

Nuclear Terrorism and Rational Choice

submitted by

Simen Andreas Ellingsen M. Sc.

for the Degree of Doctor of Philosophy

King's College London
University of London
Department of War Studies

2009

Abstract

The prospect of nuclear terrorism, terrorist acts with nuclear fission explosives, is analysed by means of rational choice theory, a methodology borrowed from economics which has hitherto not been systematically applied to nuclear terrorism. The methodology allows the formalisation and modelling of key choices faced by both the aspiring nuclear terrorist and a potential target government in order to work out best strategies under the assumptions that the players are rational and intelligent.

Four relevant decision situations are studied: The terrorist's choice of whether to embark on an ambitious and expensive nuclear project or to stay with tried and trusted conventional methods; The choice of fissile material for a terrorist bent on building a nuclear weapon: highly enriched uranium (HEU) and plutonium as fissile material; The government's choice of prioritising between branches of fissile materials safeguards (HEU versus

plutonium); and the strategic interplay between terrorist and government in the case where the terrorist has acquired a nuclear weapon and must decide whether to use it to attack, for extortion (blackmail) or to deter an attack upon his own interests.

Several key conclusions reached are of direct policy applicability. A simple decision theoretical analysis shows that heavy emphasis on HEU over plutonium in safeguards measures is justified. It is demonstrated that relative deterrence (by denial) of nuclear terrorism in favour of conventional means is possible, and the conditions for which are found. It is found, moreover, that to use an acquired nuclear weapon for blackmail or deterrence purposes is almost never preferable for a terrorist, and the best response of a government to an explicit nuclear terrorist threat is almost always forceful response.

Table of Contents

- 1 - Introduction, Literature review and methodology	9
1.1 Chapter overview	9
1.2 Definitions and terms	10
1.3 The nuclear terrorism dispute: Would they? And could they?	12
1.4 Notable recent contributions to nuclear terrorism research	17
1.5 Previous applications of game theory to terrorism	19
1.6 Conclusion	20
- 2 - Methodology: problems and strengths of rational choice theory	21
2.1 Chapter outline	21
2.2 Rational choice and terrorism	21
2.3 Game theory and decision theory: concepts and an example	22
2.4 Claimed capabilities of rational choice models	24
2.5 Criticisms of rational choice models on philosophical grounds	24
2.6 The sociologist criticism	26
2.7 Criticism on the grounds of poor performance	26
2.8 Reduction of a state to a single mind	27
2.9 Are terrorists rational?	30
2.10 Evaluating the utility of gaming	31
2.11 Conclusion	32
- 3 - Keeping fissile materials out of terrorist hands: HEU, plutonium and a prioritised response	33
3.1 Research question	34
3.2 Literature overview and outline of chapter	34
3.3 Terrorist preference: HEU	34
3.4 Making a crude Pu implosion device	35
3.5 Availability of nuclear materials: safeguards and challenges	42
3.6 Conclusions so far	46
- 4 - Safeguards: HEU vs. Plutonium – a formal analysis	47
4.1 Chapter overview	47
4.2 Introduction to the methodology and a toy model	47
4.3 The model	50
4.4 Numerical estimation of parameters	55
4.5 A numerical example for further analysis	61
4.6 Assumptions and limitations of the model	65
4.7 The gains of gaming	67
4.8 Conclusions	68
- 5 - The HEU/Plutonium choice: From the Terrorists' Point of View	69
5.1 Research question and outline	69
5.2 Acquiring the material: strength or stealth?	69
5.3 A different angle	71
5.4 The choice of fissile materials from the terrorist's point of view	72
5.5 Analysis: the price of time and the fear of failure	77
5.6 Fissile material in small batches: the hodgepodge bomb	83
5.7 The gains and weaknesses of gaming	84
5.8 Conclusion: policy implications	85

- 6 - Deterring terrorists from attempted use of nuclear weapons	87
6.1 Research question and scope	88
6.2 Chapter outline	88
6.3 The literature on terrorism and deterrence	88
6.4 The model	94
6.5 A discussion of means of deterrence in light of a simulation	96
6.6 The advantages of gaming versus a qualitative approach	105
6.7 Conclusions and policy implications	106
- 7 - Nuclear blackmail and other strategic uses of a terrorist nuclear weapon	108
7.1 Research question and chapter outline	108
7.2 Example: al Qaida	109
7.3 Principal issues of use or non-use	110
7.4 The credible threat	111
7.5 Scenarios and incentives	112
7.6 Strategic interplay of nuclear blackmail: a gaming approach	115
7.7 The merits of modelling	122
7.8 Conclusions and policy implications	123
- 8 - Conclusions and outlook	125
8.1 Conclusions from research chapters	125
8.2 Comparing policy implications of different chapters	128
8.3 The proof of the pudding: has the research approach been effective?	129
8.4 Where to from here?	131

Appendices

- A - List of seizures of attempted smuggling of fissile materials	136
- B - Numerical programme used in chapter 6	138
- C - Some mathematics	140
C.1 Functions	140
C.2 Graphs and plots	141
C.3 The derivative: the slope of the function	141
C.4 Functions of two variables and the partial derivative	142
C.5 Summation of power series	144
- D - Some thoughts on a 'Second Line of Defence'	145
D.1 Second layers in literature and policy: a brief review	145
D.2 A diagrammatic outline of the paths to nuclear terrorism	149
D.3 Value for money in second layers: the challenge	149
D.4 A simple model of nuclear terrorism countermeasures	150
D.5 Outlook: Optimum funding and organisation of the 'second line of defence'	153
D.6 Some early conclusions	154

Bibliography

Dissertations	154
Books and monographs	154
Reports (governmental and non-governmental)	156
Papers, articles, reviews and chapters in edited volumes	157
Other documents	162
Speeches, testimonies and transcripts	162
Press reports and op.ed.s	163
Web pages and software	163

List of Abbreviations

CBP	US Customs and Border Protection
CIA	Central Intelligence Agency
CSI	Container Security Initiative
C-TPAT	Customs Trade Partnership Against Terrorism
DoD	US Department of Defense
DoE	US Department of Energy
DoS	US Department of State
EU	The European Union
FBI	Federal Bureau of Intelligence
FY	Fiscal Year
GAO	US Government Accountability Office
ICBM	Intercontinental Ballistic Missile
HEU	highly enriched uranium
kT	Kilotonne (1 million kg) TNT
LEU	low enriched uranium
MOX	mixed-oxide fuel
MPC&A	Materials Protection, Control and Accountability (US DoE) National Nuclear Security Administration
NNSA	(US DoE) National Nuclear Security Administration
NPT	Non-Proliferation Treaty
NRDC	US Natural Resources Defense Council
NSG	Nuclear Suppliers Group
NYC	New York City
OTA	Office of Technology Assessment
PSI	Proliferation Security Initiative
Pu	Plutonium
SIS	Secret Intelligence Service (UK)
U	Uranium
UK	The United Kingdom of Great Britain and Northern Ireland
UN	United Nations
US	The United States of America
WW2	The Second World War

List of figures

Illustration 1.1.: Papers on nuclear terrorism in a 'would-could' diagram.	18
Illustration 1.2.: A cartoon representation of the research approach (author's sketch)..	19
Illustration 4.1: Utility function of toy game.	48
Illustration 4.2.: Model utility function with two free variables.....	49
Illustration 4.3.: Layout of safeguards game: HEU vs. Pu.....	50
Illustration 4.4.: Damage area of a 10kT bomb at Grand Central Station New York.....	57
Illustration 4.5.: Contour plot of numerical example utility function.....	62
Illustration 4.6.: Zoom of figure 4.5 with two possible scenarios and best responses. ..	63
Illustration 4.7.: Contour plot of same utility function with incremental expenditure.	64
Illustration 4.8.: Contour plot of same utility function with fixed total cost.....	64
Illustration 5.1.: The HEU/Pu game seen from the terrorist's point of view.....	73
Illustration 5.2.: Plot of the fraction on the left side of eq. (5.6).....	75
Illustration 5.3.: Expected waiting time, equation (5.19).	81
Illustration 5.4.: The deterrence time, equation (5.20).	81
Illustration 6.1.: The game of nuclear vs. conventional strategy.	95
Illustration 6.2.: Average payoff per round for different values of δ	98
Illustration 6.3.: The standard deviation of accumulated payoff per round.....	101
Illustration 6.4.: Standard deviation relative to expected payoff as functions of δ	102
Illustration 7.1.: Simple decision theoretical model of nuclear terrorist blackmail.	115
Illustration 7.2.: The general layout of the two-player nuclear blackmail game.....	116
Illustration 7.3.: The game in its simplest embodiment.....	119
Illustration 7.4.: Generalised game: response could fail, dismissal has cost.....	121
Illustration 7.5.: The game assuming $A > G$ and nonzero direct attack failure rate.....	121
Illustration 7.6.: The game of fig. 7.5 with forceful response penalty.....	122
Illustration B.1.: Structure of programme of chapter 6.....	138
Illustration B.2.: The programme structure in more detail (with pseudocode).....	139
Illustration C.1.: Examples of functions: 'Costa function' (left) and BMI (right).....	142
Illustration C.2.: A graph of a function of one variable.....	142
Illustration C.3.: A 3D graph with a single maximum.....	143
Illustration D.1.: A simple model of second layers.....	151
Illustration D.2.: Different paths to a nuclear weapon.	152

List of Tables

Table 4.1: Estimated economic impact of terrorist attacks.	56
Table A.1: List of smuggling incidents with fissile materials.....	136

Preface

First and foremost I thank my supervisors who have supported my work in innumerable ways throughout these years. Professor Peter D. Zimmerman (now Prof. Emerit.) got me interested in the subject of nuclear terrorism in the first place and persuaded me to move to London and embark on the work which was to result in this thesis. His immense knowledge and wise guidance has been a major help and inspiration, and without his encouragement this project would not have made it past the first year.

Equally heartfelt thanks go to Dr. James Acton who has invested enormous time and energy to get me through the last stages of this project. After taking over as my 'unofficial' supervisor he has worked through every page of each chapter with painstaking thoroughness, often several times, and rooted out a large number of inconsistencies, factual errors and logical flaws. Much of the honour for the completion of this work belongs to James whose combined knowledge and persistence I could not have managed without.

I received much help during the initial stages of the project from Dr. Morten Bremer Mærli, who helped me write the project proposal and get me introduced to the non-proliferation community in Norway. Much encouragement and inspiration was also received from Dr. Steinar Høybråten at the Norwegian Defence Research Establishment, for which I am grateful. Dr. Michael S. Goodman moreover provided me with extremely useful feedback on the introductory chapters of the thesis, and certain parts of the formal analysis in the thesis was inspired and aided by discussions with my office mate and fellow physicist Kristian Etienne Einarsrud.

Beatus homo, qui invenit sapientiam, et qui affluit prudentia.

I am grateful for all the support I have been given at the Department of Energy and Process Engineering at the University of Science and Technology in Trondheim, Norway, my employer during the second

half of the research period. I have been allowed the necessary flexibility which allowed me to keep working on my 'London PhD', and the understanding from my fluid mechanics teaching boss, Professor Lars Sætran, has been instrumental and enabled me to fly to London for supervision. Likewise my supervisor in Trondheim, the brilliant Professor Iver Brevik has been supportive all along, even if my rate of production has been slow over the last year from the burden of doing two PhDs simultaneously.

Finding my way through the bureaucracy of the British university system would have been a daunting task for me, had it not been for the great amount of help and service I received from Gill Woods and Helen Fisher at the Department of War Studies. I could hardly have done this without you!

Of great importance was the support and encouragement I received from three of my fellow PhD students at War Studies through many a dark hour during my year in London, my good friends Graham Gerard Ong-Webb, Valerie Arnould, and Bhavini Rama. Gerard was also the one to get me started with rational choice theory in the first place and persuaded me to take the course *Game Theory for Political Science* at London School of Economics with him.

Further thanks should go to my family for support and patience throughout, and to my friends, primarily Erik who suffered my stressful presence during a year of living in the same apartment.

Last, but most importantly, my thanks go to my wife Inger whose caring and support has kept me afloat through many a dark hour. I don't think I'd have finished this project without your indefatigable encouragement. You'll hopefully be seeing a little more of me in the evenings now this thesis is handed in.

Qui gloriatur, in Domino gloriatur.

Introduction, Literature review and methodology

Although the freezing international climate of the Cold War seems to have thawed considerably, the nuclear threat is according to many analysts, as prevalent as ever. Director General of the IAEA, Mohamed El Baradei, for one, concluded in 2005 that 'the threat of nuclear war has never been greater'¹. Much as the calculi of deterrence and massive nuclear retaliation may have become less prevalent, nuclear arms could be gaining newfound strategic importance for smaller actors: against a technologically advanced adversary it is a weapon that, even in small numbers, can immediately compensate for inferior conventional capacities. Nuclear weapons 'no longer represent the frontier of technology' Betts sums up; 'Increasingly, they will become the weapons of the weak'². Non-state actors count among those who could see such potential.

Recent events have shown with terrifying clarity that the new breed of terrorists is not averse to killing civilians in the hundreds and thousands for what they see as the ultimate cause, and the notion by Jenkins that 'terrorism is theatre' and 'terrorists want a lot of people watching, not a lot of people dead'³ has been called into question by many analysts⁴.

Yet the successful detonation of an atomic bomb by a terrorist group could dwarf even the attack on the World Trade Center on September 11, 2001⁵. Scholars,

journalists⁶ and politicians alike have seen the possibility of terrorists acquiring and detonating a real nuclear weapon, painting perhaps the grimmest picture of a terrorist attack imaginable. A coarse calculation by Bunn and co-workers⁷ estimates the consequences of a terrorist nuclear weapon of yield 10kT if detonated at the Grand Central train station, New York, on a normal day of the week. Such a bomb would reduce around one square kilometer of a city to burning rubble⁸. The number of dead is conservatively put at 500,000 and the direct economic cost to more than \$1 trillion. The indirect and long-term costs, both in money and life, will likely be much higher, and, former UN Secretary General Kofi Annan has predicted, the economic plummet following the incident will cast millions of people worldwide into hardship and poverty⁹.

The imminence of the above described threat, and how to deal with it, has been the source of debate for several decades. The aim of this thesis is to move the debate further by analysing different relevant scenarios and possible countermeasures by means of formal models and rational choice methodology.

1.1 Chapter overview

I review the existing literature on nuclear terrorism and show that although a number of important contributions have been made and progress has been significant over the last decade or so, the discussion has tended to focus on terrorists' nuclear intent and capability. I argue that a fruitful avenue for further research progress is to step beyond these questions and analyse scenarios and consequences given

1 E. Follath & G.M. Mascolo 'Al Qaida also wants the Bomb' *Der Spiegel* (February 8 2005)*

2 Richard K. Betts 'The New Threat of Mass Destruction' *Foreign Affairs* 77:1 (1998) p.27

3 Brian Michael Jenkins 'Will Terrorists Go Nuclear?' RAND paper (November 1975)* p.4

4 Including Jenkins himself: Brian Michael Jenkins 'The New Age of Terrorism' Chapter 8 of Kamien (ed.) *The McGraw-Hill Homeland Security Handbook* (New York:McGraw Hill, 2006)* p.118.

5 The total energy released in the terrorist attacks on Lower Manhattan - the kinetic energy of the two planes, the exploding aircraft fuel and the potential energy released as the two buildings collapsed - has been calculated to add up to approximately the equivalent of 0.2 kilotonnes (kT) of TNT. While some perspective is offered by this, such numbers are not directly comparable due to the very different way in which the energy is released. While the yield of the Hiroshima

bomb was 65 times the energy release of September 11, relative difference in damage could be considerably larger than this. BBC 'The Destructive Forces Unleashed' (September 18 2001)*

6 E.g. the op.ed. in New York Times the very morning this thesis was sent for printing: Jeffrey Goldberg 'On Nov. 4, Remember 9/11' *The New York Times* op.ed. (September 9, 2008)*

7 Matthew Bunn, Anthony Wier and John P. Holdren *Controlling Nuclear Warheads and Materials* report of the Project Managing the Atom (Belfer Center for Science and International Affairs, Harvard University, 2003)*

8 See figure 4.4. The area would stretch from the southern end of Central Park to Madison Square Park, and Times Square, the UN building, the Theatre District and Madison Square Garden would all be in ruins.

9 Kofi Annan, Secretary-General of the United Nations, 'A Global Strategy for Fighting Terrorism' Keynote Address to the Closing Plenary of the International Summit on Democracy, Terrorism and Security (Madrid, 10 March 2005)*

assumptions about the nature and stature of the terrorist. Rational choice methodology is proposed as a means to achieve this end.

1.2 Definitions and terms

I define 'terrorism' as

Acts or threats of strong violence targeting civilians usually including destruction of civilian property, including planning of such acts or threats, by non-state actors with the purpose of creating a condition of fear, drawing attention, creating instability, and affecting an audience beyond the victims directly targeted.

The definition is adapted from that used by Maerli¹⁰. I define 'nuclear terrorism' as

Terrorism whose primary means is explosion by nuclear fission or fusion, or reasonably educated attempts at such.

A 'terrorist' is a person guilty of terrorism, a 'nuclear terrorist' one guilty of nuclear terrorism.

A few notes to the definitions chosen are called for. There exists no real consensus amongst analysts as to how 'terrorism' should be defined, and there are a number of alternatives to choose from¹¹. The most important problem pointed out is that the notion of terrorism is often in the eye of the beholder: 'one man's terrorist is another man's freedom fighter'¹². The definition chosen herein emphasises the violence and the dread terrorism deliberately inspires, while placing no restrictions as to what might be the overarching political motive behind the attack. Also, destruction of property alone is not defined as terrorism if it may not be perceived as a threat of violence. No particular effort has been made to delve into the deep waters of philosophy in this respect, and

like other authors before¹³ I interpret 'terrorism' in an operational sense, reserving the right to use the term somewhat pragmatically, recognising that I will inescapably - despite my best efforts to avoid bias - take 'the view from the west'.

The restriction of terrorism to non-state actors excludes e.g. 'terror balance' between states. Many protesters of recent wars have dubbed state actions 'terrorism' - I merely note that this disaccords with the definition employed herein.

Furthermore, the definition of nuclear terror is narrow and will only include true nuclear explosions. Other authors include such actions as nuclear sabotage (destruction of nuclear instalments for radiological contamination) and radiological dispersion weapons¹⁴. These means are disregarded here, for manageability, but also because the scale of the threat and the set of feasible countermeasures pertaining to each of these branches differ greatly, making a distinction natural. Much as for example a failed nuclear explosive (a 'fizzle') might have similar destructive effect as a large radiological dispersion device¹⁵, the former is included in the definition for its attempted nuclear explosion, the latter is not.

I will include only 'reasonably educated attempts' at nuclear explosives, to exclude hypothetical cases where terrorists employ, for example, non-fissile materials in the misguided belief that it will somehow cause a nuclear yield¹⁶. Furthermore I will demand for threats of nuclear explosions to be plausible before counting them as 'nuclear terrorism', to exclude obvious hoaxes, a number of which have surfaced over the years¹⁷.

The definition of nuclear terrorism includes nuclear fusion for completeness. A home-made device utilising nuclear fusion is a scenario so far-fetched it is nowhere discussed, yet were terrorists to acquire a

10 Morten Bremer Maerli *Crude Nukes on the Loose?: Preventing Nuclear Terrorism by Means of Optimum Nuclear Husbandry, Transparency, and Non-Intrusive Fissile Material Verification* PhD dissertation (University of Oslo, 2004) p.13

11 Mark Burgess 'Terrorism: The Problems of Definition' Center for Defense Information (August 2002)*. See also Thomas J. Badey 'Defining international Terrorism: A Pragmatic Approach' *Terrorism and Political Violence* 10:1 (Spring 1998) pp.90-107.

12 Paul K. Davis and Brian Michael Jenkins *Deterrence & Influence in Counterterrorism: A Component in the War on al Qaeda* (RAND, 2002)* p.67.

13 Robin Frost 'Nuclear Terrorism Post-9/11: Assessing the Risks' *Global Security* 18:4 (2004) p.398.

14 e.g. Charles D. Ferguson and William C. Potter with Amy Sands, Leonard S. Spector and Fred L. Wehling *The Four Faces of Nuclear Terrorism* (New York:Routledge, 2005)*

15 Richard L. Garwin 'The technology of megaterror' *Technology Review* 105:7 (2002) p.66

16 For example, a 'nuclear bomb recipe' was discovered on an al Qaida friendly web page prescribing the use of the non-fissile material radium for a nuclear weapon. This would not be a 'reasonably educated' attempt. Sammy Salama and Lydia Hansell 'Does Intent Equal Capability?: Al-Qaeda and Weapons of Mass Destruction' *The Nonproliferation Review* 12:3 (2005) p.636.

17 Michael Levi *On Nuclear Terrorism* (Cambridge, MA: Harvard University Press, 2007) p. 120

boosted fission or fusion weapon, then its use or threatened use would be defined as nuclear terrorism.

Finally I bypass the intricacies of jurisdiction by defining a terrorist by being 'guilty of' terrorism by simply assuming that every person in the world either is or is not guilty of the acts mentioned, independently of whether this has been or can be tried in a court of law. The fact that there exists a grey area will be disregarded as it is of little consequence to the discussion presented herein.

1.2.1 *Research questions and research approach*

The overarching research question to be treated is

In the face of a nuclear terrorist threat, what are the optimal policy countermeasures for a government seeing itself as a potential target?

The research question itself is not new. What sets the present thesis apart from previous work is the research approach which rests upon two key features: (1) explicitly assuming there exists a terrorist adversary with the intent and some capability to mount a nuclear attack, and (2) employing rational choice theory to answer the above by analysing decision processes as seen by both actors in this scenario. The former of these assumptions is not in itself novel; it has been (at least implicitly) made by authors in the past in order to make policy recommendations for nuclear terrorism. The latter point has never previously been employed to the specific threat of nuclear terrorism to the author's knowledge. Point (2) implies assuming both terrorist and government to conform to some definition of rationality, as will be discussed towards the end of the chapter.

The working hypothesis that the utilisation of formal methodology may offer a new perspective is tested indirectly by application to a number of sub-questions to that above, one per chapter through chapters 3 through 7¹⁸, and its success is briefly assessed in chapter 8.

1.2.2 *The thesis*

The thesis is organised so that each chapter has an individual research question and short literature review. Each of the chapter research questions are

18 Chapters 3 and 4 have the same research question.

subsets of the more general research question in the introduction. This choice follows naturally from the chosen research approach, in which the novelty of the work presented rests primarily with the methodology rather than the research questions themselves. It was deemed that readability could be much improved if the many and varied aspects of anti nuclear terrorism policy serving as background knowledge for our analysis were presented as they became of relevance, dispersed throughout the thesis chapters. This gives each of the gaming chapters 3 to 7 a certain degree of independence.

1.2.3 *Limitations and Scope*

While each chapter has its individual research question and scope a few general limitations may be noted. In addition to restricting myself to 'true' nuclear terrorism according to my definition, I will primarily concentrate on projects in which the terrorist, given fissile material, attempts to construct a crude device.

The option of stealing a finished device is an important worry which is not a matter of focus in this thesis for reasons of restrictions of space and time. The issue has been treated expertly before¹⁹ yet it seems probable in the light of the analysis herein that a formal methodology could be fruitfully applied to aspects of nuclear terrorism with finished military weapons as well.

Likewise, the prospect of a state deliberately sponsoring a terrorist nuclear project in some way, while a very serious concern, is not a matter of focus herein. A state which inadvertently transfers nuclear material to terrorists via lax safeguards, however, is of course included. A brief discussion of this issue is provided in chapter 6 as background to the question of deterrence of nuclear terrorism, but is otherwise not treated.

I will implicitly assume throughout that unless explicitly stating otherwise, a terrorist-made nuclear weapon using highly enriched uranium will be of a gun-type design while a plutonium weapon will be an implosion device of some sort (see chapter 3 for details about these terms).

This thesis, like the vast majority of the literature²⁰

19 e.g. Bunn, Wier and Holdren *Controlling Nuclear Warheads and Materials* and Ferguson and Potter *The Four Faces of Nuclear Terrorism*.

20 A notable and most commendable exception where US and Russia are treated as equally important targets is Graham Allison and Andrei Kokoshin 'The New Containment' in

on nuclear terrorism, is somewhat US-biased despite the fact that the United States is far from the only possible target of such an attack; some would argue not even the most likely²¹. Importantly, however, the issue is in every way a global one. The 'antiterrorist player' in the games of chapters 3 through 7 in no way needs to be interpreted as the US government, but using the US as illustration and example is convenient since such an abundance of data exists on US anti-nuclear terrorism efforts. Furthermore, the US has taken upon itself a leading rôle in the matter and has expressed its belief that it is the likely target of future attacks²², which makes it the natural example for gaming purposes. Likewise, al Qaida is sometimes used as example although the terrorist players in the gaming chapters are generic and do not need to represent a single existing group.

1.2.4 Notes on terminology and references

The term 'weapons of mass destruction' and its abbreviation, WMD, will not be used in this thesis excepting quotations. The term is commonly used²³ to collectively denote nuclear, biological, chemical and radiological weapons²⁴. In the author's opinion, only nuclear amongst these can be said to cause 'mass destruction'; biological, chemical and radiological

devices could potentially cause many *casualties*²⁵, but at a typical level of destruction less than that of a conventional car bomb. Indeed, such devices might be in form of food contamination, an aerosol or dispersal of fine powder, causing no direct destruction whatsoever²⁶. By the standard understanding, the anthrax spiked letters killing 5 in late 2001 would be a 'WMD' while the September 11 attacks killing some 3000 would not. The author believes, as others have argued before him²⁷, that the term can indeed be misleading for policy uses; lumping threats of very different scale and nature together can conceal the necessity for a prioritised and nuanced response.

I follow a convention from game theory of thinking of player 1 as male and player 2 as female²⁸. The fact that this causes the terrorist player in all gaming chapters except 7 to be female should not be interpreted as a comment on the ever ongoing battle of the sexes.

Much of the literature referred to in footnotes is available online without subscription. These are marked with an asterisk (*); internet addresses are given in the bibliography.

1.3 The nuclear terrorism dispute: Would they? And could they?

The debate over the prospect of terrorists wielding nuclear weapons goes back at least to the 1970s²⁹ and has been ongoing in ebb and flow since. A revival of the discourse took place from the mid 1990s, following terrorist bombings of the World Trade Centre in 1993, the Oklahoma City federal building in

Kayyem and Pangì (eds.) *First to Arrive: State and Local Responses to Terrorism* (Cambridge, MA:MIT Press, 2004) pp.9-20.

- 21 Besides US and Russia, Israel comes to mind. It has been named a potential target for nuclear strike by Bin Laden's second-in-command al-Zarqawi in 2004; Salama and Hansell 'Does Intent Equal Capability?' p.628. Dunlop and Smith make the case that a target on the Eurasian continent, Moscow in particular, is far more likely since it would much simplify the required smuggling operation. William Dunlop and Harold Smith 'Who Did It? Using International Forensics to Detect and Deter Nuclear Terrorism' *Arms Control Today* (October 2006) pp.6-10
- 22 See Introduction in President George W. Bush *The National Security Strategy of the United States of America* (Washington DC, 2002)* p.15
- 23 Although a precise definition of this term has yet to be agreed upon by its users. See discussion and references below.
- 24 Often bundled as CBRN or another constellation of these letters. Occasionally as ABC weapons such as in Morten Bremer Maerli 'Relearning the ABCs: Terrorists and «Weapons of Mass Destruction»' *The Nonproliferation Review* (Summer 2000), pp.108-119. There is no agreed upon definition of WMD, however - see discussion in Gavin Cameron 'WMD Terrorism in the United States: The Threat and Possible Countermeasures' *The Nonproliferation Review* (Spring 2000) p.163.
- 25 Note that contagious biological agents can potentially produce many more casualties than either chemical or radiological weapons.
- 26 Notably, cleaning up a contaminated area after radiological dispersal can cause much indirect destruction.
- 27 Most outspoken of all on this matter is perhaps the famous Amitai Etzioni *Pre-empting Nuclear Terrorism in a New Global Order* (The Foreign Policy Center, 2004)* pp.8-9. See also George Perkovich 'Deconflating "WMD"' Paper #17, Weapons of Mass Destruction Commission (October 2004)*
- 28 e.g. Martin J. Osborne *An Introduction to Game Theory* (New York:Oxford University Press, 2004) p.xv
- 29 Two of the earliest whistleblowers were Willrich and Krieger. Their reports seem simplistic in the light of the current debate with their heavy emphasis on vulnerability and little attention given to intent. Mason Willrich 'Terrorists Keep Out!' *The Bulletin of the Atomic Scientists* 31 (May 1975) and David Krieger 'Terrorists and Nuclear Technology' *The Bulletin of the Atomic Scientists* 31 (June 1975). Both authors refer extensively to Theodore Taylor's work.

1995 and the sarin gas attack on the Tokyo underground the same year³⁰.

A certain refinement of the discourse is notable from the 1970s to the 1990s. The early warnings typically came from physicists involved with the development of the highly advanced nuclear arsenal of the United States and who regarded the bombs of the 1940s and '50s as somewhat primitive³¹. Since the 1990s an understanding has emerged in almost all literature that constructing even the simplest device is not an altogether trivial affair. Also, analysis of aspects beyond technicalities, vulnerabilities and consequences progressed significantly, including terrorist motivations and policy implications. Since the attacks on September 11, 2001 the quantity of published material on nuclear terrorism, often in a bundle with other non-conventional forms of attack, has grown even further, partly reflecting a manifest policy change in the western world towards taking the threat seriously. Much progress has been made in developing more nuanced policy responses, yet novel contributions to the understanding of terrorist intent and capabilities have been somewhat slower in coming. As I will argue below, a certain circularity of the scholarly debate in this respect has been detectable over the last decade or so, a trend broken by a number of innovative approaches during the last few years.

The 'risk' of an event is quite often defined as its probability of occurring multiplied by its consequence³². The probability in turn, when speaking of an attack by a hostile actor, is roughly equivalent with the term 'threat'³³, typically defined as the product of intent and capability³⁴. As one would expect, the debate over the importance of nuclear terrorism has focussed on evaluating these three

quantities - intent, capability and consequence³⁵. While there is all but consensus of recent date that the consequences of a true nuclear attack would be enormous³⁶, two questions have remained the core of the debate: 'Would they?' and 'Could they?'.

Beyond superficial differences (such as references to the September 11 attacks in more recent papers) the disputes between experts in the late nineties are similar to those seen in the academic press in recent years³⁷. Let me quote but a couple of examples for illustration. In one of the primary papers of the 90s wave, Falkenrath warns of nuclear terrorism³⁸. He admits it is not highly probable, but that the consequences still makes the risk a high one. He faces stark disagreement from Kamp³⁹, another early voicer of opinions on the matter⁴⁰. Kamp argues, quoting Jenkins⁴¹, that terrorists do not *want* nuclear weapons, since they generally seek attention, not bloodshed. Furthermore, he argues, obtaining a weapon would be extremely difficult. He cites the troubles Iraq faced in their search for the bomb saying 'it is difficult to imagine that a small terrorist group ... would find bomb building easier'⁴². The fact that the necessary amount of fissile material needed far exceeds the seized content of any actual smuggling case⁴³ shows, says Kamp, that obtaining this material is nearly impossible⁴⁴. 'Nuclear' does not deserve a place in the triad of serious threats with 'chemical' and 'biological',

30 See e.g. preface to Richard A. Falkenrath, Robert D. Newman, and Bradley Thayer *America's Achilles' Heel: Nuclear, Biological, and Chemical Terrorism and Covert Attack* (Cambridge Mass.:MIT Press, 1998)

31 Notable examples include Theodore Taylor in John McPhee's *The Curve of Binding Energy* (New York: Farrar, Straus & Giroux, 1974), Carson Mark, Theodore Taylor, Eugene Eyster, William Maraman, and Jacob Wechsler 'Can Terrorists Build Nuclear Weapons?' in Leventhal and Alexander (eds) *Preventing Nuclear Terrorism* (Lexington: Lexington Books, 1987) pp. 55-65 and Luis Alvarez *The Adventures of a Physicist* (New York: Basic Books 1988)

32 Stanley Kaplan and B. John Garrick 'On the Qualitative Definition of Risk' *Risk Analysis* 1:1 (1981) p.12

33 Henry H. Willis 'Guiding Resource Allocations Based on Terrorism Risk' *Risk Analysis* 27:3 (2007) p.598

34 Or a somewhat more complicated function of these two. *ibid.* p. 599

35 Those who agree the threat is real have furthermore quarrelled about what policy response is best suited, a debate I shall save for later.

36 Mueller and Mueller are amongst relatively few to argue that too much attention has been given to the prospect of a Hiroshima-sized nuclear attack against a US city, which would be 'horrible though not apocalyptic' and which the the US 'can, however grimly, readily absorb'. Respectively, John Mueller and Karl Mueller 'Sanctions of Mass Destruction' *Foreign Affairs* 78:3 (1999) p. 45 and John Mueller 'Simplicity and Spook: Terrorism and the Dynamics of Threat Exaggeration' *International Studies Perspectives* 6:2 (2005) p. 208.

37 This tendency of repetition has also been recognised by Gary Ackerman ('WMD Terrorism Research: Whereto from Here?' *International Studies Review* 7 (2005) pp.140-143), whose suggestions for future directions coincide very well with the research presented herein. Daniel Gressang wrote already in 2001 that 'proponents of different visions of the terrorist WMD potential appear to have reached something of a stalemate in the ongoing dialogue' ('Audience and Message: Assessing Terrorist WMD Potential' *Terrorism and Political Violence* 13:3 (2001) p.84); his paper reviews the debate on terrorist motivations.

38 Richard A. Falkenrath 'Confronting Nuclear, Biological and Chemical Terrorism' *Survival* 40:3 (1998) pp.43-65. As the title indicates, the threats treated are not exclusively nuclear.

Kamp concludes. Two years later, Maerli argued the exact opposite: nuclear weapons are much more in line with terrorists' preference for big bangs and spectacular violence than the sneaking death of chemicals or bacteria⁴⁵.

A good decade has passed since Kamp's article, and the same arguments are still being exchanged. The partial standstill is underscored by Bunn and Wier - two of the foremost advocates for the urgency of the nuclear terrorist threat - quoting and rebutting Kamp's opinion⁴⁶ eight years after its publication⁴⁷. The intent of the terrorists to seek nuclear weapons is argued for by quoting Osama bin Laden naming such actions a 'religious duty', by pointing to the *fatwa* issued by a leading Islamic cleric in 2003 declaring the killing of millions of Americans to be permissible, and reported attempts by Al Qaida to obtain highly enriched uranium (HEU). The example of Iraq used by Kamp is irrelevant say Bunn and Wier, since Iraq's technical struggles were primarily related to the attempted enrichment of uranium, a step a terror organisation would certainly sidestep. Furthermore requirements of weapon reliability and safety will be much stricter in a military setting⁴⁸.

Bunn and Wier like most of the debaters present sound arguments and thorough analysis. Their attacks on Kamp demonstrate, however, that the debate had not moved on; Kamp's relevance has not diminished. There are several more recent counter arguments downplaying the nuclear terrorist threat as well, notably by Frost⁴⁹. Frost iterates Kamp: nuclear

materials are difficult to acquire, building the device requires enormous effort and facilities, and no terrorist group wishes to even try. Pluta and Zimmerman dissent utterly, attacking Frost point for point. Nuclear materials are surely available at the right price, they argue, Frost's prescriptions for necessary equipment to build a device⁵⁰ are 'ludicrous'⁵¹, and al Qaida's intention to seek nuclear capacity is beyond doubt for much the same reasons as quoted by Bunn and Wier. Although some novel points are made, not very much is new since ten years previous: for a good decade the debate has tended to focus on pros and cons, mostly recognised by both sides but weighted differently according to the authors' convictions.

In recent years, however, a number of publications have offered new approaches, indicating that the debate is now moving on. I discuss examples below.

I have located the stance of a selection of authors⁵², primarily of journal articles, along the axes of 'would' and 'could' in figure 1.1. Most notably, the figure shows that most authors either believe terrorists both would and could build nuclear weapons, or they doubt both intent and capacity. Only May and Parachini seem to take the view that they probably would if only they could; these papers are not very explicit in this respect.

It is fair at this point to note already that while the

39 Karl-Heinz Kamp, Joseph F. Pilat, Jessica Stern and Richard A. Falkenrath 'WMD Terrorism: An Exchange' *Survival* 40:4 (1998) pp.168-183

40 Karl-Heinz Kamp 'An Overrated Nightmare' *The Bulletin of the Atomic Scientists* 52:4 (1996) pp.30-34.

41 Brian Michael Jenkins 'Will Terrorists Go Nuclear?' *Orbis* 29:3 (1985) pp.507-515.

42 Kamp 'An Overrated Nightmare' p.33.

43 This still holds true today.

44 These are but some of Kamp's points, notably.

45 Maerli 'Relearning the ABCs'

46 Kamp 'An Overrated Nightmare'

47 Matthew Bunn and Anthony Wier 'The Seven Myths of Nuclear Terrorism' *Current History* (April 2005) p.153

48 A more relevant example, technically different but often cited, is the failure of Aum Shinrikyo's attempts to obtain biological weapons despite enormous wealth, skilled personnel, well-equipped laboratories and calm working conditions. cf. William Roseau 'Aum Shinrikyo's Biological Weapons Program: Why Did it Fail?' *Studies in Conflict & Terrorism* 24 (2001) pp.289-301.

49 Robin M. Frost 'Nuclear Terrorism After 9/11' *Adelphi Paper*

45:378 (2005).

