Effects of Mass Media and Opinion Exchange on Extremist Group Formation

Steven Butler[1] and Joanna J. Bryson[1]

Department of Computer Science, University of Bath, BA2 7AY. UK

Abstract. Contemporary communication technologies are thought to facilitate the growth of the small autonomous terrorist groups indicative of 'modern' terrorism. In this study, the Animal Liberation Front provides an example of a culturally distinct organisation from which an extremist force—the Animal Rights Militia—emerges to pose the threat of violence. Agent-based modelling is used to simulate this emergence, and subsequent dynamics, under a variety of conditions. The simulation model not only implements local opinion exchange within a population, but also the polarising effect of mass media. Results show the significance of mass media, of limits on the cell size that are independent of the policing efforts, and the impact of societies with rapid population turn over, such as are found near universities.

Keywords. Agent-based simulation of social phenomena, cultural dynamics, opinion dynamics.

1 Introduction

The autonomous nature of small, terrorist networks or *cells* is indicative of 'modern' terrorism—a terrorist movement that does not need a hierarchical command structure, but instead simply the dissemination of goals, guidelines and ideologies. Modern communication technologies, the Internet in particular, easily facilitate such cultural transactions. This is a trend that can only become more widespread, and so it is important for us to understand how these cells may develop and the factors that may aid or hinder their growth.

In this study, we attempt to provide an insight into the proliferation of extremist violence from within a culturally distinct group. We create a model, the TerrorPlex, which simulates the emergence of terrorist groups and the subsequent population interaction dynamics on the individual and group level. A range of factors are explored within the model, including the impact of mass media, the nature of cell structures (particularly size limits) and the impact of population turnover, such as is found around universities.

The group chosen for this study is the Animal Liberation Front (ALF), a trans-global collective of individuals largely united by an opposition to animal testing. The ALF is not in itself a violent movement, but the assumption can be made that the Animal Rights Militia (ARM), which *has* claimed responsibility for numerous acts of violence, consists of extremist ALF members intent on recruiting into secretive groups other ALF members who exhibit an extremist bent. This arrangement is thought to be typical of terrorist organisations, many of which might be seen as a more significant threat than

the ARM. However, the ALF/ARM are relatively accessible for study, and the ARM do meet technical definitions of a terrorist group. The aims and rules of the ALF are readily available[1], enabling individuals and small groups to follow the ALF agenda whilst retaining autonomy. Guidelines even recommended size limits for activist groups.

This degree of transparency can be logically credited to the relative mildness of ARM actions. In Great Britain, Animal Rights activists are not responsible for attacks involving high-grade explosives and resulting in the death of targeted persons. From a sociological perspective, this is unlikely to change (see Taylor [10]) due largely to the emotional values that usually provide the impetus for involvement with the animal rights movement or similar issues.

Although one may face opposition when attempting to label the ARM as 'terrorist', their *modus operandi* is similar to that of the autonomous networks that are indicative of modern terrorism. We may define any sub-national group as terrorist if it performs premeditated actions that are:

a) Politically or ideologically motivated,
b) Targeted against non-combatants,
c) Violent or destructive, and
d) Executed in a clandestine manner.

The ARM is clearly encapsulated by this definition and it is hoped that the results of this study may be extended to other groups which—while operationally similar—have a tendency for more extreme violence.

Another culturally distinguishable group modelled in the TerrorPlex simulation is Pro-Test. This is the pro-animal-testing movement which emerged during 2006 in response to the campaigns waged by the ALF. In our model, any agent 'in favour' of animal testing is identified as Pro-Test, broadly reflecting the similar autonomy to be found within the ALF.

To better understand the dynamics underlying the growth of terrorist organisations, we chose to build an Agent-Based Model (ABM). After a review of ABM frameworks available, the Java distribution of RePast [9] was selected. The basic principles for the agent interaction were derived from the published ALF guidelines and other externally documented factors pertaining to clandestine and terrorist groups.

This paper provides a brief description of the simulation model, comprising of a simply definable environment and agent attributes. We shall then cover functionality of the particularly pertinent features of opinion exchange and recruitment, before detailing how a mass media influence is implemented. We then review key results, including circumstances in which the ALF 'self annihilates'. Readers at any time seeking further clarity and a more broad range of recorded experiments are directed to Butler [3]. Code for the TerrorPlex is available on demand from the web or by request from the authors.

