse the government building, may remain unknown.

Bombing and armed attacks are the dominant attack methods of choice for the lone wolf terrorist. When these attack methods are deployed in a concentrated period of time, the attacks may constitute examples of spree lone wolf terrorism. Spree lone wolves represent unique challenges to law enforcement agencies charged with their pre-emption or apprehension. Because of the time-concentrated nature of the attacks, the investigative process evolves exceedingly quickly. Even if law enforcement agencies become aware of the identity of the lone wolf terrorist, there may not be enough time to intercept him before he strikes again. Even if law enforcement become aware that a 'spree' attack is in progress, many injuries and fatalities may be inflicted before the law enforcement response effectively neutralises the spree lone wolf terrorist. Despite the difficulties presented by spree terrorism, the decision-making framework of the spree lone wolf reduces to one that is

[7] What might be required is some estimate of an expected number of injuries and fatalities from shooting sprees in confined locations, such as schools, campuses or youth camps. The injuries and fatalities inflicted by lone shooters (though not necessarily lone wolf terrorists) have sometimes been very substantial. The mean and standard deviation of the injuries and fatalities inflicted by such specialised armed attacks may represent the statistical foundation for a new attack method category.

[8] AFP, November 19 2011.

characterised by opportunities and choices. Both opportunities and choices can be further delineated within a risk-reward (expected payoff-risk) analytical framework that may begin to encompass the spree lone wolf. Once encompassed to some degree within an economic analytical framework, further insights into the behaviour of the spree lone wolf may be obtained. Ideally, these insights will provide advice that is operationally relevant to law enforcement agencies engaged in the pursuit of spree lone wolf terrorists.

THE LONE WOLF: A MODEL OF CHOICE

Building upon the economic analysis of crime and criminal behaviour (Becker 1968; Ehrlich 1973), the terrorist is usually depicted as attempting to maximise an expected utility function (Landes 1978; Sandler, Tschirhart and Cauley 1983; Enders and Sandler 1993). Factors that are salient to the particular choice problem confronting the terrorist or terrorist organisation are included within an expected utility function. Some of these factors contribute positively to utility. For example, monetary payoffs from ransoms or (non-pecuniary) concessions granted by the government contribute positively to the terrorist's utility. Some factors contribute negatively to utility. For example, the penalties imposed on terrorists who are apprehended. If, for a particular terrorist activity, the chances of positive payoffs increase or the absolute amount of the positive payoffs increase, terrorists might be encouraged to participate in that activity. If, however, the chances of failure increase or the absolute magnitude of the negative payoffs (penalties) increases, terrorists might be deterred from participating in the activity. This 'deterrence' effect, which is derived from economic analysis, operates alongside a 'substitution' effect. The terrorist who is deterred from one activity may substitute for it another type of activity. The careful exposition of the simultaneous operation of the deterrence and substitution effects is one of the fundamental contributions of defence economics to the study of terrorism and terrorist behaviour. The analysis is well-supported by empirical evidence[9].

A simple generic expected utility function for the terrorist can be constructed as:

$$U_t = E\sum_z u(z)\pi(z) \tag{1}$$

Where *E* is the expectations operator, *z* is the expected payoff of the terrorist activity, $u(z)$ is a utility function that relates the expected payoff to the terrorist's utility and $\pi(z)$ is the subjective probability

[9] For example, the fortification of US embassies reduced direct attacks on these locations. However, there was an increase in assassinations and other types of terrorist activity directed towards embassy personnel outside of the fortified areas (Sandler and Enders 2004).

of an outcome (Phillips 2005). This generic utility function can be augmented to any degree of complexity in order to capture the salient features of different terrorism scenarios. For example, Landes (1978) was interested in hijacking. He constructed the (non-generic) expected utility function as:

$$EU = (1 - P_a)U(W_j) + P_a P_c U(W_i - S) + P_a(1 - P_c)U(W_j - C) \quad (2)$$

Here the hijacker's expected utility depends upon the probability of apprehension in country *i*, P_a, the probability (given apprehension) of being convicted and sentenced to prison, P_c, the hijacker's wealth, W_i and W_j, in each of two countries, *i* and *j*, the monetary equivalent of the sentence, *S*, in country *i* and the monetary costs, *C*, associated with apprehension but no sentence (example, costs of defence). Similarly, Sandler et al. (1983) were interested in hostage-taking scenarios where terrorists made demands on the government. They constructed the expected utility function as:

$$U = U(L, D, \Pi; C^*) \quad (3)$$

The terrorist has an objective. This may, perhaps, be the promotion of a particular cause. There are two ways that this objective may be achieved. The first is by legal methods, such as lobbying or producing and distributing pamphlets. These activities have a payoff *L*. The second is by illegal or terroristic methods that involves taking hostages and making demands, *D*, of the government. The outcome of this risky activity is uncertain. The probability that the demands are met is Π and the most recent concession granted by the government is *C**.

In Sandler et al.'s model, the government also has an expected utility function:

$$G = G(C, \Omega, a_2; D^*) \quad (4)$$

The government's expected utility depends upon its concessions policy, *C*, the probability of re-election Ω, the amenities of office, a_2, and the demands of the terrorists, *D**. With these two expected utility functions in place, the economic analysis proceeds to the examination of various 'comparative statics'. This means that, holding all other variables constant, the logical outcomes of changes in one of the parameters are explored. The interaction between the terrorists and the government in a hostage-taking negotiation scenario yields a number of important comparative statics results within the Sandler et al. (1983) theoretical framework, including the 'deterrence' and 'substitution' effects.

Although important results have been derived from the application of orthodox expected utility analysis to the study of terrorism and terrorist behaviour, the results are not usually directly operationally relevant. This should not be taken to mean that the results are impractical but rather, to the extent that the results inform and guide law enforcement actions, they usually do so at a high level and do not penetrate to the operational level where such matters as the pre-emption of terrorist attacks at *particular* targets, the planning for *particular* types of attack methods that terrorists may deploy and the pursuit of *particular* individual terrorists occupy the attention of *particular* law enforcement agents. To derive directly operationally relevant results, the expected utility theoretical framework must be made to work in a different way such that computable results emerge from the analysis and the expected utility theoretical framework must be used in a different way such that inferences about individual terrorists or terrorist organisations can be drawn from a combination of the theoretical framework and the evidence that characterises the scene and nature of particular terrorist attacks. In this way, the application of the expected utility theory to the analysis of terrorism can be taken to the operational level as an additional tool that may be deployed in the investigative process.

Computable results can be obtained from a mean-variance expected utility framework. Within this framework, the terrorist's expected utility function contains just two arguments: (1) mean (expected payoff); and (2) variance (risk). The formal logic of choice operates in the same manner as the standard expected utility analysis but the analysis is less abstract in the sense that the complete choice set and the set of optimal choices is computable if the payoffs that characterise the opportunities and choices are empirically observable. Phillips (2009) shows how this approach generates an opportunity set of terrorist attack methods and attack method combinations and derives an optimal or 'efficient' set of attack method combinations that have the highest expected payoffs for each level of risk. Phillips (2009) measures the payoffs in terms of the injuries and fatalities generated by each attack method and combination. This approach has the significant advantage of equating the optimal choices of terrorists with their most dangerous choices. Critics of the optimisation process inherent within the economic analysis of choice must concede that even if the terrorist may make mistakes in his choices of attack methods, it is far better to be prepared for his most dangerous (optimal) choices than to rely on him making a mistake.