50 Here Frost is in turn quoting Friedrich Steinhausler 'What It Takes to Become a Nuclear Terrorist' *American Behavioral Scientist* 46:6 (2003) pp.782-795

51 Anna M. Pluta and Peter D. Zimmerman 'Nuclear Terrorism: A Disheartening Dissent' *Survival* 48:2 (2006) p.62

52 References include (note: positions are approximate; some authors give their position only implicitly): David Albright, Kathryn Buehler & Holly Higgins 'Bin Laden and the bomb' *The Bulletin of Atomic Scientists* 58:1 (2002) pp. 23-24; Graham Allison 'How to Stop Nuclear Terror' *Foreign Affairs* 83:1 (2004) pp.64-74; George Bunn, Fritz Steinhausler and Lyudmilla Zaitseva 'Strengthening Nuclear Security Against Terrorists and Thieves Through Better Training' *The Nonproliferation Review* (Fall-Winter 2001) pp.1-13; Bunn and Wier 'Seven Myths' *op.cit.*; Gavin Cameron 'WMD Terrorism in the United States: The Threat and Possible Countermeasures' *The Nonproliferation Review* 7:1 (Spring 2000) pp. 162-179; Richard A. Falkenrath 'Confronting Nuclear...' *op.cit.*; Frost 'Nuclear Terrorism Post 9/11' *op.cit.*; Frost 'Nuclear Terrorism After 9/11' *op.cit.*; Sigfried Hecker 'Toward a Comprehensive Safeguards System: Keeping Fissile Materials Out of Terrorists' Hands' *The Annals (of the AAPSS)* 607 (2006) pp.121-132; Kamp 'An Overrated Nightmare' *op.cit.*; Maerli 'Relearning the ABCs' *op.cit.*; Michael May 'September 11 and the Need for International Nuclear Agreements', in *Nuclear Issues in the Post-September 11 Era; Recherches & Documents No*

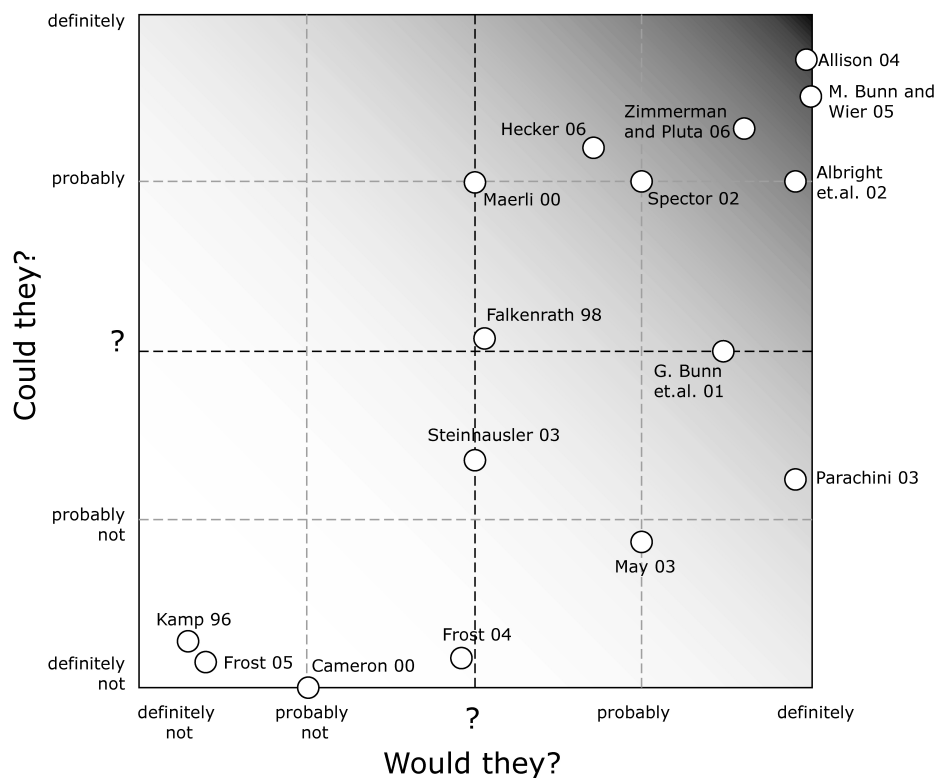


Illustration 1.1.: Papers on nuclear terrorism in a 'would-could' diagram.

academic debate over nuclear terrorism may have exhibited a certain circularity, policy with respect to nuclear terrorism has certainly moved on since the mid nineties. In 1991 the United States launched the Cooperative Threat Reduction (CTR) project⁵³. Since then some \$13 billion have been spent through that and a number of other initiatives on prevention of nuclear terrorism⁵⁴, mainly to improve the severe vulnerabilities in the wake of the collapse of the Soviet Union leaving an enormous nuclear complex at risk of theft and bribery, made worse by Russia's economic

plummet in the late nineties. While some will argue that politicians are still not doing enough⁵⁵ and that the measures taken should be prioritised differently⁵⁶, indications are that politicians and the press, particularly in the United States, are taking the threat of nuclear terrorism very seriously.

A survey of the literature on nuclear terrorism reveals that relatively few attempts have been made to approach the dispute differently. Most authors have gone at the task with a view to answering the 'would' and the 'could', and since the number of arguments for or against a given conclusion is not infinite, a certain repetition is perhaps unavoidable. It seems reasonable to speculate that a partial reason could be that many experts who see the threat as real have very reasonably aimed primarily to alert the reader to the urgency of the matter, and making novel contributions to an academic discourse might have been a secondary concern. This is exemplified by Allison's summary of what he believes the task of the nuclear terrorism researcher should be: 'Since there is

30 (Foundation pour la Recherche Stratégique, 2003)* pp. 113-124; John Parachini 'Putting WMD Terrorism into Perspective' *The Washington Quarterly* 26:4 (2003) pp.37-50; Pluta and Zimmerman 'Nuclear Terrorism' *op.cit.*; Leonard S. Spector 'The New Landscape of Nuclear Terrorism' in Barletta (ed) *After 9/11: Preventing Mass-Destruction Terrorism and Weapons Proliferation* Occasional Paper #8, Center for Nonproliferation Studies (Monterey Institute of International Studies, 2002)*; Steinhausler 'What It Takes to Become a Nuclear Terrorist' *op. cit.*

53 Often referred to as the Nunn-Lugar programme after the two US senators Samuel Nunn and Richard Lugar who proposed it.

54 Matthew Bunn *Securing the Bomb 2007* Project Managing the Atom (Belfer Center for Science and International Affairs, Harvard University, 2007)* p. 161

55 See e.g. Bunn, Wier and Holdren *Controlling Nuclear Warheads and Materials* and Bunn *Securing the Bomb 2007*

56 See discussions in chapter 3 and 4.

no established methodology, the soundest way to proceed is to ask and answer the core questions: who, what, where, when, and how?⁵⁷

Indeed, if all of Allison's questions could be answered with certainty a government would have much of the knowledge it needs. Research on the potentialities of terrorism will, however, unfortunately always be limited by the fundamental obstacle that it is impossible to know with certainty what the terrorists intend and how capable they are. Because of the fundamental obstacles to knowledge, thus, the 'would/could' debate might never be concluded⁵⁸. Such limitations should not be taken to imply that trying to infer something about nuclear terrorism by studying related phenomena is futile, and the importance of the contributions of the 'standard' literature on nuclear terrorism has clearly been paramount. It does indicate, however, that approaches which can sidestep some of the uncertainties are worth exploring.

This observation alone thus seems enough motivation to try to step beyond the debate over 'would' and 'could'. Instead of getting caught in the web of uncertainties stemming from the clandestine nature of the terrorist threat, I will ask 'assuming they *were* to try and get a nuclear bomb, then what?'. It is equivalent, one might say, to *assuming* intent and some capability in order to see where it leads us. Such an assumption is not in principle different from what previous authors have made when recommending what they see as the best policy response to a potential threat. Explicitly tossing aside uncertainties leaves more room for policy analysis, however, and I will show in the following how at least a handful of recent publications do this explicitly with success.

The thesis aspires to introduce a rational standard by which the pros and cons of different policy measures can be compared *quantitatively*. The current literature largely provides all the most important arguments for and against different policy options and priorities, and yet different analysts come to different conclusions based on different weighing of similar qualitative arguments. I seek herein to find a methodology by which to compare pros and cons more rigorously, for which I use rational choice theory, a tool borrowed from formal social science. While some will disagree (as discussed in the next

chapter) that rationality, suitably defined⁵⁹, is the best scale by which to measure policy actions and strategies, it proves to be a powerful tool in producing results and conclusions for policy application or further discussions. The analysis presented forms but a small start, but points to possibly fruitful future directions. The research approach is illustrated in cartoon form in figure 1.2.

This thesis does not profess to show that terrorists have a definite capability or intent to produce and detonate nuclear weapons. Nor does it need to. However, if the author shared the view of Frost and Kamp, the study herein would hardly be worthwhile. A brief clarification of the author's views is thus called for.

There are persuasive arguments, this author finds, that at least al Qaida wants to possess a nuclear weapon⁶⁰. The fact that they have proclaimed this is a weak indication - talk is cheap, after all, and a glimpse into the western press would tell bin Laden that the simplest way of spreading fear is using words like 'nuclear'. Stronger indications are al Qaida's meetings with leading Pakistani nuclear scientists⁶¹ and reported failed attempts to purchase HEU⁶² and acquire knowledge about nuclear weapon design⁶³. While such incidents do prove that al Qaida have at least made serious investigations into the feasibility of the nuclear option, they do not imply beyond doubt an intent to go through with such a costly and ambitious project. Possibly al Qaida has merely been charting terrain, yet this alone signifies at least some

59 I use the oldest and arguably simplest framework of maximisation of utility. See next chapter.

60 One should bear in mind that al Qaida is today a loosely affiliated network whose extent is difficult to define and separate from a more general extremist Islamic movement. A useful overview of the recent developments of Al Qaida is provided by Bruce Riedel 'The Return of the Knights: al-Qaeda and the Fruits of Middle-East Disorder' *Survival* 49:3 (September 2007) pp.107-120

61 David Albright and Holly Higgins 'A Bomb for the Ummah' *The Bulletin of Atomic Scientists* 59:2 (2003) pp.49-55.

62 Albright, Buehler & Higgins 'Bin Laden and the bomb'

63 David Albright 'Al Qaeda's Nuclear Program: Through the Window of Seized Documents' *Policy Forum Online* The Nautilus Institute, *Special Forum* 47 (November 2002)*. See also The Commission on the Intelligence Capabilities of the United States Regarding Weapons of Mass Destruction *Report to the President of the United States* (Washington D.C., 2005)* and the 9/11 report concludes that al Qaida 'has tried to acquire or make weapons of mass destruction for at least 10 years'. The National Commission on Terrorist Attacks Upon the United States *The 9/11 Report* Paperback ed. (New York: St. Martins, 2004) p. 545

57 Graham Allison 'The ongoing failure of imagination' *Bulletin of the Atomic Scientists* 62:5 (2006) p.37

58 Except, of course, were a nuclear terrorist attack to occur.




Illustration 1.2.: A cartoon representation of the research approach (author's sketch).

intent, without going so far as does Zimmerman and Lewis, indicating that 'price per murder' might form the basis for parts of al Qaida's calculus⁶⁴. It is important to note, however that the analysis in the gaming chapters of the thesis involves a generic terrorist organisation and applies equally well to other groups with nuclear intent, both present and future, whether they exist today or not.

The consequences to the attacked party being so devastating, even a small probability will give a great risk. The magnitude of the effects certainly justifies efforts to try to quantify the risk, thus, which depends crucially on the probability of various scenarios. Only when some way is found by which to arrive at quantitative estimates for the threat can direct comparison be made with other risks so that responses may be prioritised and optimised.

The construction of a gun-type HEU weapon may not be such a daunting task as commonly believed (see discussion in the following chapter), and I believe a capable group would have a modest but certainly non-zero probability of success. Bunn tentatively puts the probability of a nuclear terrorist attack over the next decade at 29% and the expected cost for the target country at \$100 billion yearly, while arguing 'even a risk dramatically smaller than that estimated ... would justify a broad range of actions to reduce the

threat⁶⁵. Given an estimate of 500,000 people perished in such an attack⁶⁶, maybe a million will be seriously injured and a similar number be homeless. In addition it is reasonable to assume that at least five times the number of casualties, 2.5 million, will be deprived of close friends or relatives. While it is dubious to quantify death, grief and strain in this manner, some perspective is offered by noting that the relatively high 29% estimate gives roughly an expected 1.7 million people dead, injured, homeless or deprived of close family members and loved ones in this manner by 2018, a figure well over a hundred times that of the September 11 attacks, estimated similarly. The possibility of such carnage must be *extremely* remote to justify negligence, I conclude.

1.4 Notable recent contributions to nuclear terrorism research

As mentioned some notable examples of relatively recent date deserve mention, where authors have approached the problem differently and moved beyond some of the uncertainties involved in all research on terrorism.

First it is fair to note at this point the number of

64 Peter D. Zimmerman and Jeffrey G. Lewis 'The Bomb in the Backyard' *Foreign Policy* (November/December 2006) p.34

65 Bunn 'A Mathematical Model...' p. 103.

66 Bunn, Wier, and Holdren *Controlling Nuclear Warheads and Materials* p.16.

articles and think-tank reports which have dealt with more specific policy issues in connection with nuclear terrorism⁶⁷. These analysts, one might say, are making the same assumption I am herein⁶⁸ that the threat from nuclear terrorism *is* real and must be dealt with in one way or another. These have made very important contributions of direct applicability to policy-making without delving too far into the 'would' and 'could'. Some papers, furthermore, deal only with terrorist motivations without considering capabilities in any detail⁶⁹, while others again pay particular attention to specific related topics such as the role of fissile materials security in Russia⁷⁰ or elsewhere⁷¹, or HEU-fuelled research reactors⁷². While not necessarily 'scholarly' in nature, the importance of this literature should not be downplayed.

Perhaps the definitive proof that academia has now started to stride beyond the 'would' and 'could' is the recent book by Michael Levi⁷³. He wastes no time debating terrorists' wish to acquire nuclear weapons, knowing that this has been treated by numerous authors before, but skips straight to an analysis of best policy responses in the face of such a threat. Levi offers a whole new way of thinking about defences against nuclear terrorism, in particular so-called second layers of defence, i.e. defences beyond the securing of stocks of nuclear materials. By not reiterating all the many arguments why nuclear terrorism is a possible threat (other than as is called

for for policy analysis) the novelty of Levi's book is striking and represents a forward leap of the academic debate on responses to nuclear terrorism in this author's opinion⁷⁴. His sidestepping the entire 'would/could' debate may be the feature which allows his analysis to reach as far as it does. While the novelty of the book lies primarily in its treatment of 'second layers', not a matter of focus in this thesis⁷⁵, Levi's book will be referred to throughout the present work.

Franck and Melese remain (to the author's knowledge) the only ones to have applied formal methodology to the terrorists' choice of opting for non-conventional weapons⁷⁶. Their paper, reviewed in chapter 5, provides interesting analysis somewhat different from what is attempted herein. In his doctoral thesis⁷⁷ Bunn attempted to quantify the risk of nuclear terrorism, suggesting numerical estimates of great utility for the research presented herein, and providing very concise numbers which will doubtlessly spur new and interesting debate⁷⁸.

Zimmerman and Lewis avoided the uncertainties by posing the question: if terrorists were to attack the US with a nuclear weapon, how would it be most practical to do it and what would it cost?⁷⁹ They too assume intent and a basic capability in order take a step further. Their paper is among the clearest example of the benefits of not having to deal with motivations and capabilities: by asking new questions, new answers are found. Among the conclusions of the paper are that a terrorist project to build a nuclear weapon from highly enriched uranium is probably most easily done inside the target country (assumed to be the United States). In this case requirements of infrastructure such as availability of electricity are easily met, and since

67 Perhaps foremost among these is the series of reports from Managing the Atom, primarily the *Securing the Bomb* series. A few other examples of good policy oriented reports of direct interest to us are Etzioni *Pre-empting Nuclear Terrorism ...* (*op.cit.*) *, Brian Finlay and Andrew Grotto *The Race to Secure Russia's Loose Nukes: Progress since 9/11* (The Henry L. Stimson Center, 2005)*; Brian D. Finlay and Elizabeth Turpen *25 Steps to Prevent Nuclear Terror: A Guide for Policymakers* (The Henry L. Stimson Center, 2007)*

68 'Assumption' may not be entirely accurate in the case of these reports in that most often the imminence of the terrorism threat is argued for.

69 E.g. the very thorough if somewhat dated analysis of terrorist motivation to employ unconventional weapons by Bruce Hoffman 'Terrorism and Weapons of Mass Destruction: An Analysis of Trends and Motivations' Rand paper #8039-1 (1999)*

70 e.g. Jon B. Wolfsthal and Tom Z. Collina 'Nuclear Terrorism and Warhead Control in Russia' *Survival* 44:2 (2002) pp. 71-82

71 e.g. Rose Gottemoeller with Rebecca Longworth 'Enhancing Nuclear Security in the Counter-Terrorism Struggle: India and Pakistan as a New Region for Cooperation' Carnegie Endowment Working Paper #29 (August 2002)*

72 e.g. William C. Potter 'Nuclear Terrorism and the Global Politics of Civilian HEU Elimination' *The Nonproliferation Review* 15:2 (2008) pp. 135-158

73 Levi *On Nuclear Terrorism*

74 See also Simen A. Ellingsen 'Strengthening the Second Line of Defense' (book review) *The Nonproliferation Review* 15:2 (2008) pp. 399-402.

75 but reviewed in appendix D.

76 Raymond E. Franck and Francois Melese 'Exploring the Structure of Terrorists' WMD Decisions: A Game Theory Approach' *Defense & Security Analysis* 20:4 (2004) pp.355-372. I perform a similar exploration in chapter 6.

77 Matthew Bunn *Guardians at the Gates of Hell: Estimating the Risk of Nuclear Theft and Terrorism - and Identifying the Highest-Priority Risks of Nuclear Theft* PhD dissertation (Massachusetts Institute of Technology, 2007)*

78 cf. also Matthew Bunn 'A Mathematical Model of the Risk of Nuclear Terrorism' *The Annals of the AAPSS* 607 pp.103-120 (2006)

79 Zimmerman and Lewis 'The Bomb in the Backyard' pp.32-39.

export control laws are inapplicable almost all the necessary equipment (except the uranium) may be bought from the online second-hand marketplace Ebay for a modest sum.

The reports from Project Managing the Atom should be mentioned at this point. Since the early 2000s this project, a part of the more far-reaching Nuclear Threat Initiative, has produced yearly reports on the security of nuclear materials, penned by Matthew Bunn and colleagues⁸⁰. The reports focus on policy options to deal with possible nuclear terrorism, as well as surveying progress on US threat reduction programmes, and are valuable sources for their data collection and analysis, and policy recommendations inferred from it. In the same vein may be mentioned a few reports from war-game type research where participants have played out scenarios following a nuclear terrorist attack as a means to analyse its results and best responses⁸¹.

Finally, although of a strictly qualitative nature, Dunn's analysis of whether al Qaeda can be deterred from use post acquisition of nuclear weapons is another good example of how interesting contributions may stem from making assumptions to bypass the uncertain⁸². His paper is examined more closely in following chapters.

1.5 Previous applications of game theory to terrorism

Given the enormous abundance of literature on the subject of terrorism, even its nuclear variant, surprisingly little has been done in the nexus of terrorism and game theory or other rational choice methodology. The majority of work done is due to Professor Todd Sandler and various co-workers⁸³,

who have published widely. Their work has been drawn upon indirectly as inspiration for the models used in later chapters, yet none of Sandler's formal methodology papers deals specifically with nuclear or non-conventional terrorism⁸⁴. Among the themes treated by Sandler and co-workers are the ways by which hardening targets deflects attacks onto other targets⁸⁵, how lack of cooperation between countries in the face of international terrorism leads to overemphasis on defence over pre-emption thus deflecting attacks onto the countries least able to defend themselves⁸⁶, whether or not a state should negotiate with a terrorist group seeking concessions⁸⁷, the rôle of private sponsors of terrorism⁸⁸, and, of some peripheral relevance to our chapter 6, strategic choice between large and small terrorist attacks as a means to achieve ends against a government adversary and a government's balance between defence and intelligence in the face of different scale attacks⁸⁹.

A handful of other authors have recently applied game theory to terrorism, and while surely valuable, are of limited use to us since they too consider special aspects of terrorism to the side of that considered here⁹⁰. The literature on conventional terrorism and game theory is encouraging more generally, in that it demonstrates the utility of such a methodology in this field.

80 The latest one is Bunn *Securing the Bomb* 2007.

81 Charles Meade and Roger C. Molander *Considering the Effects of a Catastrophic Terrorist Attack* Technical Report (RAND, 2006)*; Ashton B. Carter, Michael M. May, and William J. Perry *The Day After: Action in the 24 Hours Following a Nuclear Blast in an American City* Workshop Report (The Preventive Defense Project, Harvard and Stanford Universities, 2007)*; Ashton B. Carter, Michael M. May, and William J. Perry 'The Day After: Action Following a Nuclear Blast in an American City' *The Washington Quarterly* 30:4 (2007)

82 Lewis A. Dunn 'Can al Qaeda Be Deterred from Using Nuclear Weapons' *Occasional paper #3* (Center for the Study of Weapons of Mass Destruction, National Defense University, July 2005)*

83 Indeed, an article reviewing the field four years back refers almost exclusively to Sandler's papers; Todd Sandler and Daniel G. Arce M. 'Terrorism & Game Theory' *Simulation and*

Gaming 34:3 (2003) pp.319-337.

84 The recent paper: Kate Ivanova and Todd Sandler 'CBRN Incidents: Political Regimes, Perpetrators, and Targets' *Terrorism and Political Violence* 18 (2006) pp.423-448, does not utilise rational choice methodology.

85 Sandler and Arce 'Terrorism & Game Theory'.

86 Todd Sandler 'Collective Action and Transnational Terrorism' *The World Economy* 26:6 (2003) pp.779-802; Daniel G. Arce M. and Todd Sandler 'Counterterrorism: A Game-Theoretic Analysis' *Journal of Conflict Resolution* 49:2 (2005) pp.183-200; Todd Sandler and Kevin Siqueira 'Global terrorism: deterrence versus pre-emption' *The Canadian Journal of Economics* 39:4 (2006) pp. 1370-1387.

87 Todd Sandler, John T. Tschirhart, and Jon Cauley 'A Theoretical Analysis of Transnational Terrorism' *The American Political Science Review* 77:1 (1983) pp.36-54; Harvey E. Lapan and Todd Sandler 'To Bargain or Not to Bargain: That is the question' *American Economic Review* 78:2 (1988) pp.16-20.

88 Kevin Siqueira and Todd Sandler 'Terrorists versus the Government: Strategic Interaction, Support, and Sponsorship' *Journal of Conflict Resolution* 50:6 (2006) pp.878-898

89 Daniel G. Arce and Todd Sandler 'Terrorist Signalling and the Value of Intelligence' *British Journal of Political Science* 37:4 (2007) pp.573-586.

90 Notably: Terrorist motivations for suicide attack were analysed formally by Pittel and Rübbecke 'What Directs a

In summary, while the above mentioned exploits combining terrorism research and formal methodology are valuable, the application of formal methodology to terrorism is a discipline in its infancy, and should be explored further both in general and for the special case of nuclear terrorism, to which rational choice theory has hardly been applied at all. The literature pertaining more specifically to each of my models is reviewed briefly in the relevant chapters.

1.6 Conclusion

A review of the large body of academic literature published on nuclear terrorism shows that novel insights have been slow in coming over the last decade or so. The same arguments have been repeated with no view to a consensus emerging in the debate over the intent and capability of terrorist groups to acquire nuclear weapons, although the world of policy has moved forward during this period and there is now a broadly shared understanding in western governments that nuclear terrorism is a serious threat.

An important factor proposed as explanation for relative slowness of academic progress is that compared to the number of pages published on the issue of nuclear terrorism, relatively few attempts have been made to apply different approaches and methodologies to the subject. Great uncertainties stemming from the lack of precedents of nuclear terrorism and the clandestine nature of terrorism in general make certain claims about motivations and capabilities impossible, a fact which in itself warrants the employment of alternative approaches to explore different aspects of the potential for nuclear terrorism and policies to defend against such a threat.

Given that many western policy makers appear to have been persuaded of the imminence of the nuclear terrorist threat, it seems natural that the academic debate now moves from focussing on intent and capability to a greater emphasis on policy responses. I propose that a research approach which assumes terrorist intent and some capability and uses rational choice theory to explore different scenarios could be a fruitful way of obtaining novel insights and inform policy-making.

Rational choice theory is chosen due to several

appealing properties for my purposes. As detailed in the next chapter, it has the normative power to devise a 'rational ideal' against which real actions may be compared, suitable for analysis of policy choices in the face of an uncertain threat such as terrorism. Furthermore, the use of modelling involves a different mode of argument which enables the researcher to step beyond the uncertainties surrounding the premises of the discourse (such as the question of the intent and capability of terrorists to 'go nuclear') and analyse the implications of the assumptions instead⁹¹. A closer discussion of the powers and limitations of rational choice theory is provided in the following chapter.

Terrorist?' *Defence and Peace Economics* 17:4 (2006) pp.311-328, and detection of terrorists by Basuchoudhary and Razzolini: 'Hiding in plain sight: using signals to detect terrorists' *Public Choice* 128 (2006) pp.245-255.

91 The assumptions themselves must in turn be qualified in some way.

Methodology: problems and strengths of rational choice theory

Rational choice theory, of which game theory and decision theory are special cases, has been the subject of heated discussion for decades. Its proponents have shown more zeal than that of perhaps any other methodology in social science, with matching harshness amongst its critics from most branches in the humanities. Bennett summarises: 'Interest in the use of Game Theory ... has alternated between periods of high hopes and great expectations and periods of disfavor. The one constant has been controversy'¹. For this reason, it is necessary to be sure that the points raised against this methodology do not derail the enterprise set out in the present thesis.

In the previous chapter I argued that the clandestine nature of terrorism with its many uncertainties sets natural limits for the certainty with which questions of terrorist capabilities and intentions may be determined from available data. It seems reasonable to assume, therefore, that terrorism is a field of research well suited for alternative approaches whereby some uncertainties may be sidestepped through the explicit use of (arguably) plausible assumptions whose consequences are subsequently analysed. The process of modelling a situation from the social sciences in a way that is both simple enough to allow transparent analysis and complex enough to capture the critical characteristics of the situation is a formal way of arriving at just such a set of explicit assumptions with the obvious advantage that once a model is arrived at, the analyst has all the tried and trusted tools of mathematics at her disposal.

2.1 Chapter outline

I start the present chapter by further motivating the use of rational choice theory in the field of terrorism. Thereafter I go through a classic example from game theory in order to demonstrate what this methodology could look like and introduce some key

concepts drawn upon in later chapters. I go on to review some of the main criticisms raised against rational choice theory: those which may be classified as 'philosophical' and criticisms on grounds of poor performance. Finally I discuss the issue of terrorist rationality and introduce a definition of rationality which allows me to assume terrorists 'sufficiently rational' for gaming purposes.

Above and beyond the present chapter discussing the methodology in general, each modelling chapter contains a concluding section evaluating the utility of formal theory in the context of that chapter's research enterprise. More detailed and context specific discussion of strengths, limitations and cautious use of modelling are contained in these sections, and summarised in a corresponding section of the concluding chapter.

2.2 Rational choice and terrorism

Rational choice theory, borrowed from economics, is tailored to inform us on the strategic thinking in situations where different actors have conflicting interests or where a choice must be made under uncertainty. From this alone it is reasonable to presume *a priori* that such a methodology may be suitable to the field of terrorism. Claim Sandler and Enders, two of the most active researchers in the nexus of terrorism and rational choice theory²,

Economic methodology is particularly well-suited to provide insights in studying terrorism. Economic analysis can account for the strategic interactions among opposing interests - e.g., the terrorists and the authorities, or between two targeted countries. Rational choice models ... can be applied to ascertain how terrorists are apt to respond to policy-induced changes to their constraints. The same methods can be used to analyze how governments react to terrorist-induced changes to their policymaking environment.

In short, the interaction between terrorist and its target government exhibits all the hallmarks of strategic interplay which makes it natural to assume the use of rational choice to be a fruitful approach.

As seen in the preceding chapter, rational choice

1 Peter G. Bennett 'Modelling Decisions in International Relations: Game Theory and Beyond' *Mershon International Studies Review* 39 1995 p. 19.

2 Todd Sandler and Walter Enders 'An economic perspective on transnational terrorism' *European Journal of Political Economy* 20 (2004) p.302

theory has been employed in the field of terrorism with success in the past to analyse strategic interactions between terrorists and government (such as hostage taking and negotiations³), terrorist decision making (such as choice of target and size of attack⁴) and government responses (deterrence vs. pre-emption⁵). The merits of game theory in the field of terrorism are celebrated in two recent reviews⁶.

2.3 Game theory and decision theory: concepts and an example

As an example of what game theory can look like, and to introduce and exemplify the key concepts drawn upon in the chapters to come, I will use the most classic of all games, the prisoner's dilemma⁷.

Here is the situation: Two criminals are arrested suspected of the same serious crime which they committed together. The police do not have sufficient evidence to convict them, however, and needs a witness to testify. The two are therefore kept in separate cells and asked to testify against each other. If one prisoner testifies and the other does not, he who testifies goes free but the other receives a long imprisonment. If neither testifies, however, both are imprisoned for a much shorter time for minor offences. If both testify, both are imprisoned for a long time, but receive some rebate for co-operating.

I will use *game theory* to analyse this situation. The first concept to be introduced is the *payoff function* or *utility function*, which depends on the outcome of the game. This function defines the rationality of the players: rational play is defined as consistently choosing the strategy which maximises the payoff function. The outcome of the game in turn depends on the set of strategies employed by the players. In this

game each player has two different strategies, so there are four different outcomes of the game; we will call the strategies of testifying and staying silent D and C respectively (for 'defect' and 'co-operate'). Say the short, medium and long imprisonments are respectively 1, 6 and 9 years. Let us assume that the players extract a negative payoff equal to the number of years in prison. Then the situation looks like this:

	C	D
C	-1,-1	-9,0
D	0,-9	-6,-6

Per convention we assume that player one has rows and player two columns and the numbers in the cells represent the payoff for each player in each outcome of the game.

Now see it from the situation of player 1, say. He does not know what the other player is going to do. If player 2 co-operates (that is, she keeps her mouth shut), player 1 is better off defecting, because it will give him a payoff of 0 instead of -1. If on the other hand player 2 defects, player 1 is better off defecting as well since it will give a payoff -6 instead of -9. Thus the best strategy for both players is to defect even though the *collective* best strategy would have been to co-operate.

The reason this game has become a classic is probably that its applications as an analogy to real situations are almost endless. Consider for example the arms race of the Cold War and attempts by both the United States and the Soviet Union to disarm. Pledges notwithstanding, disarmament proved difficult and both countries still have enough nuclear weapons to destroy the Earth. The prisoner's dilemma is a simple way to explain such difficulty: if one player disarms (co-operates) the other can achieve military supremacy by defecting and keeping the weapons anyway. The solution would seem to be to introduce dialogue into the game and establish confidence (this will be a more complex model).

Other more trivial examples abound in which the socially best solution is not obtained because people act selfishly or do not communicate. If everybody waiting for their luggage by the baggage reclaim band in the airport were to take one step back, for example, all the shoving could be avoided and everyone could see when their bag came along. Unless such social behaviour is somehow enforced, however, everybody has an incentive to defect from such a co-operative

3 Harvey E. Lapan and Todd Sandler 'To Bargain or Not to Bargain: That is the question' *American Economic Review* 78:2 (1988) pp.16-20.
 4 Daniel G. Arce and Todd Sandler 'Terrorist Signalling and the Value of Intelligence' *British Journal of Political Science* 37:4 (2007) pp.573-586
 5 Sandler, Todd and Kevin Siqueira 'Global terrorism: deterrence versus pre-emption' *Canadian Journal of Economics* 39:4 (2006) pp.1370-1387
 6 Sandler and Enders 'An economic perspective on transnational terrorism' and Todd Sandler and Daniel G. Arce M. 'Terrorism & Game Theory' *Simulation & Gaming* 34:3 (2003) pp.326-329.
 7 The prisoner's dilemma is analysed in any introductory book in game theory. For example Robert Gibbons *A Primer in Game Theory* (Hemel Hempstead: Harvester Wheatsheaf, 1992) pp.2-3

strategy because they can gain an even better position by moving right next to the band.

A game theoretical analysis will typically aim to establish the *equilibria* of the game. The simplest type, the Nash equilibrium, is defined as a profile of strategies $\{s\}$ so that no player i may do better by deviating from s_i as long as all other players j employ strategy s_j . In our example, $\{s\}$ would be one of four options: (C,C), (C,D), (D,C) and (D,D); in more complex games the number of strategies can be much larger. In the simple prisoner's dilemma there was a single Nash equilibrium $\{s\}=(D,D)$, that is, assuming all players know all other players will play according to this strategy profile, no player has an incentive to deviate from $\{s\}$. The social optimum (C,C) which maximises the total payoff of both players however, is unstable because knowing the other player followed this strategy, each would have an incentive to defect. (More interesting analysis can be obtained by repeating the game indefinitely. Then each player has the option to 'punish' the other player for bad behaviour in the past. Such games are beyond the scope of this section).

Game theory in its simple form is little changed since it originated with such thinkers as Nash and Von Neumann⁸. Principally we can summarise a standard game theory exercise thus: Two or more players face choices among a set of alternatives that may lead to one of a set of possible outcomes. Each player is assumed to order these outcomes transitively⁹ by preference represented by a *payoff* function or *utility* function and act *rationally*, that is, so as to consistently maximise his or her *expected* payoff. Whilst game theory is characterised by the interplay of strategies of several players (the best choice of strategy depends on which strategy is employed by the other players), *decision theory* is its simpler counterpart in which a single player faces a choice under uncertainty.

I will make use of decision theory only in its arguably simplest embodiment in this thesis, the peculiarities and details of which are presented in detail in chapter 4. There exist several alternative formulations and more specialised branches of decision theory. Some of which emphasise the fear of loss rather than the motivation for gain¹⁰ and others,

like 'statistical decision theory' and 'Bayesian decision theory', employ more advanced mathematical tools¹¹.

In the simple decision theory I draw upon, the player will somehow assign a probability to each outcome; the probabilities formally represent the player's *beliefs* about the state of affairs. The *expected utility*¹² of a strategy is the sum of the payoffs from the various outcomes that can come of employing the strategy, multiplied by the probabilities of the respective outcomes. This form of decision theory is sometimes termed 'utility theory' or 'expected utility theory'¹³. As before a rational player is defined as a player who is consistent in playing according to his or her preferences. I will argue that it is reasonable to confine rationality to lie in this consistency only, not in the preferences themselves which are left unrestricted, as is common in the game theory literature¹⁴.

I have chosen to eschew the temptation to attempt the employment of more complete and analytically powerful versions of decision theory (and game theory for that matter) herein. Partly this is a question of taste, partly for manageability, yet I see at least two good reasons why simplicity is commendable in the current setting. The goal of the formal analysis in this thesis is exploratory and explanatory. Therefore, the simpler the theory that is applied, the more transparent and easy to analyse are the results which are obtained from it. In trying to explore or explain a phenomenon, it is little use employing a methodology which is so complex that it becomes unnecessarily difficult to interpret the results in the real-world setting I wished to scrutinise in the first place. Secondly, all else equal, the easier it is to follow an argument, the stronger and more fruitful it is. Surely an easily understood argument will be more convincing and accessible to more readers and therefore also more effective in stimulating debate. This is not to say that my decision theoretical efforts in the following chapters avoid somewhat advanced mathematics altogether, but the basic principle of

Decision Theory (New York: Courier Dover, 1988)

11 See e.g. James O. Berger *Statistical Decision Theory* (Berlin: Springer, 1985)

12 'Utility' and 'payoff' are used interchangeably in this thesis.

13 For a review of expected utility theory compared to alternative formulations of decision theory, see Mette Wik 'Individual Decision Theory under Risk: Deficiencies and alternatives to expected utility theory' Department of Economics and Resource Management, Norwegian Agricultural University, Discussion Paper #D-23 (1996)*

14 e.g. Osborne *An Introduction to Game Theory* section 1.2

8 the latter, incidentally, an important figure in the Manhattan Project to design the first nuclear weapons..

9 A transitive ordering of A, B and C with respect to a logical operator \leq means that if $A \leq B$ and $B \leq C$ then $A \leq C$.

10 e.g. Herman Chernoff and Lincoln E. Moses *Elementary*

utility maximisation is very easily grasped even if the formal steps between the game and its consequences are beyond some readers' maths knowledge.

2.4 Claimed capabilities of rational choice models

The proponents of rational choice theory will count a number of properties of rational choice theory as indications of its superior capabilities (almost none of which have been undisputed). The theory of rational choice claims to be not only of explanatory use, but also normative, providing a 'best response' to real situations. The ideal thus upheld is one which celebrates rationality over emotions, consistency over whimsicality.

Further strengths claimed have been listed in a highly useful review by Ward¹⁵ including:

- Modelling forces the user to be explicit about the assumptions behind the argument.
- The act of modelling itself necessitates a discussion of which factors are important and which are not.
- If applied correctly, formal theory is undeniably logical.
- Even if empirical data deviate from conclusions reached through a rational model, it creates a norm against which real actions may be judged.

As to the first bullet point, it is reasonable to be cautious: explicitness about *some* assumptions is necessary, but partial concealment is still possible, and modelling by no means exempts the author from the ethics of research. Quite often assumptions inherent in the structure of the model itself are not commented explicitly, identifiable only by the trained eye. A model could be said to provide a framework which makes it easier to bring assumptions forth however, which is in itself a commendable trait.

The third bullet might be the primary argument for the application of formal methodology in this context. The use of a model 'moves' the argument to a different domain. In a 'standard' argument, a number of premises are laid out and it is argued that they imply a certain conclusion. The critic may show that the

premises are false or incomplete or alternatively that the implication is invalid and that a different conclusion should have been reached based on the premises. Using a model, however, the logical implications are beyond doubt. The critic must instead turn to the assumptions behind the model and show that the model itself is flawed in some way, or that the outcome of the modelling does not imply what is claimed. Formal methodology does not inherently make stronger arguments, only different arguments.

Thus, a formal model can play the rôle of fresh blood to a discourse which has become stale and repetitive, providing a primary rationale for its proposed application to nuclear terrorism herein.

Finally, the normative ability is of some importance to us, since there exist no directly relevant empirical data to draw on; no nuclear terrorist attack has been attempted to date. Whichever way the problem is turned, thus, some level of speculation will ever be present. The strength of the normative approach was demonstrated by Zimmerman and Lewis¹⁶ ('*assuming* intent and capability how *should* the educated terrorist go about the project?') and will be made use of herein.

2.5 Criticisms of rational choice models on philosophical grounds

The fundamental assumptions of rational choice theory, that players seek to maximise personal utility and have the rational capacity, time and emotional detachment necessary to choose the right course of action, have been the target of the most fundamental, roughly 'philosophical' criticism. Ward names several general modes of these, some of which I will treat briefly in turn¹⁷:

1. From sociologists: Rational models downplay social structure and holistic explanations
2. From psychologists: In real life, actions are often not rationally motivated
3. From mainstream political scientists: the assumptions of rational choice are implausible, invariably causing faulty predictions.

The psychology criticism, closely related to that from political scientists, seems reasonable: who

15 Hugh Ward 'Rational Choice' in D. Marsh and G. Stoker (eds.) *Theory and Methods in Political Science* 2nd edition (Hampshire:Palgrave MacMillan, 2002) pp.65-89. Ward's excellent treatment is a recommended supplement to this chapter since I am forced to omit many points for the sake of brevity.