[1] The Animal Liberation Primer by Animal Liberation Front Information Services is available from: http://www.animalliberationfront.com/ALFront/ALFPrime.htm [Accessed 21 June 2007].

2 The TerrorPlex Model

We present an ABM by first specifying the environment, then the agents' characteristics or *state*, and finally the agents' behaviour [2].

2.1 Environment

During runtime, the model is viewed as a two-dimensional grid, but is in fact a three-dimensional torus, or 'doughnut' shape. The world inhabited by the agents is not the conventional physical-proximity model used to such great effect by pioneers such as Epstein and Axtell [5]; it is a metaphysical domain intended to acknowledge the pervasiveness of modern communication technologies. The metaphysical arrangement represents an environment where each agent has a number of neighbours with which it communicates regularly. The proximity is not necessarily representative of a physical closeness, but rather a psychological familiarity. In this way, the simple model accommodates a range of social interactions, with neighbouring agents potentially representing school friends, work colleagues, family and household members, as well as Internet chat room or notice board associates.

The effect of the Internet in particular is highly contested as either a moderating or polarising force. It was noted that, whilst it would be difficult to identify cases where the Internet has resulted in an individual turning to violence, it would be impossible to gauge how many have been turned from violence by the same medium. We handle the impact of mass media differently from the social network, as will be discussed below under "agent behaviour".

2.2 Agent Characteristics

Agents within the simulation are objects. As such, they are endowed with a number of attributes which constitute something akin to a 'personality', plus other attributes used for housekeeping functions. Personality attributes are randomly generated numeric values that fall between 1 and 100. These are identified as Law; an agent's inherent respect for law and order, and Pacifist, which represents the agent's pacifism. We may call these *fixed parameters*, since once set, these values remain constant for the duration of an agent's lifespan. Law and Pacifist play an important role when used in conjunction with a measure of personal isolation for discerning an agent's suitability for recruitment into an extremist group. Broadly then, an agent's tendency towards violent action can be calculated like so:

$$Tendency = \frac{Isolation}{Law \times Pacifism} \tag{1}$$

What we refer to as 'housekeeping' attributes are necessary for controlling the simulation and include, for example, those which identify whether an agent is an ALF member, describe an agent's location, and maintain a record of any groups to which an agent belongs. In this model, an agent's position is another example of a fixed parameter, whereas details of membership can be classed as *dynamic*, since these may vary throughout the course of a simulation.

Table 1. Agent attributes

Attribute	Type
x	integer
y	integer
law	integer
pacifist	integer
isolation	integer
myGroup	integer
lifetime	integer
myDuration	integer
isProTest	boolean
isALF	boolean
isActive	boolean
CIS	int array

Table 2. Cultural thresholds

No. of 1s	Cultural Type	Colour
< 2	Pro–Test	Blue
2 - 7	Neutral	Green
> 7	ALF	Red

Cultural Identity To describe the views of an agent in relation to animal testing, we employed a Cultural Identity String, or CIS. The CIS is an integer array consisting of a binary pattern, which is at the core of opinion exchange within the model. For each agent, this array is randomly generated when the agent first enters the simulation. Three conditions may arise from the content of this full dynamic parameter. With respect to the animal testing debate, an agent which has a large majority of 0s in its CIS is designated as Pro-Test, a large majority of 1s means ALF, and a mixture of values within upper and lower bounds indicate that an agent is Neutral.

2.3 Agent Behaviour

On each cycle of the TerrorPlex simulation, agents engage on a course of behaviour dependent on their cultural identity. Behaviour includes the exchange of opinion and, for some agents, recruitment. For example, ALF agents maintain their cultural identity by not updating their own opinion in accordance with their neighbours. However, if their tendency towards violent action increases to a pre-defined threshold then they become activist agents. Agents undergo an increased propensity for violence if they become more isolated or if the *terror gradient* increases. The terror gradient is calculated periodically and is a measure of the difference in the prevalence of terrorist activity between two points in time. Neutral agents possess a more substantial behaviour repertoire. They may assimilate the cultural identity of neighbours, incorporate mass media opinion and potentially develop an affiliation with Pro-Test or the ALF.