The generic mean-variance expected utility function may be expressed as:

$$U = f\left(E_R, \sigma_R^2\right) \tag{5}$$

Equation (5) says that the terrorist's expected utility is a function of the expected payoffs, *R*, and the variance or risk of those payoffs, σ_R^2. If the payoffs to terrorist attacks are measured in terms of injuries and fatalities[10], each terrorist attack method has an expected payoff approximated by the mean of the injuries and fatalities generated empirically by the deployment of the particular attack method. The risk that the terrorist must bear in order to obtain a particular expected payoff is reflected by the empirically observed variance of the payoffs. Bombs sometimes do not explode on time or at all. Terrorists are sometimes apprehended during an attack. Fewer (or more) people sometimes gather at the target location than anticipated. Police or security agents sometimes discover a bombing device before it detonates. And so on. The variance reflects the degree to which the actual outcomes diverge from those which were expected. When payoffs are measured in terms of injuries and fatalities, the variance reflects the degree to which the actual injuries and fatalities that are inflicted by a terrorist attack method may diverge from the number of injuries and fatalities that the terrorist expected. The actual number may be higher or lower than expected.

A mean-variance analysis based on the injuries and fatalities that empirically characterise RAND's attack method categories permits the delineation of the terrorist's opportunities and provides the foundation for an analysis of the terrorist's optimal choice. The complete opportunity set is simply the complete set of expected payoff-risk 'pairs' associated with every feasible attack method and attack method combination. Because of the statistical structure of the payoffs to each attack method, specifically the correlation structure of the payoffs, the terrorist's set of opportunities will be a convex set in expected payoff-risk space. The imperfect correlation that characterises the injuries and fatalities generated by each attack method ensure that a linear opportunity set is never observed and a concave opportunity set is impossible because such a geometrical property would be associated with more-than-perfectly-positive correlation. The terrorist's opportunity set must look something like this:

[10] This should not be a controversial approach. Violent terrorism inflicts injuries and fatalities. Terrorists, unless they are completely naive, expect injuries and fatalities to result from their violent terrorist actions. It is a small step, in light of both the words and deeds of terrorists which make it very clear that inflicting a maximum amount of human tragedy is immediate objective, to work from the position that terrorists attempt to maximise injuries and fatalities. The mean-variance expected utility approach would work equally well if the terrorist was supposed to 'optimise' injuries and fatalities which, in the case of assassination, could mean one single fatality.

Figure One The Lone Wolf's Opportunity Set

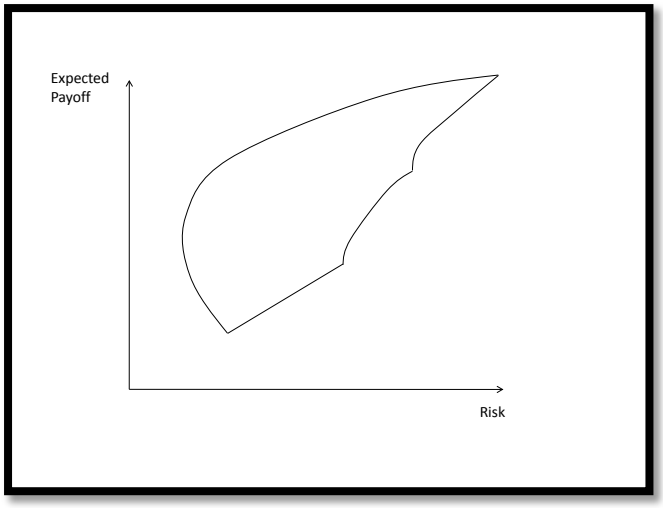

The optimal and most dangerous choices for the terrorist are those that have the highest expected payoff at each level of risk. These can be computed by solving a quadratic programming problem or, equivalently, applying Markowitz's (1952) critical line algorithm. The quadratic programming problem involves finding the resource allocation weighting for each attack method contained within a combination that yields the highest expected payoff for a given level of risk. The complete set of these 'efficient' combinations will be the uppermost boundary of the complete opportunity set. Each attack method or attack method combination contained within the efficient opportunity set has an expected payoff that is higher than any other combination at that particular level risk. The efficient set of attack methods and combinations is delineated by a dashed line superimposed over the appropriate region of the complete set of opportunities:

Figure Two The Lone Wolf's Efficient Opportunity Set

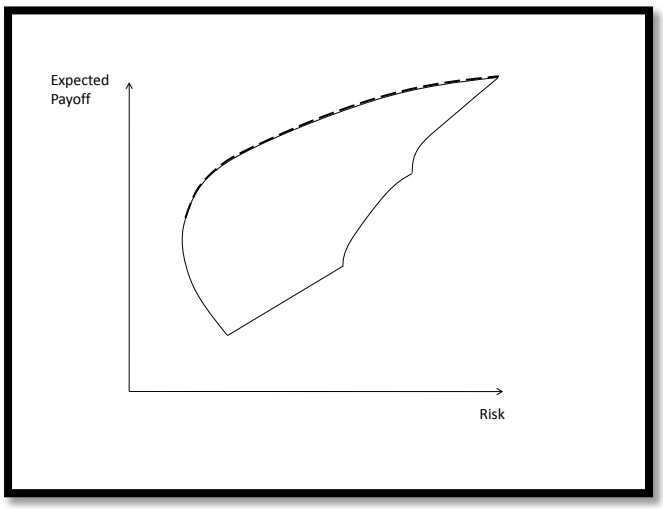

This statistical or algorithmic process permits the calculation of the opportunity set for the terrorist. The process will work with any set of payoffs that have a semblance of statistical structure. Once the opportunity set is delineated, the terrorist's choice from that opportunity set may be analysed. As a model of individual choice in the purest sense, this analytical framework is particularly well-suited to the analysis of the choices of lone wolf terrorists. The lone wolf must choose from the opportunity set. Though he may wish to deploy an attack method that inflicts a very high level of human tragedy without bearing any risk, there is no feasible opportunity to do so. He cannot escape from the risk-reward payoff. The lone wolf will choose an attack method or combination of attack methods that yields the maximum payoff at the desired level of risk. A very risk averse lone wolf will choose attack methods or combinations that have a lower expected payoff and a lower risk. A less risk averse or more risk seeking lone wolf will choose attack methods or combinations that have a higher expected payoff and higher risk. The choices of these two 'types' of lone wolves can be depicted by the set of indifference maps superimposed over the opportunity set.

Figure Three The Lone Wolf's Optimal 'Most Dangerous' Choices

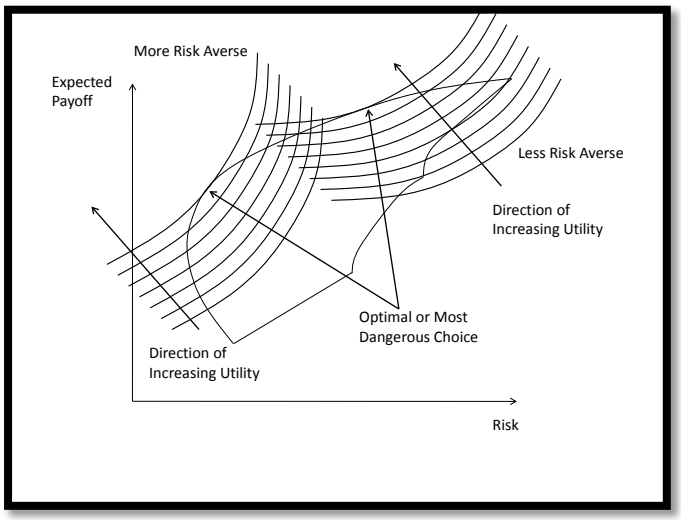

Mean-variance expected utility analysis approximates full expected utility analysis. It frequently does so even when payoffs are not normally distributed or when agents are not described by quadratic expected utility. For mean-variance analysis to hold as an approximation, the terrorist need only think in terms of a range of equally likely gains and losses centred on zero (Tobin 1958, p.74). Mean-variance analysis is also able to generate computable results, which makes it a much more suitable foundation for the provision of operationally relevant investigative advice than the strictly logical-theoretical conclusions of state preference theory (Tobin 1969). If 'quadratic' utility is used as the specific functional form for the generic equation (5), any preference ordering of attack methods produced by the quadratic utility function will certainly conform to the well-known von Neumann and Morgenstern (1944) axioms. For this reason, the quadratic utility function is usually used as the

foundation for the analysis of choice within the mean-variance framework. Although there are some shortcomings that must be recognised when using this type of expected utility function in certain settings, it will, in any case, approximate a broad class of expected utility functions (see Elton et al. 2003, p.232; Kroll, Levy and Markowitz 1984; Levy and Markowitz 1979).