16 Peter D. Zimmerman and Jeffrey G. Lewis 'The Bomb in the Backyard' *Foreign Policy* (November/December 2006) pp. 32-39.

17 Ward 'Rational Choice' p.71

couldn't say he has made a choice at some point which was impulsive and not very well thought through? Non-rational motivations such as revenge, spite, care, and fairness often form basis of human actions, the psychologists say, and they are doubtlessly right¹⁸. Particularly, the omission of altruism is criticised. For my purposes salvation is threefold. First, there is a rule of thumb that the higher the stakes the more rational the players¹⁹; nuclear terrorism arguably is a prospect with high stakes on both sides. Secondly, and perhaps most importantly, it is easy to model preferences other than selfishness by letting players extract utility from altruistic action, say. Indeed, this is *necessary* in order to model suicide terrorism from the point of view of the individual bomber, whose actions cannot easily be motivated by narrow self interest alone²⁰. Thirdly, in the absence of clinical tests of the mental abilities of, say, leading al Qaida figures, a normative approach using assumptions that might not be entirely accurate is far better than nothing.

Barry has argued that there is a logical problem with models where preferences are not dictated by self-interest²¹. His argument is that theories become all but untestable from empirical data since some combination of self-interest and altruism will always give the right prediction. Replacing 'altruism' by 'blood thirst' to model fanatical terrorism, the criticism presumably remains the same. While Barry's point is valid, it is arguable that this is not so much a problem of methodology as of unethical research: tailoring assumptions in order to produce a certain conclusion is never acceptable, regardless of discipline. The key, this author thinks, is explicitness of assumptions so that conclusions may be tested against alternative starting points.

A related criticism is that of *post hoc* modelling tailored to explain what has already happened. Again, this author believes it ultimately an ethical question. Rational choice theory *has* come up with many novel

hypotheses²², and papers using game theory will often try to show how their conclusions are non-trivial and cannot easily be arrived at via other paths. The criticism only partly applies to the research herein, since in the case of nuclear terrorism there is no empirical data to be posterior *to*.

As a last point in this section I mention the following question raised by some: why should maximisation of expected utility be the criterion for rationality? This question is reviewed by Broome²³ who emphasises two sub-questions. Firstly, is it necessarily so that 'utility' is an arithmetic quantity which can be quantified? This is a deeply philosophical question which was addressed by such thinkers as von Neumann and Morgenstern who proposed an axiomatic utility theory in which the ordering of preferences, not the utility itself, was the central point. While this conceptual problem could be serious in certain contexts, an argument why it probably isn't for our purposes is the relative ease with which one can arrive at reasonably plausible ways to quantify expected utility in the particular games considered. Moreover, the very effort of modelling the terrorist's expected utility function creates fertile ground for investigating how different types of terrorists behave within the same game, and although no such model will ever be certain to yield accurate predictions, it provides ample opportunity for variation whenever the goal is not prediction but exploration and explanation.

Secondly, why should a rational player maximise the *expected* utility in particular? From a pragmatic point of view this objection is probably far more serious. The chief problem with maximisation of expected utility is that it implies neutrality towards risk. This will be further elaborated in chapter 6 and also elsewhere, but in simple terms stem from the fact that two options which have the same expected utility can represent widely different levels of risk. Imagine you were given the choice between \$10,000 with 1% probability or \$100 for sure. Both options have an expected utility of \$100, but which alternative is chosen depends strongly on how much risk you are willing to take. Someone desperate for money will

18 More on the psychology criticism: See Jon Elster 'Emotions and Political Economy' *Journal of Economic Literature* 36 (1998) pp.47-74, a reply to Matthew Rabin 'Psychology and Economics' *Journal of Economic Literature* 36 (1998) pp.11-46.

19 e.g. Jon Elster 'When Rationality Fails' in Schweers Cook and Levi (eds.) *The Limits of Rationality* (Chicago: University of Chicago Press, 1990) pp.40-41.

20 excepting, perhaps, the case where the afterlife is expected to be preferable to the present.

21 Brian Barry *Sociologists, Economists and Democracy* (Chicago: Chicago University Press, 1978) pp.33-37.

22 One of the few to contest this is Stephen M. Walt 'Rigor or Rigor Mortis?: Rational Choice and Security Studies' *International Security* 23:4 (1999) pp.5-48. Further discussed below.

23 John Broome 'Should a Rational Agent Maximize Expected Utility?' in Schweers Cook and Levi (eds.) *The Limits of Rationality* (Chicago: The University of Chicago Press, 1990) pp. 132-145

probably take the \$100 whereas someone to whom this sum is insignificant might take the uncertain option. What the use of expected utility theory means for my purposes, thus, is that account of the terrorist and government's attitudes to risk must be taken separately.

Arguably, expected utility theory is the simplest principle upon which to analyse games, and has been chosen for that reason. In a field so early in its infancy as rational choice theory analysis of nuclear terrorism, the soundest approach is arguably to take the easiest route possible and establish simple benchmark results upon which future efforts can be based. Efforts to explicitly incorporate the effects of risk attitudes would seem a natural extension for the future. Note, however, that the concept of 'utility' allows, to some extent, for taking preferences of risk into account in many cases. If the player is very risk averse, say, it is sometimes possible to manually decrease the utility of options which are always risky.

2.6 The sociologist criticism

The main criticism from sociologists is that a model will often inherently assume a certain social structure and take this for granted. This is certainly true, and something to be aware of: devising a game means explicating how communication flows, which player decides what and in which order. Again, one may argue that it may not be a fundamental flaw of the methodology but rather the usage of it, and disagreeing with a model employed and its assumptions is a valid form of criticism. Naturally, any game will vastly *simplify* the social structures at play, but inferring from this that nothing fruitful can come of explicitly disregarding some complexity I find much too strong a claim. Taking a step beyond the swamp of complexity by means of assumption is just the trait that makes rational choice theory powerful, but in research as in life everything comes at a price.

The level of simplification, a constant source of criticism²⁴, is always a challenge in modelling and is ultimately a question of balance. Simplicity, as it were, is the strength and the weakness of rational choice models: a simple model makes for simple and transparent analysis, but may miss important aspects of the situation modelled; a complex model could capture more of the complexity of the real world, but the analysis is more opaque and in extreme cases the

very behaviour of the model becomes inexplicable²⁵. Accuracy is necessary if the aim is precise *prediction*, but also requires precise input data; simplicity is opportune if the aim is as herein, *understanding*²⁶. Users of game theory furthermore are typically pragmatic: 'the proof of the pudding is in the eating', Osborne says, 'if a model enhances our understanding of the world, then it serves its purpose'²⁷.

Further criticism of the assumption of rational players, also from sociologists, is of relevance²⁸. Some points are:

1. actors employ a paradigmatic filter which biases use of information
2. actors make limited efforts to search for available options
3. actors are under pressure to appear consistent even at the cost of failure

It would be wrong to doubt the validity of such objections, but perhaps their importance for our purposes is not great. The clandestine nature of terrorist groups poses specific challenges, and irremovable uncertainty makes it necessary to employ some imperfectly justified presumption about how decisions are made. One need not dismiss the importance of such factors as listed above, however; only choose to explicitly ignore them to see where it leads. It might bring us forward, or in the least provoke new debate. On the whole it falls in with the whole debate over simplifications involved in any modelling procedure.

2.7 Criticism on the grounds of poor performance

Walt, a former rational choice theorist turned sceptic, criticises formal theory in an important paper by measuring its performance with respect to three criteria: (1) logical consistency and precision, (2) originality and (3) empirical validity²⁹. While the last of these points mainly concerns descriptive models to explain historical trends and events *post hoc*³⁰, and the

25 *ibid.* p.39

26 see also Nigel Gilbert and Klaus G. Troitzsch *Simulation for the Social Scientist* 2nd edition (Buckham: Open University Press, 2005) p.18

27 Osborne *An Introduction to Game Theory* p.7

28 Ward 'Rational Choice' p.74. Outtake of Ward's full 'list'.

29 Stephen M. Walt 'Rigor or Rigor Mortis?: Rational Choice and Security Studies' *International Security* 23:4 (1999) pp.5-48

30 Walt makes the point that the employers of formal models have a history of paying little attention to whether the

24 Bennett 'Modelling Decisions...' p.27

first regards consistent use of the theory, merely a point to take on board, the second is of great importance to us.

Nuclear terrorism is a field of research which has not generally inspired originality as we have seen. Novelty, however, may not only lie in the research questions set out or the answers given to these, but also in the methodology applied in the process. Walt, however, argues that formal theory has not been able to serve this purpose.

Walt points to two fundamental problems with the originality of formal approaches, the first of which is what he terms 'methodological overkill'³¹. Formal theory, he says, has shown a tendency to show what everybody knows in a way no one can understand. Long and mathematically complicated calculations will often lead to commonsensical conclusions, Walt says, providing no new understanding at the end of the day. One cannot argue that this is not a potential weakness of formal theory. If a reader must command advanced calculus to follow the argument, the power of the analysis is somewhat diminished in that fewer readers can access the content.

Arguably, this is a problem of how a tool is applied, ethics and sound judgement. Exceptionally complex calculations only seem justified if the task itself is correspondingly complex or somehow novel and uncharted. It is hard to attribute this trait to the methodology itself, however, since the same is in principle true for every theory: opaque arguments *never* serve the purpose of increasing understanding be they formal or otherwise, something Walt recognises³². This author agrees with Walt on this point and has sought to keep models and analysis as simple as feasible.

Walt's second problem is what he calls 'old wine in new bottles'³³. Many efforts using formal theory, Walt argues, have practically re-invented previous hypotheses using new methodology and sometimes terminology. While it will appear from Walt's examples that some authors have indeed overlooked some previous publications in their field and invented terminology equivalent with existing work, Walt nonetheless comes across as unreasonably harsh on this point.

There is a need, this author believes, to distinguish between 'trivial' and 'not groundbreaking'. Although a

hypotheses proposed are supported by empirical data.

31 Walt 'Rigor or Rigor Mortis?' p.23

32 *ibid.* p.21

33 *ibid.* p.26

conclusion reached by formal means corresponds well with previous efforts, the very fact that it was reached by a completely different mode of argument typically means it was no waste of time. Rather, the hypothesis is thereby strengthened, adding weight to one side of the debate³⁴.

To be sure, conclusions like 'countries which have the most to gain and the least to lose from going to war will most often do so' adds little to common sense, yet the fact that formal analysis has not generally led to counter-intuitive conclusions should be reassuring if anything. To flip the argument around: If formal theory did *not* tend to produce the same conclusions as other approaches *that* would have been a sign that something was rotten about it.

When, if one adopted Walt's strict demands for originality, would commencing a formal analysis *ever* be justified? Upon undertaking a research enterprise, one does not as a rule know what the conclusion is going to be – if one did, it would be poor use of time at best, unethical at worst. Using Walt's criterion, only a fraction of the research done by means of formal methodology – that which turned out counter to previous work – would ever be published, and the effort would hardly be worthwhile. It will seem that the only practical solution would be to abandon this methodology altogether, a conclusion much more extreme than Walt's own.

2.8 *Reduction of a state to a single mind*

In two notable papers³⁵ and a now classic monograph³⁶ focussing on the example of the Cuban Missile Crisis, Graham Allison makes the case that assuming nations may be reduced to a single economical mind, while occasionally very useful, is often an overly simplistic model of the real process of decision making. In the words of Allison (who in more recent years has become a principal debater in

34 Notably, it is common in the theoretical natural sciences that following a seminal result, several publications will follow, showing the same result in a different manner. Indeed, the novel result is often not widely accepted to be correct until it has been through this process.

35 Graham T. Allison 'Conceptual Models and the Cuban Missile Crisis: Rational Policy, Organization Process, and Bureaucratic Politics' RAND paper #3919 (1968)*; Graham T. Allison 'Conceptual Models and the Cuban Missile Crisis' *The American Political Science Review* 63:3 (1969) pp.689-718.

36 Graham T. Allison *Essence of Decision: Explaining the Cuban Missile Crisis* (Boston: Little, Brown, 1971)

the field of nuclear terrorism) and co-worker Halperin³⁷,

[T]his simplification – like all simplifications - obscures as well as reveals. In particular, it obscures the persistently neglected fact of bureaucracy: the «maker» of government policy is not one calculated decision-maker, but rather a conglomerate of large organizations and political actors who differ substantially about what their government should do on any particular issue and who compete in attempting to affect both governmental decision and the actions of their government.

This form of decision making, contrasting the 'single mind' picture typically employed, is what Allison terms 'bureaucratic politics'³⁸.

In my analysis in chapters 4 through 7 I make just such a simplification; both government and terrorist is assumed to be reducible, as an adequate approximation in the context of the chapter's research question, to a single mind. In particular the scenario studied in chapter 7, where a government must decide on a course of action in the face of a nuclear blackmail threat, fits well with the type of confrontational scenario typical for the use of Game Theory in the days of the Cold War which Allison cautions against. The notion that a government may adequately be reduced in this manner and that the same is true for a terrorist organisation are assumptions with different motivations and require different kinds of caution. I will therefore discuss them separately, beginning with a government.

2.8.1 A government as a single mind

Firstly, it is of importance to recognise that the example primarily used by Allison to make his point, the Cuban Missile Crisis, criticised the single-mind assumption being used to furnish an *explanation* of events *after* they had already happened. As Allison himself notes, making recommendations for future action is a 'related, but logically separable enterprise'³⁹. The distinction is partly one between the

normative and the descriptive domain: how *should* a government act in a given situation, and how *were* decisions reached in practice. I will argue that in order to make a general recommendation, only the first, normative, question must be answered, but in putting the general recommendation to action in a specific scenario involving specific actors, the second requires analysis as well.

There are two different perspectives to consider in this case. The first is general, independent of which government is in question, and *external* to the decision making process. As an independent researcher considering general and hypothetical scenarios it is natural to choose this perspective. The other perspective is specific to a particular government at a particular time, and is *internal* in that it has information about the various sub-state actors who interact in the forming of the final policy decision. This would be the perspective of a President or Prime Minister, say, who would prefer a certain course of action, but is facing pressure from several weighty actors with interests beyond or even conflicting with the interests of the hypothetical 'single actor' state.

Because the decision making process will be individual to each country, an analysis which aspires to apply to many states is limited to making few and general assumptions. As a first approach in a search for a policy recommendation, the reduction of the decision process to a single mind is then natural. It is probably true that the decision making processes in industrialised countries are sufficiently similar that a more detailed study of internal considerations could be performed without excessive loss of generality, and one could always decide upon a particular nation state to study in order to make the analysis concrete. No doubt these would both be valuable enterprises which form a next step beyond the analyses herein.

Assuming Allison is right, the policy recommendations made in this thesis will, from the perspective of a policy maker in a particular country, in some sense have the form of *goals* for the outcome of the actual decision process. The policy maker wishing to follow the recommendation must in turn recognise who the important sub-state actors are, what preferences they have, and try to manoeuvre in such a way that the policy eventually enacted is the desired one. The fact that a recommendation must in practice be accompanied by a plan for putting theory into practice under the actual circumstances does not, however, invalidate the recommendation itself.

Put differently and a little simplistically, it is all

37 Graham T. Allison and Morton H. Halperin 'Bureaucratic Politics: A Paradigm and Some Policy Implications' *World Politics* 24:2 (1972) p.42.

38 This criticism is somewhat related to the sociology criticism mentioned above.

39 Allison 'Conceptual Models...' *Am. Poli. Sci. Rev., op. cit.* p.689, footnote.

right for the independent researcher with the *external* perspective to assert that a government *should*, in its own interest, act as though it were a coherent and rational player. The government's adversaries, after all, will also be largely external: they will often observe only the decisions made and attribute them to the state as a whole, regardless of how they were reached.

In fact, in the normative case the argument can be turned on its head. The fact that real decision making is a jumble of conflicting interests and agendas should make it all the more important that relatively disinterested experts make solidly argued general recommendations which are independent of petty interests of conflicting organisations.

Such arguments offer only partial salvation and cannot change the fact that there is a contrast between how governments make decisions and the model I employ, of a single-mind government. The dichotomy of perspectives, moreover, is perhaps not quite so clear cut as portrayed. Adopting Osborne's pragmatic view⁴⁰, however, the important question is a pragmatic one: what is the simplest way of modelling which can tell us something useful? If the simplest model can give rise to useful recommendations, introducing more complexity is an unnecessary effort and could just obscure things. Should it be found to be too general, abstract or far removed from reality to be of use, however, more realism must be introduced. Opinions will certainly vary from researcher to researcher. I have chosen to try and establish general frameworks herein, in the hope that applicability has not suffered overly from this.

2.8.2 *A terrorist group as a single mind*

In a sense there may be more serious reasons to question the reduction of terrorist groups to a single mind because while the analyst can recommend certain actions to the government in a strictly normative sense, there will always be a descriptive element in questions concerning the terrorist's decision making. As is discussed several places in the thesis, the ideal question to answer when it comes to a terrorist adversary is not so much how they should think as how they actually *do* think.

The motivation for making the modelling assumption that the terrorist may be approximated as a single mind is therefore quite different than is the

case for a government. For any given government in the industrialised world, there exists a wealth of information about how decisions actually come about, who the interested sub-state parties are and so on, and the single mind assumption is a large *reduction* in complexity. For a terrorist adversary (and especially a generic one such as assumed herein), the opposite is the case: one typically has only sporadic and incomplete indications of how decisions are actually made in a given terrorist cell, and instead of throwing away information one has available, the single-mind assumption fills holes where important information is missing.

One can obviously envisage similar processes within a terrorist cell as those which Allison describes within a nation state. Different members can have different personal agendas and decisions may not be made in unison. The potential target government, however, is again expected to be external to all this, and if it learns of the actual course of events at all, it is typically long after the decision was made.

Considering a particular terrorist organisation, one would quite obviously do well not to ignore the information one does have, and should it be known, say, that a particular cell is plagued by internal strife, this may be something the government can make use of. Furthermore, should evidence surface that a given terrorist organisation is making decisions which do not conform to the expectations based on a rational and single-minded set of assumptions, the government should take this into account. Yet herein lies a key rationale for my approach: when rational and single-minded *benchmarks* have been established it is possible to compare, and doing so may bring the government a little closer to answering the question of how the terrorist group *really* thinks.

The frameworks I establish in the research chapters must therefore be used with caution where terrorist single-mindedness is an important assumption, as it is in chapters 5, 6, and, in particular, 7. The games and resulting equations developed in this thesis are tools for the policy maker in order to approach the corresponding questions in a systematic way, and like other tools they should be applied correctly; the consequences of erroneously assuming the terrorist to act in a certain way should be analysed as well and the conclusion used to decide the bounds within which the theory may safely be applied.

Note therefore in particular in chapter 7, how the normative value of the recommended course of action for the government in that chapter is strengthened by

⁴⁰ See quote above; 'if a model enhances our understanding of the world, then it serves its purpose'. Osborne *An Introduction to Game Theory* p.7

the fact that should the terrorist not act according to the assumptions made, the prescribed course of action is nevertheless expected to be the best available. Also in chapters 5 and 6 the consequences should the assumptions used to model the terrorist decision making be faulty do not preclude the use of the model as it is. The analysis in chapter 5, if anything, is normative for the terrorist actor, and deliberating in line with the model is arguably the best the terrorist can do in her own interest. In chapter 6 we find that the best means of deterring a terrorist from attempting acquisition of nuclear weapons are also the means which will hinder her from succeeding, should she not be deterred.

2.9 Are terrorists rational?

Psychologists and sociologists alike argue, as we have seen, that rationality is not always a good representation of the working of the human mind at the best of times. A human who devotes his endeavours to terrorist activities, one might think is even further from the rational ideal. At one level one could argue that since little is known about the psyche of specific terrorist leaders, some assumption about their reasoning must always be made, and that rationality is always one interesting assumption to explore, whether it represents realistic terrorists or not. For the policy maker this is hardly entirely satisfactory, however, because real policy must defend against real threats, thus for them analysis of defences against a hypothetical enemy is not very useful. It is arguably not reasonable to insist that a real actor must adhere perfectly to the assumptions of rationality in order for such a theory to capture some of the player's characteristics, and even less so if the theory's goal is instead to set a rational ideal that the player *should* live up to in his or her own interest, yet it is worth posing the question whether terrorist organisations can be seen as sufficiently rational for a rational choice theory analysis of their choices to connect sufficiently well with reality to be of use.

There exists a considerable body of literature on the question of whether or not terrorism may rightly be explained rationally. The key question that concern most authors would be whether a *homo economicus*⁴¹

41 or *homo oeconomicus*. A construct sometimes attributed to John Stuart Mill, frequently used in economic theory of a 'person' motivated solely by accumulation of wealth, avoidance of unnecessary labour and who has the intelligence to take optimal action towards these goals. Joseph Persky 'The Ethology of Homo Economicus' *Journal of Economic*

could ever justify the apparently irrational action that terrorism is, suicide terrorism in particular. The focus has shifted between the leaders of terrorist groups and the individual bomber, and conclusions have varied throughout the spectrum from assigning terrorists complete self-interested rationality⁴² to being guided by altogether different forces⁴³.

I will briefly survey some key papers from the debate on this topic and how it leads to a useful distinction between different definitions of rationality. It is concluded that rational choice models may be used with terrorist actors if assumptions of self-interest and rational expectations are relaxed, and it has already been argued that none of these assumptions are strictly necessary for modelling.

One of the most referred to supporters of the notion that terrorists are rational is Pape. He argues that terrorists are quite plainly rational even in the strictest sense of *homo economicus*, and the argument is mainly that 'terrorism is rational because it works'. Among Pape's five principal findings⁴⁴ are that suicide terrorism is strategic, not random, that it is specifically designed to make democracies make specific concessions and that the reason that suicide terrorism has been steadily rising over past 20 years is that terrorists have learnt that it pays. The rationality Pape advocates is perhaps best attributed to the *leaders* of terrorist groups resorting to suicide tactics.

Pape is harshly criticised by Abrahms⁴⁵, who upholds that his primary conclusion that 'terrorism works' is faulty. Terrorist acts have on the whole done more harm than good to the position of the interest groups behind them, he argues, and shows how the relative successes of marginalised groups that Pape credits to terrorism can be adequately explained by other factors.

The most important contribution of Abrahms' critique for our purposes may be to credit the terrorists a *certain* rationality⁴⁶. 'Substantial rationality is concerned with the consequences of the decision, while procedural rationality makes no claim that the

Perspectives 9:2 (1995) pp.221-231.

42 Robert A. Pape 'The Strategic Logic of Suicide Terrorism' *American Political Science Review* 97:3 (2003) pp.343-361.

43 For references to literature applying various models, see Pape 'Strategic Logic..' page 343.

44 Pape 'The Strategic Logic...' pp.344-345.

45 Max Abrahms 'Are Terrorists Really Rational? The Palestinian Example' *Orbis* Summer 2004 pp.533-549

46 *ibid.* p.546

actor correctly perceives the environment⁴⁷. Terrorists, argues Abrahms, may be said to be procedurally rational but not substantially so.

The question thus raised is whether 'does it work' is the appropriate criterion by which to judge terrorist rationality. The answer leads us to the question of rational expectations – is it reasonable or not for the terrorists to expect to gain what they hope from their violent actions – a criterion pertaining to the strict *homo economicus* sense of rationality normally assumed in economy. Could it not suffice that the terrorist has a *belief* that violence will somehow have a beneficial outcome for himself; a belief which does not have to be justified by historical evidence. In fact, in any game theoretical model, the deep beliefs upon which the players base their preferences are prerequisites for the model and hence lie outside the model itself. They must therefore be argued for in non-rational and non-formal terms.

It is natural therefore to specify more closely just what 'rationality' is demanded for modelling purposes. After carefully surveying the debate, Caplan⁴⁸, shows that three different interpretations of rationality appear to co-exist which I will refer to as the weak, moderate and strict interpretations respectively⁴⁹:

- *Weak sense*: Anyone who uses means to achieve ends is rational by definition. This is typically applied by high-level theoreticians.
- *Moderate sense*: Rationality equates with consistently maximising a well-behaved utility function. This requires an ordered and transitive set of outcomes.
- *Strong sense*: Rationality is based on narrow self-interest and/or rational expectations. Implying that (1) a person who makes the same mistake repeatedly is irrational and (2) preferences must be based on 'economical' self interest. This is in essence the *homo economicus*.

Regardless of its justification, for our modelling purposes it is arguably natural to use the moderate sense in the case of terrorists. The weak sense is useless for rational choice theories, and a strong

assumption of rationality as argued before is not necessary for the purpose of modelling, hence there is no practical reason for us to make such strict demands.

When, moreover, the applicability of the strong sense of rationality to terrorists is in some doubt while the moderate (or procedural) sense appears less problematic, I conclude from this that a terrorist organisation or group, to the extent that it can be mimicked by a single rational mind, may be modelled as rational. However, care must be taken to be explicit about just how the preferences of the terrorist are modelled since unlike *homo economicus* who comes with a pre-defined set of preferences, the moderately rational terrorist must have his preferences specified for each new game.

2.10 Evaluating the utility of gaming

When evaluating the rôle of a model it is important to recognise what a model should and should not be expected to do. Towards the end of each of our research chapters I have included a short section evaluating the gains from using formal theory towards answering that chapter's research question. Here I establish the foundations upon which such an evaluation should rest: in particular the goals which gaming should achieve beyond what could have been readily done with purely qualitative means.

Importantly, a model can never produce conclusions which are not, in principle, already inherent in the assumptions which go into designing it. Hence there will always seem to be a certain element of circularity about this mode of argument if the model is taken to be a premise of the analysis. Rather, the model should in a sense be seen as a conclusion, the conclusion of the process of modelling, a qualitative process of applying judgement to decide how best to represent the essence of a complicated problem in a condensed mathematical form.

Hence the model and the mathematical methods of analysing it, are no more (or less) than tools which allow the analyst to analyse and visualise the implications of the assumptions made in the modelling process. It would be wrong to expect a paintbrush to produce an artwork, and somewhat similarly, gaming does not create *fundamentally* new knowledge. If one could say that the artwork *exists* (in a potential sense) from the moment the artist has planned it in his mind, and that painting it is but a mechanical means of applying the idea to canvas, it is

47 The idea is from Herbert Simon 'From Substantive to Procedural Rationality' in Latsis (ed.) *Methods and Appraisal in Economics* (Cambridge: Cambridge University Press, 1976) 130-131

48 Bryan Caplan 'Terrorism: The relevance of the rational choice model' *Public Choice* 128 (2006) pp.91-107

49 *ibid.* p.93.

somewhat analogous to the act of gaming. The creation of the artwork corresponds to the creation of the game with the judgement and qualitative analysis it entails, and actually gaming it out is only a way of extracting the implications of the game in a more explicit form.

Thus, even if a problem could be well analysed qualitatively, for certain problems a game has the potential to extend these arguments in two ways; to explicate implications and conclusions which are buried deep amongst the available evidence, but which one may not easily realise; and to formulate succinct mathematical results which contain a wealth of information which it would be ineffective to try to lay out in words.

Just like other tools, game theory and decision theory are effective for some tasks but will certainly be ineffective for others. To use the same analogy again, painting may be superior to words in capturing beautiful scenery, say, but painting a book would not be as useful as to describe its contents. Likewise, some research questions invite the use of formal methodology for their analysis, others do not. Obviously, apart from being of general interest, the research questions analysed in the research chapters of this thesis have been chosen because they were expected *a priori* to work well with gaming. Therefore, when I conclude that the gaming efforts herein have generally proved rather successful, this is perhaps not so surprising.

When evaluating the value of doing the exercise of gaming, the obvious question is whether the same conclusions could be reached by qualitative arguments alone. But a second criterion of success would be the derivation of simple formulas which make intuitive sense, which form a tool for reducing a complicated problem to the limited task of estimating a few parameters, and which contain precise information about how these parameters interact in a given rational decision. If a sufficiently simple game can be devised, such formulae will tend to follow almost automatically, but it is up to the judgement of the analyst and users of these formulae to decide whether the model does in fact capture all essentials.

2.11 Conclusion

By the strictest standards on rationality, upheld by some and applied by many, there is reason to doubt whether suicide terrorism, even from the perspective of terrorist leaders whose lives are not typically willingly sacrificed, may be explained by rational

choice. I argue that by limiting rationality to connote consistency in choosing the preferred strategy amongst several while imposing no restrictions on the preferences themselves, terrorists are *rational enough* for modelling purposes. Such a definition of rationality, however, necessitates further assumptions about the way a terrorist actor extracts utility in various outcomes, which must be made explicit and appropriately canvassed.

I review and discuss a number of criticisms of rational choice theory voiced over the years. A discussion of what is reasonable to expect from a model is also given for future reference in the later evaluation of use of formal methodology in the research questions herein.

While noting the number of limitations of rational choice theory, it appears I am on reasonably safe grounds in my application of formal theory to nuclear terrorism within the framework of the assumptions and simplifications made use of.

Keeping fissile materials out of terrorist hands: HEU, plutonium and a prioritised response

The three devices that made up the first generation of nuclear weapons - Trinity, 'Little Boy' and 'Fat Man' - represented two very different weapon concepts. Out of necessity the active material that provided the enormous amount of energy released in the three explosions came from two different elements, one of which, plutonium, was discovered only a few years earlier. The development of the nuclear bomb during the Second World War was a race against time to get the new weapon ready soon enough to play a part in the war, and since uranium needs to go through a slow and complicated process - *enrichment* - before it is usable in a weapon, enough material could not be produced for more than a single bomb. Given the rushed circumstances of their creation it should surprise no-one that the first two types of nuclear weapons are also considered the most primitive and the two designs it has been argued that a terrorist organisation could possibly build¹.

Today the situation is wholly different; neither highly enriched uranium (HEU) nor plutonium is in short supply. Thousands of tonnes of them are stored in a range of countries all over the world, held in very different locations both private and military². It is no longer necessary for a proliferator to enrich his own uranium or to produce his own plutonium in a reactor and extract it from the very radioactive spent fuel by complicated chemical reprocessing which was the only possibility during World War 2 (WW2)³; he could obtain the materials covertly or overtly, from one of the many states holding stores of it.

All the different measures that are in place to

hinder *fissile materials*⁴ from falling into the wrong hands are commonly dubbed *nuclear safeguards*. There is much to say about safeguards, and it has many aspects, as will be elaborated later in section 3.5, where I also define the term more carefully. What is intuitively clear, however, is that safeguards like all government activities cost money, and when there are more tasks than there is money to do them (and this is nearly always the case) one must prioritise. Hence the question: what to safeguard first?

To answer this I will make use of a formal model from rational choice theory. Before embarking on the actual modelling, however, an introduction of the two elements is called for. Our question leads naturally to another: which element poses the greatest proliferation threat? This question, and the more general one of the prospects of terrorists building a crude nuclear device from illicitly acquired fissile material, has been studied extensively in the past⁵. I will discuss the question in the context of the choice between HEU and Pu here, a discussion which serves the additional purpose of providing an introduction of the key technical concepts of nuclear weapons.

Potential proliferators may be divided into two groups: state actors and non-state actors. The latter is often equated with 'terrorists' since they are the one type of non-state actor which is typically imagined to have any interest in acquiring nuclear weapons. I will argue that, to a proliferating state it would not matter very much if their supply of nuclear materials were uranium or plutonium. To a non-state actor with very limited capabilities, however, it could make a world of difference⁶; there are several reasons as will be discussed below why the terrorists would prefer HEU over plutonium⁷, the latter posing a significantly greater technical challenge.

1 A classic treatment of this question is J. Carson Mark, Theodore Taylor, Eugene Eyster, William Maraman and Jacob Wechsler 'Can Terrorists Build Nuclear Weapons?' in Leventhal and Alexander (eds.) *Preventing Nuclear Terrorism* (Lexington: Lexington Books, 1987)* pp. 55-65.

2 e.g. International Panel of Fissile Materials *Global Fissile Material Report 2007* (IPFM, 2007)*

3 Plutonium is still being manufactured from spent reactor fuel in reprocessing plants in several countries.

4 i.e. materials that can sustain a fission chain reaction, including enriched uranium and plutonium.

5 See e.g. Charles D. Ferguson and William C. Potter 'Improvised Nuclear Devices and Nuclear Terrorism' Paper #2, Weapons of Mass Destruction Commission (June 2004)* and Matthew Bunn and Anthony Wier 'Terrorist Nuclear Weapon Construction: How Difficult?' *The Annals of the AASPP* 607 (2006) pp. 133-149

6 A detailed study of terrorist weapon construction with a suboptimal choice of nuclear materials is found in Michael Levi *On Nuclear Terrorism* (Cambridge, MA: Harvard University Press, 2007) pp. 66-97

7 In this chapter I will distinguish only between HEU and Pu, not between the various forms in which the two materials may be obtained, which is also of importance. While overlooking the importance of such distinctions, this simplifies the analysis considerably. For a more complete, qualitative analysis, see *ibid*.

It turns out that with only very moderate assumptions one can show (within the limitations of the model employed) that US plutonium safeguards measures are probably overfunded at present, and HEU measures almost certainly underfunded. First, however, an introduction of the two elements is called for in which I argue qualitatively that with regard to proliferation to a non-state actor HEU constitutes a graver threat than plutonium.

3.1 Research question

The research question to be discussed both in this chapter and the next is

Based on the threat of proliferation of nuclear weapons to terrorists, to what extent is there reason to prioritise HEU safeguards measures over plutonium safeguards?

The present chapter is a qualitative treatment of the multifaceted problem of prioritisation of safeguards between the two different fissile elements. The next chapter draws upon the qualitative analysis found herein to model the situation and draw conclusions with important policy implications.

3.2 Literature overview and outline of chapter

The numerous authors dealing specifically with nuclear non-proliferation to non-state actors take somewhat varying positions as to whether one fissile material should be given priority over the other. On the one side are Ferguson and Potter whose list of 'urgent priorities' against nuclear terrorism lists 'Put HEU first' (rather than plutonium and radiological sources) right at the top⁸. Maerli's doctoral thesis, another example, focuses on HEU⁹. The reports from the 'Managing the Atom' project at Harvard¹⁰ on the

8 Charles D. Ferguson and William C. Potter with Amy Sands, Leonard S. Spector and Fred L. Wehling *The Four Faces of Nuclear Terrorism* (New York: Routledge, 2005)* pp.324-336

9 This is in accordance with the thesis' research questions. Morten Bremer Maerli *Crude Nukes on the Loose?: Preventing Nuclear Terrorism by Means of Optimum Nuclear Husbandry, Transparency and Non-Intrusive Fissile Material Verification* Dissertation (Dr. Philos.) (University of Oslo, 2004) pp.77-81.

10 e.g. Matthew Bunn, Anthony Wier and John P. Holdren *Controlling Nuclear Warheads and Materials*, and Matthew Bunn and Anthony Wier *Securing the Bomb 2006*, reports of the Project Managing the Atom (Belfer Center for Science and International Affairs, Harvard University, 2003 and 2006 respectively)* see also Matthew Bunn 'Preventing a Nuclear

other hand, speak mainly of 'fissile nuclear materials' as a group, in effect treating HEU and Pu as equals to terrorists although the added technical challenge of using plutonium are duly discussed¹¹, and the same authors hold the opinion that the prospect of a terrorist plutonium bomb is 'very real'¹². In one of the most cited books on nuclear terrorism Allison likewise treats HEU and plutonium as essentially equal for 'build it yourself' purposes¹³. The same is true of Von Hippel, another prominent proponent of safeguards against nuclear terrorism¹⁴.

On the question of prioritisation of safeguards efforts between the isotopes, existing literature gives us all the pros and cons, but rarely attempts to weigh them against each other by any quantitative means.

3.2.1 Chapter outline

The chapter is structured as follows. First I discuss qualitatively what separates HEU and plutonium from the terrorist proliferator's point of view. This chapter contains much analysis and information which will be referred to in later chapters and acts partly as a pre-study for the game introduced in chapter 4.

3.3 Terrorist preference: HEU

Physicist Luis Alvarez, Nobel Prize laureate and an important part of the physics team at Los Alamos during the Second World War, remarks¹⁵

with modern weapons-grade uranium, the background neutron rate is so low that terrorists, if they had such material, would have a good chance of setting off a high-yield explosion simply by dropping one half of the material onto the other half. Most people seem unaware that if separated U-235 is at hand it's a trivial job to set off a nuclear explosion, whereas if only

9/11' *Issues in Science and Technology* (Winter 2005)* pp.55-62 where no distinction is made between HEU and Pu.

11 Bunn *et al.* *Controlling Nuclear Warheads and Materials* p.28

12 Bunn and Wier 'Terrorist Nuclear Weapon Construction' p.143.

13 Graham Allison *Nuclear Terrorism: The Ultimate Preventable Catastrophe* (New York: Times Books, 2004) pp.92-98

14 Frank Von Hippel, 'Recommendations for Preventing Nuclear Terrorism' *Journal of the Federation of American Scientists* 54:6 (2001)* pp. 1-10

15 Luis Alvarez *Adventures of a Physicist* (New York: Basic Books, 1988) p.125. See footnote 20.

plutonium is available, making it explode is the most difficult technical job I know Given a supply of U-235, however, even a high school kid could make a bomb in short order.

No doubt the simplest nuclear bomb to build and successfully detonate is by a good measure the *gun design* using HEU¹⁶. This design, which uses the principle that Alvarez describes, simply fires one subcritical¹⁷ lump of active material into another; the pieces might be shaped as a 'nut' and fitting 'bolt' where the 'bolt' is fired inside a cylinder¹⁸, somewhat resembling the interior ballistics of a navy ordnance cannon. The 'bolt' hits the 'nut' mounted in the opposite end, the total mass is now supercritical and a nuclear explosion occurs; this is the simple principle of a gun design.

This design has rarely been used for military purposes due to its relative inefficiency. The United States tested only a handful of these weapons¹⁹ (The first of which was 'Little Boy', the device used against Hiroshima in 1945) and the only other known military devices using a uranium gun design were built by South Africa and dismantled when the country gave up their nuclear arsenal in the early 1990s. Needless to say, although 'inefficient' by the standards of nations that sport thermonuclear weapons, the images from Hiroshima demonstrate with all possible clarity that the successful detonation of such a device would dwarf even the most atrocious grand-scale slaughter at the hands of terrorists to date.

While it is certain that Alvarez exaggerates for dramatic effect²⁰ (one should remember that this was

written at a time when the general consensus was that building a nuclear weapon required a multi-million dollar 'Manhattan'-scale project) and that even a HEU gun will be a very ambitious undertaking for an inexperienced technical team, few knew better than the late Alvarez the difficulties of designing the *other* type of weapon, the plutonium implosion device: he was the very man who invented the mechanism to detonate the explosives which set off such a bomb²¹, whose design was realised in the Trinity test device and the bomb 'Fat Man' which caused more than 70,000 people to perish in Nagasaki in August 1945²². The implosion weapon, the simplest workable design using plutonium, requires advanced command of high explosives to produce a very precise spherically converging shockwave by detonating a number of shaped charges very nearly simultaneously.