Opinion Exchange Opinion exchange within the model is achieved on a local basis with net effects propagated—where applicable—throughout the agent population. The mechanism for these exchanges is known as tag-flipping, described by Epstein and Axtell [5]. An agent eligible for interaction compares the value held at a randomly selected location within its CIS with the corresponding value in each eligible neighbour's CIS. If the majority of neighbouring agents share the same value, then no change is made.

However, if most neighbours have a different value for the particular CIS location, then the agent flips its own value to match the neighbouring majority.

Deffuant et al. [4] provide a model for opinion exchange between agents in which a threshold governs whether an exchange is made. If two agents have an opinion that differs by a margin exceeding the threshold, then no exchange of opinion is possible. They call this 'bounded confidence'. A threshold, per se, is not used here but rather the social divisions or cultural thresholds that are implied by the makeup of agent's CIS string. These divisions are essential aspects of our model. Further to dictating an agent's behaviour, they may act as barriers to communication, so that neutral agents may—given certain parameters—identify ALF and Pro-Test agents as beyond the scope of their cultural interactions. Similarly, ARM agents are *always* omitted from opinion exchanges as the secrecy of such individuals would be closely guarded.

Recruitment Gould [6] states that the process whereby individuals become involved in movements through people whom they know, is one of the "most frequently cited facts about social ties and activism". It will come as no surprise then, that the recruitment of ALF agents into activist groups by ARM agents plays a pivotal role in our simulation. A very small proportion of activist agents come to exist in the model under their own volition (see above). ALF agents whose tendency towards violent action reaches a pre-defined 'conversion' threshold become ARM agents. These individuals are responsible for recruiting ALF agents into their clandestine groups, and in turn the recruited become the recruiters. A more complete model could provide mechanisms for group norms and polarisation, but that is beyond the scope of the present work[2].It should be noted at this point, that although ARM agents must actively recruit ALF agents, there is no such requirement for an agent to become culturally Pro-Test or ALF. This is justifiable since Pro-Test members and (non-ARM) ALF sympathisers do not operate illegally, and so are not obliged to exercise the degree of secrecy that the ARM must do to avoid prosecution.

For an agent to be recruited into an ARM group, several conditions must be met:

1. There must be communication (physical or virtual) between both parties.
2. There must exist mutual trust between these two, although not necessarily between the new member and all existing group members.
3. The person being recruited must be known to share the same ideologies as those of the group.
4. Based upon the person's ideologies and other attributes, it should be deemed likely that they would wish to become part of the group.
5. The group must want a new recruit.

These rules are met within the simulation as follows:

Communication may exist between agents who are classed as 'neighbouring' in the radial catchment area within the social space defined by the environment and described

[2] Berry et al. [1] present an interesting paper in which entire cliques are eligible for conversion to terror when their 'disgruntlement'—calculated as an average over all members—exceeds some threshold.

earlier. Parameters for the catchment—either von Neumann or Moore neighbourhood methods (see below) and the length of the radius—are controlled through the interface.

Mutual trust is established between two agents after having been neighbours for a pre-defined number of simulation cycles. This may appear to be a rather crude approach to a very complex notion, but one must recall the nature of the simulation space; agent proximity represents an emotional connection. Therefore, it should not be too unreasonable to assume that a prolonged period of emotional familiarity may engender a degree of trust between two parties.

As human beings are adept at assessing others' values through contextual, conversational and visual clues, ARM agents are able to identify those neighbours who share ALF ideologies by direct access to the variable state.

The likelihood that an agent could be recruited is represented by their tendency towards violent action. This is a personal attribute, but if we have already said that mutual trust exists, then we can suppose that such knowledge has been accumulated by the recruiting party. An existing ARM agent will access the '*tendency*' value of a potential recruit and if it falls below a set threshold then the ALF agent is eligible for recruiting.

ARM agents accept suitable recruits until the membership quota of their cell (if any) is met.

2.4 Mass Media

With careful definition of the agent space, we have implicitly managed of one aspect of modernity—communication technologies—by not restricting ourselves to an environment of physical proximity. Also to be considered is another phenomenon unique to recent history; mass media. The effects of mass media cannot be handled as conveniently as modern communication devices, but demand an explicit solution. Indeed, without such an approach, we would be unable to exercise any convenient control over the behaviour of mass media within our model and thus analyse recorded variations.

Mckeown and Sheehy [8] build upon the bounded confidence model [4] by introducing a second mechanism for mass communication. This mechanism involves the interaction of a fixed opinion, representing a media source.