Following Tobin (1958), the quadratic expected utility function can be expressed as:

$$U(R) = (1+b)R + b\sigma_R^2 \tag{6}$$

The lone wolf ranks his preferences for attack methods and makes choices on the basis of risk σ_R^2 and expected payoff *R*. He will not just rank the attack method with the highest expected injuries and fatalities as the best attack method. He will consider the risks involved. The result is a preference ranking for attack methods that depends on how the lone wolf feels about the risk-reward ratio of each attack method. If the lone wolf happens, for a given payoff, to prefer a smaller variance of payoffs (reward) to a larger variance, then $-1 < b < 0$. Like the Becker, Ehrlich and Sandler et al. expected utility functions, the quadratic function depicts a terrorist obtaining utility from some factors. In this case, the two relevant factors are reward and risk, measured by the injuries and fatalities expected to be generated by a terrorist attack and the standard deviation of those injuries and fatalities. Because additional reward is 'good' and adds to the lone wolf's utility *R* enters positively into the function. Because additional risk is 'bad' and detracts from the lone wolf's utility, σ_R^2 enters negatively into the function for a risk averse lone wolf[11]. This approach not only simplifies the expected utility analysis that is assumed to underlie the lone wolf's choices but also allows computable results to be generated rather than purely logical results. Because we can empirically observe the two moments of the distribution of injuries and fatalities generated by terrorist attack methods, we can analyse the lone wolf's choices as if they were made on the basis of this distribution and reach computable numerical conclusions about those choices.

Mean-variance analysis provides the foundation for an economic analysis of the opportunities and choices of the lone wolf terrorist. It does so in a manner that is more suitable to generating operationally relevant computable results. The lone wolf terrorist chooses the most dangerous attack method or combination of attack methods at the level of risk he wishes to bear. If he is very risk averse he will choose attack methods or combinations that inflict a low amount of human tragedy at a low level of risk. If he is more risk seeking he will choose attack methods or combinations that inflict a higher amount of human tragedy. He cannot escape the risk-reward trade-off. If he chooses attack

[11] We can deal with a situation where risk is good. This will be discussed when risk preferences are considered.

methods with a higher expected payoff, he bears greater risk. By using the empirical data for each attack method, Phillips (2011) developed an analysis that aimed to identify the attack methods that would be chosen by lone wolf terrorists across a range of risk levels. Our task here is to contribute further in this direction by applying similar techniques to the analysis of the 'spree' lone wolf. The objective is to encompass within economic analysis aspects of the behaviour of spree lone wolves who are willing to suddenly 'plunge' all of their resources, including time, into violent terrorism even though they had not previously engaged in violent terrorism. The analysis should permit us to draw additional inferences about the types of behaviour that can be expected from such a spree lone wolf and identify some of the circumstances that may precipitate a spree. This analysis will take us some way closer to providing operationally relevant advice to law enforcement and security agencies tasked with pre-empting or pursuing this most dangerous of lone wolf terrorists.

THE SPREE LONE WOLF: OPPORTUNITIES AND CHOICES

Like the serial lone wolf and the transnational terrorist, the spree lone wolf is faced with opportunities and the necessity to make a choice. The RAND categories of attack methods, each with their expected payoffs and risks and the possibility that they may be combined represent his opportunities. Fundamentally, his opportunities are expected payoff-risk 'pairs'. A serial lone wolf, a transnational terrorist or a spree lone wolf, if he chooses optimally, chooses the most dangerous attack method or combination of attack methods at his chosen level of risk. One important difference, perhaps the most important difference, between the spree lone wolf and the serial lone wolf is the time-concentrated nature of the spree lone wolf's engagement in violent terrorism. Within the expected utility model of choice that we have been developing, a spree lone wolf's most distinct characteristic is his willingness to suddenly and all at once 'plunge' all of his resources, including time, into violent terrorism when he had not previously been engaged in violent terrorism. Only a lone wolf with a particular type of risk preference will be observed to suddenly and completely 'plunge' all of his resources into violent terrorism. What is more, the spree lone wolf who is willing, under certain circumstances, to plunge all of his resources into violent terrorism will be extremely sensitive to changes in the expected payoffs to violent terrorism. Even a small adjustment to the risk-reward trade-off may precipitate a spree of terroristic violence.

Start with a generic lone wolf terrorist. He has no special properties. His immediate objective in engaging in violent terrorism is to inflict an amount of human tragedy. He chooses his attack method or combination of attack methods such that the injuries and fatalities that are inflicted by his violent

terrorism are a maximum for the level of risk that he bears[12]. A higher expected payoff is accompanied by higher risk. Within a mean-variance analytical framework his opportunities can be depicted as a convex set of expected payoff-risk pairs. His most dangerous choices will be at the uppermost boundary of the opportunity set. His choices will reflect his preference for risk. If he is very risk averse his choices will be located in the lower risk range of opportunities. If he is less risk averse his choices will be located in the higher risk range of opportunities. A lone wolf has limited resources, including time. To the extent that his choices involve engaging in violent terrorism he has less time and other resources to allocate to non-violent forms of expression or other legitimate activities such as working a regular job. It is the sudden and complete plunge into violent terrorism that distinguishes the serial lone wolf and the spree lone wolf. Whereas the serial lone wolf's behaviour can be completely explained as the optimising behaviour of a risk averse individual requiring increments in expected payoffs in return for bearing additional risk, the 'plunging' behaviour of the spree lone wolf—zero allocation of resources to violent terrorism then, suddenly and all at once, complete allocation of resources to violent terrorism—is his defining characteristic. The analytical tools with which to encompass this aspect of the spree lone wolf emerge from Tobin's (1958, p.76) classical treatment of risk preference.

As we have seen, the opportunities of the lone wolf are a convex set of risk-reward pairs in expected payoff-risk space. The convexity of the opportunity set reflects the positive nature of the risk-reward trade-off and the gains that may be obtained by combining attack methods. Specifically, combination yields increases in expected payoffs without increases in risk. If the lone wolf may exhibit varying degrees of specialisation in violent terrorism, the efficient opportunity set will be expanded or contracted according to the percentage of his total resources, including time, which he devotes to violent terrorist activities. The efficient opportunity set for the non-specialist lone wolf will lie within the efficient opportunity set of the lone wolf who devotes a greater percentage of his resources to violent terrorism. The efficient opportunity set can be computed for different levels of specialisation. The solution process for Markowitz's (1952) critical line algorithm involves computing the optimal resource allocations across different attack methods when the total resource allocations (weightings) to the different attack methods may not exceed 1.00 or, equivalently, a resource allocation of 100 percent. If the lone wolf does not specialise in violent terrorism, the total resource allocation must be less than 1.00. The Markowitz approach applies equally well in such scenarios where the total resource allocation is 5 percent, 10 percent, 18 percent, 75 percent and so on.

[12] Note that if the lone wolf wishes to inflict a small number of fatalities, he can choose an attack method that has a lower expected payoff and lower risk. However, even with a less overtly violent attack method such as hostage-taking, injuries and fatalities may be high. The risk associated with each attack method encompasses the possibility that the actual number of fatalities generated by an attack will higher than expected.