3.4 Making a crude Pu implosion device

HEU and plutonium, as was duly demonstrated during the last World War, can both be used for nuclear explosives. But for sharing this vital trait, there are important physical differences between the two materials that set them apart both in their commercial and military uses and the proliferation concern their continued existence, and indeed production, causes.

fission yield is approximately

$$E \approx 500 \left(\frac{\nu_0 v \tau}{d_0} \right)^{3/2} M e_m .$$

Here ν_0 is the ideal value of the effective number of neutrons per fission (the fizzle will start when the pieces are close enough together for the effective neutron number to become positive), v is the velocity of the falling slab as it approaches the still slab, τ is the average time between two spontaneous neutrons in ²³⁵U, d_0 is the distance the falling piece travels from fission starts until the mounting pressure pushes the system apart and fission stops, M is the total mass of the two pieces, e_m is the energy released per kg of uranium which is fissioned. Serber uses the quantities $\nu_0 = 0.3$, $\tau = 10^{-8}$ s, $d = 10$ cm and $e_m = 20,000$ kg TNT per kg fissioned HEU. Furthermore we assume $M = 50$ kg and with simple Newtonian mechanics find that a piece which is dropped from a height of 2m will obtain a velocity of about 6m/s (see any undergraduate physics textbook). With these numbers one finds a yield of 0.4 kg of TNT, not a big explosion, but enough to kill whoever drops the piece (note that this is a rough calculation). This is probably an overestimate, since the falling piece has so little momentum it will likely be pushed off of the other before impact by the pressure from early fission. The likely effect seems to be that no explosion occurs, but the neutron radiation will probably suffice to kill or seriously injure

-
- 16 Bearing in mind, of course, that design of nuclear weapons is subject to strict military secrecy, the author cannot pretend to possess knowledge of the plethora of designs that have been successfully attempted over the past six decades.
- 17 When an amount of uranium or plutonium is large enough to just sustain a chain reaction, it is called critical. A smaller amount than this is subcritical, and a greater supercritical.
- 18 One could equally well fire the 'nut' onto the 'bolt' of course as was done in 'Little Boy'. E.g. John Custer-Mullen *Atom Bombs: The Top Secret Inside Story of Little boy and Fat Man* (Self-published, 2004)
- 19 See list of US nuclear tests at Carey Sublette's *The Nuclear Weapon Archive* Online: <http://www.nuclearweaponarchive.org/Usa/Tests/index.html>
- 20 Actually, it is quite simple to demonstrate that Alvarez' claim is not accurate. Serber (reference below) gives a formula for calculating the yield of a fizzle given a number of assumptions (these may not be fulfilled here, but we are interested only in an order of magnitude calculation, for which Serber's method should suffice). The formula for the

3.4.1 The gun and the implosion device explained

Nuclear weapons and reactors work because of two properties of nuclear fission which makes a *chain reaction* possible. When a heavy nucleus such as that of ^{235}U or a Pu isotope²³ splits in twain (or *fissions*) neutrons are emitted, numbering on average 2.52 and 2.95 for the two elements respectively²⁴. The second property is that when these neutrons in turn smash into a nearby ^{235}U or Pu nucleus, this causes the second nucleus to fission as well. Thus more neutrons are produced and the chain reaction is started. Clearly, for a chain reaction to escalate, the average number of neutrons per fission must be larger than 1, which would just suffice to keep the reaction rate constant. Many of the produced neutrons do not produce further fission, however: they may escape from the system or be absorbed by the fissile nuclei or another type of nucleus present, hence one speaks of the *effective* number of neutrons per fission after subtracting the one neutron necessary to keep the reaction going and the fraction of neutrons which do not further the chain reaction. If the effective neutron number is positive, a chain reaction starts, otherwise the reaction will die out.

As explained above, a system where the chain reaction is allowed to increase exponentially is called supercritical. If the chain reaction goes on for a sufficiently long time, the chain reaction releases enormous amounts of energy and an explosion occurs. The system is quickly blown apart by the extreme temperature and pressure, however, and the bomb becomes subcritical again. The trick is therefore

anyone in the room (this claim is not verified here). Alvarez' point remains valid, however: the principle of a uranium gun is very simple. It seems likely that, given two pieces of HEU metal a 'high school kid' could, if not make a bomb, at least manage to wreak some havoc. Robert Serber *The Los Alamos Primer* (Berkeley: University of California Press, 1992) section 17.

- 21 The exploding bridgewire detonator, see Lillian Hoddeson, Paul W. Henriksen, Roger A. Meade and Catherine Westfall *Critical Assembly: A Technical History of Los Alamos during the Oppenheimer Years 1943-1945* (Cambridge: Cambridge University Press, 1993) pp.171-173
- 22 Richard Rhodes *The Making of the Atomic Bomb* (New York: Simon & Schuster Paperbacks, 1986) p.740
- 23 All plutonium nuclei (^{239}Pu , ^{240}Pu , ^{241}Pu , and ^{242}Pu) can in principle be fissioned. J. Carson Mark 'Reactor-Grade Plutonium's Explosive Properties' Nuclear Control Institute (1990)*.
- 24 Serber *The Los Alamos Primer* p.20.

to create conditions under which the chain reaction can escalate for as long a time as possible before the device disassembles.

If one were to build a nuclear bomb using plutonium, the gun method described above would not work well. Both in HEU and Pu there is a certain activity of *spontaneous* fission, i.e. heavy nuclei fission on their own account (or so one might think of it) at a certain rate emitting neutrons which can potentially start the chain reaction too soon.

Because of the ever-present background of neutrons from spontaneous fission, therefore, the two pieces of HEU in a gun design need to be slammed together sufficiently quickly - if they approach each other too slowly, neutrons from spontaneous fission in one lump will start inducing fission in the other and vice versa, commencing the explosion before the two pieces are properly in place. While most neutrons are still lost to the gap between the pieces, the fission still generates enough heat and pressure to blow the device apart before the explosion can begin properly. The result is a very small bang by nuclear standards - a *fizzle*. HEU emits spontaneous neutrons at a fairly low rate (about a hundred million per second for a critical mass²⁵), so a well designed gun can shoot the two pieces completely together between two spontaneous neutron emissions. The same assembly speed is not nearly fast enough when the material is plutonium, however, whose spontaneous fission activity for a critical mass is much higher. A more complex design is called for.

The way military plutonium weapons work is not to slam individual bits together, but to compress a 'pit' of plutonium metal into a denser state by making clever use of high explosives²⁶. Whilst the initial state is subcritical, shock compression of the fissile material forces the nuclei closer together to form an easier target for the neutrons to hit²⁷. The increased density makes the pit supercritical, a neutron generator, or *initiator*, starts the chain reaction at just the right moment and the bomb goes off. This is the implosion device.

This is not as straightforward as it sounds; getting the spherically converging shockwave right, in fact, was among the greatest challenges to the Manhattan project designing the first generation of nuclear

25 *ibid.* p. 12

26 For further discussion of the implosion design see Mark 'Reactor-Grade Plutonium's Explosive Properties'

27 The critical mass is inversely proportional to the density squared. Serber *The Los Alamos Primer* p. 27

weapons²⁸. It was found that simply packing a shell of explosives around the plutonium core would not do: even a pressure wave differing very slightly from spherical shape would flatten rather than compress the plutonium core, dramatically decreasing the yield of the weapon to a *fizzle*, or no yield whatever. The solution that was found was to use so-called explosive lenses: shapes of faster and slower high explosives that produce a pressure wave of a desired shape.

At the time of the Manhattan project, implosion and shaped charges had hardly been utilised before - now more than 60 years have passed, and the knowledge of spherical shock waves is widely used in several branches of engineering²⁹. Half a century of research on shock wave compression physics has provided any competent engineer with access to a good library just about what he or she needs to get a certain understanding of an implosion system³⁰. Computer software and calculating power have furthermore progressed immensely even since the 1970s by which time 'the era of new concepts in [US] nuclear weapons design had virtually come to an end'³¹. Finally, terrorist demands on reliability and predictability of yield will be much lower than that of any military state, so precision requirements can be relaxed somewhat³². This said, however, obtaining the skill and experience required to make even a finished blueprint into a working device could still be an important obstacle³³, as discussed further in chapter 6.

There is also the question of transporting the device after it has been successfully built. While a gun

design will have the shape of a cylinder, a crude implosion device will be roughly spherical and considerably more bulky. The measurements pertaining to the two nuclear bombs used in war might give us an idea of what the terrorist should expect to move around. 'Little Boy', the uranium gun, was 3.05m long, 71cm diameter and weighed something over 4000kg³⁴. 'Fat Man' (the implosion device) was, as the name indicates, more bulky with its length, diameter and weight roughly 3.25m, 155cm and 4550kg respectively³⁵. To be sure, these bombs had bulletproof armour and were designed to be aerodynamic, hence were much heavier than would be necessary for a delivery, say, by car. Yet a terrorist organisation will not have all the equipment and expertise of the US military ordnance workshops at hand and will be unlikely to be able to, nor indeed attempt to, optimise their design with respect to weight and size. A mass of some 2-3000 kg for the HEU gun and 3-4000 for the implosion device might be a reasonable guess³⁶, thus it is clear that such a device is a challenge to smuggle due to its sheer size and weight. And the Pu device significantly more so than the uranium gun, although this difference is probably much less important than those pertaining to device construction treated above.

3.4.2 HEU implosion weapon?

There is also the possibility of making an implosion weapon using HEU. To the best of the author's knowledge, all currently stocked military HEU weapons are of implosion type. As explained, a gun design would be much easier to realise, yet given HEU, the implosion design is always a possibility. The major advantage of an implosion type weapon would be that less HEU metal is needed. Therefore it is conceivable that a terrorist could consider this weapon, in particular when she has too little material for a gun and few prospects of acquiring more³⁷. Combinations of uranium and plutonium could also

28 See e.g. Hoddeson *et al.* *Critical Assembly* chapters 8,9 and 14.

29 See the preface of R. F. Trunin *Shock Compression of Condensed Materials* (Cambridge: Cambridge University Press, 1998) pp. ix-xi.

30 See e.g. the proceedings of the annual conferences of the American Physical Society Topical Group on Shock Compression of Condensed Matter from 1981 to present (for a list of references, see www.shockphysics.org), as well as textbooks such as Trunin *Shock Compression of Condensed Materials*. It is probably no coincidence that although the publisher's notes on the back flap of the latter volume predicts it will be of interest to 'condensed matter physicists, material scientists, earth scientists and astrophysicists', the only library in the author's native Norway found to hold a copy was that of the Norwegian Defence Research Establishment.

31 'Nuclear Weapons Technology' Section 5 of *Militarily Critical Technologies List Part II* (Washington DC: US Department of Defense, 1998)* p. II-5-3

32 Bunn and Wier 'Terrorist Nuclear Weapon...' p.142

33 Brian A. Jackson 'Technology Acquisition by Terrorist Groups: Threat Assessment Informed by Lessons from Private Sector Technology Adoption' *Studies in Conflict and Terrorism* 24 pp. 183-213 (2001)

34 Surprisingly reports differ slightly. 4100kg is reported in E.F. Newley 'Development of the Nuclear weapon' in T.I. Williams (ed) *A History of Technology* vol. 7 (Oxford: Clarendon Press, 1979) pt.1 p.276.

35 e.g. Newley p. 276.

36 The uranium gun weapons made by South Africa weighed as little as 1 tonne, yet it is doubtful whether a non-state group could achieve this. Peter D. Zimmerman 'Proliferation: Bronze Medal Technology is Enough' *Orbis* (Winter 1994) p.77

37 For a given mass of HEU an implosion design would also be able to give a greater yield, yet it is unlikely that a well informed terrorist will risk this if the gun design is an option.

be used, should the terrorist possess a quantity of both. It is therefore worth keeping this option in mind, although we shall not focus on it in the following.

3.4.3 *Acquiring and transporting the material*

The specifics of nuclear material safeguards, commercial and military uses and storage are many and tangled. Providing any completeness on this question is beyond the scope of this thesis, but an introduction to materials acquisition which the reader may wish to refer to is given in the next section, and more extensively in the literature on this specific topic³⁸. In summary, available literature gives no reason to believe that one material should be significantly easier to come by than the other.

It is highly questionable whether the sparse record of radioactive materials actually seized in transit is any guide as to whether there exists an illegal market for fissile materials or not. The list of smuggling cases, rather short once all instances involving non-fissile or non-weapons usable materials are removed, paints a picture of a disorganised market for nuclear materials with amateurish sellers and no visible buyers. Most of the nuclear brokers appear to be opportunist, seeking quick profit from what they can get their hands on rather than responding to a real demand from a market³⁹.

The cases involving smuggling of HEU are both more numerous and also more serious (larger quantities involved) than those involving Pu, of which there are only two: 6.15g of nearly pure ²³⁹Pu was seized from the garage of a minor criminal in Germany in 1994 and the same year 363g of mixed-isotope Pu was seized at Munich airport as the result of a sting operation⁴⁰. A list of the known seizures of

fissile material trafficking is provided in appendix A. Given the small number of cases, the relative emphasis on HEU in smuggling seizures might not even represent a statistically significant trend, much less a real preference among potential buyers of fissile materials.

The radiation from plutonium might be a little harder to shield from detection than that of uranium. However, the needed material is neither large⁴¹, difficult to shield nor particularly dangerous if handled with care.

The most common way of detecting the presence of fissile material is using a gamma-ray detector⁴². Uranium and plutonium nuclei⁴³ are unstable and decay with time to other elements, and in each decay radiation is emitted, most often including at least one gamma quantum whose energy is specific to the decay process. The decay process, in turn, is specific to the nucleus that decayed, providing (in principle) a means to identify which nuclei are present and in which quantity. Typical γ -energies from ²³⁵U (the prominent isotope in HEU) are lower than that of relevant plutonium isotopes⁴⁴, and since as a rule of thumb more energetic radiation is more penetrating, a given γ -activity of radiation from a lump of Pu is in principle somewhat more easily detectable than the same activity from HEU. Furthermore a quantity of plutonium will have a higher activity than a similar mass of HEU.

Nonetheless assuming the plutonium is highly enriched in the isotope ²³⁹Pu, radiation levels are so low that although more shielding will be called for, it is debatable whether the difference in radiation detectability is important (for less optimal forms of plutonium the radiation levels can be much higher due to the presence of light elements). The chances that an item be randomly detected in transit this way are remote in both cases if the smuggler is clever⁴². For example, an experiment conducted by ABC News in

38 There exists an extensive literature that the reader may refer to. A good starting point is Richard L. Garwin and Georges Charpak *Megawatts and Megatons: A Turning Point in the Nuclear Age?* (New York: Alfred A. Knopf, 2001). Other excellent resources include the series of *Securing the Bomb* reports from the Project Managing the Atom and Brian Finlay and Andrew Grotto *The Race to Secure Russia's Loose Nukes: Progress since 9/11* (The Henry L. Stimson Center, 2005)*.

39 For an in-depth but somewhat dated discussion, see e.g. chapter 2 of Rensselaer W. Lee III *Smuggling Armageddon* (New York: St. Martin's Press, 1998). A more updated 'synopsis' is Lee's 'Nuclear Smuggling: Patterns and Responses' *Parameters* 33:1 (2003) pp.95-111.

40 Lists of smuggling cases are given in numerous sources, e.g. the *Nuclear Threat Initiative Research Library* webpage:

http://www.nti.org/e_research/e3_special_nucltrafficking.html

41 A critical mass of a HEU or plutonium metal sphere is an object the size of a softball or tennis ball respectively.

42 See appendix D.

43 Several uranium nuclei have very long lifetimes, however, the longest being ²³⁸U whose half-life is comparable to the age of the universe.

44 Gamma signature energies of ²³⁵U is approximately 186 keV, while both ²³⁹Pu and ²⁴⁰Pu (the prominent isotopes of Pu) emit γ -photons of energies around 640 keV. David Spears (ed.) *Technology R&D for Arms Control* (Office of Nonproliferation Research and Engineering, US Department of Energy 2001)* pp. 46 and 33 respectively.

co-operation with the Natural Resources Defence Council in 2002 smuggled 6.8 kg cylinder of depleted uranium (slightly less radioactive than HEU) by air from the United States to Vienna, thence by train to Istanbul crossing four extra-EU borders and by container ship back to the States. Although the container containing the cylinder was one of the few that were x-rayed upon arrival in the US, the uranium was not detected⁴⁵. An assembled implosion device using plutonium, supposedly easier to detect than HEU, is deemed 'almost undetectable with passive methods' in a careful study by Fetter and co-authors⁴⁶. Passive detection *en route* to construction site, it thus seems, is likely only subsequent to a tip or intelligence report.

With regards to acquisition and transportation, thus, I do not expect the differences between the two substances to incline terrorist preferences notably in either direction. We look therefore to the next phase of a proliferation project.

3.4.4 Designing and building the device

Comparison of the requirements for building a gun-type versus an implosion type nuclear device has been expertly treated by many authors⁴⁷ and I will not reiterate them but focus on what additional requirements a plutonium project brings as compared to the uranium gun.

Apart from weapons design, uranium as a material is significantly easier to handle than plutonium. For one, it is less radioactive; in fact a lump of HEU metal could be handled by hand without much of a health risk. But the greater danger from plutonium is not primarily from bulk radiation but inhalation of airborne particles⁴⁸. In order to form the weapon parts, the metal must be cast and then ground, milled or otherwise machined. Metal dust is produced in the process and inhaled plutonium dust is lethal even in milligram quantities because of its α -activity⁴⁹.

45 Christopher E. Paine 'Preventing Nuclear Terrorism', testimony for the Hearing of the National Security, Veteran Affairs, and International Subcommittee of the House Government Reform Committee (September 24, 2002)

46 Steve Fetter *et al.* 'Detecting Nuclear Warheads' *Science & Global Security* 1 (1990) p.246

47 e.g. Bunn and Wier 'Terrorist Nuclear Weapon...' and Ferguson and Potter *The Four Faces ...* pp.131-149.

48 If the plutonium is 'reactor-grade', that is, it has a large proportion of the isotope ²⁴⁰Pu as well as contamination of light elements, the bulk radiation can also pose a notable health risk.

49 R. H. Clarke *et al.* 'The environmental safety and health

Uranium and plutonium are corrosive and are therefore often not stored in the pure metal form most usable for weapons⁵⁰. If the material is stored as oxide powder or an alloy, chemical processing is needed to reduce oxide to metal or separate the metal from the alloy; in any such process, plutonium demands more care and special equipment to protect the worker from harm⁵¹.

An overall point is that of extreme health hazards when working with plutonium with amateur equipment due to the toxicity of plutonium when inhaled or otherwise taken in⁵². This requires some special equipment such as glove boxes and at least some simple remote-handling equipment. It should be noted, however, that the Afghan drug industry, with which al Qaida is reported to sustain connections, employs similar equipment for chemically converting opium poppy seeds to heroin, due to the toxicity of airborne heroin⁵³.

While some will argue that a terrorist who is willing to blow himself up in the attack anyway will be unconcerned about hazards to his health, one must remember that a weapons project of this type will take many months, more likely several years⁵⁴. Inhaling even milligram quantities of plutonium dust will result in death or disability in a shorter time than this. Furthermore, while the quick death in a spectacular bang has evidently been attractive to some terrorists, the slow and painful death by radiation sickness

implications of plutonium' *Journal of Radiological Protection* 16:2 (1996) pp.91-105.

50 Siegfried Hecker 'Towards a Comprehensive Safeguards System: Keeping Fissile Materials out of Terrorists' Hands' *The Annals of the AAPSS* 607 (2006) p.123

51 Mark *et al.* 'Can Terrorists Build Nuclear Weapons?' *op.cit.* p.58 and 61.

52 The toxicity of plutonium has often been much exaggerated. For a sober analysis, see Bernard L. Cohen 'Hazards from Plutonium toxicity' *Health Physics* 32:5 (1977) pp. 359-379. Cohen estimates approximately 1 eventual fatality per 18 grammes dispersed in a city and puts the lethal dose of inhaled plutonium dust at 200 microgrammes. Both uranium and plutonium are also chemically toxic due to being heavy metals.

53 Bunn and Wier 'Terrorist Nuclear Weapon...' p.138

54 Pluta and Zimmerman expect the building phase to 'exceed one year' whilst a project management analysis by Harney *et al.* indicate surprisingly that as much as 4 years could be required for a state proliferator to weaponise stolen HEU (bearing in mind that military weapons have other and tougher requirements than terrorist ones). Anna M. Pluta and Peter D. Zimmerman 'Nuclear Terrorism: A Disheartening Dissent' *Survival* 48:2 p.62 and Robert Harney *et al.* 'Anatomy of a Project to Produce a First Nuclear Weapon' *Science and Global Security* 14 (2006) p.169.

might not be as appealing⁵⁵, and finally, while pawns may be sacrificed, nuclear weapon scientists are hard to come by, and a single death by accident could be a major setback. Hence, whereas little regard for the health of the workers could speed up the process by avoiding the stringent safety regulations enforced in any legal facility, it could just as easily derail the project before a workable bomb is finished.

Work on explosives will also involve dangers. If an implosion device is made, extensive work with high explosives must be performed, presumably more dangerous than the relatively slow-burning propellants used in a gun design. Even if terrorists have extensive experience with explosives, accidents have been widespread in many terrorist organisations. For example, it has been estimated that the Irish Republican Army lost approximately 120 members due to accidents with explosives between 1970 and 1996⁵⁶.

Further complications arise because the stable form of plutonium at room temperature (the alpha phase) is brittle and very difficult to machine, necessitating some metallurgical solution - several such are well known but not easy to manage⁵⁷.

There can be no doubt that designing either weapon type is not a trivial matter and demands much effort by personnel with specialised skills. In their authoritative paper Mark *et al.* prescribe the following requirements⁵⁸

The preparation of [weapon design drawings] requires a large number of man-hours and the direct participation of individuals thoroughly informed in several quite distinct areas: the physical, chemical, and metallurgical properties of the various materials to be used, as well as the characteristics affecting their fabrication; neutronic properties; radiation effects, both nuclear and biological; technology

concerning high explosives and/or chemical propellants; some hydrodynamics; electrical circuitry; and others.

This was written in the 1980's however; nowadays computer drawing and simulation software requiring no more than a standard desktop computer can vastly simplify many of the above tasks and much of the information previously reserved for experts is now available on the internet and open sources⁵⁹. Furthermore one must again emphasise the difference between a terrorist and a military project⁶⁰, relaxing very significantly what 'thoroughly informed' means for each of the points mentioned.

It is tempting at this point to mention the 'Nth Country' experiment started 1964 at the Lawrence Livermore Laboratory in which two recent physics PhDs were asked to make a detailed nuclear weapons design based only on open and published sources⁶¹. 30 months after it commenced a blueprint was finished for a 20kT implosion device. A jury of experts deemed it workable beyond question. It should be noted that the two students, Dobson and Selden, opted for the implosion design because a gun design would be too easy⁶². Published resources that designers will find helpful have grown enormously in number since the 60s, and as Garwin and Charpak say, 'It should not be assumed that terrorists or other groups wanting to make nuclear weapons cannot read.'⁶³ One must be careful in extracting too much from this particular experiment for, while it shows that designing even an implosion device is not an impossible undertaking, the story of the Nth Country experiment neglects the practical problems of obtaining the necessary data for the design as well as building the actual device. The participating students could ask that experiments 'be performed' and were subsequently given the data. Many of these experiments would presumably involve advanced equipment such as flash X-ray cameras, pindomes⁶⁴

55 A point elaborated by Adam Dolnik 'Die and Let Die' *Studies in Conflict & Terrorism* 26 (2003) p.29

56 Brian A. Jackson 'Technology Acquisition by Terrorist Groups: Threat Assessment Informed by Lessons from Private Sector Technology Adoption' *Studies in Conflict & Terrorism* 24 (2001) p.193

57 The plutonium used in the first weapons was stabilised in the so-called delta phase by alloying it with a small concentration of gallium. Richard D. Baker, Siegfried S. Hecker and Delbert R. Harbur 'Plutonium: A Wartime Nightmare but a Metallurgist's Dream' *Los Alamos Science* (Winter/Spring 1983)* pp.142-151

58 Mark *et al.* 'Can Terrorists Build Nuclear Weapons?' p. 58

59 Point made by Pluta and Zimmerman 'Nuclear Terrorism ...' *op.cit.* p.63

60 A military weapon must be safe, reliable and predictable. Terrorist weapons need not be either. e.g. *ibid.* p. 61.

61 Dan Stober 'No Experience Necessary' *The Bulletin of the Atomic Scientists* 59:2 (2003) pp.56-63.

62 *ibid.* p.57

63 Garwin and Charpak *Megawatts and Megatons op.cit.* p.348

64 A device looking a little like a rolled up hedgehog, whose pins measure pressure. It is perhaps the simplest tool to measure the achieved symmetry of implosion. See Nuclear Transfer and Supplier Policy Division, Department of Energy (US) *Annex 3 of the Handbook for Notification of Experts to Iraq (UNSC*

and detectors for neutrons and gamma rays, none of which is easily obtainable⁶⁵.

In 1977 a Princeton undergraduate designed an implosion weapon from unclassified information for a thesis. His professor, who was knowledgeable about weapon design, gave him the best mark, and the US government classified the paper⁶⁶. While such events are notable, it is important to bear in mind that the path from blueprint to actual assembly poses many additional and significant hurdles.

By far the hardest part will be creating the converging shockwave in an implosion device. The propelling of the moving HEU piece inside the uranium gun is relatively straightforward (also literally) and the challenge is basically one of making a charge of the right strength so that the pieces combine fast enough, but without the propellant damaging the casing and tamper nor sending the uranium piece flying through the opposite end before a reaction can start. With enough natural or depleted uranium available and access to a secluded test field, this can be tried with a realistic system until the right balance is found⁶⁷. The same is not the case for a plutonium implosion system. Although one can find stand-in materials for plutonium, accounts of the extensive implosion testing at Los Alamos during WW2 show that a trial and failure procedure, even with unlimited access to explosive lenses (which are non-trivial and time-consuming to cast and shape, and not easy to buy in large numbers without detection), will take exceedingly long and more likely will never work at all. A considerable theoretical as well as experimental effort will be required⁶⁸.

3.4.5 Terrorists and fizzling: other plutonium options?

A large number of papers and books dealing with nuclear terrorism will tell the reader that a gun design cannot be used with plutonium, and that implosion will only be effective if the shockwave produced by

Resolution 1051) (DoE, 1998)* §11.4-§11.5

65 This point is emphasised in Michael Levi's *On Nuclear Terrorism* pp.74-76

66 Allison *Nuclear Terrorism: The Ultimate Preventable Catastrophe* pp.87-88

67 Peter D. Zimmerman and Jeffrey G. Lewis 'The Bomb in the Backyard' *Foreign Policy* (November/December 2006) p.37

68 See e.g. chapter 8 of Hoddeson *et al. Critical Assembly* to see how some of the world's most talented physicists and engineers struggled to obtain an acceptably symmetrical implosion and diagnose it.

the high explosives is highly symmetrical. But what if the terrorist built a plutonium gun nonetheless, or gave up trying to get the implosion right and decided to detonate a very crude implosion, maybe without explosive lenses?

The answer seems to be that the explosion would probably be small, but the terrorist could be lucky⁶⁹. If the terrorist got really lucky, a sloppy implosion system might be able to yield perhaps a hundred tonnes of TNT. With less luck, a few tonnes might be achieved, and in the other lower end of the scale (where the plutonium gun is found), only a little more than the energy of the high explosives themselves⁷⁰. The explosion, if small on a nuclear scale, could still be a most serious terror incident. Says Paine, a researcher for the US Natural Resources Defence Council (NRDC), 'If plutonium were used in a crude gun assembly device the yield most likely would be substantially less than a kiloton, but it could be larger than the explosion that destroyed the Federal Building in Oklahoma City.'⁷¹ The Oklahoma explosion equalled approximately two tonnes of TNT.⁷²

What one *would* get, especially if some fission were achieved, was a large dirty bomb⁷³, in itself highly dangerous. If the wind and weather was right, the radioactive cloud could travel far and make large areas uninhabitable with contamination and the public fear and panic it would inspire would likely be very considerable, even if the body count might not rival that of previous conventional attacks.

So why not, then, go for plutonium after all, being the most toxic of the fissile options, and make a massive radioactive poison bomb? Lewis A. Dunn gives a tentative theory why such dirty weapons may not be very attractive to the terrorists often deemed

69 For a comprehensive discussion of implosion design fizzle yields, see Mark 'Reactor-Grade Plutonium's Explosive Properties' *op.cit.* The discussion assumes a Trinity-type implosion system is used to implode reactor-grade plutonium, somewhat different from the question at hand, but the discussion might give some idea of the probabilities involved. Notably Mark concludes that 'not even the best weapons-grade plutonium is of any interest in connection with a gun-type assembly system.' (p. 5) I will not engage in a further calculation of fizzle yields for crude designs herein.

70 Office of Technology Assessment *Nuclear Proliferation and Safeguards* (OTA, 1977)* pp.141-142.

71 Christopher E. Paine 'Preventing Nuclear Terrorism', *op.cit.*

72 Richard A. Falkenrath, Robert D. Newman, and Bradley Thayer *America's Achilles' Heel: Nuclear, Biological, and Chemical Terrorism and Covert Attack* (Cambridge Mass.:MIT Press, 1998) p.155.

73 Richhard L. Garwin 'The technology of megaterror' *Technology Review* 105:7 (2002) p. 66

likely to employ them. He has studied al Qaida⁷⁴ throughout its history and attempts to see patterns in their choice of terror targets and methods⁷⁵. He too claims puzzlement that no radiological device has yet been used, but points to the trend for al Qaida to attempt to cause what he terms 'spectacular effects and visually pleasing destruction'. Whilst a successful nuclear detonation would probably satisfy the criterion of 'visually pleasing', a failed weapon might not. On this basis, he concludes that radiological weapons are inconsistent with al Qaida's past targeting practice⁷⁶ and one might extrapolate this to include a nuclear device in the case where the terrorists herself is fairly certain it will yield a fizzle or less.

Dunn also points to al Qaida's preference for operations of some complexity, often involving simultaneous attacks on several places at once, such as the failed 'Bojinka' plot to bring down 6 or more aircraft simultaneously in 1995, the 1998 bombing of US Embassies in East Africa, the four hijackings in the September 11 attacks and the Madrid train bombings in 2004. 'Simultaneous attacks' says Dunn, 'demonstrate the organization's operational sophistication to members, potential recruits, other outsiders and opponents'⁷⁷. Beyond doubt, a successful nuclear attack will demonstrate sophistication. The pressure to appear sophisticated in the eyes of their audience and get their one shot right could create a significant risk aversion.

While these arguments may not persuade anyone that e.g. al Qaida would not attempt usage given access to plutonium, it indicates that they might not be interested in investing great resources in a guaranteed fizzle, especially given that there is a preferable alternative, and if so will save their money until HEU is within reach. This choice as seen from the terrorist's point of view is detailed in chapter 5.

3.5 Availability of nuclear materials: safeguards and challenges

As the last part of the theoretical background, I will

74 In this chapter as in other chapters, al Qaida will be frequently used as example of a terrorist organisation.

75 Lewis A. Dunn 'Can al Qaeda be Deterred from Using Nuclear Weapons?' Center for the Study of Weapons of Mass Destruction, National Defense University *Occasional Paper* #3 (2005)* pp.9-10.

76 *ibid.* p.11

77 *ibid.* p.15

briefly go through the challenges of nuclear safeguards and the availability of nuclear materials to the terrorist in search of them. I will use the term safeguards in a broad sense, and define it as any specific measure to secure weapons usable fissile materials and nuclear weapons or weapon components from theft, or any other means of unauthorised removal, from its legal owner. In line with the focus of this thesis, by 'safeguards' I will often refer only to fissile materials. Safeguards measures, forming what is often termed the 'first layer of defence' against nuclear terrorism, thus include such measures as installation of security equipment at storage sites, hiring and training of personnel, implementation of improved security routines, relocating stores to fewer sites, improving measurement and accounting procedures, verification and transparency, as well as measures to dispose of excess fissile materials⁷⁸. By 'specific measures' it is implied that anti-terror and crime reduction efforts in general are not included.

In a notable paper⁷⁹, Allison sets out what he sees as a simple and workable doctrine to halt what he terms 'the ultimate preventable catastrophe' in a corresponding book⁸⁰. Allison introduces the 'doctrine of "Three No's": no loose nukes, no new nascent nukes, and no new nuclear states'⁸¹. The doctrine is based on a simple 'matter of physics: without fissile material, you can't have a nuclear bomb'⁸².

While this is an undeniable truth, critics have questioned whether Allison's doctrine provides much help in practice. Potter, Ferguson and Spector argue that Allison's view fails to grasp the vastness and complexity of the problem of safeguarding nuclear warheads and materials, and that more nuance is called for. Of interest to us is their argument that a sharp differentiation between plutonium and HEU is necessary, where the latter should be emphasised over the former. Allison, according to Potter and co-authors⁸³,

78 Such as by downblending HEU to lower enrichment levels for use in reactors, or storing plutonium as mixed oxide or in a matrix with high-level nuclear waste. See Frans Berkhout *et al.* 'Disposition of Separated Plutonium' *Science and Global Security* 3:3 (1993) pp. 161-213.

79 Graham Allison 'How to Stop Nuclear Terror' *Foreign Affairs* 83 (2004) pp.64-74.

80 Allison *Nuclear Terrorism: The Ultimate Preventable Catastrophe*

81 Allison 'How to Stop Nuclear Terror' p.65. By 'nascent nukes' is meant fissile, weapons-usable material.

82 *ibid.* p.64

83 William C. Potter, Charles D. Ferguson, and Leonard S.

...fails to make a crucial distinction between highly enriched uranium ..., which terrorists may already have the capacity to turn into the simplest [improvised nuclear device]..., and plutonium, which is much more difficult to turn into a weapon. Prior to September 11, when states presented the main proliferation challenge, it made sense to treat HEU and plutonium as roughly equivalent dangers. Today, however, when nonstate actors constitute a far greater nuclear threat, priority must be given to rapidly securing, consolidating, and eliminating the vast global stocks of HEU.

It is likely that Allison deliberately overlooks details in order to cut through to his most important point, that of the paramount importance of improved safeguards worldwide, a general principle which none of his critics disputes. Whilst this treatise will not attempt to assay the innumerable facets and subtleties of nuclear materials safeguards, I shall very briefly go through some of the aspects that make keeping these materials safe a unique challenge. Hecker gives five reasons why Allison's proclaiming that fissile materials should be kept as safe as the treasures in the Kremlin Armoury and the gold in Fort Knox⁸⁴, inspires an erroneously simplistic image of the real challenge.⁸⁵

1. *Existing inventories are thousands of times what is needed for one bomb:* Approximately 1.9 million kg of HEU and 0.49 million kg of separated Pu⁸⁶ exist worldwide. The

Spector 'The Four Faces of Nuclear Terrorism and the Need for a Prioritized Response' *Foreign Affairs* 83:3 (2004) pp.130-132. Note that whether non-state actors form a 'far greater threat' is a disputed issue not to be addressed here.

84 Allison 'How to Stop Nuclear Terror' p.64

85 Siegfried Hecker 'Towards a Comprehensive Safeguards System: Keeping Fissile Materials out of Terrorists' Hands' *The Annals of the AAPSS* 607 (2006) p.122-125. Figures and facts in the below list is from Hecker where nothing else is specified.

86 This figure differs slightly from Hecker's figure. Obtained from David Albright and Kimberly Kramer 'Global Stocks of Nuclear Explosive Materials: Summary Tables and Charts' in *Global Stocks of Nuclear Explosive Materials* (Institute for Science and International Security (ISIS), September 2005)*. ISIS' periodically updated reports on global stockpiles are probably the most authoritative account available in the open literature, although limited by secrecy and lack of transparency in many countries. The estimates of the International Panel on Fissile Materials are 1.7 and 0.50 million kg of HEU and Pu respectively. IPFM *Global Fissile Material Report 2007* (IPFM, 2007)* p.10 and 14.

quantities needed for a crude nuclear weapon may be roughly 50kg and 15kg respectively⁸⁷, much less than even the uncertainties with which physical inventories are accounted for.

2. *Fissile materials exist in many different forms:* Uranium and plutonium are highly reactive metals that oxidise quickly and are therefore not stored as pure metal bars like the gold of Fort Knox, but e.g. as oxide or nitride powders, uranium hexafluoride gas or in alloys with other metals like aluminium⁸⁸. For weapons and reactor use, metallic or ceramic (oxide) forms are required respectively, so to become fuel elements or weapons parts the fissile materials undergo chemical processing. In the weapons case, casting, and machining are also necessary. For plutonium, moreover, all of these processes must be performed with remote equipment in specially sealed containers due to its radioactivity, toxicity and special chemical properties. Accounting for fissile materials gram for gram throughout all these processes is clearly a daunting task.
3. *Fissile materials exist in many and diverse locations:* Separated plutonium and HEU exist in state owned enrichment, fuel processing and reprocessing facilities and in storage facilities, as well as in transport between these. Furthermore, HEU is still used to fuel some 140 research reactors in 40 countries, many of which private and with extremely varying security instalments⁸⁹.
4. *Fissile materials are difficult to measure and handle:* Keeping track of materials means measuring the weight and enrichment level of the material, a difficult task even in idealised laboratory conditions. Plutonium requires special handling for health reasons, and moreover exists in seven different metal forms, all of different density. Gamma-ray and neutron detection are typically used to measure mass and enrichment level for

87 These figures assume that crude technology is available, and that a gun design is used for HEU and implosion for Pu.

88 Ferguson and Potter *The Four Faces of Nuclear Terrorism* p.120

89 Figures differ slightly. Hecker reports 120, Bunn 147 and the International Panel on Fissile Materials (IPFM) reports 140. Hecker 'Towards a Comprehensive ...' p. 124; Matthew Bunn *Securing the Bomb 2007*, report of the Project Managing the Atom (Belfer Center for Science and International Affairs, Harvard University, 2007)*; IPFM *Global Fissile Material Report 2007* (IPFM, 2007)* p.13.

plutonium and HEU, but the technique is not perfect because a lump of such heavy metal will absorb much of its own radiation, hence the measurement strictly only applies to the *surface* layer of a metal piece while its bulk may hide a different isotopic composition. Spoofing measurements is thus possible (hiding HEU inside a layer of depleted uranium, for example). Signature radiation of ²³⁵U is so weakly penetrating, moreover, that its presence is easily shielded from detection⁹⁰.