Because our study concerns the two opposing forces of ALF and Pro-Test, it is sensible that we incorporate two corresponding media opinion sources. These are essentially two CIS elements, but deliberately and statically biased in opposite directions. The TerrorPlex simulation by default asserts a probability of 0.7 that the binary values will be 0 for one media source, and a probability of 0.3 for the other source. Wherever possible, randomising functions are used. When the mass media functions are enabled, an agent will incorporate one of these media sources into the tag-flipping exercise as though it were a neighbouring agent. Whilst a neutral agent will liberally select at random a media source to assimilate, ALF and Pro-Test agents will not be so indiscriminate. These agents will ignore the media source which does not represent their cultural identity, thus facilitating the reinforcement of ideologies that is found to act within groups.

To add realism to our model, we have employed an innovative approach to mass media representation by harnessing the power of feedback. To paraphrase Wardlaw [11],

when terrorism becomes institutionalised, the perceived significance of terrorist acts is reduced. Therefore, it may follow that the effect of media bias upon an individual will be less if violent acts are routinely reported. Indeed, if violence becomes commonplace, it may not be reported at all. With this in mind, a 'terror gradient' has been implemented whereby the media presence in an agent's tag-flipping routine is weighted in accordance with the relative change in terrorist activity from one period to another. A significant change in activity will herald a more significant effect upon the routine. For our purposes, we shall say that there is a direct correlation between the number of activist agents and the prevalence of terrorist acts.

2.5 Simulation Execution

At run-time, an initial representation of the model is created. Data channels are defined for both real-time graphical output and the recording of simulation data for analysis. These channels are updated upon each cycle. Elements recorded to disk are the number of ARM, ALF, Pro-Test and Neutral agents, plus tag-flipping activity, the number of ARM groups and the terror gradient.

Every aspect of the model has a default value within the code. However, the majority of these may be changed through the user interface. Before execution then, the user can easily change parameters which include environment size, population size, whether mass media is present, and the range of an agent's view. One may also express a preference for either a von Neumann or Moore neighbourhood method [7]. The von Neumann method searches for neighbouring agents in the North, South, East and West directions only, so that given a range of one, the maximum number of neighbours would be four. Moore, however, includes the diagonal directions North-East, South-East, South-West and North-West, potentially yielding eight neighbours when given the same range.

Table 3. TerrorPlex simulation parameters

Parameter	Type	Entry
Population	Integer	Direct
GridSize	Integer	Direct
RestrictGroupSize	Boolean	Check–box
MassMedia	Boolean	Check–box
NeighbourMethod	Integer	Drop–down List
Range	Integer	Direct
IgnoreOtherCultures	Boolean	Check–box
Respawn	Boolean	Check–box
StepLimit	Integer	Direct

The three-dimensional torus environment is generated and projected onto the screen as a grid, and this is populated with a randomly positioned agent population. Each agent object is instantiated with its attributes as described earlier.

A schedule is defined which asserts that each agent shall undergo one of several behavioural routines (also described earlier) once every cycle. Which routine is selected will depend upon the cultural identity of an agent and any relevant parameters. The program does not iterate through a sequential list of agents, but rather a list that is randomised at the beginning of each cycle. The schedule also ensures that the terror gradient is calculated at regular intervals.

3 Results

The experiments we conduct provide a substantial dataset supporting interesting real-time observations. The most noteworthy results can be seen in tests for which the *Respawn* option is active. Respawn regularly replenishes the agent population by removing those who have reached their assigned lifespan and reintroducing to the simulation an equal number of new agents. Respawned agents are created with their attributes being randomly-assigned under the same rules as at simulation startup. When this option is not in use, the population stagnates after relatively few iterations of the simulation and does not allow us to enjoy the dynamics of changeable social interactions (see Figure 1(d)). Notably, in human societies we also see more social change and sometimes radicalism come from areas such as university towns where there is a steady but not overly rapid turn over of population. If turn over is too rapid, trust cannot be established.

Some expected simulation behaviours were recorded, that include an excellent example of clustering. This occurs where neutral agents are free to interact with ALF and Pro-Test agents, but the 'respawn' option is disengaged. Under these conditions and an absence of mass media presence, we see that the three cultures achieve a similar abundance within the population in distinctive clusters. In examples where the mass media mechanism is applied, clustering is still evident, but the number of neutral agents will dwindle until none remain.