Using the RAND data for injuries and fatalities generated by each attack method, the Markowitz critical line algorithm can be applied to scenarios where the lone wolf specialises to different degrees. By specialisation we mean the allocation of resources, including time, to violent terrorism—either single attack methods or combinations of attack methods—rather than specialisation in sense of a single attack method. At each level of risk and each level of specialisation, the investigator may determine which attack method combination[13] yields the highest expected payoff. Formally, the opportunity sets at various levels of specialisation are computed by solving the quadratic programming problem with resource allocation constraints of 0.10, 0.25, 0.50 and 0.75:

$$\min \sigma_C^2 * = \sum_{i=1}^{n} \sum_{j=1}^{n} w_i w_j \rho_{ij} \sigma_i \sigma_j$$

Subject to the constraints:

$$\sum_{i=1}^{n} w_i = 0.10; 0.25; 0.50; 0.75$$

$$w_i \geq 0$$

Our very non-specialist lone wolf who allocates only 10 percent of resources to terrorism has significantly less opportunity for exploiting the gains from combination. As a consequence, his efficient set of opportunities is situated relatively farther to the South-West in expected payoff-risk space. Our 'mid range' specialist lone wolves allocate (a) 25 percent; (b) 50 percent; and (c) 75 percent of resources to terrorism. As the allocation increases—as the degree of specialisation increases—more opportunities become available to generate higher expected payoffs at each given level of risk. This is the exploitation of the underlying statistical structure of the attack methods. The benefits from combination begin to become apparent when the expected payoffs at each level of risk are compared at different levels of specialisation. Where the very non-specialist lone wolf could not form an attack method combination with an expected payoff in excess of 1.00 fatalities and injuries, the 25 percent allocation is mostly characterised by attack method combination opportunities with expected payoffs in excess of 1.00 fatalities and injuries. Furthermore, even at the same level of risk there are feasible combinations at the 25 percent level of resource allocation that yield higher fatalities and injuries. Of course, a further increase in specialisation to 50 percent resource allocation to

[13] We work, for the moment, on the basis that the lone wolf can form a combination constituted by more than one attack method.

terrorism yields a similar increment in the expected level of human tragedy for each level of risk, as too does a further increase to 75 percent. Each time, the opportunity set moves further to the North-West in expected payoff-risk space.

The efficient opportunity sets for each level of specialisation, including complete specialisation in violent terrorism, emerge geometrically as:

Figure Four The Lone Wolf's Efficient Opportunity Sets (Computed with RAND Data)

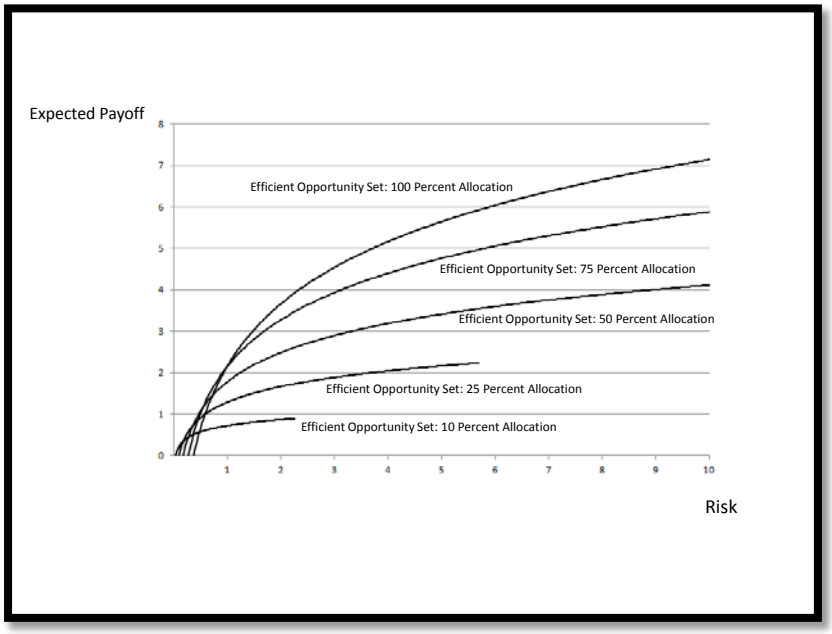

At all of these levels of specialisation—10 percent, 25 percent, 50 percent, 75 percent and 100 percent—familiar patterns emerge to characterise the lone wolf's opportunities to inflict human tragedy. A combination of assassination and bombing dominates the lowest risk combinations across all levels of specialisation whereas the higher risk combinations are dominated by bombing, hostage-taking and unconventional attacks across all of the levels of specialisation. Arson is all but ruled out whilst armed attacks feature in a large proportion of the attack method combinations across all levels of specialisation. Armed attacks do not dominate the attack method combinations at any level of risk but they are present. When full combination of the attack methods is available to the lone wolf, the statistical structure of the attack method combinations diminishes or enhances the expected prevalence of any particular attack method. Fundamentally, the statistical structure highlights the gains that may accrue to the lone wolf by combining attack methods and reveals the structure of the opportunities available to the lone wolf at each level of risk. As the lone wolf becomes more specialised, his opportunities for inflicting human tragedy increase at each level of risk.

The optimal and most dangerous choices of the lone wolf may be depicted as the maximisation of an expected utility function of the form:

$$U(R) = (1+b)R + b\sigma_R^2 \tag{6}$$

In the case of risk aversion $-1 < b < 0$. This does not mean that the lone wolf will not bear risk but only that he requires increments in expected payoffs for bearing additional risk. He maximises his expected utility by weighing up the risk-reward trade-offs presented to him by the opportunity set. If he is very risk averse, he will choose a lower level of specialisation and a combination of attack methods that has a lower expected payoff and lower risk. The indifference curves associated with this type of lone wolf are very steeply concave in expected payoff-risk space. A possible solution for this lone wolf is the point of tangency of the highest indifference curve with the boundary of one of the non-specialist efficient opportunity sets:

Figure Five The Risk Averse Lone Wolf's Optimal Choice

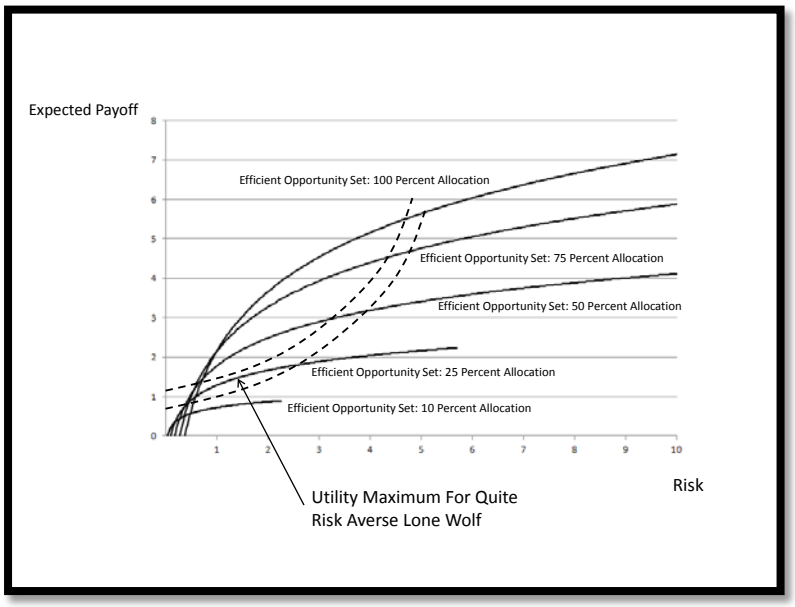

If the lone wolf is less risk averse, he may choose a higher level of specialisation and attack method combinations that have a higher expected payoff and higher risk. A possible solution for this lone wolf is the point of tangency of the highest indifference curve with the boundary of one of higher-specialisation efficient opportunity sets:

Figure Six The Less Risk Averse Lone Wolf's Optimal Choice

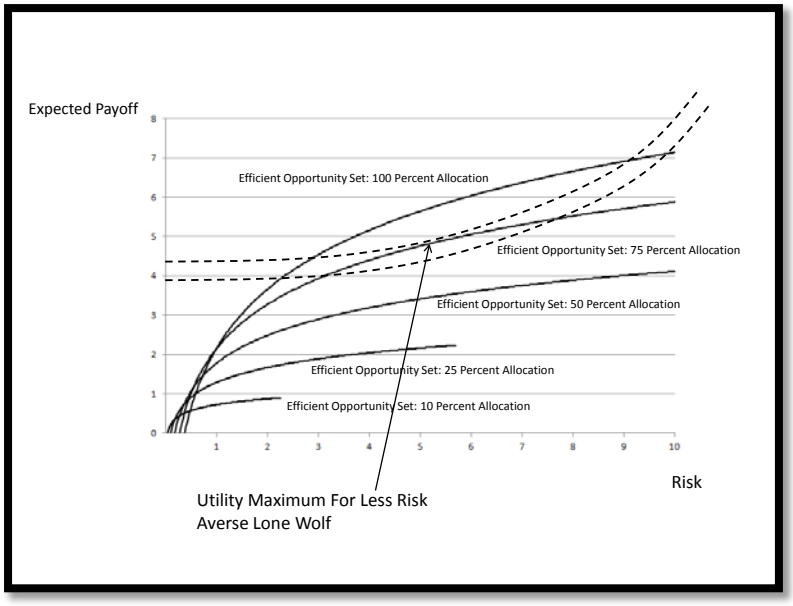

It is possible for the risk averse lone wolf to specialise in violent terrorism. A very risk averse lone wolf specialist who devotes 100 percent of his resources to violent terrorism will select those attack methods and combinations with lower expected payoffs and lower risk. A less risk averse lone wolf specialist who devotes 100 percent of his resources to violent terrorism will select those attack methods and combinations with higher expected payoffs and high risk. Possible solutions for these lone wolves are represented by the points of tangency of the highest indifference curve for each type of lone wolf with the boundary of the 100 percent specialisation efficient opportunity set:

Figure Seven The Lone Wolf's Optimal Choices with 100 Percent of Resources Allocated to Violent Terrorism (Lower and Higher Risk Aversion Cases)

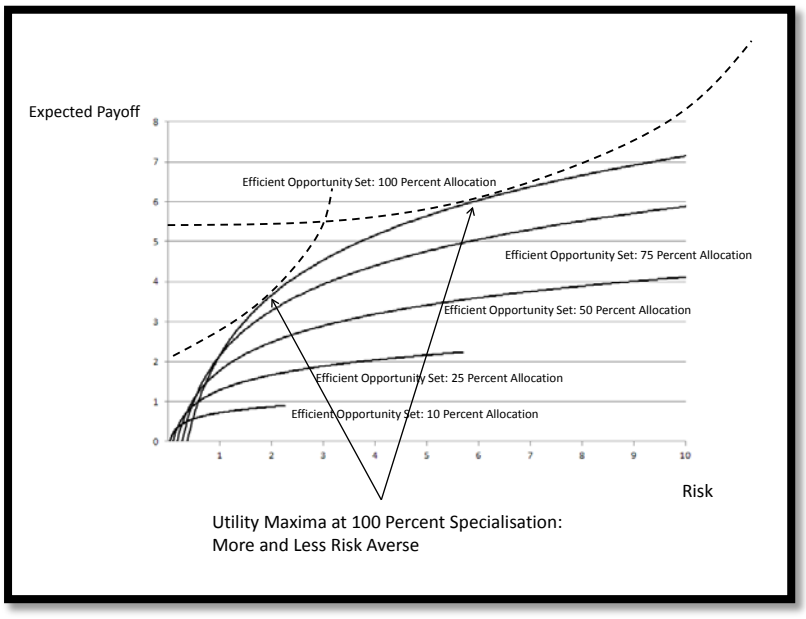

In order to encompass the spree lone wolf within this analytical framework, it is necessary to move beyond the treatment of risk aversion and consider risk seeking behaviour. For transnational terrorists and lone wolves, risk aversion accurately reflects the theoretical logic of choices and is in accordance with empirical fact. That is, if all terrorists were risk seeking we would only ever see attacks of the riskiest kind. A risk seeking individual always seeks the point of maximum expected payoff and maximum risk. A risk averse individual, with a particularly low level of risk aversion, may also find that his utility is maximised at the point of maximum expected payoff and maximum risk. It is clear, however, that this is not an empirical feature of terrorism that is regularly observed. Risk aversion encompasses both this less common feature of terrorism and the more common cases where terrorists do not engage at the point of maximum risk. Risk aversion will explain nearly all terrorist behaviour. A risk seeking lone wolf may, however, be introduced to the analytical framework. He will maximise the same quadratic expected utility function:

$$U(R) = (1+b)R + b\sigma_R^2 \tag{6}$$

But in the case of risk seeking $0 < b < 1$. The indifference curves for the risk seeking lone wolf must reflect the fact that he is willing to give up expected payoff in order to bear more risk. To reflect this defining characteristic of the risk seeking lone wolf, the indifference curves for the risk seeking lone wolf will be downward sloping in expected payoff-risk space. As Tobin (1958, p.78) points out, the point of maximum risk will always be the point at which utility is maximised for a risk seeking individual. Within the analytical context that we have been constructing, the risk seeking lone wolf will always be at the point of maximum risk in the 100 percent-specialisation efficient opportunity set. This is easily proved by inspection of the relevant geometry:

Figure Eight The Risk Seeking Lone Wolf's Optimal Choice

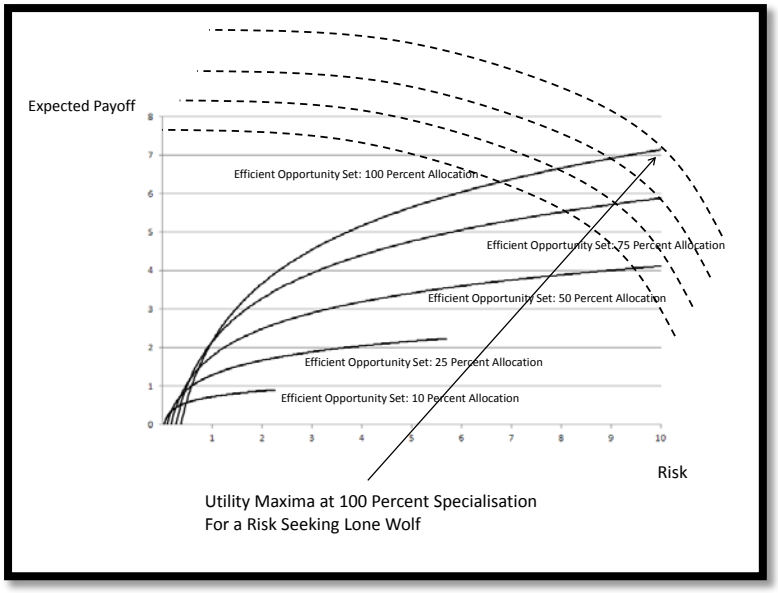

Of the two types of lone wolves, the risk averse lone wolf is the more likely to deploy the types of attack methods—armed attacks and bombings—that are associated with the lone wolf activity that has been empirically observed. The risk averse lone wolf is also the more likely to be less than completely specialised in violent terrorism. He is more likely to devote at least some of his time and resources to other activities. The more or less risk averse lone wolf typifies the serial lone wolf terrorism that has been empirically observed and encompasses many of the idiosyncrasies exhibited by serial lone wolf terrorists. To understand the spree lone wolf it would appear that we must look beyond risk aversion to, perhaps, risk seeking behaviour. Unfortunately, we do not find the defining characteristics of the spree lone wolf reflected by the risk seeking lone wolf. The risk seeking lone wolf is more likely than the risk averse lone wolf to devote all of his resources, including time, to violent terrorism but the actual attack methods that he will choose will be dominated by the riskiest attack methods: unconventional attacks and hostage-taking. It seems that the exact type of behaviour—no resource allocation to violent terrorism then, suddenly, 100 percent allocation to violent terrorism, particularly armed attacks or bombing—that most distinguishes the spree lone wolf is not yet reflected in either of the cases that we have examined.