5. *Military secrecy hampers safeguards and transparency*: Nuclear weapons design and stockpile specifics are amongst the most closely held military secrets. Verification of stockpiles performed by foreign inspectors⁹¹ therefore necessitates the use of so-called non-intrusive measurements (verification of reported weight and enrichment level without revealing secrets such as weapons design and geometry), which is inherently less accurate than if access were more direct. In some secret nuclear weapons sites, foreign inspection personnel are not granted access at all, as is also the case in worrisome stockpiles such as those of Pakistan, India and North Korea⁹².

For these five reasons, Hecker says, 'simply locking up all of the materials is not a feasible course of action. Many states do not even know what "all" is'⁹³. A sixth reason is that nuclear materials are tradable and traded commodities and therefore a large amount of such materials will be in transit between buyer and seller at any given time, much increasing the difficulty of keeping detailed track of it all⁹⁴.

90 The best overview of technology for verification measurements of fissile materials is perhaps David Spears (ed.) *Technology R&D for Arms Control Office of Nonproliferation and Arms Control* (US Department of Energy, 2001)*.

91 The government in question, presumably, has every access to its own weapons and materials.

92 Recent developments show some promise that North Korea may again permit UN inspections of its nuclear complex (note that unlike the Pakistani weapons, the nuclear warheads of North Korea and India use plutonium, not HEU). Also Pakistan has recently confirmed that it is co-operating with the United States on nuclear security and accounting, but very little information is currently available as to the nature of this co-operation. Bunn *Securing the Bomb 2007* p. 59.

93 Hecker 'Toward a Comprehensive Nuclear Safeguards System' p.125.

94 Point made by Peter Zimmerman, private communication.

Today's safeguards, sadly, are often assessed to be clearly inadequate to deal with the risk of proliferation of fissile materials to terrorism⁹⁵. Whilst most of the focus is directed to Russia, both in literature (such as the reports from *Managing the Atom*) and international programmes (the United States are currently funding several bilateral programmes to safeguard and eliminate stocks of fissile materials in Russia totalling more than one billion dollars annually⁹⁶), the amount of material needed for one bomb is so small it can be obtained from almost any of the more than 40 countries in the world keeping such materials⁹⁷. Bunn and Wier sum up the situation outside Russia thus 'There are no binding global standards for nuclear security, and in practice the security at sites where the essential ingredients of nuclear weapons are located ranges from excellent to appalling'⁹⁸. The existing legislation, the *Convention on the Physical Protection of Nuclear Materials*⁹⁹ is, according to critics, 'vague in its requirements, applies primarily to international transport of materials, does not cover military materials at all, and has no provisions for verification or enforcement'¹⁰⁰.

Since the fall of the Soviet Union, Russia's nuclear weapons complex has been considerably downsized, leaving thousands of nuclear experts, who used to be considered the nation's finest, unemployed or on unsustainably small salaries. The complaint was put succinctly by the head of the Russian nuclear inspection agency some years ago, saying 'Highly qualified specialists who work in secret nuclear towns

95 See e.g. Hecker 'Toward a Comprehensive...', and Matthew Bunn *Securing the Bomb 2007* (Belfer Center for Science and International Affairs, Harvard University, 2006)*

96 e.g. Anthony Wier and Matthew Bunn *Funding for U.S. Efforts to Improve Controls Over Nuclear Weapons, Materials and Expertise Overseas: Recent Developments and Trends* report from Project *Managing the Atom* (Belfer Center for Science and International Affairs, Harvard University, 2007)*, table 1.

97 Notably, most research reactors do not store enough fissile material for a weapon, so to obtain an adequate amount, thefts would likely include several such sites.

98 Matthew Bunn and Anthony Wier *Securing the Bomb 2006* Project *Managing the Atom* (Belfer Center for Science and International Affairs, Harvard University, 2006)*, p.19.

99 International Atomic Energy Agency *Convention on the Physical Protection of Nuclear Materials With amendment* (2005)*

100 Matthew Bunn, John P. Holdren and Anthony Wier *Securing Nuclear Weapons and Materials: Seven Steps for Immediate Action* report from Project *Managing the Atom* (Belfer Center for Science and International Affairs, Harvard University, 2002)* p.57

earn less than the cleaning women who work in the Moscow subway¹⁰¹. This downsizing, concludes Rensselaer Lee III, an authority on nuclear smuggling, has 'catastrophically corroded employee well-being and morale, greatly increasing the risk of nuclear theft'¹⁰².

This fact is underlined by a study of the actual smuggling attempts that have been interdicted¹⁰³. These incidents tend to be spurred by the economic needs of opportunistic personnel seeking quick profits, only to fall victims of police stings (indeed, German police has been criticised, by Lee among others, for encouraging the theft in the first place, creating the very market they claim to hunt down and in the process victimising employees in a desperate situation who would not otherwise have perpetrated the misdeed). Noting the ease with which police operations have been able to persuade employees to commit theft for money, it is unlikely that terrorist organisations will find it more difficult, especially if they have ties to Russian organised crime. Such a group would also be willing and able to exercise types of pressure on employees that the police never could.

The situation in Russia has improved somewhat since the 90s, with at least temporary employment provided for many former nuclear workers¹⁰⁴. The fact remains, however, that a large number of Russians holding exactly the expertise Al Qaida would require in order to turn stolen material into a weapon are still in dire economic need, unable to support their families. A report from the Carnegie Endowment for International Peace of 2001 reported that some 200 scientists in Russian nuclear cities proclaimed themselves willing to 'work for anyone and do anything'¹⁰⁵. While Russia's economic upturn of recent years has improved the conditions in the country's

nuclear sector, it is likely that should an economic recession occur, conditions will again turn to the worse. Russia's economic boost has been based largely on high oil prices, a cheap rouble and low domestic labour cost, all fluctuating conditions that are not guaranteed to last¹⁰⁶.

Nonetheless, Hecker only rates Russia the fourth greatest proliferation threat today, after Pakistan, North Korea and HEU-fuelled research reactors worldwide¹⁰⁷. Hecker's point is a valid one: one must not be blinded by the vastness of the Russian nuclear materials stockpile, for the amount of materials needed for a bomb is comparatively so minute it could be obtained from a number of countries. Pakistan already hosted the notorious A.Q. Khan, a grand-scale broker in nuclear technology¹⁰⁸, and some of its tribal areas in the north are dominated by militant Islamic groups sympathetic of al Qaida. Whilst Pakistan's weapons and materials might be physically well protected, the threat from armed attacks and inside theft is cause for much concern. Given the recent political turmoil in the country, it is not hard to envision a worst case scenario in which an extremist Islamic faction comes to power, perhaps by a *coup d'état*, creating the world's first Islamic fundamentalist nuclear weapons state¹⁰⁹. A recent paper by Luongo and Salik downplays such concerns, yet admits there are many challenges left, and their sense of security seems to depend heavily on the unity and stability of the Pakistani military¹¹⁰.

North Korea has already proven its readiness to sell missile technology, even assembled missile systems. Recent press reports have indicated that North Korea may have supplied a plutonium producing nuclear reactor to Syria¹¹¹. While perhaps still a big step, covert wholesale of excess separated plutonium, one may speculate, could be a possible

101 Rensselaer W. Lee III *Smuggling Armageddon: The Nuclear Black Market in the Former Soviet Union and Europe* (New York: St. Martin's Press, 1998) p.37. Lee in turn quotes the German magazine *Die Woche*.

102 Lee *Smuggling Armageddon* p.35.

103 For a list of smuggling cases with some details, see e.g. the *Nuclear Threat Initiative Research Library* webpage: http://www.nti.org/e_research/e3_special_nucltrafficking.html. Based on the International Atomic Energy Agency's *Illicit Trafficking Database* available to members from <http://www-ns.iaea.org/security/itdb.htm>, a synopsis of which is found in appendix A.

104 Bunn and Wier *Securing the Bomb 2006* pp.84-85.

105 Jon Wolfsthal 'Surveying the Nuclear Cities' *Bulletin of the Atomic Scientists* 57:4 (2001) pp.15-17. Also Valentin Tikhonov *Russia's Nuclear and Missile Complex: The Human Factor in Proliferation* (CEIP, 2001)*

106 Pluta and Zimmerman 'Nuclear Terrorism: A Disheartening Dissent' pp. 59-60.

107 Hecker 'Towards a Comprehensive Safeguards System' p.130.

108 See Gordon Corrae *Shopping For Bombs* (London: Hurst & co., 2006)

109 Many have voiced this worry, e.g. Ferguson and Potter *The Four Faces ...* p.125.

110 Kenneth N. Luongo and Naeem Salik 'Building Confidence in Pakistan's Nuclear Security' *Arms Control Today* (December 2007)*

111 Glenn Kessler 'N. Korea, Syria may be at work on nuclear facility' *Washington Post* (September 13, 2007)* p. A12; 'Report: North Korea provided technical assistance to Syria to build nuclear reactor' *International Herald Tribune* (February 18, 2008)*

next step for a regime desperate for money¹¹².

Based on this it may not be so surprising of Pluta and Zimmerman to go so far as to conclude that 'It seems certain that at some price nuclear explosive material is available to well-funded terrorists.¹¹³ Assuming this is true, what would the optimal strategy be for a terrorist organisation with such ambitions? Much as al Qaida has proven to be well funded (the attacks on September 11 alone are estimated to have cost them some \$400,000-\$500,000¹¹⁴), the purchase of fissile material will doubtlessly be a large investment. Zimmerman and Lewis find that if al Qaida were to choose its method of attack by the appalling index of price per person perished, a nuclear weapon would likely be the most cost-efficient by far. But as they point out, 'spending \$5-10 million to kill 100,000 people is a bargain only if you have \$5-10 million to spend in the first place.¹¹⁵ Such a sum would not be spent lightly, and it is likely that Bin Laden and his co-conspirators would think hard and well before making their choice of what to buy. The old rule of thumb 'the higher the stakes, the more rational the players'¹¹⁶ seems to be applicable.

3.6 Conclusions so far

This chapter has analysed and compared the nuclear terrorism threat from highly enriched uranium (HEU) and plutonium. Our main conclusion thus far is that insofar as the terrorist has a choice she would choose HEU over plutonium for her nuclear project. As quantitative estimates go (with reference to the next chapter) I conclude from the analysis herein that the *a priori* probability of a terrorist attempt to construct nuclear weapons will involve HEU rather than Pu is greater than 50%.

Most importantly, a HEU gun-assembled nuclear weapon is much simpler to design and build than a plutonium device, which would need to be realised by an implosion mechanism making advanced use of explosives. Even a project to make a rather crude plutonium implosion would exceed a uranium gun project in operational complexity and its probable

yield would still be significantly smaller than that of the gun due to a high probability of predetonation.

Since even a fizzle yield will be a very significant terrorism incident, predetonation is only a concern to a terrorist who is either motivated to cause maximum damage or has a fear of appearing incompetent in the eyes of potential supporters, or both. I argued that there is empirical evidence in favour of the latter of these notions, making for a risk aversion which could incline preferences towards the safer of the two options: HEU. The notion that 'maximum damage' be a primary measure of success in the terrorist's calculus is perhaps a little simplistic, yet is arguably a useful assumption for policy making as demonstrated and discussed in chapters 4 through 6.

Furthermore the properties of plutonium pose significant hurdles which a project using HEU would avoid. Plutonium is more radioactive and poisonous and processing it calls for some extra equipment such as glove boxes, inert atmosphere casing and remote handling devices. Moreover, plutonium is a difficult material to work with, being extremely brittle in metal form.

The qualitative analysis presented in this chapter, which synthesises research by a number of experts over the years, paves the way for the model introduced in the next chapter where a formal decision theoretic analysis will be undertaken.

112 Hecker 'Towards a Comprehensive Safeguards System' p.130.

113 Pluta and Zimmerman 'Nuclear Terrorism ...' p.60.

114 National Commission on the Terrorist Attacks Upon the United States *The 9/11 Report* (St. Martin's Paperbacks, 2004)* p.249

115 Zimmerman and Lewis 'The Bomb in the Backyard' p. 39.

116 See e.g. Jon Elster 'When Rationality Fails' in Schweers Cook and Levi (eds) *The Limits of Rationality* (Chicago: University of Chicago Press, 1990) pp.40-41.

Safeguards: HEU vs. Plutonium — a formal analysis

In the previous chapter I argued that there are very good reasons to believe that given the choice, a terrorist will opt for highly enriched uranium over plutonium. Furthermore a nuclear attack which fizzles, while probably sufficient to dwarf most terror incidents to date, is still far preferable to the targeted government than if the bomb had the yield of the Hiroshima attack, obtainable with the simplest type of nuclear weapons.

4.1 Chapter overview

The research question to be addressed in this chapter was introduced at the beginning of the previous chapter where it was partially treated in a qualitative way. In this chapter I will use the conclusions of chapter 3 to devise a model with a view to analysing what policy implications follow from them. The chapter starts with a presentation of the key concepts of decision theory drawn upon in the analysis.

Two potentially powerful results (equations (4.10) and (4.12)) are derived and their interpretation and significance are discussed. A demonstration of the applicability of these formulas is given thereafter, when some estimates for the real numbers behind the algebraic symbols is provided in the case of current safeguards efforts undertaken by the United States primarily in the former Soviet Union. A further development of the theory makes way for a succinct criterion to determine situations where full priority should be given to HEU over Pu. Finally, the main assumptions and simplifications inherent in the formal model used are discussed to assess the robustness of the conclusions drawn from the analysis.

4.2 Introduction to the methodology and a toy model

Before embarking on the full analysis, let us consider a 'toy model' in order to explain the basic principles of the methodology employed in this and

several subsequent chapters.

The concept of a *game* in rational choice theory was introduced in chapter 2 and we quickly recapitulate the basics as well as introduce the conceptually simpler *decision theory*¹. Each game has one or more players, each with a set of options for how to play the game. The word 'game' becomes natural since just as in most board games, our formal game also has a sequence of one or more simultaneous or subsequent *moves*. In chess, for example, moves are consecutive and alternate between two players. A plan of how to move under different conditions is called a strategy. When there are two or more players, we call the theory employed *game theory*, and in this case the best action by a player will depend on which strategies the other players employ. For game theory models the analysis is typically a quest for so-called equilibria of strategies, that is, sets of strategies (collections of one strategy per player) that are *best responses* to each other. When all players play according to a set of strategies which form an equilibrium, no player has any incentive to deviate from their equilibrium strategy². In this way, especially when a game has only one equilibrium, the equilibrium concept has been found to predict and describe the behaviour of actors in certain real scenarios as our example of the prisoner's dilemma demonstrated.

To define rigorously what is meant by 'the best' outcome for a player, rational choice theory introduces the concept of a *utility function*, which is a measure of each player's *utility* or *payoff*. Per definition, a rational player will try to maximise his expected utility³. The utility function will depend on which *outcome* the game has (any interesting game has more than one possible outcome), and decides each player's preferences with respect to the possible outcomes. In a single game of chess the preferred outcome (highest value of the utility function) is obviously to win the game, a draw is in between and a loss is the least preferred outcome. We will denote the utility function with the symbol *U*.

1 There exists a large number of introductory textbooks in game theory and rational choice theory. For an easy-to-read introduction, refer to Robert Gibbons *A Primer in Game Theory* (Hemel Hempstead: Harvester Wheatsheaf, 1992)

2 The equilibrium described here is the so-called Nash equilibrium. This concept has been generalised in many ways to extend to more complex games. See Gibbons *ibid*.

3 There are alternatives to utility theory which will not be used in this thesis. See discussions in chapter 2 and John Broome 'Should a Rational Agent Maximize Expected Utility?' in Schweers and Cook *The Limits of Rationality* (Chicago: The University of Chicago Press, 1990) pp. 132-145.

A game like chess is said to be a game of *perfect information*, which means that both players know at any stage of the game exactly what the history of the game is - who has done what. An example of a game of imperfect information could be a game of rock-paper-scissors; each player must now respond to the other player's action before having observed what it is.

Chess is also a game of *complete information*: both players have all relevant information about the game and can put themselves in the other player's stead and ask 'what would I have done if I were he/she?'. Formally we may say that each player knows the other player's possible strategies and corresponding payoff function⁴. In card games, however, the situation is different. Here, each player only knows their own cards and hence cannot be sure what the opponents' best responses are to his or her own actions. Mathematically, the players must now create utility functions by assigning probabilities to the different possible hands the other players may have and revise the utility function in the light of new information for each step of the game. Most real players will perform this process based on experience rather than arithmetics, but professional card players will often play by such explicit calculations of odds.

In the case where there is only one player, the theory is called *decision theory*. The goal is still to devise a utility function and maximise this. Decision theoretical models are most interesting in the case where a decision must be made under uncertainty, that is, with imperfect or incomplete information about the environment in which the decision must be made. Where there is a lack of information, the player must form a *belief* about what the probability is for different possible *states* of the world in which the decision is made.

Let me first regard a very simple 'toy model' of a game in which a government actor must decide how much money to spend on anti-terrorist activities. His utility function will have the form

$$U = -C - C_T \cdot p_T$$

where U is the utility, C is the (absolute) cost of antiterrorism measures, C_T is the (believed) cost of a terrorist attack and p_T is the (believed) probability of such an attack. Both C_T and p_T can depend on C . The utility function will then look something like that shown graphically in figure 4.1. Note that payoffs are

⁴ Specifically, to know the other player's payoff function one needs to know what the other player's possible strategies are.

negative for all C as they should in an antiterror game. To the left of the ideal value, the slope of the graph is positive, to the right it is negative corresponding to underfunding and overfunding respectively. The expenditure C which our government player can expend is the sole variable, so the utility function is only a function of C .

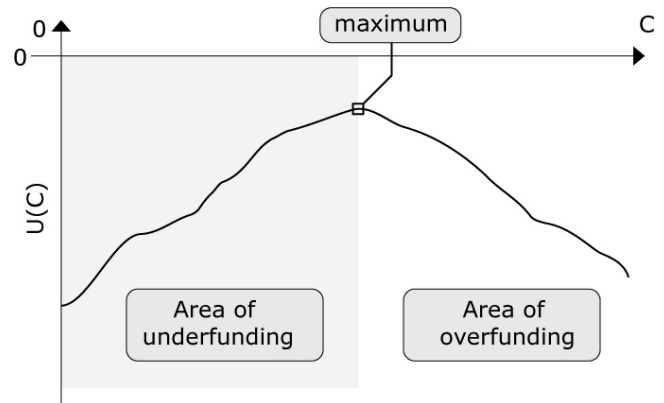


Illustration 4.1: Utility function of toy game.

As seen in the figure, the utility function in this case will arguably have a single maximum and fall off in both directions from this maximum. The reason for this assertion is that if the government has spent very little on anti-terrorism measures, a relatively large gain in security can be obtained for a small sum of money. On the far side of the maximum, too much is already spent and spending more is a waste of money.

As an analogy, consider a house full of valuable things, but whose door does not have a lock. Imagine that local thieves in the neighbourhood have been known to break into doors without locks with an 80% probability but only 20% if they need to force the lock. Imagine furthermore that the thieves are expected to steal \$1,000 worth of goods. If a lock costs \$100 it is clearly a good investment, since it decreases the expected loss from thievery from 80% \$1,000 = \$800 to 20% \$1,000 = \$200, that is for an expenditure of \$100 the owner of the house gets \$600 worth of security. If, however, the house is already a fortress and the probability of burglary is negligible to start with, one extra lock makes no difference, and the \$100 is a waste of money.

4.2.1 Improvement by estimate of slope: method of steepest ascent

The ideal solution would be for the government with the utility function of figure 4.1 to find the

maximum on the curve and spend this sum. In real life, however, this strategy is often impossible: In order to estimate the shape of $U(C)$ for all values of C the government must estimate a set of probabilities pertaining to all the unknowns involved in the decision. It is normally possible, with some sound judgement, to come up with estimates of what such probabilities might be under the conditions one observes today, but in order to determine what $U(C)$ looks like for *all* values of C , the player must also be able to estimate these probabilities as they *would be* under very different, hypothetical conditions. In the situations I shall consider uncertainties become so large so quickly that looking for the maximum in this manner is not helpful: even if we managed to allocate it, its position would be so uncertain that the result would most likely be useless.




Illustration 4.2.: Model utility function with two free variables.

There is an alternative approach, however, which is not as ideal but much more powerful. Assume now that the government does not spend all of the money in one go, but a little at a time. This models how governments typically work, with a new budget to be approved every year. For each round, then, the government can re-assess the parameters that underlie the utility function and find out whether to spend more money. The way to determine if enough money has been spent or not is then to measure the *slope* of the curve at status quo: if the curve is pointing upwards, maximum is not yet reached and more should be expended. In optimisation theory this is called the *method of steepest ascent*⁵. If the curve is flat

5 Or 'descent' depending on definition of the utility function. For a more advanced introduction, see e.g. chapter 9 of

or has negative slope, however, enough has been allocated and the government player has come as close to the equilibrium as he could. While this procedure could miss the ideal value by a fair bit, it has the major advantage that only estimates of present parameters are needed, so the approach can be employed with much better confidence.

In the following section we will consider a game with two free parameters rather than one. The generalisation is immediate and is shown in figure 4.2. The obvious analogy now is that of a mountaineer wanting to reach the top of a hill, the higher he gets, the higher the payoff. He needs now to find the right *direction* to walk in (previously we could only choose between back and forth) in order to approach the top of the hill. The ideal solution is if he can see the summit and head straight for it, but imagine he is walking in fog so that visibility only extends a little way in every direction. This corresponds to a government having very vague information about how the world would have been in states different from the actual situation today. The best plan in this case is instead to measure the slope of the ground and walk in the direction where the terrain rises the most steeply⁶.

There is one sense, however, in which the mountaineer analogy is not good when the game is like that above where a level of expenditure is to be decided on. While the mountaineer can move equally well in both directions along the east-west and north-south axes, a total expenditure can only be increased. There is no way to un-spend money which is gone⁷. In the hill analogy it is as though the mountaineer could only choose from directions in between, say, due east and due north. If, say, the walker finds himself to the north west of the peak already, the best he can do is walk due east until he reaches a maximum along the east-west axis. That is as high as he can come. Analogously, if the government finds it has already spent too much on some project, the maximum payoff which was available at the start can no longer be obtained. The best strategy is simply to cut funding to this project.

There is a subtlety when interpreting the kind of

Hubertus Th. Jongen, Klaus Meer, and Eberhard Triesch *Optimization Theory* (New York: Kluwer, 2004).

6 Mathematically this direction is given by the *gradient* of the utility function, as we will discuss briefly later. See also appendix C.

7 In special cases it might be possible to sell equipment and get most of the money back, but this would be the exception. As a rule sunken cost cannot be retrieved.

game whose payoff function is depicted in figure 4.2 in terms of cost and benefit. In the above I have implicitly assumed that the payoff function is a static quantity given once and for all, which is to say that any change in the state of the world comes about as a consequence of the player's choice of action. If U represents a mountain, this is a good approximation on the time scale in which humans operate. In politics, however, this is not necessarily so. The player could find that as he takes steps to try and close in on the top of the hill, it is as though the hill itself moves bringing him either higher or lower in the process. I will return to the implications of the assumption of static payoff later in this chapter.

4.3 The model

With the analysis presented in chapter 3 and the methodology outlined above I now have all the necessary considerations to introduce game theoretical model. The model⁸ is shown in figure 4.3. It is at first a generalisation of the above since it has two players; we may call them an antiterrorist ('player A' or just 'A') and a terrorist ('player T' or 'T'). I will show how the game can eventually be reduced to a choice theoretical problem seen from the point of view of the antiterrorist player.

The antiterrorist is most readily interpreted to be a government⁹; there is no need to specify a specific state, yet we assume the government in question regards itself as the likely target of the attack (the calculus changes slightly if this is not the case). The 'angle' symbol, \angle , denotes a continuum choice node, a square a discrete choice node and a circle is either a chance node or an end node.

The game starts by player A deciding on the costs c_u and c_p to be spent on anti nuclear terror activities; the former is the sum to be spent on anti uranium terror efforts and the latter is the sum dedicated to the plutonium branch¹⁰. Note that the dimension of T_u, T_p, C, c_u and c_p is money (\$, £, etc.), whereas all other symbols are non-dimensional. For convenience in the following we define the total cost $C = c_u + c_p$. In the following we assume c_u and c_p to be *independent* and

free variables, but will later examine what the consequences would be if C were fixed and player A were only free to distribute the total sum among the two branches.

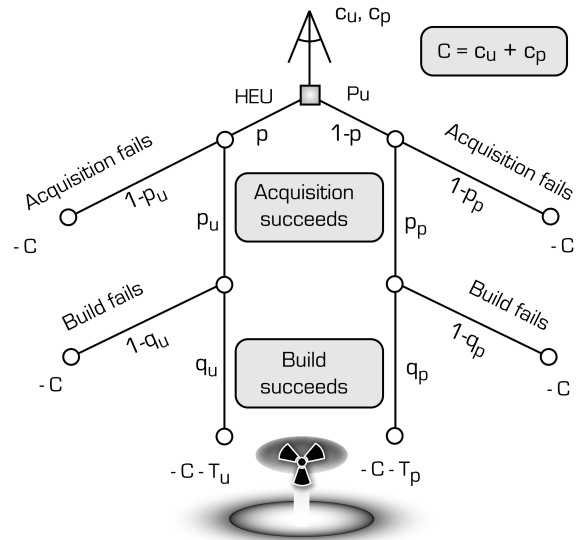


Illustration 4.3.: Layout of safeguards game: HEU vs. Pu

Much in the way of the natural sciences I do all my calculations with symbols rather than numbers - this allows me to reach very general results since the values behind the symbols may still be varied in the end. I will return to a discussion of exactly how they fit the real picture later.

Next, the terrorist, player T, chooses whether to opt for HEU or Pu to achieve her nuclear ambitions. Nature in turn decides whether player T manages to acquire the necessary material (probability p_u and p_p for HEU and Pu respectively) and finally whether she succeeds in building a working device and using it in a successful attack (successful design, construction, transportation and detonation are considered a single achievement for simplicity - probabilities q_u and q_p of success).

I will focus on player A and therefore only specify the payoffs (i.e. utility) for this player in this chapter¹¹. There are six different outcomes of the game. Four of these produce no terror attack, and are assumed to give the same payoff for A, equal to the cost of the safeguards effort¹², $-C$. In the case of successful attack,

8 inspired by one used by Sandler and Arce. Todd Sandler and Daniel G. Arce M. 'Terrorism & Game Theory' *Simulation & Gaming* 34:3 (2003) pp.326-329.

9 Generalisations are possible, but not considered herein.

10 Strictly speaking these two categories are not mutually exclusive since some facilities such as weapons assembly and disassembly factories store both materials. This is not too problematic, however, as will be explained.

11 I assume, of course, that the preferable outcome for the terrorist is a successful nuclear attack.

12 Surely there are many antiterrorism expenses other than nuclear safeguards, but to the extent they do not affect p 's and q 's they merely add a constant to all payoffs which is of no consequence to our analysis.

there is an additional cost. For generality I distinguish between attacks with the two materials: if the bomb is made with HEU, the cost of an attack is T_u and if the material is Pu, it is T_p .

The game is one of incomplete information; in particular, player A is assumed not to know for sure what player T's perceived payoff function is for the two options. Based on whatever information he has available, however, he has the *belief* that player T chooses HEU with a probability p and plutonium with a probability $1-p$. The game is viewed from the perspective of player A, and so the quantities p_u, p_p, q_u, q_p, T_u and T_p are player A's estimates (not necessarily equal to T's). Explicitly, what happens is that by spending c_u and c_p , A is able to shift the values of p, p_u, p_p in a favourable way before the terrorist gets the chance to move. Thus, A needs to estimate both what the present values are and how different spending levels are likely to change this.

It is worth emphasising that it is up to the beliefs of player A whether he considers it likely that a bomb will produce a significant yield, with correspondingly large values for T_u and T_p , or if he thinks the bomb has a significant probability of being a fizzle but still produce a sizeable bang and considerable radiological contamination. However, one must demand that he does *not* count an attack causing a fizzle as a 'failure' at the second chance nodes, since an estimated payoff $-C$ would be erroneous in this case; the cost of a grand scale radiological attack will also be huge, although small relative to a true nuclear explosion¹³.

If player A thinks it overwhelmingly likely that a terrorist produced implosion device will produce a fizzle, say, whilst a HEU gun-assembled device is more likely to give a several-kiloton yield, this should be reflected in a value of T_p smaller than T_u . Formally one might introduce separate damage estimates for fizzle and non-fizzle, say T_f and T_n , and a probability for non-fizzle, say ρ , so that $T = \rho T_n + (1 - \rho)T_f$ (different values of ρ for HEU and Pu is understood).

13 Player A could in principle assume that player T would never detonate a weapon unless she *knew* it would work, an unrealistic assumption since a small scale enterprise would have few chances of establishing such confidence in the performance of its device before it was tested. Even at the Trinity test, the first test of an implosion device, the nuclear physicists at Los Alamos were uncertain whether the weapon would work as they hoped. The anecdote goes that physicist *extraordinaire* Enrico Fermi annoyed his anxious colleagues the night before the test by jokingly taking bets on all conceivable outcomes of the test, including whether it would ignite the atmosphere. See e.g. Richard Rhodes *The Making of the Atomic Bomb* (New York: Simon & Schuster, 1986) chapter 18.

T_u and T_p are thus estimated by averaging over different possible yields (one could generalise this to different target choices as well). We will perform such an estimate in a later section of the chapter.

The game in figure 4.3 may be reduced to a decision theoretical problem, because since A has so little knowledge of player T's reasoning, T is fully represented by the single variable p . By introducing probability estimates not only regarding Nature's moves (p_u, p_p, q_u and q_p) but the move by the terrorist as well (p), and assuming that all player A can do after investing his resources is wait and pray, the terrorist threat as modelled is mathematically equivalent to a partly preventable natural disaster. The simplification that the q 's are constants is discussed briefly later in this chapter and at length in chapter 5.

4.3.1 Mathematical assumptions

I will need some assumptions about how the probabilities p, p_u, p_p, q_u and q_p vary with the variables c_u and c_p . My first assumption is that the probabilities of successful acquisition of material depend only on the amount of resources channelled to the branch in question, meaning that

$$p_u = p_u(c_u) \tag{4.1}$$

$$p_p = p_p(c_p) \tag{4.2}$$

and that these are decreasing functions of their argument, that is: the more resources are poured into a branch by player A, the smaller the probability that player T will get hands on the necessary material becomes:

$$\frac{dp_u}{dc_u} \leq 0 \quad \text{and} \quad \frac{dp_p}{dc_p} \leq 0 \quad . \tag{4.3}$$

Equations (4.1) and (4.2) are an approximation. In practice they imply that I neglect possible synergy effects between the two safeguards branches. While samples of the two elements are rarely protected by the same piece of equipment as discussed below¹⁴, it is plausible that less palpable measures such as security culture could spill over between the two branches. There is reason to believe that in the light of the large overall uncertainties pertaining to antiterrorism policy decisions, this approximation is unproblematic.

14 Charles D. Ferguson and William C. Potter with Amy Sands, Leonard S. Spector and Fred L. Wehling *The Four Faces of Nuclear Terrorism* (New York: Routledge, 2005)* pp.120-124.

In practice it takes a certain amount of funding to keep safeguards at *status quo* so they do not fall into neglect and disrepair. This causes subtle problems which have to do with the assumption of constancy of the utility function as discussed at the end of section 4.2.1. We have not worked out U yet but it is clear that it must depend on p_u and p_p hence if A were to do nothing we could be expected to 'drift' further away from the optimal value since, left to their own devices p_u and p_p would increase with time at some rate. Clearly, significant funds must be allocated to countering this¹⁵. On the other hand, other countries could choose to take action to improve safeguards on their own, causing the opposite effect. As I will return to, however, my model does not in fact have a 'time' and a more complex model would be called for if one wished to fully incorporate such effects. The easiest is to simply assume that C does not include costs to retain *status quo*, only those that go to improvement, and that any safeguards measures taken by other states are well known and worked into the estimates of the probabilities of the model, hence into U itself.

Player T's believed probability of choosing HEU is assumed to depend on both variables

$$p = p(c_u, c_p) \quad (4.4)$$

in a way so that increasing c_u decreases p and increasing c_p increases p :

$$\partial_u p \leq 0 \quad \text{and} \quad \partial_p p \geq 0 \quad (4.5)$$

where I have introduced the shorthand notation ∂_u and ∂_p meaning the partial derivative with respect to c_u and c_p .

Notably, there is no option left for player T to abandon her nuclear ambitions altogether, hence the total level of deterrence, C , cannot change her mind, however high. This very notion has been used to define fanaticism in a game theoretical setting: a fanatical player is one who gains a net benefit even if the mission fails¹⁶. Much as this is a simplification, I will simply assume for now that T is bent on nuclear weapons however low her chances of success and return to the nuances of this question in chapters 5 and 6.

15 The US Department of Energy spends in the range of \$1.5 billion on domestic safeguards yearly. Presumably the majority of this sum is spent on sustaining *status quo*. Matthew Bunn *Securing the Bomb 2007*, report of the Project Managing the Atom (Belfer Center for Science and International Affairs, Harvard University, 2007)* p.58

16 Sandler and Arce M. 'Terrorism & Game Theory' p. 321

My final assumption of dependences on variables is that q_u , q_p , T_u and T_p may not be influenced at all by player A. First of all this simplifies calculations considerably. Secondly, the feasibility of a 'second line of defence' is not on trial in this chapter¹⁷, so assuming it true for simplicity is permitted to the extent that it does not affect other conclusions of the chapter. Notably, I *do not* assume that intelligence and other second line measures do not play *any* rôle (their expected success is reflected in q_u and q_p), merely that their chances of success are out of player A's hands, on a different budget so to speak. Also, measures by player A are assumed not to affect the performance of the bombs produced. I furthermore assume that q_u , q_p , p_u and p_p are all statistically independent.

4.3.2 The utility function U and its derivatives

I shall play the game as seen from the perspective of player A, locating the ideal values of some pair of variables that will maximise his expected payoff. The expected payoff or *utility function* of player A, denoted U , is the sum of the payoff of each outcome of the game multiplied by the probability of that outcome. With the assumption of statistical independence I readily find this value equal to

$$U = -C - pp_u q_u T_u - (1-p)p_p q_p T_p \quad (4.6)$$

and its partial derivatives are

$$\begin{aligned} \partial_u U = & -1 + (-\partial_u p)(p_u q_u T_u - p_p q_p T_p) \\ & + p \left(-\frac{dp_u}{dc_u} \right) q_u T_u \end{aligned} \quad (4.7)$$

$$\begin{aligned} \partial_p U = & -1 - (\partial_p p)(p_u q_u T_u - p_p q_p T_p) \\ & + (1-p) \left(-\frac{dp_p}{dc_p} \right) q_p T_p \end{aligned} \quad (4.8)$$

A notation is introduced to keep track of the sign of the various terms: Where a minus is placed inside the differentiation parentheses, such as $(-\partial_u p)$, this indicates that the derivative itself is non-positive so the entire parenthesised expression is non-negative.

17 See appendix D for more on 'second layers', in which a similar game is used in a brief analysis of this assumption.

4.3.3 An important quantity: the asymmetry cost

Prior to further investigation I argue that based on the preceding qualitative discussion, the factor $(p_u q_u T_u - p_p q_p T_p)$ may safely be assumed to be positive: Based on the discussion from previous sections it should be obvious that while p_u and p_p are roughly in the same order of magnitude, q_u per expectation exceeds q_p , probably by a margin which is not small¹⁸. Also, because I deem an untested implosion device more likely to produce a fizzle than a gun design, T_u should be larger than T_p per expectation, hence $p_u q_u T_u$ should exceed $p_p q_p T_p$. The factor is a recurring one, so I define the shorthand notation $\Delta T \equiv p_u q_u T_u - p_p q_p T_p$ and assume $\Delta T > 0$.

Theoretically, whenever $\Delta T \neq 0$, it signals an asymmetry of the safeguards situation in that the two fissile materials are not equivalent. $\Delta T > 0$ means the threat from HEU is in a sense intrinsically¹⁹ greater than that from plutonium whereas $\Delta T < 0$ would have signalled the opposite. The conclusion $\Delta T > 0$ is therefore a mathematical expression of the conclusions of the preceding chapter, and I will refer to it as the asymmetry cost.

If one agrees that $T_u > T_p$ and $q_u > q_p$ independently of expenditures, the only region of the c_u, c_p plane where ΔT could possibly be negative is a region where $p_u \ll p_p$, presumably located in a region of very large values of c_u and comparatively small c_p . An idea of how U might look like as a function of c_u and c_p is found in figure 4.5; the area where ΔT is negative is found in a slim band of values towards the bottom right of such a figure where $c_u \gg c_p$. Our estimates in the case of the United States later in the chapter indicate that its current situation is far from this region.

The reader should note that the above argument that the asymmetry cost ΔT is positive is the only plausibility argument in the symbolic analysis, and will turn out to be of crucial importance for the conclusions we are able to draw at the end of the day.

18 Expressions such as 'not close to' and 'much larger than' are used in our discussion. While such terms are not rigorously defined and must be understood in the context they are put, their meaning should be unproblematic. As a rule of thumb, 'much larger than' might mean the difference in magnitude is at least one order, yet again this depends on the expression in question. For example 0.5 might typically be considered 'larger than and not close to' 0.2, and 'much larger than' 0.01.

19 Here this means independently of the choices and preferences of the terrorist in question.

4.3.4 The slope of the utility function with respect to plutonium cost

The slope of the utility function along the c_p axis is positive if $\partial_p U > 0$ and negative if $\partial_p U < 0$. In the former case it is opportune to spend more on plutonium safeguards, since the benefit is greater than the cost, in the latter case the security obtained in return for the expenditure is not worth the price. We find that measures to improve plutonium safeguards (remember we assume *status quo* is automatically retained) should not be funded further if

$$(1-p)q_p \left(-\frac{dp_p}{dc_p} \right) < \frac{1+(\partial_p p)\Delta T}{T_p}, \quad (4.9)$$

given some value of c_u as yet unspecified. Whether it holds or not can only be determined by a numerical estimate which I will undertake shortly.