Tests also provided a few surprises. It was shown that placing limits on the size of activist groups was absolutely necessary for avoiding a saturation of activist agents within the ALF community. Emergent terror cells in the clustering examples mentioned above are very sparse when the default group size restriction of 5 is imposed. If the restriction is lifted however, ARM membership increases to include nearly the entire ALF contingent. We had originally assumed that the terror cells would be limited in size by the inherent difficulties involved in recruiting suitable candidates. In real societies, cell size is also limited by policing, since the larger the number of members, the higher the probability that the cell is discovered. Adding this probability of discovery into the TerrorPlex would be interesting further work.

The polarising effect of mass media can be observed throughout the tests as expected. Polarisation of opinion is evident without mass media, but at a much-reduced rate. An unforeseen phenomenon though, is that of a self-annihilating ALF. What we witness is the initial growth of both ALF and Pro-Test groups. This may also be accompanied —in the case of unrestricted group size—by a proportional increase in ARM activity or, as we see with low-value size restrictions, a steady increase in ARM numbers. In either situation, and for other restriction values between these two extremes, there

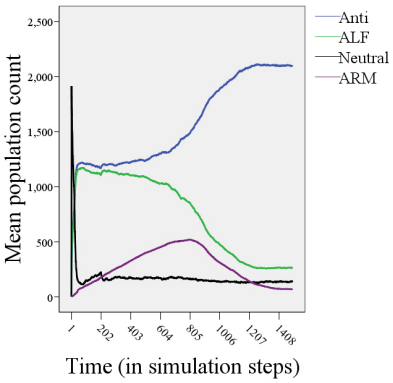

(a) mass media, respawn, cell limit of 5

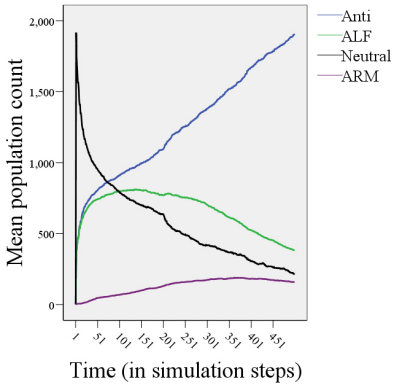

(b) no mass media, respawn, cell limit of 5

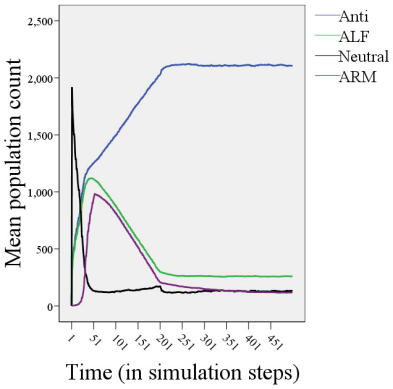

(c) mass media, respawn, no cell limit

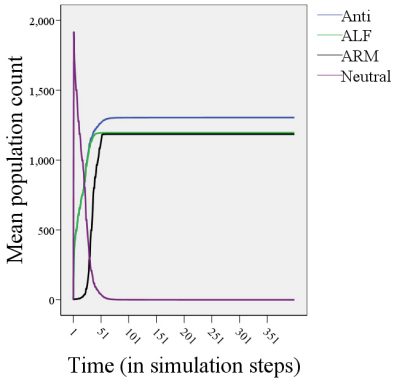

(d) mass media, no respawn, no cell limit

Fig. 1. Plots showing means for 15 runs in each condition. *Anti* here represents Pro-Test. *Mass media* determines whether the two mass media influences are available to agents (see main text.) In the *Respawn* condition, each agent is randomly assigned a lifespan of between 0-200 cycles, otherwise agents enter the simulation only at its beginning. *Cell limit* indicates the size at which terrorist cells no longer recruit even if compatible agents are found.

is a clearly observable phenomenon (see Figure 1). After the initial polarisation into ALF and Pro-Test groups, there is a period of 'power sharing' between the two cultural identities. This gives way to a divergence of fortunes however, where ALF membership diminishes suddenly, with Pro-Test agents increasing in number at an equal rate to gain almost complete consensus. This reduction in ALF members is closely related to ARM membership and, in fact, is easily explainable by the stipulations that we have made with respect to the transfer of cultural identity.