There is a special type of risk seeking lone wolf that exhibits the characteristic behaviour of the spree lone wolf. This type of lone wolf is willing to absorb additional risk in return for increases in expected payoffs. However, unlike the risk averse lone wolf, this type of lone wolf displays diminishing marginal rate of substitution of risk for expected payoffs. At higher levels of risk he requires very little additional expected payoff in order to be enticed to bear an additional increment of risk. He views the risk-reward trade-off in a different way from our other two types of lone wolf terrorist. At very high levels of risk, it will take only a very small increase in expected payoff to entice this wolf to

bear a lot more risk. This makes him very dangerous. When the expected payoffs to violent terrorism increase, this type of lone wolf may suddenly switch all of his resources, including time, to violent terrorism. This type of lone wolf will be observed to suddenly switch all of his resources into the risky violent attack methods or combinations of attack methods that are available to him where previously all of his resources were devoted to other, non-violent activities. All at once he will 'plunge' all of his resources into the available attack methods. All at once he will engage in a spree.

To encompass this type of behaviour within our analytical framework requires us to think of the efficient sets as dynamic rather than presented statically on the printed page. As the correlation structures that characterise the relationship of each attack method with the others change through time, the opportunity set may move up or down. It will always be a convex set unless the attack methods all exhibit perfectly positive correlation with each other but it does not always have to be stationary. Law enforcement actions will affect the risk and reward of some attack methods, making others more or less desirable and terrorists may innovate. When the risk-reward trade-off is altered, the serial lone wolf may adjust his choices in light of the new opportunities but providing his preference for risk remains the same, his new choices are unlikely to be very much different from his old choices. The most likely outcome is that for the serial lone wolf is that his choices now yield higher (or lower) rewards for a given level of risk. The spree lone wolf is different. Before the alteration of the risk-reward ratio he lay hidden from law enforcement. He did not allocate any resources to violent terrorism. Now the change in the risk-reward ratio brings him forth. His first attack is not the beginning of a serial campaign. His first attack is a spree of violence characterised by the plunging of all of his resources into violent terrorism. Such behaviour is encompassed by indifference curves that are upward sloping in expected payoff-risk space to reflect the special property that Tobin (1958) calls 'plunging'. In our context: no resources allocated to the violent terrorism then, following a change in the risk-reward trade-off, all resources 'plunged' into violent terrorism.

Figure Nine The Spree Lone Wolf's Optimal Choice

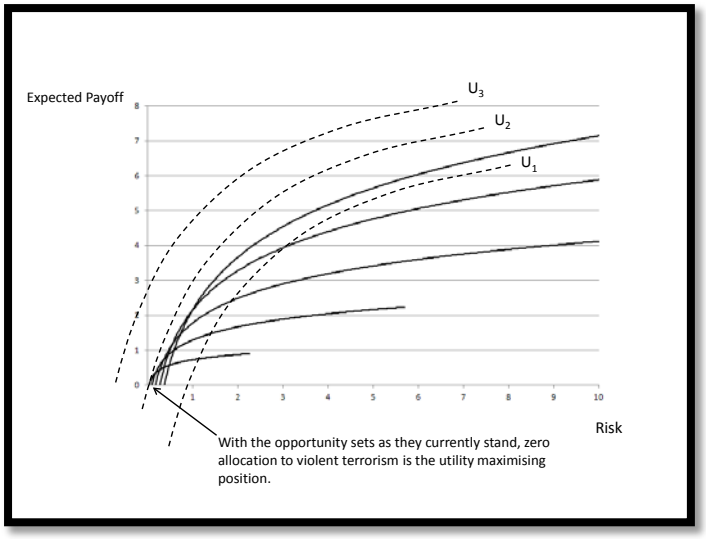

With the opportunity sets as they stand, the utility maximising position for the lone wolf is zero allocation of resources to violent terrorism. However, if he were to perceive an innovation in the risk-reward trade-offs available to him such that the efficient opportunity sets all move upwards such that the uppermost efficient set (100 percent specialisation) intersects an indifference curve higher than U_2, the lone wolf will 'plunge' all of his resources into violent terrorism. The plunging of all resources, including time, into violent terrorism may manifest itself as a spree. Where law enforcement had been previously unaware of the existence of a potential lone wolf threat because of his avoidance of violent terrorist actions, the spree lone wolf now emerges suddenly. He does not emerge as an individual who allocates some resources to terrorism and some resources to other activities. The spree lone wolf suddenly and all at once plunges all of his resources, including time, into violent terrorism. The spree of violence will likely continue until he is apprehended or until he perceives that the risk-reward trade-off no longer favours his violent terrorist actions. Due to the time-concentrated nature of the spree, it is unlikely that the lone wolf's perceptions will change. He must therefore be apprehended in order for the spree to end.

THE PLUNGE AND THE SPREE

The plunging of all resources, including time, into violent terrorism might be accepted by some economists as an appropriate analogy for a spree without further justification. However, the metaphorical mapping between 'plunging' and 'spree' behaviour does have a stronger foundation that provides further justification for this conclusion. The analysis is formally a single-period analysis. Because the single analytical period may align with any period of real chronological time, it is usually implicitly assumed within the analysis that opportunities and choices take place over some period of chronological time. This implies that some 'time diversification' characterises the choices of the lone

wolf. Time diversification refers to a property of the statistical structure of the payoffs to risky activities. Specifically, the expected variance is lower at longer horizons than at short horizons. For any particular series of attack method deployments, the variance of the actual injuries and fatalities may be higher after the first two attacks of the series than after six or seven attacks. In this case, the allocation of resources to terrorism over time rather than all at once may yield an accumulated amount of human tragedy at a lower level of risk. The 'plunger' shuns time diversification and pursues the risk-reward trade-off that prevails at the single point in time that he decides to 'plunge'. For the spree lone wolf, the actual chronological time period that corresponds to the analytical time period is one that is condensed or concentrated into a very small period. The spree occurs not only because the spree lone wolf all at once allocates all of his resources to violent terrorism but because he also shuns the statistical property of violent terrorism that is reflected by the term 'time diversification'.

THE SPREE LONE WOLF: SINGLE ATTACK METHOD ONLY

If the lone wolf cannot combine attack methods or if some attack methods are not available, the spree lone wolf, in plunging all of his resources into violent terrorism, can also be expected to plunge all of his resources into the single available attack method that yields the highest expected payoff at the level of risk that he wishes to bear. Both serial lone wolves and spree lone wolves have generally chosen armed attacks or bombing as their preferred attack methods, which leads us to suspect that combining attack methods is not always feasible[14]. If lone wolves can only choose a single attack method, the opportunity set reduces to a set of risk-reward 'pairs' that are formed on the basis of the expected payoffs and risks of each of the attack methods and the degree of specialisation in violent terrorism that characterises the lone wolf. The lone wolf's choices when no combination of attack methods is possible may be analysed by determining which degree of specialisation and which attack method yields the highest expected payoff when only one attack method can be chosen. Formally, at different levels of risk solve the quadratic programming problem:

$$\max E(R_C) = \sum_{i=1} w_i E(R_i) = R*$$

Subject to a target level of return variance (risk), σ_C^2 *:

$$\sigma_C^2* = \sigma_i^2$$

[14] This was assumed by Phillips (2011) on the basis that a single individual offender may find it difficult to allocate resources across more than one attack method during a single period of time.