It is not so hard to get a qualitative picture of just what the above inequality says. On the left hand side is found the *product* of arguably the three most important quantities for a policy maker to estimate when deciding on the amount of resources to put into the plutonium branch:

- q_p : The likelihood that, given plutonium in some form, the terrorist will be able to build some device and use it. Conventional wisdom would indicate that this is a small number compared to unity.
- $1-p$: The perceived probability that a terrorist will opt for plutonium over HEU. Putting together notions that terrorists would not be interested in a fizzling weapon and a low estimated value of q_p , this number should also be fairly small compared to unity even after the investment C^{20} . We argued in chapter 3 that $1-p$ is at least smaller than 50%.
- dp_p/dc_p : This number represents player A's power to do anything about the threat from plutonium terrorism. More precisely it is the rate at which player A, by pouring resources into the Pu branch, is able to decrease the chances of terrorist plutonium acquisition; a measure of 'value for money' of plutonium safeguards. A recent evaluation of the progress of US safeguards and materials elimination projects in Russia suggests that this value (be it the Pu or HEU branch) may

20 Discussed extensively in the next chapter.

be small (compared, for example, with $1/C$), concluding that '...for most of these [safeguards and elimination] programs, progress is constrained more by limited cooperation with foreign partners and bureaucratic impediments than it is by lack of funds.'²¹ A higher value of c_p , thus, may not make a great difference to p_p .

The right hand side of (4.9) is at first a little harder to relate to real life quantities, but one will notice that since $\Delta T > 0$, the numerator of the fraction must be larger than 1 and the entire fraction hence larger than T_p^{-1} . In other words:

If Player A deems that subsequent to expenditures c_u and c_p ,

$$(1-p)q_p \left(-\frac{dp_p}{dc_p} \right) < \frac{1}{T_p} \quad (4.10)$$

then the planned anti-plutonium terrorism improvement measures are *overfunded*²².

If player A finds that (4.10) holds true before the sum c_p is spent, the ideal value is $c_p = 0$ unless better value for money can be achieved (i.e. a higher value of $-dp_p/dc_p$). If he finds that (4.10) holds true *after* spending c_p , he can conclude that he should have spent less.

If, coupled with the above reasons for believing the left hand side of (4.10) small, player A also believes an attempt at Pu-bomb terrorism will result in a dispersion weapon rather than a nuclear explosion, T_p should be scaled considerably smaller than T_u and the inverse cost on the right of (4.10) might not be so small. This should further increase the plausibility of the truth of the inequality, yet it should be noted that the more player A believes an implosion bomb to fizzle, the easier it will become to build a weapon according to A's expectations, i.e. a larger value of q_p . Thus a shift of beliefs in this direction will increase the value of T_p^{-1} , but in (4.10) this should be partly but not entirely cancelled by a corresponding increase in the value of q_p (*cetera paribus*).

It should be emphasised already that even if (4.10)

holds true at present this needs not stay true forever. When HEU safeguards have been sufficiently strengthened it could well be opportune to turn to plutonium once more. This will be further explained towards the end of our analysis.

The expression (4.10), used as final result for its simplicity, is a sufficient but not a necessary criterion for the conclusion. The weaker, *necessary*, demand is (4.9). Note that if $(\partial_p p)\Delta T$ is in the order of magnitude of unity or more, the stricter demand (4.10) may be *much* stronger than necessary to conclude the overfunding of plutonium, as planned or at present.

A striking feature of (4.9) and (4.10) is the absence of the absolute quantities c_u and c_p themselves. The conclusion in other words, is that it does not matter exactly how much money is spent (or at some point in time has already been spent), only the extent to which spending *more* will *improve* the situation. While sunken cost is clearly a fallacy, player A can do nothing more about it than add it to his book of lessons learned. This disables arguments of the sort 'we have spent so much on this project already, so it would be a waste if we didn't complete it'. In other words, whatever is already spent is important only in that it has brought us to status quo; independently of how the present state of affairs was reached, the question is how best to improve from here.

4.3.5 The slope of the utility function with respect to HEU cost

In exactly the same way we use (4.7) to determine the slope of U with respect to c_u . This time, as will be clear, it is more natural to formulate an underfunding criterion, i.e. an inequality whose truth implies that $\partial_u U > 0$. We find that more should be spent on HEU safeguards improvement if

$$pq_u \left(-\frac{dp_u}{dc_u} \right) > \frac{1 - (-\partial_u p)\Delta T}{T_u} \quad (4.11)$$

Recognising that the right hand side of (4.11) is smaller than T_u^{-1} , I get a more elegant but stronger criterion similar to (4.10) obtained previously:

If Player A deems that subsequent to expenditures c_u and c_p ,

$$p \left(-\frac{dp_u}{dc_u} \right) q_u > \frac{1}{T_u} \quad (4.12)$$

then HEU efforts are *underfunded* at the

21 Matthew Bunn and Anthony Wier *Securing the Bomb 2006*, report of the Project Managing the Atom (Belfer Center for Science and International Affairs, Harvard University, 2006)*, p. viii.

22 Note already (this is elaborated further later) that this is with respect solely to nuclear terrorism. There could be other good reasons why plutonium should be kept safe.

current total deterrence level.

By comparing (4.11) to (4.9) one may make an interesting observation: while the right hand side of (4.9) is greater than $1/T_p$, the right hand side of (4.11) is *smaller* than $1/T_u$. This is why the natural criterion for HEU is for underfunding while that for Pu was for overfunding, a property which may be traced back to my assertion that ΔT is presently positive²³. This property makes for an important inclination towards prioritising HEU over Pu: The reader will be able to verify that if (4.12) were found to be true, this makes a much more robust claim about HEU than (4.10) does about plutonium if it is likewise true because unless $(-\partial_{ip})\Delta T$ and $(-\partial_{ip})\Delta T$ are much smaller than unity, (4.12) is much stronger compared to (4.11) than (4.10) is compared with (4.9). Indeed if $(-\partial_{ip})\Delta T > 1$, HEU efforts are *sure* to be underfunded according to the model, since the right hand side of (4.11) is then negative while the left hand side is positive²⁴!

4.4 Numerical estimation of parameters

The two inequalities (4.12) and (4.10) form a powerful toolkit for a government worried about nuclear terror attacks against it to divide its resources between the two branches. The potential power of this result, however, may only be drawn upon if the government is able to estimate with some confidence numerical values for the parameters and variables that form these conditions. Let me therefore briefly discuss the various quantities involved so as to tie the so far algebraic analysis to the present day political picture. It turns out that a very rough estimate allows me to draw some conclusions about the US safeguards programme in the former Soviet Union, demonstrating the power of the tool developed.

It is reasonable to assume that a government such as the United States or Britain will have access to considerable amounts of data exempt from the public domain and employ experts able to estimate the numbers in question with much greater authority than I may using data from open sources only. The inequalities (4.12) and (4.10) may thus be regarded as tools for analysis and synthesis of intelligence.

However extensive the amount of intelligence data

23 In turn a consequence of the relatively greater belief in a terrorist HEU project to succeed.

24 Note that a big increase in HEU expenditures would eventually bring $(-\partial_{ip})\Delta T$ below unity again by decreasing p and p_u .

available however, there will always be a considerable element of judgement and 'gut feeling' involved in estimating parameters. Numbers like q_u and q_p , for example may not be determined from previous cases, since there are none. Hence uncertainties will inevitably be large, and while some of the parameters could with sufficient effort be estimated quite accurately, I deemed there would not be much to gain from this. The numerical estimation is therefore somewhat rough. Surprisingly, some fairly robust conclusions and interesting analysis may nonetheless be drawn from this analysis.

4.4.1 The cost of a nuclear terrorist attack

The damage inflicted by an attack depends on the choice of target and means of delivery; it is not obvious that even a successful attack using a nuclear device is devised to maximise carnage. Given the large number of deaths on Manhattan, it is easy to forget that two of the four aeroplanes hi-jacked on September 11 2001 were destined for targets that could *not* have produced similar body counts (the Pentagon and probably Capitol Hill), but were powerful symbols. Hence the estimate for a value of the T 's due to Bunn and companions discussed below, assuming the weapon be detonated at Grand Central, New York at a busy hour, although 'conservative' given this assumption²⁵, may not represent the real choice of target, hence might be an overestimate. Yet how to estimate such numbers, at the end of the day, comes down judgement.

The value of human life is not easily be measured in dollars, and yet, when estimating the relative costs of different evils, be it earthquakes, wars, traffic accidents or cigarette smoking, this is what a government must do. Resources are always limited and avoiding death tolls from dramatic events altogether is wishful thinking. Governments must be concerned with the continued existence of the nation and choose the lesser of evils where necessary. In the following I will specify how I proceed to put numbers to lost lives and property.

We will very roughly estimate T_u and T_p using numbers from the September 11 attacks for comparison. Bunn and co-workers provide an

25 Matthew Bunn, Anthony Wier and John P. Holdren *Controlling Nuclear Warheads and Materials* report of the Project Managing the Atom (Belfer Center for Science and International Affairs, Harvard University, 2003)* pp. 15-19.

estimate for this value which in the author's view is a sound one²⁵. The scenario treated is one in which a bomb with an explosive yield of 10kT is detonated at Grand Central train station, New York on a normal working day. This is exactly the scenario that was allegedly reported to US intelligence in October 2001 and only determined to be a false warning weeks later²⁶. The estimate presented is half a million people dead and a direct cost to the United States alone amounting conservatively to at least one trillion dollars. Upon redoing the calculations somewhat less conservatively (details in the following) I find a more realistic estimate of approximately \$3 trillion. For perspective, this is approximately the value of the entire US federal budget for fiscal year 2008, of \$2.9 trillion²⁷. Including the many indirect expenses from after-effects on economy, the total cost would 'inevitably be several times the direct cost', Bunn and co-workers claim²⁵ although their notion might be too strong since some long term effects are accounted for.

Table 4.1: Estimated economic impact of terrorist attacks.

	September 11	10kT bomb	1T fizzle
Lost physical capital	22	300	15
Lost human capital	9	1500	0.7
Lost GNP	50*	500	50
Medical treatment	8*	100	8
Economic revitalisation	5*	5	5
Cleaning & Decontamination	0.6	20	5
Long term lost GCP	58	580	80
Total:	153*	3005	164

Numbers in \$billions, approximated to nearest integer value. Lost Gross City Product (GCP) calculated over first 3 years after attack. Lost GNP does not include lost GCP. *Deviates from Comptroller estimate, commented below.

The scenario described is something of a 'worst case': the bomb has a yield similar to the Hiroshima bomb and the target is chosen to have the gravest effect on the targeted state. A more detailed analysis was performed at the RAND think-tank of a scenario

in which the port of Long Beach, Los Angeles was the target of a 10kT bomb as it arrived in the harbour by container ship²⁸. With the great uncertainties involved, which scenario is used matters little and the RAND numbers coincide fairly well with mine²⁹; the report is nonetheless an excellent introduction to the far-reaching consequences of such an attack.

Let me also estimate very roughly the effects of a fizzle, also in central New York. I will choose a yield at the lower end of the scale of nuclear yields, at 1 tonne of TNT³⁰. To help me I have the calculations of Bunn and collaborators as well as the estimates of the fiscal impact of the 2001 attacks on Lower Manhattan by the New York City (NYC) Comptroller³¹. The Comptroller estimates the direct cost to NYC from the attacks to amount to between \$82 and \$94 billion and upon adding the impact beyond NYC and after-effects, I argue that the likely cost of a fizzle would be in the same order of magnitude. Costs of the three different scenarios are categorised in table 4.1³².

The area affected by a 10kT blast at Grand Central station delivered on the ground is visualised in figure 4.4 based on the online 'nuclear weapon effects calculator' provided by the Federation of American Scientists³³. If the terrorists were to deliver the weapon by air, for example in a suicidal detonation on board a hired aeroplane, the affected area would be considerably larger.

28 Charles Meade and Roger C. Molander *Considering the Effects of a Catastrophic Terrorist Attack* Technical Report (RAND, 2006)*

29 *ibid.* p.7. The numbers, totalling \$1 trillion, do not include long-term effects.

30 It is highly improbable that a design conceptually able to produce several kilotons could possibly produce a yield this low (Peter Zimmerman, private communication). Our estimates include, however, designs which are all but guaranteed to fizzle, such as the plutonium gun and other more creative designs which the terrorist might think of.

31 William C. Thompson, Jr., Comptroller New York City *One Year Later: The Fiscal Impact of 9/11 on New York City* (New York: Office of the Comptroller, 2002)*

32 Numbers from Thompson p. 1 and Bunn *et al. Controlling Nuclear Warheads and Materials* p. 18, though the latter holds only tentative numbers. No attempt has been made to compensate for changing dollar value since 2002.

33 Lucas Royland 'Nuclear Weapon Effects Calculator', online from the Federation of American Scientists <http://www.fas.org/main/content.jsp?formAction=297&contentId=367>. Based on Samuel Glasstone and Philip J. Dolan (eds) *The Effects of Nuclear Weapons* 3rd edition (Washington D.C.:GPO, 1977)*. The high quality photo is from Google Earth[®]. The Effects Calculator does not offer a photo of New York; the circles were fitted by comparing a Google Earth[®] photo of

26 Massimo Calabresi and Romesh Ratnesar 'Can We Stop the Next Attack?' *TIME Magazine* (March 11 2002)*

27 Office of Management and Budget, The White House *Budget of the United States Government FY2008* (Washington D.C.: The White House, 2007)*

Physical capital

The physical damage of the September 11 attacks were estimated to \$22 billion by the Comptroller (the historical and cultural value of a landmark such as the World Trade Center can, of course, not be measured in dollars). For the 1T fizzle I put this at somewhat less. The blast itself will probably cause much less damage than that of September 11, being similar in size to the bomb that destroyed the Murrah building in Oklahoma in 1995, but much additional damage will be caused by buildings which must be demolished due to high contamination levels. The physical damage of the 10kT blast is set, probably rather conservatively, to some 13 times that of September 11th 2001, based on figure 4.4.

Lost human capital

Lost human capital for the 10kT bomb is found by extrapolating the Comptroller's number for the approximately 3000 dead on Manhattan to 500,000 dead. The Comptroller's number, in turn, is based on the value of the workforce as indicated by the sum of salaries of those dead. The number must be lowered somewhat, since the area affected would include areas where average income is lower than on lower Manhattan. On the other hand, there would be a large number of wounded and disabled in addition to casualties, probably more than making up for this in terms of lost working force.

The number of dead in the 1T fizzle attack is likely to be much lower than that of 2001: the blast will presumably be similar in size to that which destroyed the Oklahoma City Federal building, an incident which took 168 lives and injured 800. Using these numbers to extrapolate the September 11 figures, I end up at some \$0.7 billion (some account is taken for disability following injury).

Medical treatment, economic revitalisation and cleaning

Costs of medical treatment for the attack is found on the basis of the governmental compensation and

Washington D.C. at exactly the scale used in figure 4.4 to the photo of D.C. in the effects calculator. The circles were then transferred to the photo of New York, centered at Grand Central Station. The uncertainties following this procedure are probably much smaller than those of the calculation itself, which does not take into account the screening by geography and buildings. (Google allows use of data from Google Earth for all non-commercial educational purposes; please adhere to Google's conditions if reprinting figure 4.4).

benefit payments following September 11, provided by a comprehensive RAND report³⁴. These totalled \$8 billion following the September 11 attacks. For the 1T fizzle the same number is used: While the number of 800 injured is not large, the number of wounded from the September 11 attacks was small compared to casualties. A nuclear attack, however, small, will probably cause widespread fear and anxiety, and calls on medical services to test for possible effects of radiation will possibly be out of proportion with the actual health risk posed by the fallout.



Illustration 4.4.: Damage area of a 10kT bomb at Grand Central Station New York.

Bomb assumed delivered by car, based on the Nuclear Weapon Effects Calculator by Lucas Royland³³. Blue circle: homes completely destroyed and stronger commercial buildings severely damaged due to high pressure blast wave. Red circle: Widespread fires due to intense heat. Yellow circle: Moderate damage to buildings, risk due to flying debris caused by blast wave.

Cleaning costs are extremely hard to estimate, and the numbers must be taken as utterly tentative (no numerical estimates are given by Bunn *et al.*). Economic revitalisation to individuals and businesses following September 11 covered by the government amounted to around \$5 billion³⁴. This may be regarded as something of a voluntary undertaking by the government, and while the need for measures to restore normality will be huge following a true nuclear attack, the calls on the government's purse will be so enormous in such an event that such efforts likely cannot be given funding proportional to the damage done. Expenditures for revitalization will aim to mitigate the long term economic effects,

34 Lloyd Dixon and Rachel Kaganoff Stern *Compensation for Losses from the 9/11 attacks* (RAND, 2004)*. Tables on p.132 and 135

represented by the lost GNP, hence could, if implemented well, pay off in the long run.

Lost Gross National and City products

Whilst after September 11 airports were closed, cargo freight over sea and land continued; after a nuclear attack, all ports will probably be closed for a while after, having enormous effect on the economy. The Comptroller's estimate, at that, includes only NYC, whilst the rest of the nation suffered indirect costs as well. The \$500 billions tentatively include lost GNP, including indirect costs. According to Bunn and associates, this is conservative, yet their claim that the direct cost (which already includes almost all of the long term effects, save behavioural shift in the rest of the country) should be multiplied several times to obtain the indirect might seem an overcompensation.

The Comptroller does not include the indirect costs of lost national product beyond that of NYC, yet only a very tentative assessment is needed in this already very rough calculation. Adding another \$50 billion to the above sum should cover national losses in airline traffic and other effects³⁵, in which case I end at a very rough total cost of \$140 billion for the September 11 attacks.

A further \$580 billion is estimated to be lost from Gross City Product (GCP) over the following years. The Comptroller estimate for September 11 of \$58 billion is used for the 1T fizzle as well, whereas an ostensibly very high \$580 billion is used for the 10kT bomb. There are several reasons why this should be a high sum in my opinion. The majority of the cost of property damage after September 11 was covered by insurance³⁴, but following such a major catastrophe as the 10kT bomb, there will be an acute shortage of insurance money³⁶. With possible widespread bankruptcy among insurance companies reconstruction is likely to be delayed, and the economic consequences could persist for much longer than the three years by which New York City was back on track after the devastating attacks.

The loss in GCP for the 1T fizzle is somewhat higher than for September 11, since a larger area must likely be evacuated for clean-up and for a longer time. However, since the evacuated workers will not

generally be injured, it is likely that solutions will be found to mitigate these effects by moving business to temporary locations. Long-term (after first year) effects are assumed to be like the effects suffered by New York from 2002 onwards so that the difference is due to the first year after the attack.

The total cost

Notably, the expenses from the U.S. 'Global War on Terror' on Afghanistan and Iraq³⁷ that followed the attack are not included in table 4.1, nor is this cost included in any of the estimates for antiterrorism expenses elsewhere. The costs of the military conflicts engaged under this banner will have cost the United States alone some \$737 billion by the end of 2008, surpassing the cost of the Vietnam War, measured in 2006 dollars³⁸. An assessment of the helpfulness of these wars in abolishing terrorism is beyond the scope of this thesis. Some perspective on the sums of money involved may be offered, however, by noting that this is almost 100 times what Congress has spent on non-proliferation programmes in the former Soviet Union since 1992³⁹.

Using the model from page 51 I put T_p somewhere between $T_f = \$164$ billion and $T_n = \$3005$ billion, and write $T_p = \rho T_n + (1-\rho)T_f$ where ρ is the probability of the bomb exploding with 10kT yield. Since improvised devices using HEU and Pu will have similar yields in the fizzle and non-fizzle cases, T_n and T_f are used for both, with a higher value for ρ in the HEU case. How to estimate ρ in each case is a question of judgement. I will choose figures which in my opinion (based on considerations in chapter 3) are certainly not underestimating the threat from plutonium terrorism and also not overestimating the HEU threat, since we will be using the numbers to evaluate (4.10) and (4.12). I put ρ at 20% for Pu and 70% for HEU, giving the estimates to one significant figure

$$T_p = \$7 \cdot 10^{11} \quad \text{and} \quad T_u = \$2 \cdot 10^{12} \quad . \quad (4.13)$$

35 The airline industry's losses 2001-2003 have been estimated to amount to some \$23.7 billion by the Air Transport Association 'Airlines Outline War Impact' News Release (ATA, 2003)*. It is unknown how authoritative this account is.

36 Meade and Molander *Considering the Effects of a Catastrophic Terrorist Attack* p. xvii

37 Noting that the connection between Iraq and the September 11 attacks is highly questionable in the least. The alleged connection nonetheless formed an important part of the rhetoric.

38 Stephen M. Kosiak, Center for Strategic and Budgetary Assessments 'The Global War on Terror: Costs, Cost Growth and Estimating Funding Requirements' Testimony before the Senate Budget Committee (June 6 2007)* p. 2.

39 US Government Accountability Office *Weapons of Mass Destruction: Nonproliferation Programs Need Better Integration* GAO-05-157 (2005)* p. 5

The difference between the two is approximately a factor 3. Note that both of these estimates are lower than an alternative estimate of Bunn's of \$4 trillion⁴⁰.

4.4.2 Rough estimate of value for money in US safeguards programs

The other quantities in equations (4.10) and (4.12) are even harder to establish than the costs. Especially probabilities may be extremely hard to establish with confidence, even for an intelligence organisation like CIA or SIS, with their access to classified information.

Regarding the quantities $(-dp_u/dc_u)$ and $(-dp_p/dc_p)$. For very tentative numbers, consider US grants to safeguards efforts, primarily in former Soviet states, in fiscal year 2005⁴¹. During this year, the United States spent some \$1.5 billion on safeguards and security in DoE's own installations⁴² and a further \$1.08 billion abroad⁴³, chiefly in Russia. For simplicity I assume that the efforts on US ground were just adequate to maintain *status quo*, hence I do not count it as part of C. The efforts in Russia, however, surely had a much larger impact. During this period, a further 3% of fissile material in Russia was upgraded to a 'comprehensive' level of safeguards, whereas a further 3% received 'rapid' safeguards arrangements⁴⁴.

Such figures, however, paint an overly simplistic picture. Equipment and alarm systems are no good if personnel can be paid to bypass them, and the corruption in Russia is cause for much concern and a prerequisite assumption that Russia is able to maintain *status quo* by themselves might be flawed⁴⁵ (indeed, Bunn and Wier warn against the use of such numbers as a direct measure of progress⁴⁶). Furthermore, safeguards are only as good as the weakest link, provided the terrorist knows where the weakest link is to be found. By the end of fiscal year 2006 an estimated 70 buildings in Russia have yet to receive rapid upgrades⁴⁷ and with sufficient reconnaissance a terrorist proliferator could discover

which buildings these are and target these rather than those which are more secure. As an upper bound for these slopes, 3%/\$1 billion might be reasonable, yet the indication is that this is an overestimate.

The difference between the slopes for uranium and plutonium may also be perceivable; the US-Russia Plutonium Management and Disposition Agreement for example, has yet to be implemented⁴⁸, and once commenced can make a real and relatively cheap difference. A similar effort, the blending down of HEU, is under way and its achievements are notable⁴⁹. Global stockpiles still amount to thousands of tonnes, however, whilst one bomb requires merely tens of kilograms. The above estimates being so rough, thus, I approximate the rate of change of p_u and p_p with c_u and c_p to be equal.

Whilst $(-p_i) = 3\%/\$1 \text{ billion}$ ($i = u, p$) is a high estimate, let us keep it for further calculations. Importantly, however, Bunn and Wier conclude that this ratio could be much higher with improved cooperation.

With the above estimated numbers inserted, the 'overfunding criterion' (4.10) reads approximately

$$(1-p)q_p < 0.05 . \quad (4.14)$$

As the left hand is a probability squared, 0.05 is not particularly small and would be obtained, for example, if $(1-p)$ and q_p were both $\approx 22\%$. On the basis of this very rough estimate it is difficult to conclude whether this is true or false at present, merely that it may not be dismissed on grounds of being unlikely; in truth there is serious reason to believe the inequality holds true. If instead $(-p_i) = 1\%/\$1 \text{ billion}$, an equally likely estimate, the corresponding $(1-p)$ and q_p must each be 39% for the inequality to be false, a *very* generous estimate indeed for their true values! The indication is thus that plutonium efforts could be overfunded at present, but the numbers are too rough for a definite conclusion.

Equation (4.12) is estimated the same way. I use the same rate of change of acquisition probability $(-dp_u/dc_u) = 3\%/\$1B$, as I used for p_p with c_p above and T_u from (4.13) to find the sufficient criterion for HEU efforts to be underfunded as

$$pq_u > 0.02 . \quad (4.15)$$

If one assumes, as argued above as well as in the next chapter, that p is at least greater than $1/2$, an

40 Matthew Bunn 'A Mathematical Model of the Risk of Nuclear Terrorism' *The Annals of the AAPSS* 607 (2006) p.106.

41 I use the US as example here as elsewhere.

42 Bunn and Wier *Securing the Bomb* 2006 p.47

43 *ibid.* p. viii

44 *ibid.* p. 54

45 e.g. Anna M. Pluta and Peter D. Zimmerman 'Nuclear Terrorism: A Disheartening Dissent' *Survival* 48:2 (2006) p. 56-60.

46 Bunn and Wier *Securing the Bomb* 2006 p. 15

47 Bunn *Securing the Bomb* 2007 pp. 64-66

48 *ibid.* p.99-101

49 *ibid.*

unreasonably small value of q_u is required for this to be false. Using, say, $p = 0.6$, q_u would need to be no more than about 3%, and probably even less if the less strict criterion were applied. Even using a smaller rate of change, like $(-dp_u/dc_u) = 1\%/\$1 \text{ billion}$, a value of q_u of 5-10% is still required. Notably, Bunn judges q to be 28% not specifying whether Pu or HEU is used⁵⁰. Assuming the probability that the material is HEU is $p < 1$ the quantity q_u will be somewhat larger than this since a nonzero probability of Pu will tend to decrease the chances of a successful weaponisation of the fissile material.

While my numerical calculations indicate that plutonium efforts may well be overfunded, it seems certain from the above calculations that HEU efforts are underfunded, implying an ideal value of c_u significantly larger than c_p ⁵¹. It is important to note that the value for money numbers used are specific to US efforts, thus the result strictly speaking applies only to the US programmes. It is reasonable to assume, for example, that Russia and other countries of proliferation concern could achieve far more per dollar domestically.

Our 'local' method of measuring slopes must be thought of carefully. Consider safeguards efforts in Russia for example. At the start of the Materials Protection, Control & Accounting (MPC&A) programme⁵², for example, there was a large number of facilities to secure, and by the 'weakest link' reasoning above, each new site which was secured increased the overall security only very little. Now that the number of sites left to secure is relatively small, each site will close a relatively large gap compared to the number of buildings left to secure and the instantaneous value for money is much better. Still the increased slope would not have been possible had the period of little increase in security not been undertaken. At the outset, thus, a calculation like that above could have shown that HEU efforts were overfunded, but seen as a whole, the MPC&A project has clearly been worthwhile.

This problem is not a fundamental weakness of the model, but of the way the derivatives are calculated. When interpreted as *mean* derivatives of an entire project like MPC&A and its many sub-projects,

inequalities (4.10) and (4.12) give good information whether upgrade programmes give good enough value for money. While 3%/\$1 billion might be a gross overestimate of the immediate security payback from money spent, when comparing the cost of the entire programme to the security gain the programme is expected to yield, such a figure might still be reasonable. Since 1993 through fiscal year 2008 the United States has spent/budgeted around \$11 billion dollars to safeguards abroad, primarily in the Soviet Union⁵³. An average value of 3% per billion dollars and a rough assumption that half of the funding was spent on HEU measures would indicate that p_u has been reduced by 66% over this period, which is a high number, but maybe not unreasonably so.

Further precautions must be taken as well, for the numbers employed are tentative and the model very simple. A fair amount of judgement must be applied when estimating the numbers, and other authors would doubtlessly have chosen different figures. Nonetheless, the above values could be changed considerably without changing these tentative conclusions, and so this author finds them convincing.

4.4.3 The effect of changing stockpiles

As also discussed later, when making such estimates as I have above one should notice that there could be external reasons why parameter values can change with time due to other reasons than safeguards improvement. I already discussed the fact that a significant budget post will be to maintain protection at status quo, which we assume is not included in the model. Of real-world trends which will be of significance, however, the most notable is perhaps the different rate of growth or reduction of the global stockpiles HEU and separated plutonium.

While a number of measures have been under way for years to reduce the global stockpiles of HEU⁵⁴, and the production of HEU has virtually stopped in most countries⁵⁵, neither is true of plutonium. Efforts to dispose of excess Russian and US plutonium have made no progress⁵⁶ and civilian separation of plutonium continues in several countries, at a total rate of approximately 10 metric tonnes a year⁵⁷. If this

50 Bunn 'A Mathematical Model ...' p.106. 28% is the product of figures given for successful assembly and delivery respectively.

51 If Pu efforts are indeed overfunded at present, of course, it implies an ideal value of c_p to be zero.

52 A programme by the US Department of Energy to safeguard fissile materials primarily in the former Soviet Union.

53 Based on Figure A-1 of Bunn *Securing the bomb 2007* p.154.

54 International Panel on Fissile Materials *Global Fissile Material Report 2007* (IPFM, 2007)* p.25.

55 *ibid.* p.11.

56 *ibid.* p.41.

57 *ibid.* p. 7.

trend continues, it seems likely that the day will come when illicitly obtaining plutonium is significantly easier than acquisition of the terrorist's material of choice, HEU.

The model introduced in this chapter can account for this change to a large extent, when realising that its intention is to give advice to a government at a particular point in time. Assuming we find ourselves in a future situation where plutonium is much more easily available than HEU, however, the above estimates and plausibility arguments for parameter values will not all be valid any more. Inequalities (4.9) and (4.11) will be, since they derive directly from the model, which does not in itself assume anything about the values of the parameters.

For equalities (4.10) and (4.12) to be stricter versions of these, however, one must require that the asymmetry cost ΔT is positive, and this key conclusion, which I argued is all but certain to hold true today, could change with increasing Pu availability. The premise for the guaranteed positiveness of ΔT was that a HEU project has a greater probability of success than one using plutonium, whereas the availability of the two materials is approximately the same. Clearly the former statement will not change with changing stockpiles, but the latter will, and could eventually balance out the asymmetry or even reverse it. In fact, thinking about such a scenario demonstrates the power of devising a general model because while our conclusion of HEU overfunding and possible Pu underfunding could change with changing conditions, the inequalities (4.9) and (4.11) will be equally valid.

On the other hand, if the goal were to analyse the long term effect of changing stockpiles, rather than be a policy making tool at a particular point in time, a different model which explicitly incorporates how the environment changes with time could no doubt shed further light on the matter. Such a model, if it were to take account of the terrorist's decision as well, would likely have to be more complex and make more assumptions, since the timeline of varying stockpiles, which is comparatively well known and documented, must run alongside the timeline of an unfolding terrorist plot which is not known. Ways could probably be found to salvage this now familiar asymmetry of available information. A reasonable approximation could probably let probabilities associated with clandestine terrorist activity to be measured per time unit and incorporate the changes

in stockpile and stockpile protection in a more detailed and time-dependent way. Indeed, this could be a valuable extension of the research reported herein for the future.

On a rough level, however, the model in this chapter and its resulting inequalities can give a good idea about how prioritisation of safeguards would be affected by changing stockpiles, simply by seeing what happens to the inequalities (4.9) and (4.11) upon varying the relevant parameters: the probability of the terrorist choosing Pu over HEU, and the 'value for money' along the plutonium branch. It will be a rough tool only, however, because the model is fundamentally instantaneous and cannot be used to account for changes over time in a fully consistent way.

4.5 A numerical example for further analysis

Given the surprising conclusions so far it would be interesting to get some visual impression of what the utility function (4.6) could reasonably look like in a plane where the axes are c_u and c_p . To do this, however, I must make rather more specific assumptions than what we have so far, and primarily for the sake of illustration let me assume the following form of p , p_u and p_p :

$$p = 0.9 - 0.45 \frac{c_u}{c_u + c_p}$$

$$p_u = 0.1 + \frac{0.5}{c_u/2 + 1}; \quad \text{and} \quad p_p = 0.1 + \frac{0.5}{c_p/2 + 1}.$$

Here, c_u and c_p are per \$billion. We use like before $q_u = 0.7$ and $q_p = 0.2$ and T_u, T_p as given by (4.13).

It is important to note that the above examples are for demonstration only, and the absolute numbers used have no significance in themselves. They fulfil the mathematical assumptions set out at the beginning of the chapter and use numbers in the order of magnitude of those estimated in the previous section. Regardless of whether the reader finds these particular functions satisfactory examples, the resulting graph gives important insights into how to interpret the analytical results of previous sections.




Illustration 4.5.: Contour plot of numerical example utility function.

The utility function which follows from these assumptions is shown as a contour plot in figure 4.5. The example yields the maximum (Nash equilibrium) choice⁵⁸ at around $c_u = \$30$ billion, and $c_p = \$10$ billion; a total cost $C = \$40$ billion. If player A finds the solution unreasonably expensive, he may try to employ diplomatic tools to improve cooperation with countries of proliferation concern in order to change the behaviour of p_u and p_p into falling off more steeply. This would move the maximum to a lower value of C . In the opposite case, if cooperation were to ground to a full stop, more money would make no difference, the slopes of p_u and p_p would fall to zero and the point optimum value moves to infinity. Although based on input data which are so uncertain they border on the arbitrary, the figure shows some interesting properties which we will use for the analysis in the following.

4.5.1 Could Pu measures possibly be overfunded?

While the conclusion that HEU measures should be boosted considerably is relatively uncontroversial and conforms with the assertions of several analysts as we have seen, the notion that plutonium measures could be overfunded will no doubt cause surprise. I noted that my numbers are not precise enough to come to a definite conclusion on the matter, but indicate the definite possibility that plutonium measures could be overfunded. How to interpret this result?

58 A Nash equilibrium is a strategy profile $\{s_i\}$ such that no player i can do better by deviating from their strategy s_i , given that all other players stick to $\{s_i\}$.

There is at least one intuitive reason to doubt whether the overfunding conclusion is correct, namely the 'weakest link' argument. If there are still unsecured buildings in the world holding stores of plutonium large enough to fuel nuclear weapons, how can Pu measures then possibly be underfunded before every building is secure?

The way to address this question is, I believe, twofold. It is important as pointed out above, to look beyond next year's budget when estimating the security reward which follows from a safeguards project. While securing one more building might not seem worthwhile, securing all the remaining buildings could still very well be. This could lead to an underestimate of the relevant slope, although as we have argued, it seems unlikely that I have been guilty of doing this in the present context.

If even after carefully considering the whole enterprise and its effect one finds that (4.10) is still satisfied, the conclusion is that the present level of safeguarding together with the perceived difficulty of building a plutonium implosion device together provide security enough. The probability of success of a terrorist plutonium plot, and hence the expected damage from plutonium terrorism, is sufficiently small that further spending to decrease it is not justified. This is an uncomfortable conclusion which must certainly be regarded in much greater detail than I have above to be taken at face value. The fact that better plutonium safety obviously is within reach makes the conclusion all the harder to swallow.

And yet, a conclusion of plutonium overfunding is perhaps not as problematic as it first appears, as I will explain. Regard figure 4.5 once more. Forget about the actual numbers on the axes and the sums of money spent so far (the numbers used to create the plot are so uncertain that these numbers are of little significance anyway). The question is then: where are we at present? In fact⁵⁹ if we assume that the actual expenditure of something like \$5 billion along both axes gives an idea of where we are, the situation in the figure matches our assertion. Thinking of the contour plot as a map, we find a steep uphill in the c_u direction as we should, and a slight downhill along the c_p axis. But notice now that \$5 billion might still be below the ideal value! In fact, starting from the origin $c_u = c_p = 0$ we see that the direction of steepest ascent is along the c_u axis for some time until around \$9 billions where it suddenly becomes opportune to spend money on

59 And not entirely accidentally, since the numbers used were chosen to match our previous estimates.

plutonium. Only very little at first, and then more until the summit is finally reached at something like $c_p = \$5-6$ billion.

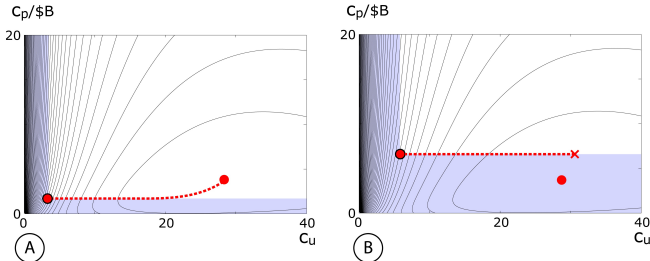


Illustration 4.6.: Zoom of figure 4.5 with two possible scenarios and best responses.

The blue areas are forbidden since sunken cost cannot be retrieved. The black-rimmed circle to the left in each image is present situation, the red circle to the right is optimum value. A: the plutonium expenditures are not yet past optimum value, but slope is initially negative in c_p direction. B: more than the optimum value is already expended. Best approach is to improve HEU measures until maximum is reached in the c_u direction.

The important implication is that while plutonium efforts seem to be overfunded at present, they need not stay so if sufficient funds are allocated to HEU. So long as the plutonium slope is estimated to be negative, funding should be directed at HEU efforts. If the maximum has already been passed as in figure 4.6B, one will find the best value one can reach when the slope along the HEU direction becomes zero. If the ideal plutonium value is *not* yet reached as in figure 4.6A, the plutonium slope will turn positive again at some point.

The policy recommendations that follow from the above analysis should therefore be to heavily prioritise HEU measures and secure sufficient funding for this branch. The above analysis is not a sufficiently steady fundament upon which to recommend the complete cessation of plutonium safeguards improvement measures, but plutonium projects must not be allowed to get in the way of the more important safeguards projects involving HEU, for example during budget negotiations. In case of a shortness of resources it seems justified to halt plutonium projects temporarily in order to ensure that improvement of HEU security continues.

4.5.2 Special case: Fixed and insufficient total spending level

Before concluding I mention briefly an interesting modification of the model in which the government at the start of a fiscal year has a fixed sum ΔC to spend

on safeguards measures that year (the period can of course be shorter or longer, the point being that the government can take one step towards improving the situation each period). This reflects the fact that in a typical democratic state governmental expenditure is bounded by a budget passed by a parliament each year.

Let me furthermore assume that the current situation is one in which one is quite far from the optimum value of $U(c_u, c_p)$ (compared to what can realistically be achieved in one period), and that the allocated ΔC is quite far from enough to bring the situation from *status quo* to the maximum value of U or the best obtainable value as in figure 4.6B. Such a model is attractive for modelling reality reasonably well, and because it is an immediate and moderate generalisation of what we have done so far.

Standard multivariable calculus says (see appendix B) that the *direction* in the c_u, c_p -plane in which the value of U rises *most steeply* is given by the vector⁶⁰ we define as

$$\vec{C} \equiv [c_u, c_p]_{\text{optimal}} = \nabla U, \quad (4.16)$$

where square brackets denote a set of vector components, the first along the c_u axis and the second along the c_p axis, and the vector operator $\vec{\nabla}$ ('nabla') is defined as

$$\nabla \equiv [\partial_u, \partial_p] = \hat{c}_u \partial_u + \hat{c}_p \partial_p$$

where a hat denotes a unit vector⁶¹. This readily gives us a simple formula for optimal spending if the incremental spending ΔC can only take us a bit of the way to the final optimum: a distribution between HEU and Pu branches so that the sum ΔC increases the expected utility maximally⁶². This is illustrated in figure 4.7. In fact there are more rigorous ways of solving this problem, but the formalism of steepest rise is intuitive and a good approximation.