For reasons of personal security, ARM activists exist outside of the normal social context. As such, they are omitted from the cultural view of other agents. A 'neutral' agent, when assimilating its neighbours' culture, will not involve an ARM activist, but *will* involve an ALF member. If there is a situation where the ARM population has increased to a majority of the ALF population, then an agent involved in cultural exchange will have mostly other neutral agents and Pro-Test agents to engage with. This imbalance quickly leads to a simulation space dominated by Pro-Test agents. The proliferation of ARM activists has effectively led to the annihilation of the entire ALF movement.

As mentioned, tests have also demonstrated the importance of size restrictions upon ARM groups, showing that they have a tendency to balloon in size where restrictions are not in place. Herein lies a surprising and interesting relationship: The sustainability of the ALF population actually increases with the application of more stringent restrictions on ARM group size. Far from aiding the growth of the ALF, the presence of ARM agents causes a decline by increasing the probability of media attention, but concurrently removing culturally influential ALF members (those that have now turned to activism). In the TerrorPlex simulation, the ALF population is only able to compete with Pro-Test if the number of activist agents is kept low.

4 Conclusion

Beginning with just a simple definition of agent behaviour, we have been able to create a fertile simulation environment which enables us to explore emerging culturally complex scenarios for modelling terrorist recruitment and other sorts of ideological social affiliative behaviour.

Local opinion exchange—enhanced with the concept of bounded confidence—has been developed with the addition of polarised mass media sources. This mass media influence has been further advanced with a dynamic terror gradient which enables cultural changes within the agent population to be fed back into the model.

The element of feedback has, in turn, led to some fascinating results that include an example of how autonomous terrorist networks can, through their own prosperity, cause a cultural imbalance which leads to the eventual demise of not just the terrorist faction but the entire movement.

References

[1] Berry, N., Ko, T., Moy, T., Smrcka, J., Turnley, J., and Wu, B. (2004). Emergent Clique Formation in Terrorist Recruitment. In *Agent Organizations: Theory and*

Practice, Session on Organisational Models. The AAAI Press, Menlo Park, California. American Association for Artificial Intelligence workshop.

[2] Bryson, J. J., Ando, Y., and Lehmann, H. (2007). Agent-based models as scientific methodology: A case study analysing primate social behaviour. *Philosophical Transactions of the Royal Society, B — Biology*, 362(1485). in press.

[3] Butler, S. (2006). A simulated exploration into the growth of modern terrorist networks. Technical Report CSBU–2006–18, University of Bath, UK. Honours Undergraduate Dissertation, with Joanna J. Bryson.

[4] Deffuant, G., Neau, D., Amblard, F., and Weisbuch, G. (2000). Mixing Beliefs Among Interacting Agents. *Advances in Complex Systems*, 3:87–98.

[5] Epstein, J. M. and Axtell, R. (1996). *Growing Artificial Societies*. The Brookings Institution, Washington, DC.

[6] Gould, R. V. (2003). Why do Networks Matter? Rationalist and Structuralist Interpretations. In Diani, M. and McAdam, D., editors, *Social Movements and Networks*, pages 233 – 257. Oxford University Press.

[7] Hegselmann, R. and Flache, A. (1998). Understanding complex social dynamics: A plea for cellular automata based modelling. *Journal of Artificial Societies and Social Simulation*, 1(3). Available from: http://www.soc.surrey.ac.uk/JASSS/1/3/1.html.

[8] Mckeown, G. and Sheehy, N. (2006). Mass Media and Polarisation Processes in the Bounded Confidence Model of Opinion Dynamics. *Journal of Artificial Societies and Social Simulation*, 9(1). Available from: http://jasss.soc.surrey.ac.uk/9/1/11.html.

[9] North, Collier, Vos, Najlis, and Maciorowski (2004). Repast Revolution: An Overview of New Repast Developments. In *Social Dynamics: Interaction, Reflexivity and Emergence*.

[10] Taylor, B. (1998). Religion, Violence and Radical Environmentalism: From Earth First! to the Unabomber to the Earth Liberation Front. *Terrorism and Political Violence*, 10(4):1–42.

[11] Wardlaw, G. (1982). *Political Terrorism*. Cambridge University Press, Cambridge, UK. Australian Institute of Criminology.