And the constraints:

$$\sum_{i=1} w_i = 1$$

$$w_i \geq 0$$

The levels of risk at which this quadratic programming problem was solved are: 0.1, 0.2, 0.3, 0.5, 1.0, 1.5, 2.0 ...42.5. The degrees of specialisation and attack methods that correspond to the solution of the quadratic programming problem at these levels of risk are presented in the tables. At the lowest levels of risk, assassination dominates across a range of degrees of specialisation. For example, if the lone wolf bears a very small standard deviation (risk) of 0.5 fatalities and injuries per attack, a 71.5 percent degree of resource allocation (specialisation) is the optimal choice. As the level of risk that will be borne increases, assassination no longer remains an optimal choice and bombing emerges as the optimal choice of attack method. For example, at a 4.5 level of risk, a 74.9 percent degree of specialisation in bombing is optimal. At still higher levels of risk, the lone wolf will find that it is optimal to switch from bombing to hijacking and then to armed attacks. Finally, hostage taking and unconventional attacks come to dominate. At each level of risk, there is a single attack method and a degree of specialisation that produces the highest expected payoff.

For the spree lone wolf, the relevant question is whether single attack methods will ever yield an optimal solution to his utility maximisation problem and, if so, whether a 100 percent resource allocation to armed attacks or bombing such has been observed empirically is ever to be expected on the basis of the mean-variance expected utility analysis. When the relevant quadratic programming problem is solved, it emerges that as risk increases particular attack methods emerge as the optimal choice at particular levels of specialisation. At very low levels of risk, the lone wolf who can choose only a single attack method will allocate some resources to assassination. His optimal level of specialisation in violent terrorism, which in this case is equivalent to his level of specialisation in assassination, increases as the level of risk he wishes to bear increases. He never becomes a complete specialist—100 percent resource allocation—in assassination because at increasing levels of risk higher degrees of specialisation in assassination become dominated by lower degrees of specialisation in bombing. As the risk that he is willing to bear increases, his degree of specialisation in violent terrorism increases. He becomes a specialist in violent terrorism, which is equivalent to specialisation in bombing, at levels of risk (standard deviation) approach 6.00. At higher levels of risk, bombing is superseded by hijacking, in which near specialisation may be observed, and at higher still levels of risk hijacking is superseded by armed attacks. At levels of risk (standard deviation) that approach

15.00 the lone wolf becomes a specialist in violent terrorism and allocates all of his resources to armed attacks. At the appropriate levels of risk both bombing and armed attacks represent solutions to the spree lone wolf's utility maximisation problem: maximise injuries and fatalities at 100 percent specialisation in violent terrorism when only a single attack method can be chosen.

Table Three The Lone Wolf's Optimal Choice: Single Attack Method, Varying Specialisation Level

Risk (Standard Deviation)	Degree of Specialisation	Attack Method	Expected Payoff
0.1	0.143	Assassination	0.22
0.2	0.286	Assassination	0.441
0.3	0.429	Assassination	0.66
0.4	0.572	Assassination	0.883
0.5	0.715	Assassination	1.10
1.0	0.166	Bombing	0.97
1.5	0.249	Bombing	1.46
2.0	0.333	Bombing	1.96
2.5	0.416	Bombing	2.45
3.0	0.499	Bombing	2.94
3.5	0.582	Bombing	3.43
4.0	0.666	Bombing	3.92
4.5	0.749	Bombing	4.41
5.0	0.832	Bombing	4.90
5.5	0.915	Bombing	5.39
6.0	0.999	Bombing	5.88

Table Three The Lone Wolf's Optimal Choice: Single Attack Method, Varying Specialisation Level (Continued)

Risk (Standard Deviation)	Degree of Specialisation	Attack Method	Expected Payoff
6.5	0.609	Hijacking	2.38
7.0	0.655	Hijacking	2.56
7.5	0.702	Hijacking	2.74
8.0	0.749	Hijacking	2.93
8.5	0.796	Hijacking	3.11
9.0	0.843	Hijacking	3.29
9.5	0.890	Hijacking	3.47
10.0	0.936	Hijacking	3.66
10.5	0.983	Hijacking	3.84
11.0	0.743	Armed Attacks	3.95
11.5	0.777	Armed Attacks	4.13
12.0	0.810	Armed Attacks	4.31
12.5	0.844	Armed Attacks	4.49
13.0	0.878	Armed Attacks	4.67
13.5	0.912	Armed Attacks	4.85
14.0	0.945	Armed Attacks	5.03
14.5	0.979	Armed Attacks	5.21

If a lone wolf can choose just one single attack method and cannot combine attack methods, his optimal opportunities become a set characterised by a single attack method and a particular level of specialisation at particular levels of risk. For the most part, complete specialisation emerges as an optimal choice only at small number of points. At very low levels of risk, complete specialisation is not expected at all. At the mid-range of risk, we do observe complete specialisation in violent terrorism—100 percent allocation of resources to violent terrorist attack methods—with 100 percent resource allocation to bombing at 6.00 standard deviation of fatalities and injuries. At the mid-to-high ranges of risk we observe almost complete specialisation in hijacking at a 10.5 standard deviation of fatalities and injuries and almost complete specialisation in armed attacks at 14.50. At the very highest levels of risk, complete specialisation will be observed in hostage taking and, ultimately, in unconventional attacks. A lone wolf seeking to bear the highest levels of risk will find that these highest levels can only be reached through a complete specialisation in the riskiest attack methods. A plunging spree lone wolf who can only choose bombing or armed attacks will find these attack methods to solve his utility maximisation problem at their respective levels of risk. The analysis agrees with empirical observation.

CONCLUSIONS AND IMPLICATIONS FOR LAW ENFORCEMENT

The spree lone wolf terrorist presents special challenges to law enforcement. If defence economics is to contribute to the pre-emption, pursuit or understanding of the spree lone wolf, the first step is to attempt to encompass his most characteristic behaviour within an economic analytical framework. The time-concentrated nature of the spree lone wolf's terrorist actions is the most distinct feature of his behaviour. Unlike the serial lone wolf, the spree lone wolf all at once 'plunges' all of his resources, including time, into violent terrorism. Where before he was not engaged in risky violent terrorism he is now suddenly fully engaged in violent terrorist actions. What is more, he shuns the 'time diversification' that characterises the serial lone wolf's actions and directs all of his resources towards a time-concentrated spree of terrorist violence. Within a mean-variance expected utility framework, the type of behaviour that characterises a spree lone wolf is reflected by a utility maximising individual whose utility maximisation problem is solved by devoting either zero resources to violent terrorism or, following an innovation in the risk-reward structure that characterises the available attack methods, suddenly plunging 100 percent of his resources into violent terrorism. This sudden change is unique to the 'plunger' and will not be exhibited by either a risk averse lone wolf or an (ordinary) risk seeking lone wolf[15].

The distinguishing feature of the spree lone wolf emerges within the mean-variance expected utility analysis. If spree behaviour—suddenly plunging all resources, including time, into a time-concentrated spree of violent terrorism—emerges within the analytical framework it next remained to inquire whether an answer as to which attack method the spree lone wolf will plunge his resources into could also be extracted from the economic analytical framework. Empirically, lone wolves (and transnational terrorists) have been observed to favour armed attacks and bombing. This choice is consistent with formal utility maximisation within the mean-variance framework. If the spree lone wolf can choose his level of specialisation, the amount of resources he allocates to violent terrorism, and tries to inflict a maximum of human tragedy at a given level of risk when only one attack method from the RAND categories can be chosen, the formal solution of the relevant quadratic programming problem reveals that complete specialisation in bombing or armed attacks is a solution to this utility maximisation problem at the respective levels of risk. The lone wolf terrorist who suddenly and all at once plunges his resources into violent terrorism in the form of armed attacks or bombing is reflected

[15] Both the risk averse and the risk seeking lone wolf may allocate 100 percent of resources to violent terrorism. This will always be the case for the risk seeking lone wolf and will sometimes be the case for a risk averse lone wolf with a low level of risk aversion. The 'plunger' or 'spree' lone wolf, however, goes from 0 percent allocation to 100 percent allocation all at once.

within our economic analytical framework. This reflection yields further insights into the nature of the spree lone wolf.