In the case of the United States, there is some doubt, as we have seen, whether $\partial_p U$ is positive or not; if Pu safeguards and elimination measures are overfunded in the sense we explained above, $\partial_p U$ is negative and the vector prescribed by (4.16) points into the forbidden area. No vector component can be

60 In this context we can think of a vector as simply an arrow pointing in a direction on our 'map'.

61 A unit vector has length 1 and no dimension.

62 In general, this is only *exactly* true if ΔC is *infinitesimally* small (it is then denoted dC), but it is a good estimate of optimal division if U is well behaved if ΔC is sufficiently small.

allowed to be negative since sunken cost cannot be retrieved. In this case, since $\partial_u U$ is certainly positive, the optimal solution is to spend the entire sum ΔC on HEU safeguards, so that $\vec{C} = [\Delta C, 0]$.

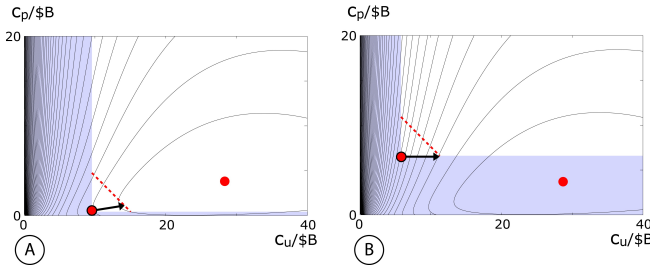


Illustration 4.7.: Contour plot of same utility function with incremental expenditure.

The incremental change ΔC can take the government player from the initial position (red circle with black rim) to any point on the red dotted line. The vectors point to the optimal division of ΔC in both cases.

If $\partial_p U$ is found to be positive, however, (4.16) yields the ideal distribution as⁶³

$$c_u = \frac{\Delta C \cdot \partial_u U}{\partial_u U + \partial_p U} \quad \text{and} \quad c_p = \frac{\Delta C \cdot \partial_p U}{\partial_u U + \partial_p U}. \quad (4.17)$$

One could also write out more specific expressions by inserting equations (4.7) and (4.8).

4.5.3 Fixed total spending level

We can push our analysis even a little bit further as an extension of the above. In this section I will have to assume that the reader has a basic command of differential algebra. Assume now that the total spending level C is fixed. As shown in figure 4.8 the possible values lie on a straight line in this case. We have chosen C somewhat below \$20 billion for illustration.

A more rigorous way of finding the best obtainable value is now to measure the slope as we move along this line either down (towards higher c_u) or up (towards lower c_u). Note that in the former case dc_u is positive, in the latter it is negative. While planned costs which are not yet expended can be increased or decreased at will, once more the available final value is limited by what is already spent, shown as usual as a forbidden blue area in figure 4.8.

Along the red dotted line, c_u and c_p are constrained by the condition $c_u + c_p = \text{constant}$, which implies

$$dc_u = -dc_p.$$

Moving along the red dotted line by varying c_u and c_p by dc_u and dc_p , the variation of the utility function is

$$dU = \partial_u U \cdot dc_u + \partial_p U \cdot dc_p = (\partial_u U - \partial_p U) dc_u = [(-\partial_u p) \Delta T + (\partial_u p) \Delta T + \Delta T'] dc_u \quad (4.18)$$

where I have used (4.7) and (4.8) and defined the quantity

$$\Delta T' \equiv p \left(-\frac{dp_u}{dc_u} \right) q_u T_u - (1-p) \left(-\frac{dp_p}{dc_p} \right) q_p T_p. \quad (4.19)$$

We could have substituted for c_u instead and obtained the same but with opposite sign.

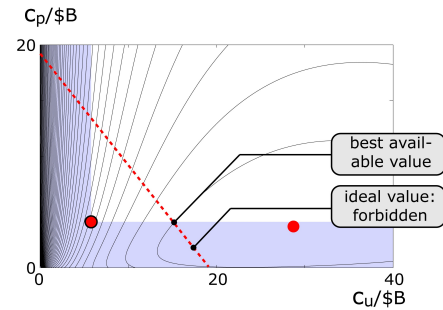


Illustration 4.8.: Contour plot of same utility function with fixed total cost.

The optimal point along the dotted line is where $dU = 0$, but this point could lie in the forbidden region as in figure 4.8. The alternative is to move along the line in a direction so that dU is positive until either a maximum or a forbidden boundary is reached. What is clear from (4.18) is that if the expression in the square brackets is positive, dc_u must also be positive for dU to be positive, which implies that when

$$(-\partial_u p) \Delta T + (\partial_u p) \Delta T + \Delta T' > 0$$

the optimal value lies at a higher value of c_u (correspondingly: a lower value of c_p). A sufficient criterion for this to hold true is that ΔT and $\Delta T'$ are both positive. Considering the definition of $\Delta T'$, it is clear that this quantity is positive roughly in the same region of the $c_u c_p$ plane where ΔT is positive: everywhere except a slim area in the lower right corner of the contour plots where $c_u > c_p$ by a large margin⁶⁴.

63 Note that ΔC is the sum of the lengths of each vector component, not the length of the vector.

64 The reasons for this are the same as argued in section 4.3.3 combined with the conclusion from chapter 3 that $p > 1-p$. If the rate of change of p_u and p_p with respect to c_u and c_p are thus

A new and very succinct criterion for priorities is thus found:

If per estimate ΔT and $\Delta T'$ are both positive after the entire sum of C is spent on HEU safeguards, this is the optimal division of C .

An important corollary is that

If ΔT and $\Delta T'$ are both currently positive, incremental spending should be allocated solely to HEU efforts until this is no longer the case.

There is reason to believe, as we have seen, that the latter situation holds true for the United States today.

While holding no new information compared to that found before, these criteria arguably sum up the essence of safeguards priorities in the face of the threat of nuclear terrorism and could be just as powerful and probably as applicable to the real political picture as the free variable criteria (4.10) and (4.12).

4.6 Assumptions and limitations of the model

Every model is a simplification. Indeed, this is the very strength of formal social science; it allows the analyst to cut through the jungle of complexity which is almost always present, and reach a conclusion which may in turn be tested empirically. Of course, with nuclear terrorism no such empirical data exists; I find myself, therefore, in the normative rather than descriptive domain, where rather than hypothesising about mathematical regularities of human interaction, the analysis aspires to create a rational ideal for future action.

I have gone from a qualitative discussion of the problem to devising a model and thence to using formal logic to show what the model implies and draw conclusions by attaching it once more to the real world by numerical estimates. What remains for my argument to be on as firm a ground as possible is to examine the assumptions upon which the model rests.

In doing so, I must consider what criterion should be applied to judge the appropriateness of a simplification. As we argued in chapter 2, the appropriate requirement of a simplification is that it

somehow aids the understanding of the problem considered. There are at least two somewhat related ways a simplifying assumption can achieve this: by explicitly ignoring complicating minor effects which are not essential, and by making way for simple and transparent analysis (mathematical or otherwise) with clear cut arguments.

The more unrealistic the simplification, the more reason there is to justify it. On the other hand, says Bennett⁶⁵ the more empirical data one has to compare with, the more complicated the model can be, since one has then better data by which to test it. In the strict sense, I have *no* empirical data on nuclear terrorism, hence a very simple model seems *a priori* like a reasonable first approach. A simple model, while a crude reproduction of reality, requires fewer explicit assumptions about how the world behaves, assumptions which in the absence of empirical results will be somewhat *ad hoc*.

The second criterion is an assay of how *fundamental* the approximations made are. By fundamental is meant, quite literally, whether the assumption forms a modelling fundament so that changing or replacing it will change qualitatively the behaviour of the model (or possibly make it collapse). Clearly, if our analysis were to fall apart at the alteration of an unrealistic assumption, it would be of questionable soundness. A problematic assumption is thus one which is both unrealistic and fundamental.

4.6.1 Time and time ordering

The model of figure 4.3 contains only a pseudo-time. Much as moves are ordered consecutively, one will notice that the model as presented does not necessitate any true passing of time - we may as well imagine that player A moves, player T moves an infinitesimal time later, and the game is over in no time whatsoever.

This is quite deliberate, motivated by the fact that one can only have a very vague idea of the real timing of the process such as portrayed. Each player knows their own timing, but that of the other player is unknown: Player A must defend against an attack that may already be on the way, in the planning stage, or which may never be attempted at all. It is deemed that this trait, as well as simplifying calculations, reflects reality well. Introducing an explicit time would increase the complexity of the model, but

in the same order of magnitude, $\Delta T'$ is sure to be positive except in a slim region $c_i \gg c_p$ where the latter two assertions no longer hold.

65 Peter G. Bennett 'Modelling Decisions in International Relations: Game Theory and Beyond' *Mershon International Studies Review* 39 (1995), p.39

might not do much in terms of improving the realism of the model, hence the chosen approach seems justified.

4.6.2 *The terrorist has yet to obtain the material*

Arguably it is a sensible stance for A to assume that player T has not yet acquired the fissile materials she needs, given our assumption that the game is beyond his control as soon as fissile material has been acquired⁶⁶. If I were to allow player A a chance to alter the q 's as well (as discussed in appendix D), the assumption becomes fundamental and must be tested. Here, the worst A can do by assuming that T is not yet beyond the acquisition stage is spending money defending against a threat that might not be real, whilst the opposite assumption could lead to a decision to do nothing when a major catastrophe could have been avoided. Hence it seems reasonable that A sets the level of deterrence *prior* to the terrorist's choice and subsequent materials acquisition.

4.6.3 *The division into HEU and Pu branches*

Plutonium and HEU are not regularly stored in the same facilities⁶⁷; HEU of terrorism concern may typically be found in research reactors, enrichment facilities and in naval fuel (fresh and lightly irradiated) whilst plutonium is found primarily in reprocessing plants, in storage for conversion into mixed-oxide fuel or in civilian reprocessing plants. However, there are instalments (for example nuclear weapons assembly/disassembly facilities and weapons component storage sites) where both materials are kept, and the safeguarding of these places does not fit cleanly into one of the two branches of my model.

Mathematically, the formalism only demands that no dollar may be counted in both categories. For the formalism to be consistent, this can be ensured in any arbitrarily chosen way, for example by splitting in half wherever both materials are safeguarded by the same equipment; deciding how exactly to make the split ultimately becomes a question of judgement. If the motivation for using the above model was, say to check if plutonium efforts are worthwhile, one may

66 This simplification is treated below

67 Ferguson and Potter *The Four Faces...* pp.120-124.

place in the Pu category only the funds that go into safeguarding *solely* plutonium. In conclusion, the exclusiveness of the categories can always be ensured, and the shortcoming is not fundamental.

4.6.4 *Non-deterability of terrorists*

It has been assumed herein that player T may neither opt not to attack at all, nor choose conventional means instead of nuclear. The simplification might well be an unrealistic one and I have devoted chapter 6 to this question alone⁶⁸. Here, however, one should remember that the model is devised so as to be beheld by the government player, and since I am not testing *whether* to defend against a nuclear threat or not, but *how*, it would not do to assume the threat is anything but real. One may believe, like Kamp and Frost⁶⁹, that the threat is *not* real and question the relevance of testing defences against it, but this is a different debate. With reference to chapter 6 the question discussed herein might be what happens if the deterrence schemes outlined in that chapter fail or are not implemented.

A major problem with deterrence of terrorism as discussed in chapter 6 is that it is very hard to verify whether a deterrence campaign works or not. It should therefore be accompanied by a safeguards programme as a backup. Furthermore our analysis in that chapter shows that safeguards form an important element of relative deterrence by decreasing the terrorist's perceived chances of obtaining fissile materials, so there is no conflict between believing in the possibility of relative deterrence and giving full priority to safeguards programmes.

As a general note all the different games presented in this thesis could be devised to be played simultaneously to try and capture more of the dynamics between the players. Safeguards priorities, priorities between safeguards and second layers of defence, terrorist choice of nuclear or conventional means and in the former case, choice of nuclear material could then be analysed in a single scenario. While this would certainly be more realistic, it would also be mathematically vastly more complicated. I deem therefore that playing one part of the game at a

68 Deterrence is simply defined as measures which persuade an adversary that a certain cause of action is not in his or her own interest. Thus relative deterrence from nuclear to conventional means of terrorism can be achieved (I argue in chapter 6) by changing important parameters of the terrorist player's cost-benefit analysis.

69 See section 1.3.

time will suffice mainly for the reason that the signalling between government and terrorist is of a weak and sporadic character, so modelling the situation so that each player plays in isolation where the other player acts as a part of the environment (and is able to influence that environment) is likely to be a good approximation. The price to be played in complexity for undertaking the more realistic game is therefore unlikely to pay off in terms of better understanding in this framework⁷⁰.

4.6.5 *The rôle of a 'second line of defence'*

I have already mentioned the simplification that player A cannot affect the passing of matters beyond hindering T's acquisition of nuclear materials (This shortcoming is analysed briefly in appendix D where a model similar to that in figure 4.3 is used to discuss the relative emphasis between a first and a second layer of defence). I will not delve into the matter here. Although it would not be difficult to incorporate a second-layer model into that used in this chapter, it was deemed wiser not to do this for the sake of keeping the two discussions apart and not complicate calculations unnecessarily. One may think of it as though a budget has already been split into first and second layer efforts so that C is the sum to be spent on safeguards; then in turn the above model yields how best to split C .

4.7 *The gains of gaming*

In this chapter I have considered a question which has been treated in the past by qualitative means. A natural question to ask is what was gained by formulating the problem in the form of a game and going through the calculations.

Indeed policy conclusions reached herein for the special case of the United States - the underfunding of HEU and possible overfunding of plutonium - were reached before by Ferguson and Potter⁷¹ by qualitative arguments alone. Moreover, it is clear that these conclusions hinge upon the qualitative analysis of chapter 3, where it was argued in a qualitative way why HEU is of greater proliferation concern. For the sake of these conclusions alone, it is not perhaps obvious that the gaming effort was more than a

detour on the way to a conclusion which was implied in rough terms already from the previous chapter. It is true, as demonstrated by the book of Ferguson and Potter, that in order to make qualitative recommendations about the prioritisation between the two fissile materials, qualitative arguments suffice. On the other hand, other experts who have surveyed the same evidence as do Ferguson and Potter, have *not* necessarily reached the conclusion that a heavy emphasis on one material should be qualified⁷².

In view of the discussion of criteria for evaluating the gains of formal analysis of section 2.10, it is clear that the main outcome of gaming is not these qualitative conclusions in themselves, which stem from the application of roughly estimated numbers in a particular context as a way to demonstrate the usage of the model. The primary product of devising a game and analysing it is the most general set of inequalities, (4.9) and (4.11). These inequalities reduce the large and complicated problem of safeguards priorities to one of estimating a few parameters, and are valid independently of what the values of these parameters may be⁷³. Beyond simply inserting numbers, they can be used in a variational sense, such as I did in section 4.4.3 where we could use these formulae to analyse the effect of fissile material stockpiles changing in time.

While such results are potentially powerful, they must be used with caution. It is a weakness as well as a strength of a model that it never captures the full picture, and before taking such general formulas at face value, the user should understand what assumptions were employed to produce them, and assess whether these match the special case under evaluation. The mathematical model is never equivalent with the real problem, only a condensed representation of it, and judgement has been made on the side of the analyst in order to devise it.

Therefore, in order to apply it in practice with confidence, the user must either fully understand the model and its premises, or have a high level of trust in someone who does. Powerful tools though they are, when applied wrongly such formulae are of limited use at best, misleading at worst. This necessity for special training is perhaps the greatest weakness of

70 It could nonetheless be an interesting undertaking for future research.

71 Ferguson and Potter *The Four Faces...*

72 See e.g. Matthew Bunn and Anthony Wier 'Terrorist Nuclear Weapon Construction: How Difficult?' *The Annals of the AAPSS* 607 (2006) 133-149

73 The simplified and stronger criteria (4.10) and (4.12), notably, are not, because they depend on the assertion that ΔT is positive, a plausibility assumption based on real-life estimates.

the methodology from a pragmatic point of view.

A final and related limitation to bear in mind is what we might call the information problem. Because the game and subsequent analysis in this chapter have assumed a *generic* government at some unspecified point in time, I have assumed to have significantly less information than a real government will have when facing the same question. Before making use of the equations derived the real government must ask itself the additional question: does any of the extra information I have make it necessary to use a different gaming representation than the general one? There seems no way of solving this problem in general; it is another point to bear in mind.

With these cautions in mind, however, the modelling in this chapter has yielded some simple and intuitive formulae and some striking results. The concept of the asymmetry cost, which proved to be a crucial quantity, fell out of the analysis almost automatically while it would not be immediate to recognise qualitatively that just this quantity is of such prominence. Another property of the scenario which is not easy to detect qualitatively is the tendency of the system to lean heavily to one side depending on the sign of the asymmetry cost; the form of the inequalities amplifies the asymmetry. Because we argued on intuitive grounds that this cost was positive, the system automatically inclines steeply towards HEU⁷⁴.

It is never possible to say with certainty what conclusions could and could not have been reached with a different methodology, since the logic which I formulate mathematically could, at least in principle, be formulated in words. That question can thus best be answered subjectively by each researcher; could I have reached this conclusion otherwise? That HEU safeguards should be given higher priority is certainly a conclusion one could reach qualitatively, but the notion that plutonium safeguards may be overfunded was one which ran counter to my intuition and was not so easy to swallow until the more detailed quantitative analysis of section 4.5.1 had been performed, providing a deeper understanding of how such a result could be possible.

4.8 Conclusions

My calculations indicate that with respect to

74 This is particularly true if the government's decision influences the terrorist's choice of material. Then the criterion (4.12) is a much tougher criterion compared to (4.11) than (4.10) is compared to (4.9).

terrorism, a strong emphasis on HEU in safeguards efforts is soundly justified in the US safeguards programmes abroad. There is reason to suspect from our analysis that US efforts to throttle proliferation of plutonium by improving safeguards are being overfunded at present and should be put on hold or at least not be allowed to get in the way of more urgent HEU efforts. The rough numerical evaluation indicates such a conclusion, but uncertainties are too large to establish certainty at this stage.

It seems certain from our analysis that HEU funding should be increased dramatically since the security benefits of investing in HEU safeguards are significantly greater than the corresponding costs. This conclusion is in line with those arrived at by Ferguson and Potter⁷⁵ by a different methodology.

While expenditures to further improve plutonium safeguards could decrease the expected utility at present, this is not necessarily a permanent conclusion. After a further improvement of HEU safeguards, the estimated terrorist preference could shift sufficiently in the direction of plutonium as to make investment in security upgrades opportune for this branch once more.

Another conclusion that may be drawn is the importance of building trust and good international relations between countries of proliferation concern and those in a position to pay for the lessening of such concern. If investment is to be beneficial, adequate value for money must be ensured. The importance of the rate of improvement with increased expenditure is made clear by the prominence of the quantities dp_u/dc_u and dp_p/dc_p in concluding results.

What has already been spent is important only in that prior efforts have brought us to *status quo*, hopefully with a few lessons learned. The absence of the absolute cost of safeguards and stockpile elimination from our criteria for optimal operation, however, renders clear that as future strategies go, threat reduction per dollar is the relevant number, whereas for example a weighing of accumulated cost of safeguards versus estimated cost of an attack is not. Roughly speaking, as long as one dollar invested in a safeguards branch lowers the nuclear terrorism threat by more than one dollar equivalent, safeguards spending should be increased for that branch.

75 Ferguson and Potter *The Four Faces...* pp.324-336.

The HEU/Plutonium choice: From the Terrorists' Point of View

I have considered the threat from HEU and Pu as a government sees it. We will now turn the situation around and view the same choice as seen from the view of a terrorist actor. It may seem daunting to attempt to understand the motives and deliberations of a given terrorist leader, and I shall make no attempt to enter into the field of terrorism psychology. Rather, I follow the scheme laid out in chapter 2 and consider a generic terrorist whom I assume to be *rational* and *intelligent* in the way that these terms are understood in rational choice theory: 'rational'¹ means the terrorist *consistently* chooses the action that is best for her; 'intelligent' means, in layman's terms, that she is *able* to work out what the best option *is*, based on the available information.

The two previous chapters have considered a question of very direct policy implications, namely the division of resources between HEU and plutonium safeguards as seen by the government paying for such an enterprise. The question treated in this chapter of a somewhat more theoretical nature and lends support to the more policy oriented analyses in the rest of the thesis more indirectly. While the question treated links the analysis below directly to the two preceding chapters, it will become clear that theoretical implications which result from the formal analysis hint at important realisations about terrorist actions and preferences in general, to which I will make frequent references in the next chapter. While policy implications from this chapter may not flow so directly from the analysis herein as they do in other gaming chapters of the thesis, the realisations and conclusions drawn are nonetheless of great relevance to a government trying to understand a rational terrorist adversary.

5.1 Research question and outline

The research question to be discussed in this chapter is

¹ A discussion of the proper understanding of this term is found in section 2.9.

What are the conditions determining whether a rational and intelligent terrorist who seeks to maximise the damage done to a target country would choose plutonium or uranium for a project to build a nuclear weapon, and how can a government opponent affect such a decision?

5.1.1 Chapter Outline

As a general backdrop to our analysis complementing that already presented in chapter 3, I discuss summarily an aspect of terrorist acquisition of nuclear materials which has not been treated in the previous, namely the choice of whether to mount an overt attack on a nuclear facility or to approach with stealth, for example through bribing insiders to theft.

I then go on to devise a model with which to analyse the terrorist's choice of fissile material and use the model to compare the two arguably most important candidate strategies: the opportunistic strategy of taking whatever material becomes available first, and a more patient strategy of waiting until the preferable material, HEU, becomes available. A standard multiplicative system of discounting is introduced to model the cost of waiting. I thereafter discuss the meaning of the price of lost time and link it to the concept of time failure, reflecting the fact that the probability of unforeseen events derailing the terrorist efforts increases with the time spent on the project. In relating the terrorist's fear of failure mathematically to the notion of discounting, however, an anomaly appears which makes way for an important realisation about absolute deterrence of terrorists.

Finally, a brief discussion about reasonable limits of terrorist intelligence (as defined in rational choice theory) is provided, and a discussion of the possible extension of the model to allow for the terrorist to make a 'hopscotch' bomb out of small batches of different fissile material. Finally the policy implications of the analysis are discussed.

5.2 Acquiring the material: strength or stealth?

The discussion in this section may not be a direct necessity to the formal analysis below, but forms a backdrop which may help in connecting the formal to the real situation. In applying the gaming conclusions to policy, qualitative discussions of this sort are also

necessary to determine how best to make use of the knowledge gained from the formal procedure. The reader is also encouraged at this point to review the sections on the intricacies of nuclear safeguards and security, presented summarily in section 3.5.

When a terrorist organisation wishes to acquire fissile materials, it has the choice between two main approaches: to mount a bank robbery style attack on a facility and force their way to the desired weapons components, or to use stealth to acquire the material undetected, most likely by use of insiders at facilities storing such materials, for example by threats and/or bribes.

From the point of view of whoever is in charge of the logistics of the terrorist organisation post acquisition – smuggling, building, transport and detonation – there can be no doubt that undetected acquisition will be preferable by far. As soon as the disappearance of several kilogrammes of HEU or plutonium is detected, efforts to track down the perpetrators of the theft and hinder use of the materials will almost undoubtedly be initiated immediately. If the theft is discovered within hours, the thief must expect that border crossings and other ways out of the country where the material was stolen will suddenly be more crowded than usual with inspectors bent on recovering the dangerous contraband.

Perhaps as importantly (as will be discussed in chapter 6), the knowledge that police and intelligence are trying very hard to track the material down will add considerably to the stress of the group of technicians actually building the weapon, increasing the probability of errors and even accidents derailing the project. It is likely that this alone is enough to make terrorist leaders contemplating the best route to the bomb want to keep the existence of the ambitious and expensive terrorist nuclear proliferation project secret for as long as possible, preferably until the bomb goes off.

While it is true that terrorist organisations have mounted armed military style attacks on public targets in the past, the incidents most often cited have not been part of a larger plot whose success depended in part on staying clandestine. Bunn, Wier and others of Project Managing the Atom have repeatedly warned about the threat from groups of armed terrorists against nuclear instalments quoting such incidents as the plane hi-jackings of September 11, 2001, the seizures of a Moscow theatre in 2002 and the

school in Beslan in 2004², all of which carefully planned and coordinated attacks by a group of terrorists. All of these incidents, however, had a violent ending which, apart from the tragic killing of innocents, precluded further near-term terrorist activity for any of the perpetrators involved.

In a recent book of journalistic nature, Langewiesche makes the case that such tactics are not preferable if the goal is to steal material which must subsequently be brought to a safe place for further construction³. Clearly it is not unthinkable that a well organised and well equipped group could succeed in such a spectacular act of robbery, yet it seems unlikely to ever be preferable as long as the alternative of using insiders is there. Langewiesche brings up the questions the terrorist leader must ask when planning the attack: Once the group is inside, what then? If they do not know exactly where to find what they are looking for, locating it and gaining access to it will take time and probably involve pressing employees for information. The Russian nuclear cities for example, storing the greater share of Russia's fissile materials, each contain hundreds of buildings so getting past the outer perimeter will only be one small step. Even assuming the terrorists know exactly where to find what they need once inside, the operation will take at least minutes, time enough for guards to organise themselves to try and keep the intruders from escaping. Even if they do escape the plant, safety is still far from secured. With police or the military on their tail, crossing borders will be much more difficult than if no suspicion were raised. Reaching a border will in most cases take hours, ample time for the nearest border controls to be alerted and reinforced.

Knowing exactly what materials the terrorists have stolen⁴, government antiterrorists will know what else the group needs in order to weaponise the uranium or plutonium and a number of incidents which would normally not attract attention, such as purchase of special workshop equipment, could now cause red flags to go up. If it has good records, the government will know the isotopic composition of the stolen material for identification purposes, and what

2 e.g. Matthew Bunn *Securing the Bomb* 2007 report from Project Managing the Atom (Belfer Center for Science and International Affairs, Harvard University, 2007) p.13

3 William Langewiesche *The Atomic Bazaar* (New York: Farrar, Strauss and Giroux, 2007) pp.48-50

4 This, of course, requires the government to have good accounting procedures in place, which may not always be the case.

chemical and metallurgical processing is needed to make the needed components. While random detection of smuggled uranium and plutonium is very difficult, it is made somewhat easier by knowing exactly what one is looking for, especially if one has some idea about where to look. Even very small quantities of airborne uranium and plutonium dust can be detected using sampling aeroplanes, for example⁵.

If a plot using an insider goes askew, on the other hand, the terrorist can often vanish and resurface for another attempt somewhere else. The insider who was bribed or coerced into helping could be given so little information as to be of little use to the police. If the theft succeeds it will presumably be detected at some point, but with luck and a clever ploy terrorists could have gained enough time to escape to safety or even finish the attack. For these reasons, Langewiesche concludes, and this author agrees, that 'a rational bomb-maker would abandon any idea of commando heroics'⁶.

The choice between strength and stealth clearly matters to the probability of success in the subsequent construction and attack. For my modelling purposes I will use only one probability of success however (as in the previous chapter). Although the above arguments indicate that the rational terrorist will try to acquire the material undetected, it is not necessary to explicitly assume this, only bear in mind the importance of the manner of acquisition. For numerical estimates, an assumption would be necessary.

5.3 A different angle

Although the situation in this chapter may be thought of as the same as that modelled in the previous chapter, the analysis I will employ is formally rather different. I argue in this section that such a different approach is natural and fruitful.

Says Bennett, 'Of all the limitations of the basic game model, arguably the most fundamental is the assumption that all the players see the same "game"'⁷. In the case where the situation to be modelled is

highly asymmetric the assumption that a single game can fruitfully portray the relevant options and preferences of both players could well be too strong.

Terrorism, under study here, is an extremely asymmetric conflict: the government on the one side has access to vast resources, military forces, police and intelligence services, whilst the terrorist must manage with limited funds and a comparatively tiny workforce most often recruited more for their ideology rather than for their skills. Also the information situation is very different for the two players; whilst the terrorist's success depends on all her plans and moves remaining clandestine, the government is to a large extent a public player - his plans and information (provided by intelligence services and the like) may be kept secret, but major moves by the government will be monitored by the press. Budgets will be scrutinised by the political opposition and many of its actions will be physically visible to the terrorist: the deployment of troops or border guards or the installation of bulky radiation detection equipment in key positions are not invisible operations and the terrorist is likely to be watching.

In the previous chapter it was possible to reduce the game theoretical model to a decision theoretical one, involving in practice a single player. This could be done by assuming that all relevant information that the government player could obtain to aid his choice would only be received *after* the choice had already been made. The cost of doing this, however, was that the previous model could not be used to analyse the factors influencing the choice of material the *terrorist* had to make.

The scarcity of information makes it natural to once again model the situation as a one-player game, set in an environment of imperfect information. A truly game-theoretical model where strategies could be played against each other would perhaps have been ideal in some ways. There are paths one could possibly pursue to achieve this, yet the situation at hand is not one that is easily modelled as a two player game. Terrorism differs from conventional warfare and even asymmetrical guerrilla warfare in the extreme sparsity of information exchange; communication between adversaries is sporadic and of a stochastic nature and might be best simulated in a computer programme. Such a project would likely be fruitful and a natural continuation of the analysis presented herein, but the simplicity, transparency and forcefulness of the reduced, decision theoretical approximation is appealing enough to justify it as a

5 Carson Mark, Theodore Taylor, Eugene Eyster, William Maraman and Jacob Wechsler 'Can Terrorists Build Nuclear Weapons?' in Leventhal and Alexander (eds.) *Preventing Nuclear Terrorism* (Lexington:Lexington Books, 1987)* p. 60

6 Langewiesche *The Atomic Bazaar* p.50

7 Peter G. Bennett 'Modelling Decisions in International Relations: Game Theory and Beyond' *Mershon International Studies Review* 39 (1995) p. 30.

first approach for my largely exploratory purposes.

5.4 The choice of fissile materials from the terrorist's point of view

When studying the choice of material for a nuclear programme seen through the eyes of the terrorist (player T), I will largely disregard the government's perspective, reducing once more to a 'decision theoretical' model. The interactions between the players is of a sporadic and random nature, and I will therefore see the two players as non-communicating as an approximation and in turn discuss how the government might influence the situation indirectly by affecting the values of some of the parameters of the game.

In the first round of the game, shown in figure 5.1, player T is given an opportunity to obtain a sufficient quantity of fissile material, and Nature decides whether it is HEU or separated plutonium⁸. HEU is picked with a probability $1-\rho$ and Pu with probability ρ , $0 \leq \rho \leq 1$. Player T can now choose to take the opportunity (accept) or wait for the next round (decline).

If T decides to decline and wait, the waiting period will induce discounting, a system to model the extent to which harvesting a payoff now is preferred to receiving it at a later time. Similar concepts exist in the economic literature, but the standard multiplicative discounting system typically employed in game theory⁹ is somewhat different, which should not confuse the reader. Here, for every round player T waits, all payoffs (whether positive or negative) are discounted by a factor $\delta \in [0,1]$. Waiting two rounds, payoffs are discounted by a factor δ^2 and so on. The discount factor represents the *impatience* of player T (a nuanced interpretation of discounting in this particular case is an important part of the analysis in the following).

There are a number of reasons for the terrorist to be impatient to get started with a planned nuclear project. The longer she waits, the larger the chance

that her plans might be discovered. Assuming that a certain level of searching activity will be required in order for the opportunities for obtaining fissile material to keep appearing, each new round means exposure and risk of detection. Furthermore, safeguards on storage places for nuclear materials may improve with time, causing perhaps an expectancy that getting the materials in the future might be more difficult than it is now. All in all there should be ample reason for δ to be smaller than unity.

A second reason for choosing a multiplicative model is that I wish my model to reflect my assumption from the previous chapter that player T is bent on going nuclear and will not change her mind with time. Doing nothing, as I will argue, will give a payoff of zero, and so discounting must then be devised so that payoffs that are initially positive will not drop below zero with time. The multiplicative system does this and hence seems a good choice. It turns out, however, that I run into trouble, necessitating a reconsideration of the assumption that player T cannot be deterred with time. But I will leave this for later.

Once player T has accepted an opportunity, Nature decides whether the construction and attack succeeds – probability q_i , $i = u,p$ – or fails – probability $(1-q_i)$. It is important to note that now q_u and q_p are as estimated by player T, not necessarily equal to the values of the same name in the previous chapter, which reflected the government's belief. The same goes for all quantities whose value cannot be determined beyond doubt, for example the damage inflicted on the antiterrorist/government (player A), T_u or T_p , which the terrorist may estimate differently than the government.

In the case where the attack succeeds, the terrorist receives a payoff $-C_T + vT_i$ where $i = u,p$. C_T is the cost of the operation for player T, and the reward for the success is vT_i , proportional to the damage inflicted on player A. The proportionality constant, as we will see, has an interesting interpretation. In the case of failure, player T receives an extra cost of failure, φ . This reflects possible costs in addition to the mere building of the bomb, this could be e.g. loss of credibility, loss of funding or loss of property or personnel if the project ends in an accident. It is likely that this quantity is small compared to C_T , but greater than zero. One could generalise the game in such a way that the terrorist could try again after a failed building attempt, a possible future development not undertaken here.

8 Of course, the success rate of the project will furthermore depend on what *form* the fissile material is in, that is, how much processing it needs before it is usable in a weapon. Plutonium from very radioactive spent reactor fuel, for example, is separable in principle, but will in practice almost certainly be useless to the terrorist. I disregard spent reactor fuel throughout this thesis (lightly irradiated spent HEU fuel from research reactors, however, is another matter altogether).

9 e.g. Martin J. Osborne *An Introduction to Game Theory* (New York: Oxford University Press, 2004) pp.421-422

I have assumed for simplicity that the cost of building a uranium bomb is equal to the cost of building a plutonium bomb. While this is perhaps a coarse approximation, it simplifies the analysis and may be good enough as a first approach.

The assumption that the terrorist's reward for a successful attack is proportional to the damage inflicted has implications that cannot be ignored. Adopting the distinction introduced by Arce and Sandler¹⁰ between political and militant terrorists – the former using terrorism as a means to obtaining a place at the negotiation table but is eventually interested in political reform, the latter bent on incurring damage on the adversary until concessions are made – such a reward system implies that player T falls heavily on the militant side. An example to illustrate the distinction¹¹ is that of Spain, which has suffered attacks from the Basque separatist organisation *Euskadi ta Askatasuna* (ETA) for many years, and which was also victim of an Islamic terrorist attack on Madrid train station in 2004. Whilst ETA's attacks were discriminate, seeking to avoid mass casualties, the train station attack was the direct opposite, designed to maximise the carnage and escalate violence. It is fairly safe to assume that any terrorist organisation with the slightest interest in detonating a nuclear device in a populated area will be bent on mass casualties and be firmly in the militant camp. The signals sent through a terrorist attack by such immensely destructive means are several – the demonstration of technical sophistication to equal that of the adversary might be one – yet a main motivation must remain the perverse will to kill in the hundred thousands¹². A model where reward is proportional to the harm inflicted, while possibly simplistic, seems to mirror such preferences reasonably well.

The present model, of course, is not well suited to explain why most terror organisations are *not* interested in acquiring nuclear weapons¹³. If the

payoff system were modelled differently, moreover, for example in terms of bargaining power or concessions in the wake of the attack, a spectacular bloodbath might very well defeat the purpose entirely, effecting harsh retaliation rather than negotiation power¹⁴. It is important to be aware of these aspects of the model: it is designed to describe a logical process of decision making based on utterly fanatical preferences. In so doing it is necessary to strongly limit the effect of actions and strategic considerations which more moderate actors might perceive as deterrents¹⁵.




Illustration 5.1.: The HEU/Pu game seen from the terrorist's point of view.

The relevant terrorist strategies in the game of figure 5.1 are the following: An opportunist strategy (O) where the first opportunity is accepted whatever the fissile material, a uranium strategy (U) where the player waits until an opportunity for uranium comes up, and a plutonium strategy (P) where plutonium is likewise awaited. More complicated strategies are thinkable, of course, but one may show that with a multiplicative discounting system, at least one of the three simple strategies will always have a higher payoff than such a combined strategy¹⁶.

Notice furthermore that, unlike in the model in the previous chapter, player T now has the option to *not*

10 Daniel Arce M. and Todd Sandler 'Terrorist Signalling and the Value of Intelligence' *British Journal of Political Science* 37:4 (2007) pp.573-586.

11 *Ibid.* p.2

12 A possible exception could be if the nuclear weapon is detonated in a remote area as a warning shot to create fear and possibly lever the government into concessions. However, such a plan would demand that the terrorist holds (or credibly claims to hold) at least one more nuclear weapon, and the delivery of the second will be very much complicated by the increased security guaranteed by the first blast.

13 See e.g. chapter 2 of Charles D. Ferguson and William C. Potter with Amy Sands, Leonard S. Spector and Fred L.

Wehlin *The Four Faces of Nuclear Terrorism* (New York: Routledge, 2005)*

14 Morten Bremer Maerli 'Relearning the ABCs: Terrorists and «Weapons of Mass Destruction»' *The Nonproliferation Review* (Summer 2000), p.110.

15 A 'deterrent' in this thesis is understood as anything that can persuade an actor that a certain cause of action is not in his or her own interest, not restricted to threats of retaliation as was the typical Cold War interpretation of the term. See chapter 6 for a discussion.

16 More on this below. I will not show this rigorously here, but it will be fairly obvious from the following analysis where more detailed waiting strategies are discussed.

try to go nuclear at all, by declining every opportunity forever. The payoff for this strategy is 0 in this model: the strategy does not have its own end node, but the result is nonetheless evident. If one thinks instead of a strategy where player T waits for N rounds and then takes whatever is available, the payoff will be discounted by a factor δ^N . As N goes to infinity, the payoff becomes 0 (assuming $\delta < 1$; in the special case $\delta = 1$ one can still *define* this strategy to have payoff 0). This constitutes a fourth strategy that I denote (N) for 'no attempt'.

As player T's financing goes, I will not be very specific in this chapter, but merely assume that she has a budget pot of C_T designated to the nuclear project and that project alone. Note that the 'zero level' of finances is set as *status quo*.