There are several important results that may be of importance to law enforcement:

- The spree lone wolf will be the most sensitive to changes in the risk-reward trade-off. A risk-seeking lone wolf will already (and always) be at the point of maximum risk and maximum expected payoff. A risk averse serial lone wolf will adjust to changes in the risk-reward trade-off by increasing or decreasing his engagement in violent terrorism or choosing different attack methods and combinations but he will never 'plunge' or engage in a spree of violence in response to changes in the risk-reward trade-off.
- The spree lone wolf is unlikely to discontinue his spree of violence once it has started. There are no features of his profile that would act to mitigate a spree of violence once it has commenced. For the spree lone wolf to decide to discontinue his spree of violence, he must perceive that the risk-reward trade-off has readjusted such that his utility maximising position is once again characterised by a zero allocation of resources to violent terrorism. Given that the spree takes place in a very short period of time, there is unlikely to be much opportunity for the spree lone wolf to reflect on the changing risk-reward trade-off once the spree is underway. For the spree to cease, he must be apprehended.
- A spree must necessarily be a very violent action. Within the mean-variance framework that has been described, a spree occurs when the lone wolf suddenly and all at once plunges all of his resources into violent terrorism. This is precipitated by a change (or perception of a change) in the risk-reward trade-off that characterises violent terrorism. If the lone wolf's perception is matched by a real upward shift in the efficient opportunity set, every attack method and combination now is expected to inflict a higher amount of human tragedy.
- Armed attacks and bombing represent utility maximising solutions for the spree (and serial) lone wolf when he is constrained to choosing a single attack method rather than a combination of different attack methods. If the lone wolf wishes to bear less risk, bombing will be the single attack method of choice. If more risk is desired, armed attacks will be the attack method choice when only one attack method can be chosen. Of course, this agrees with the empirically observed choices of lone wolf terrorists.

The spree lone wolf that emerges within the mean-variance expected utility framework is one that does not engage in violent terrorism at all when the opportunity sets presented to him do not offer a high enough ratio of risk to reward. He allocates no resources to violent terrorism under such circumstances. During this time, he may use legitimate or non-violent methods to pursue whatever

ends—political, religious, racial *etc*—motivate him. Or he may harbour his thoughts and motivations behind a mask of indifference and give no indication to anyone that he has another dimension to his character. Perceiving a change in the risk-reward ratio that characterises violent terrorism, he suddenly and all at once allocates all of his resources, including time, to violent terrorism. If he can choose only a single attack method, bombing, hijacking, armed attacks, hostage-taking or unconventional attacks will represent utility maximising choices at various levels of risk. We know that lone wolves have historically favoured armed attacks or bombing and so we may expect the spree lone wolf to centre his attention on those types of attack methods. The spree is a time-concentrated complete allocation of resources to violent terrorism.

For law enforcement agencies charged with the task of pre-empting or pursuing the spree lone wolf, the advice that emerges from the analysis is as follows. As law enforcement agencies already know, all lone wolves are difficult to pre-empt because they are almost invisible before they strike. The spree lone wolf presents a special challenge in this regard because his first strike will be a spree of violence undertaken on the basis of a complete 'plunge' of resources into a time-concentrated deployment of attack methods. It is clear that the only way to pre-empt a spree lone wolf is for governments to provide resources adequate enough for law enforcement and security agencies to collect and analyse the very scant information trail produced by a lone wolf. Remember, the lone wolf is not a part of a network that may be infiltrated. But he must acquire weaponry and materials from somewhere and he may actively seek to further some cause or other by legitimate means before suddenly turning to violent terrorism. If the lone wolf was always completely invisible before he strikes, attempts at pre-emption would be futile. If we conclude that the lone wolf is *almost* invisible before he strikes, law enforcement must be equipped with the resources to penetrate those aspects of the lone wolf's actions that necessitate him to reveal fleeting glimpses of his existence. Once he strikes, law enforcement confronts an individual who will not cease his spree of violence until he is apprehended or otherwise compromised.

Although the challenges presented to law enforcement by lone wolf terrorism are substantial, the analysis shows that the spree lone wolf's actions must be precipitated by an upward movement in the opportunity set (increase in the risk-reward ratio). Effective actions by law enforcement agencies serve to increase the risk and decrease the rewards of violent terrorism (decrease the risk-reward ratio). Effective law enforcement and vigilant security at public locations serve to ensure that the rewards available to violent terrorism do not rise more than proportionately to the risks. This represents the most effective check on the spree lone wolf. What must be guarded against at all costs is the unfavourable rebalancing through some security initiative of the deterrence-substitution effects that characterise terrorism. Defence economists have shown that security initiatives deter but also encourage substitution. In the case of spree lone wolf terrorism, the most dangerous implication is that

a targeted security initiative shifts law enforcement resources such that a particular type of terrorist action at a particular location becomes riskier without increasing the associated rewards (decrease in the risk-reward ratio) whilst simultaneously increasing the risk-reward ratio of an alternative attack method. If the spree lone wolf senses this set of circumstances, the consequences are clear.

References

AFP, 2011, "Norway Killing Spree was Plan B," November 19, 2011.

Becker, G.S. (1968), "Crime and Punishment: An Economic Alternative," *Journal of Political Economy*, Vol.78, pp.169-217.

Ehrlich, I. (1973), "Participation in Illegitimate Activities: A Theoretical and Empirical Investigation," *Journal of Political Economy*, Vol.81, pp.521-565.

Elton, E.J., Gruber, M.J., Brown, S.J. and Goetzmann, W.N. (2003), *Modern Portfolio Theory and Investment Analysis*, John Wiley & Sons, New York, New York.

Enders, W. and Sandler, T. (1993), "The Effectiveness of Anti-Terrorism Policies: Vector-Autoregression-Intervention Analysis," *American Political Science Review*, Vol.87, pp.829-844.

Instituut voor Veiligheids en Crisismanagement (2007), *Lone Wolf Terrorism*, Instituut voor Veiligheids en Crisismanagement, Netherlands.

Kroll, Y., Levy, H. and Markowitz, H.M. (1984), "Mean-Variance Versus Direct Utility Maximisation," *Journal of Finance*, Vol.39, pp.47-61.

Landes, W.M. (1978), "An Economic Study of U.S. Aircraft Hijackings 1961–1976," *Journal of Law and Economics*, Vol.21, pp.1-31.

Levy, H. and Markowitz, H.M. (1979), "Approximating Expected Utility by a Function of Mean and Variance," *American Economic Review*, Vol.69, pp.308-317.

Markowitz, H. (1952), "Portfolio Selection," *Journal of Finance*, Vol.7, pp.77-91.

Phillips, P.J. (2005), "The 'Price' of Terrorism," *Defence and Peace Economics*, Vol.16, pp.403-414.

Phillips, P.J., (2009), "Applying Portfolio Theory to the Analysis of Terrorism: Computing the Set of Attack Method Combinations From Which the Rational Terrorist Group will Choose in Order to Maximise Injuries and Fatalities," *Defence and Peace Economics*, Vol.20, pp.193-213.

Phillips, P.J. (2011), "Lone Wolf Terrorism," *Peace Economics, Peace Science and Public Policy*, Vol.17, Article 1.

Sandler, T., Tschirhart, J.T. and Cauley, J. (1983), "A Theoretical Analysis of Transnational Terrorism," *American Political Science Review*, Vol.77, pp.36-54.

Sandler, T. and Enders, W. (2004), "An Economic Perspective on Transnational Terrorism," *European Journal of Political Economy*, Vol.20, pp.301-316.

Tobin, J. (1958), "Liquidity Preference as Behaviour Towards Risk," *Review of Economic Studies*, Vol.25, pp.65-86.

Tobin, J. (1969), "Comment on Borch and Feldstein," *Review of Economic Studies*, Vol.36, pp.13-14.