5.4.1 The opportunist strategy

I commence by examining the opportunist strategy. The expected utility of this strategy is easily found to be

$$U_{(O)} = -C_T - \varphi[\rho(1-q_p) + (1-\rho)(1-q_u)] + \nu[T_u(1-\rho)q_u + T_p\rho q_p]. \quad (5.1)$$

For the terrorist to be better off trying to build the bomb than saving the money for something else (strategy (N), payoff 0), $U_{(O)}$ in (5.1) must exceed 0. This brings me to defining the following criterion for ν that may or may not be fulfilled:

$$\nu > \frac{C_T + \varphi[\rho(1-q_p) + (1-\rho)(1-q_u)]}{T_u(1-\rho)q_u + T_p\rho q_p} \quad \forall \rho \in [0,1]. \quad (5.2)$$

I will refer to (5.2) as the 'bloodlust' or 'blood thirst' criterion. When (5.2) holds for all values of ρ (as opposed to the *one* value that player T believes to be the correct one), this implies that strategies (O), (U) and (P) are all preferred to (N) since they all will have expected utility higher than 0. This follows since in all these strategies some opportunity is accepted at some point, and if building *either* option in the first round has a positive expected utility, so will it after a few rounds' waiting since discounting as implemented here does not change the sign of the payoff. If (5.2) is fulfilled, in other words, it means player T extracts so much pleasure from inflicting destruction that attempting a nuclear attack is worth the risk and cost both using HEU ($\rho = 0$) and Pu ($\rho = 1$)¹⁷.

17 Note, importantly, that player T has no other alternatives than commence a nuclear project or do nothing. The situation

5.4.2 Uranium 'waiting strategy'

Upon considering strategy (U) the relevant question regards the expected discount factor. The probability that an opportunity for HEU first comes up in round n is

$$p_n = \rho^{n-1}(1-\rho);$$

the discount factor in this case is δ^{n-1} so the expected discount factor in the uranium case is found from standard statistics to be

$$\begin{aligned} \delta_{(U)} &= \sum_{n=1}^{\infty} \delta^{n-1} p_n = \sum_{n=1}^{\infty} \delta^{n-1} \rho^{n-1} (1-\rho) \\ &= (1-\rho) \sum_{n=0}^{\infty} (\delta \rho)^n = \frac{1-\rho}{1-\delta \rho}. \end{aligned} \quad (5.3)$$

The expected payoff of the uranium strategy is then found to be

$$U_{(U)} = \frac{1-\rho}{1-\delta \rho} [\nu T_u q_u - C_T - \varphi(1-q_u)]. \quad (5.4)$$

Notice that from (5.3), $(1-\rho) \leq \delta_{(U)} \leq 1$. This implies that waiting is always preferable to giving up and opting for (N), that is, the terrorist cannot be deterred absolutely through changing the parameters δ and ρ . To understand this, consider if instead of the fraction in (5.4) one had the factor $(1-\rho)$. The reader may verify that this would be the expected payoff of the strategy 'accept in the first round if material is HEU, otherwise give up (decline forever)', and with $(1-\rho) \leq \delta_{(U)}$ this is never better than strategy (U). But then the same must be true for a strategy of the type 'wait for HEU, but if none arrives in the first n rounds then give up'. Assume T follows such a strategy and after round $n-1$ she has still had no HEU coming up. She is then in the situation described above (accept next round if HEU or decline forever) which is inferior to the strategy 'wait forever if necessary', so sticking to the former strategy cannot be preferable¹⁸.

This is interesting, since such a strategy (wait but give up after some time if no luck) is one that we would recognise from everyday life and intuitively makes sense in many cases. So why not in this case? Reasons for this may be traced back to the simplicity of the model, the 'blood thirst' criterion (5.2) and the discounting system used.

changes completely once one allows her to opt for a conventional strategy; see next chapter.

18 The term in Game Theory is that the strategy is not *subgame perfect*.

There are at least two principal reasons to decide to stop waiting for the ideal and settle for the less preferable but certain. Either (1) conditions have changed since the waiting period started; for example one could have an ultimate deadline by which something must happen, in which case the cost of waiting becomes unbearable at some point (δ becomes small), or (2) preferences can change with time. One can imagine a situation where either of these happens for the nuclear terrorist: something can convince T that if no material is obtained by a certain date the project is sure to fail. Or a new leadership of the terrorist organisation (that is, a change of T's preferences or 'type') might be of the opinion that spending the resources on conventional means or even political activities is preferable to nuclear activities.

By assuming all parameters except δ and ρ to be constant with respect to time, these scenarios are ruled so long as the bloodlust criterion is satisfied, which certifies that player T is bent on the nuclear option and will under no circumstances give up her plans. But if I allowed quantities such as q_u , q_p or φ to vary in time I could simulate both of the scenarios mentioned as reasons for preferences to change in time. The first reason to give up could be simulated by letting both q 's drop suddenly to zero in round n (guaranteed negative payoff if any option is chosen). The second, change of leadership, could be modelled by moderately altering the q 's (new leadership believes it more unlikely that build will succeed), φ (new leaders are more afraid of failure) or v (new leadership is less bloodthirsty; extracts less utility from major carnage) in round n so as to shift preferences amongst the strategies (notably, (5.2) would then not hold after round n).

Ruling out such 'wait, but not forever' strategies, thus, is possible only because of the multiplicative discount system which never changes the sign of future payoffs. They *do* become possible once a more realistic system of discounting is introduced, as we shall see.

5.4.3 Discarding the plutonium strategy

With a principally identical argument to that for (U) I deduce the expected utility for the (P)-strategy as well:

$$U_{(P)} = \frac{\rho}{1-\delta(1-\rho)} [vT_p q_p - C_T - \varphi(1-q_p)]. \quad (5.5)$$

With the arguments from the previous chapter in mind, there seems to be little or no reason to expect that strategy (P) should be preferable to (U), yet before I go on let me make a quick comparison. Strategy (U) is preferable to (P) if and only if

$$\begin{aligned} & \frac{1-\rho}{1-\delta\rho} [vT_u q_u - C_T - \varphi(1-q_u)] \\ & \geq \frac{\rho}{1-\delta(1-\rho)} [vT_p q_p - C_T - \varphi(1-q_p)] \end{aligned}$$

or

$$\begin{aligned} \frac{(1-\rho) - \delta(1-\rho)^2}{\rho - \delta\rho^2} & \geq \frac{vT_p q_p - C_T - \varphi(1-q_p)}{vT_u q_u - C_T - \varphi(1-q_u)} \\ & \equiv \frac{A_p}{A_u} \end{aligned} \quad (5.6)$$

where I have defined the quantities A_p and A_u :

$$\begin{aligned} A_p & \equiv vT_p q_p - C_T - \varphi(1-q_p); \\ A_u & \equiv vT_u q_u - C_T - \varphi(1-q_u) \end{aligned} \quad (5.7)$$

If the 'bloodlust' criterion (5.2) is satisfied, both A_u and A_p are positive. A_u , as we see, is the expected utility if T could obtain material i with certainty in the first round and accepts it.




Illustration 5.2.: Plot of the fraction on the left side of eq. (5.6).

Figure 5.2 shows a plot of the fraction on the left side of (5.6). For given values of ρ and δ , the ratio A_p/A_u must lie below the graph for (U) to be preferred to (P). If one accepts, based on the arguments laid out in chapters 3 and 4, that $A_p/A_u < 1$ then this is always the case if $\rho < 1/2$, which should be obvious from common sense. Even for higher ρ this is most likely true if the discounting is modest, like $\delta = 0.9$. This is also most intuitive: if the expected payoff using HEU

is higher than that for Pu and discounting is modest, it is better to wait for HEU than to wait for Pu even if the probability of HEU in each round is smaller than that for Pu. Indeed, if $A_p/A_u \ll 1$ (for example if A_p is only just positive whereas A_u is positive by a fair share - not an unlikely situation) then waiting for HEU is better even when discounting is harsh (e.g. $\delta = 0.1$) and the chances of HEU each round is small.

In summary, if as one expects A_p/A_u is smaller than and not close to unity¹⁹, the only time when (P) might still be preferable is when the discount factor δ is small and the probability of obtaining Pu is much larger than that for obtaining HEU. With the discussion of Pu versus HEU in the previous chapter and the assumption that player T will very roughly share this opinion, it appears I can safely disregard the strategy (P).

5.4.4 A comparison of (U) and (O)

A more interesting comparison is that between waiting for HEU and accepting the first available option. Intuitively I would expect these two strategies to be the two real options: wait until weapons-grade uranium becomes available or start right away using plutonium if that is what is at hand.

In reality, building the weapon and performing the attack will take time, and if the expected time interval between offers is shorter than the building time, a third strategy might emerge as a compromise in the case where player T is exceptionally well funded: start building Pu if that is what is available, but still be on the lookout for HEU and get it as well if the Pu bomb project is not yet far advanced. This would call for free resources possibly as large as $2C_T$ in our terminology. In the interest of keeping the degree of detail at a manageable level, I will not discuss this strategy beyond this paragraph. As a note on cost, however, much as two parallel bomb projects will be able to share some resources, the fissile material itself might well be the costliest part of the project (in the tentative calculation of Zimmerman and Lewis, fissile

material alone makes up more than 70% of the total cost C_T ; \$4 million of the total \$5.4 million²⁰). Hence if T does not really believe in the feasibility of using Pu, buying it anyway might mean wasting her one chance to get it right.

The strategy of opportunism is opportune if it brings at least as high an expected utility as waiting for HEU, that is, when

$$U_{(O)} \geq U_{(U)}.$$

This is solved with respect to the discount factor and use of the relations²¹

$$U_{(O)} = \rho A_p + (1 - \rho) A_u; \quad U_{(U)} = \frac{1 - \rho}{1 - \delta \rho} A_u \quad (5.8)$$

(A_u and A_p are defined in (5.7)) to find, quite simply, that (O) is at least as good as (U) when

$$\delta \leq \frac{A_p}{U_{(O)}} \equiv \delta_0 \quad (5.9)$$

This result is readily interpreted. If HEU comes up in the first round, the strategies prescribe identical action, so all T needs to decide is what to do in the opposite case. Consider therefore the situation where Pu has come up in the first round and T is deciding whether to accept or not. If she accepts, the expected utility is A_p but if instead she waits one more round and *then* takes whatever comes up ('just to see if HEU comes up next time...'), the expected payoff is $\delta \cdot U_{(O)}$. Just when $\delta = \delta_0$ as defined in (5.9), we see that the two expected payoffs become equal. If $\delta < \delta_0$, accepting the plutonium (opportunist strategy) is better than to keep waiting and *vice versa* if $\delta > \delta_0$. If she decides to decline the plutonium, however, and in the next round Pu comes up again, she will face the exact same choice. Hence δ_0 must be the critical value that distinguishes whether to grab the first available option or wait for HEU for as long as it takes.

Note at this point that equation (5.9) is a very general result and does not depend on the exact forms of A_p and $U_{(O)}$ which result from the specific modelling, in our case (5.7). A different payoff system would give different expected payoffs for different strategies yet so long as the structure of the game is as

19 A note is required here: I assume $A_p < A_u$ based on our arguments from this chapter and the previous that a plutonium bomb is more likely to fizzle, hence cause less damage. To assume that the *terrorist* holds this belief, however, is another matter, for our understanding of 'rationality' does not require rational expectations, as detailed in section 2.9. An argument for assuming rational expectations in this case, however, is our motivation to deal with the 'worst case' kind of terrorist, arguably the terrorist who wants to cause maximum damage *and* has good knowledge of how to do so.

20 Peter D. Zimmerman and Jeffrey G. Lewis 'The Bomb in the Backyard' *Foreign Policy* (November/December 2006) p.36. Note that this is for a HEU gun-type device; I assume in this chapter that similar numbers would apply to a project to build a crude plutonium implosion device.

21 These should be obvious from above calculations.

drawn in figure 5.1, equation (5.9) is the same. In later sections where I generalise the discount factor δ into a discount function, equation (5.9) must be revised.

The choice between an opportunist strategy and one of waiting for the better option to occur is then entirely decided by the cost of waiting balanced against the expected payoffs. This should not be unexpected since, if HEU has higher expected payoff (equal to A_u when undiscounted) than Pu (A_p undiscounted) – hence also than (O) which randomises between the two – wasted time is the *only* cost of choosing (U) over (O) in our model, just as lower expected payoff in the first round is the only cost of choosing (O) over (U).

Given our previous argument that a plutonium bomb is likely to do less damage on average than a HEU gun, it might be preferable for a government player if the terrorist could be pressured into adopting an opportunist strategy. From (5.9) and $U_{(O)}$ from (5.8), notice that there are two main parameters for player A to seek to influence:

- δ : If the terrorist player can be stressed into sufficient impatience so that δ drops below δ_0 , (O) becomes preferable to (U).
- ρ : If the probability that HEU will come up is small (that is, ρ close to unity), δ_0 approaches unity, and δ might slip beneath it.

This result is intuitive: if player T is impatient, and thinks the waiting time for HEU might be long (ρ is large), she may find the wait too long, and opt for a strategy where she takes whatever she can get. However, the result is more precise than could have been found by qualitative means alone. Given only two qualitative assumptions, that T discounts future utility by a multiplicative model and that the plutonium project can be discarded, I find in (5.9) the mathematical way in which δ , ρ and other parameters determine which strategy is preferable, thus (5.9) (together with the definitions of A_p and $U_{(O)}$) holds much more information than the mere conclusion about patience and waiting time above, which one could arrive at intuitively. For example, I may read out directly how the situation would change if other parameters, e.g. φ , ν , q_u or q_p , change, variables which are largely out of A's power to influence but whose values he may estimate and re-estimate given new information.

The direct use of the formula (5.9) and others will demand some numerical estimate of the quantities involved. In section 4.4 I did a similar exercise as seen from the government's point of view, while estimation

as performed by the terrorist adversary is more difficult and will involve an added layer of speculation. Therefore I will not include a numerical study in this chapter. Nonetheless, simple mathematical relations such as (5.9) are valuable in that they provide a method to reduce the large and multifaceted problem of terrorist preferences to one of estimating three quantities. Analytical studies of how these quantities interrelate, furthermore, provide models describing what effect may be achieved by altering different parameters of the model and, while difficult to quantify, gives at least a better qualitative understanding of how policies may affect terrorist preferences.

5.5 Analysis: the price of time and the fear of failure

It is of interest to take a closer look at the discounting system employed and see if a more 'physical' understanding can be assigned to it. The notion of discounting is an inherently elusive term, difficult to pin down and give a clear interpretation²², and it will be shown how several options are possible. Let me explore one plausible way of interpreting discounting: as the *impatience due to fear of failure*, because waiting increases the probability of unforeseen incidents²³. With this I am able to establish a relation between the discount factor δ , and the expected time τ between two opportunities for fissile materials. What emerges is an inconsistency of formalism (i.e. the interpretation is inconsistent with the standard mathematical form of discounting employed above) and with it an understanding of absolute deterrence of terrorism²⁴.

The probability that some unforeseen incident happens and derails the project will most often increase with time²⁵, and examples that come to mind

22 See e.g. discussion in Osborne *An Introduction to Game Theory* pp.421-422.

23 Another and less concrete interpretation is used in the next chapter where I *do not* link impatience directly to a fear of failure.

24 *Absolute* deterrence is used in the meaning 'detering actors from terrorist action altogether' as opposed to *relative* deterrence of some forms of terrorism (e.g. nuclear) meaning deterrence from the use of certain means of terrorism (e.g. from nuclear weapons to conventional means). See chapter 6 for further details.

25 There is a subtlety here: probability increases with the time *left to wait*, whilst the time already waited is (I assume) irrelevant. Here's an analogy: A couple is planning to have five children. Before any of them are born, there is a large

are:

- Searching for fissile material involves sticking your head out to some extent, and many of the seizures of nuclear material in transport have been the consequence of police sting operations.
- Co-conspirators possessing less zeal than their leaders may get cold feet and, given time to contemplate, decide to send an anonymous tip to player A or otherwise defect.
- A's intelligence officers, moreover, will be working constantly to try and detect the activities of player T, and the longer they search, the more probable it is that they will catch a drift.
- Pre-emptive strikes could randomly knock out the project even if no intelligence pointed to it, as could police raids and other unforeseen events. Possibly some of the key participants of the terror project will already be wanted for previous acts of violence, and must be ever watchful not to get caught

5.5.1 Time failure

The various ways by which the project may randomly derail with time listed above are not modelled individually or explicitly. As a simplification I will treat them all compiled, dubbing them collectively as sources of 'time failure' since they all increase in likelihood the longer the time horizon of the project.

Assume that there is a very small probability Δp_t that time failure will occur during the short time interval Δt . Let N be a real number greater than 1 (it is natural to think of N as an integer). The probability that a longer time $t = N \cdot \Delta t$ passes without time failure is then²⁶

probability that at least one of them will be a boy (about 97%). But after they have had four consecutive girls, the chances have shrunk to about 50%, since the probability of the sex of each child does not depend on the sex of any other child. The probability of the incident 'boy' happening at least once over the next five children is much larger than that it happens over the next *one* child, just as the chances that the project is hamstrung by some unforeseen event is larger over the next year than over the next month. The underlying assumption is that the incident that derails the project takes a short time compared to the overall time frame, which does not restrict us too much.

26 P denotes 'probability'. I assume all time intervals are independent of all others, that is, whether a time failure happens in one interval is independent of whether or not it

$$P(0 \text{ failures during } t) = f(0) = (1 - \Delta p_t)^N.$$

Using that $a^b = \exp(b \cdot \ln a)$, and $N = t/\Delta t$, I write this as

$$f(0) = e^{\ln(1 - \Delta p_t) t / \Delta t}.$$

A simple expression is thus obtained for the probability of non-failure by the time t :

$$f(0; t/\theta) = e^{-t/\theta} \quad (5.10)$$

where I have defined the quantity θ :

$$\theta = \frac{-\Delta t}{\ln(1 - \Delta p_t)}. \quad (5.11)$$

In statistics this is a special case of the so-called Poisson distribution²⁷ where the *parameter* t/θ is the *expected number* of 'time failures' during the time period t and θ is the *average time* between each time failure. $f(0; t/\theta)$ is then the probability that *exactly zero* 'time failures' have happened during a time interval t .

The time θ I will refer to as the *penetration time*, borrowing the analogy of penetration depth from physics: just as the penetration depth is about as far as (say) light is able to penetrate into a certain material, the penetration time here is about as far as player T can expect to 'penetrate' in time before it becomes overwhelmingly probable that the project has failed. The chances that the project will survive a waiting time equal to θ are $1/e \approx 37\%$.

The logarithm in (5.11) has a simple interpretation which becomes obvious if one lets Δp_t and Δt be infinitesimally small: $\Delta t \rightarrow dt$, $\Delta p_t \rightarrow dp_t$. Then one may Taylor expand the logarithm to leading order only without loss of precision:

$$\ln(1 - \Delta p_t) / \Delta t \xrightarrow{\Delta \rightarrow d} - \frac{dp}{dt}$$

and hence the penetration time can be written

$$\frac{1}{\theta} = \left(\frac{dp_t}{dt} \right)_{t=0}, \quad (5.12)$$

quite simply the reciprocal of the rate of probability of

has happened in another already.

27 See any university level textbook in statistics. The Poisson distribution is typically used to describe events that happen stochastically in time so that the number of times it happens during a time interval is proportional to the length of the time interval. The probability that an event happens x times in a time interval during which the *expected* number of events is μ is $f(x;\mu) = (\mu^x/x!) \exp(-\mu)$.

time failure per unit time²⁸. In practice, however, it is probably easier to determine the penetration time directly (from experience) than estimate this slope.

The observant reader might question whence came the restriction that the derivative be taken at $t = 0$ in (5.12). The reason is that I demanded θ to be a constant with respect to time whereas the derivative in general is not – it has a non-linearity that disappeared in the linearisation of the logarithm above. For absolute rigour therefore the slope (5.12) must be calculated at $t = 0$, but a t significantly smaller than θ is a good approximation. The reader can easily verify these statements by setting $p_t = 1 - f(0; t/\theta)$ and differentiating with respect to time.

For complete consistency, a time-based model for the probabilities q_u and q_p should be implemented as well, yet I will assume for simplicity that 'time failure' during the construction period have already been calculated into these quantities. Since I have assumed that all probabilities are statistically independent, the exponential form of $f(0; t/\theta)$ makes the two formalisms equivalent²⁹. In summary: it is unproblematic to apply time failure only to the 'waiting and searching' stage and treat the building process as instantaneous with a finite and sub-unity probability of success.

5.5.2 Relating discounting explicitly to time and the cost of failure

In order to relate $f(0; t/\theta)$ to δ it is necessary to designate a payoff for the event of time failure. One natural choice is the defeat cost, $-\varphi$, assuming if so that none of the resources C_T have yet been spent; from Lewis and Zimmerman's estimates³⁰ it is reasonable that the majority of C_T is fissile material and construction costs, so this choice of failure payoff is roughly equivalent with assuming time failure occurs prior to materials acquisition and construction. Since the following discussion is mainly of theoretical interest it does not really matter whether $-\varphi$ or $-C_T - \varphi$ or something in between is used; the reader can

simply choose a higher or lower value of $-\varphi$ if preferable. Moreover, the relation between $f(0; t/\theta)$ and δ depends on what T is waiting for. I argued that (P) is a strategy we can ignore, hence assume that the only relevant 'waiting strategy' is (U).

Assume thus that T is waiting for HEU. The payoff for choosing this strategy from the start is $U_{(U)}$, and the payoff for waiting one round and then choosing (U) is $\delta \cdot U_{(U)}$. We will now think of $U_{(U)}$ independently of the specific discounting system that leads to the equation (5.4), merely as 'the payoff of strategy (U) if the strategy is commenced now', independently of exactly how this payoff is calculated. I will assume that the expected time between two opportunities is τ . For simplicity, let us assume that *every* round takes exactly the expected time τ .³¹ There is then from (5.10) a probability $\exp(-\tau/\theta)$ that time failure will not yet have occurred after one time period. Assume that in the first round, Pu came up, so that the expected payoff for sticking to the strategy (U) (rather than opt for the plutonium) is, with our standard multiplicative system, $\delta \cdot U_{(U)}$. An alternative way to discount, without the factor δ , is by introducing the possibility of time-failure. Now there is a probability $\exp(-\tau/\theta)$ that the payoff will be $U_{(U)}$, and $1 - \exp(-\tau/\theta)$ that the payoff will instead be $-\varphi$ due to some event that derails the project during the time interval τ . We now insist that the payoff be equal in the two pictures, and get the equation

$$\delta U_{(U)} = e^{-\tau/\theta} U_{(U)} - (1 - e^{-\tau/\theta}) \varphi \quad (5.13)$$

or

$$\delta = e^{-\tau/\theta} - (1 - e^{-\tau/\theta}) \frac{\varphi}{U_{(U)}} \quad (5.14)$$

Inserting (5.4) with A_u from (5.7) into (5.14) and solving with respect to δ gives us

$$\delta = \frac{(1 - \rho) e^{-\tau/\theta} - (1 - e^{-\tau/\theta}) \varphi / A_u}{(1 - \rho) - \rho (1 - e^{-\tau/\theta}) \varphi / A_u} \quad (5.15)$$

Equation (5.14) uncovers a problem with our discount system: if τ is large enough and $\varphi > 0$, δ becomes negative! A negative δ as implemented above would cause payoffs to alternate between positive and negative values for each round, which is clearly absurd! Note moreover that (5.14) is plainly inconsistent with the way discounting is performed: if

28 Strictly, the fraction containing the logarithm in (5.11) is the exact one.

29 Say waiting takes a time t_w and building and attacking a time t_b . If then the probability of success during build was given a time-failure form, we would have $q_i = \exp(-t_{b,i}/\theta)$, $i \in u, p$ and the total probability of success would be $q_i \cdot \exp(-t_w/\theta) = \exp[-(t_w + t_{b,i})/\theta]$, i.e. the time-failure formalism applied to the total time $t_w + t_{b,i}$.

30 Peter D. Zimmerman and Jeffrey G. Lewis 'The Bomb in the Backyard' *Foreign Policy* (November/December 2006) p.36

31 In reality, τ will be statistically distributed. Introducing this adds little or nothing to our understanding, however.

it were derived assuming instead that T waits *two* rounds, for example, the left side of (5.14) would become δ^2 whilst the exponents on the right hand side would be $-2\tau/\theta$; a whole different equation. Something is obviously more subtle than meets the eye!³²

The problem can be solved both mathematically and to our intuitive satisfaction by re-interpreting the discount system. Let there be instead a discount function, $\delta(t)$, by which payoffs harvested a time t from now are discounted regardless of whether t is long or short. If t is one round, the discount function is $\delta(\tau)$, if it is two rounds, the discount function is now $\delta(2\tau)$ and so on. The discount function used before (the multiplicative model) is equivalent with a discount function

$$\delta_{\text{old}}(t) = \delta_m^{t/\tau}, \quad (5.16)$$

where δ_m is the factor by which payoffs used to be multiplied for every round that passes, which I used to call just δ (redubbed to avoid confusion). The reader should be able to verify this.

The simple equation (5.9) derived above is very general in that it does not require any detailed specification of how the payoffs A_u and A_p are calculated. However, it was derived assuming a simple multiplicative discounting system, and does not hold when such is no longer the case. When now departing from such a simple system for a more general $\delta(t)$, it is not possible to find a straightforward generalisation of (5.9) along the lines of the qualitative argument which follows that equation as I will now argue. Thinking now of discounting as an operation, the time failure understanding not only multiplies payoffs by a factor $\exp(-\tau/\theta)$, but also subtracts a term $[1-\exp(-\tau/\theta)]\varphi$. A naïve generalisation for one round would then be

$$\tilde{\delta}(\tau) = e^{-\tau/\theta} - (1 - e^{-\tau/\theta}) \frac{\varphi}{U_{(u)}}$$

but this cannot be generalised to other times than $t = \tau$ because $U_{(u)}$ is defined *at* the time when an opportunity arrives, disallowing $t < \tau$, and for longer times such as $t = 2\tau, 3\tau$ etc., such a form would presuppose that P_u came up at the first, second etc. opportunity, suppositions we may not, of course make. Moreover, such a discount function would

depend on ρ , which is not as general as it should.

One can, however, derive a simpler and ρ independent discount function by use of a simple argument. As before we will insist that

$$U_{(u)} = \langle \delta(t) \rangle A_u,$$

where $\langle \delta(t) \rangle$ is the expectation value of the discount function assuming the first opportunity to happen immediately (an average, which does not depend on t). If $\delta(t)$ is a continuous function, it follows that there exists a time t^* so that

$$\langle \delta \rangle = \delta(t^*)$$

Now imagine T has to wait a time t before the next opportunity comes up. We will require that the discount function is then $\delta(t+t^*) A_u$. But since discounting should account for the probability of time failure, it must also equal

$$\delta(t) U_{(u)} = \delta(t^*) A_u e^{-t/\theta} - (1 - e^{-t/\theta}) \varphi,$$

which gives the equation

$$\delta(t+t^*) = \delta(t^*) e^{-t/\theta} - (1 - e^{-t/\theta}) \frac{\varphi}{A_u}. \quad (5.17)$$

Making use of the obvious initial condition $\delta(0) = 1$, this implies that

$$\delta(t) = e^{-t/\theta} - (1 - e^{-t/\theta}) \frac{\varphi}{A_u}. \quad (5.18)$$

Note that (5.16) and (5.18) are equal in the special case where $\varphi=0$ and $\delta_m = \exp(-\tau/\theta)$, as the reader may verify – this is no coincidence and has interesting interpretations, as we will see. The reader may verify in any preferred way that (5.18) is a solution of (5.17).

5.5.3 The 'deterrence time'

With the new discount function (5.18) it is no longer the case that waiting for however long is always better than giving up. At some time, which I shall call the 'deterrence time' and denote τ_d , payoffs that were initially positive become negative. If player T believes she will have to wait longer than the deterrence time, thus, it is better for her to do nothing and get a payoff of zero.

So how long can player T afford herself to wait and how long should she *expect* to have to wait until HEU arrives? The expected waiting time can be calculated by summation over all rounds, just like $\underline{\delta}$ was

32 The paradox disappears in the limit $\varphi = 0$, which has interesting implications, as we shall see. Notice also that, except when $\rho=1$, there are no problems in the limit $\tau \ll \theta$.

calculated in equation (5.3) before. The probability that HEU comes up precisely in round n is $\rho^{n-1}(1-\rho)$, and after waiting $n-1$ rounds, the time passed is $(n-1)\tau$, hence by summation over all n :

$$\begin{aligned} \tau &= \sum_{n=1}^{\infty} (n-1)\tau \rho^{n-1}(1-\rho) \\ &= \rho(1-\rho)\tau \frac{d}{d\rho} \sum_{n=1}^{\infty} \rho^n = \tau \frac{\rho}{1-\rho}. \end{aligned} \quad (5.19)$$

The expected waiting time is plotted in figure 5.3 as a function of ρ ; notice that it diverges to infinity as ρ approaches unity as one would expect intuitively.




Illustration 5.3.: Expected waiting time, equation (5.19).

The expected time can now be compared with the deterrence time. I set the right hand side of (5.18) equal to zero and solve with respect to t to get the expression

$$\tau_d = \theta \ln \left(1 + \frac{A_u}{\varphi} \right). \quad (5.20)$$

Note again how general this equation is, involving only the fear of failure, the expected utility if HEU were obtained right away and the penetration time, independently of the exact form of A_u . Equation (5.20) is visualised in figure 5.4.

Using this I can compare the expected waiting time for uranium with the maximum time player T can wait until it becomes preferable to do nothing. Expected waiting time is too long if $\tau \geq \tau_d$, that is when

$$\frac{\tau}{\theta} \geq \frac{1-\rho}{\rho} \ln \left(1 + \frac{A_u}{\varphi} \right). \quad (5.21)$$

The fraction on the left hand side one recognises as the expected number of time failures per round of waiting. Note that the waiting time is *always* too long

when ρ approaches unity (fraction on the right hand side goes to zero), and likewise *never* too long if φ is zero (the logarithm diverges), although notably the divergence as $\varphi \rightarrow 0$ is slow (logarithmic).




Illustration 5.4.: The deterrence time, equation (5.20).

Thus is uncovered the *true* condition for absolute deterrence to be possible in principle: that player T has *nothing to lose*. Only if the terrorist has absolutely no fear of failure is there no chance of deterring her. As long as φ is greater than zero, however, it means that there is something that player T holds dear that is imperilled by an unsuccessful attempt. This is an important theoretical finding³³: even a wildly fanatic terrorist like the one I have modelled, whose pleasure is proportional to the damage inflicted, may in principle be deterred, for while the individual terrorist footman prepared to blow himself up may be beyond persuasion, his *organisation* is not suicidal - if its support or very existence is at risk, attacking may in principle be worse than doing nothing. Given that the aspiring nuclear terrorist will always have the alternative of choosing other means, this conclusion is primarily of importance to conventional terrorism research.

I strongly indicate, thus, that it may be impossible to devise a model based on plausible assumptions in which a rational actor cannot, even in principle, be deterred from terrorist action. This conclusion was reached as a by-product of a game designed to analyse a different question, and the question of absolute deterrence of terrorists should be analysed again with a more suitable model; this task is interesting but too peripheral to our overall research question to fit into the thesis. The generality of the

33 It has been concluded by several other authors before by qualitative arguments; notable among them are Robert F. Trager and Dessislava Zagorcheva 'Deterring Terrorists: It Can Be Done' *International Security* 30:3 (2006) .

results which lead to this conclusion in the current setting however, leads me to the hypothesis that this conclusion is quite general.

If the deterrence time is shorter than τ , of course, none of the nuclear strategies has a positive payoff. This would be the ideal situation that Allison speaks of, where fissile material is guarded 'as closely as the gold in Fort Knox'³⁴. While the preferred strategy in this chapter's game would then be (N) the implication is of course not that the terrorist should simply surrender; in reality the terrorist's alternative to a nuclear strategy is not doing nothing, but to keep to her tried and trusted tools of car bombs and other conventional means. What it does show however is the (intuitively obvious) conclusion that if such Fort Knox security could be achieved worldwide it would achieve *relative* deterrence, deterrence from a particular *means* of attack rather than attack itself, from nuclear terrorism. The question of relative deterrence is treated in the next chapter, where it will be demonstrated by a numerical simulation that safeguards is one of several policy options to achieve a relative deterrence by creating conditions so that nuclear terrorism is not the preferred option even to the terrorist bent on maximum destruction.

5.5.4 Waiting or not: the choice between (O) and (U) revisited

Now that the concepts of expected waiting time, penetration time and deterrence time are introduced, it is time to make a brief revisit to the equation (5.9) which, in the old discounting system determined the deciding value of δ above which waiting for HEU, strategy (U) became preferable to accepting the first opportunity.

We derive the result in the same way (5.9) was derived before, by first working out an expression for the expectation value of the discount function before using this to compare $U_{(U)}$ to $U_{(O)}$. The expectation value of the discount function (5.18) is readily calculated just as in (5.3) by averaging over all rounds from 0 to infinity:

$$\begin{aligned} \langle \delta(t) \rangle &= \sum_{n=1}^{\infty} \delta[(n-1)\tau] \rho^{n-1} (1-\rho) \\ &= \frac{(1-\rho)(1+\varphi/A_u)}{1-\rho e^{-\tau/\theta}} - \frac{\varphi}{A_u} \end{aligned}$$

34 Graham Allison 'How to Stop Nuclear Terror' *Foreign Affairs* 83 (2004) pp.64-65.

Note how this reduces to the form (5.3) when $\varphi=0$ and $\delta_m = \exp(-\tau/\theta)$ as it should. Since $U_{(O)} = \rho A_p + (1-\rho)A_u$ we find that an opportunist strategy (O) is preferable to the uranium waiting strategy (U) if

$$\frac{\tau}{\theta} \geq \ln \left(\frac{U_{(O)} + \varphi}{A_p + \varphi} \right),$$

which can be written in the more instructive form

$$\frac{\tau}{\theta} \geq \ln \left(1 + (1-\rho) \frac{A_u - A_p}{A_p + \varphi} \right). \quad (5.22)$$

Observe that in the following limits this relation *guaranteed* to hold true, that is, opportunism is sure to be preferable to waiting for HEU as

- $\rho \rightarrow 1$: The probability of Pu is overwhelmingly greater than for HEU;
- $\theta \rightarrow 0$: Extreme impatience, i.e. very short penetration time;
- $\tau \rightarrow \infty$: Fissile materials are becoming extremely hard to obtain at all;
- $A_u \rightarrow A_p$: The terrorist becomes indifferent between the two materials.

Note that in all of these limits the corresponding side of the inequality approaches zero or infinity in a linear fashion³⁵. Likewise, waiting for HEU is guaranteed to be preferable to opportunism when

- $A_u \rightarrow \infty$: The payoff from an HEU attack is so large that it overwhelms all other concerns (provided the same is not true of plutonium)
- $A_p \rightarrow -\varphi$: A project using plutonium is expected to imply almost certain failure.

In both of these 'waiting limits' the right hand side of inequality (5.22) diverges to infinity logarithmically, much slower than the linear tendency in the 'opportunism limits'. The difference, which is inherent in the fact that time failure incidents are distributed according to the Poisson distribution, means that it is much easier to change the terrorist's preferences towards opportunism than towards patience. Roughly put: the terrorist with a fear of failure comes with an inherent affinity for opportunism.

To be sure, all of these limits could have been arrived at by intuitive arguments alone. However, equation (5.22) contains much more information than these limits, about how these parameters *interact*, when deciding which strategy is the more opportune, also when they are not close to these limiting cases.

35 One may see this by performing a first-order Taylor expansion in the relevant small parameter.

Of special cases one could note the asymptote where $A_u \gg A_p$, in which case (5.22) reads

$$\frac{\tau}{\theta} \geq \ln \left(1 + \frac{1-\rho}{1+\varphi/A_u} \right); A_u \gg A_p .$$

This is not so unlikely a case in light of our discussions in chapter 3. If in addition the terrorist has a relatively small fear of failure, $A_u \gg \varphi$, the condition for opportunism to be preferable becomes payoff independent:

$$\frac{\tau}{\theta} \geq \ln(2-\rho); A_u \gg A_p, \varphi .$$

In the opposite case is the risk averse terrorist for whom $\varphi \rightarrow \infty$. In this case one finds, using the definitions of A_u and A_p , that (O) is preferable to (U) if

$$\frac{\tau}{\theta} \geq \ln \left[1 + (1-\rho) \left(\frac{q_u}{q_p} - 1 \right) \right]; A_u \gg A_p .$$

Now the right hand side is dictated by the probability of operational success in the two cases.

5.5.5 Limits to terrorist intelligence?

We defined above that the terrorist's rationality lies in her consistency in choosing her best option, and her intelligence equivalent to her ability to work out what the best option is. While many will agree that terrorist masterminds may adhere to some notion of rationality, the above analysis is somewhat mathematically complicated and it may seem unrealistic to assume the terrorist will actually perform such an analysis prior to making her decisions.

There are, this author believes, three reasons why the analysis as performed is valuable, and should be calculated assuming no restrictions on the terrorist's intelligence (the way such is defined).

The first is of a philosophical nature. The problem pertains to much of game theory, a theory which often becomes much more mathematically complex than what reasonably reflects the way decisions are actually made, and several attempts have been made to mitigate it. Examples of models with restrictions on either rationality or intelligence are the field of bounded rationality³⁶. The problem with such models

is that they do not necessarily help our understanding, but rather the introduction of some limitation of intelligence into the model is another *ad hoc* assumption whose consequences must be kept track of.

A second is that the problem that the consequences of ideal rationality and intelligence may be unrealistic is one which becomes important primarily in the transition from the normative to the descriptive domain³⁷. As long as results are interpreted as normative (how *should* terrorists think), there is no problem, and the corrections which may be necessary in the transition to the descriptive question (how *do* they think) is a separate question which may be approached in several ways. This transition lies beyond the scope of this thesis, but could be a very fruitful continuation of the present chapter.

The last and perhaps most compelling reason is that while the final equations drawn from modelling may be of a mathematical nature, they should aim to give rise to intuitively graspable conclusions which make sense independently of the precise form of the model and parameters chosen. If this cannot be achieved, the modelling has not been very successful whereas if it is achieved, one is lead to believe that the same conclusions could, at least in principle, be arrived at by a qualitative argument. Having reached these conclusions by a mathematical method, however, has important bonuses because it explicitly allows the analyst to keep track of all relevant parameters throughout the argument, and the equations arrived at in the end can be varied with respect to different parameters (as exemplified by the graphs in figures 5.2 through 5.4), potentially allowing the extraction of much more information than a single intuitive main conclusion.

5.6 Fissile material in small batches: the hodgepodge bomb

In the modelling above I have assumed that all the necessary fissile material can be obtained in one batch. This is of course not the only possible scenario: it may be just as likely that the terrorists will need to obtain fissile material from several places in order to acquire enough for a bomb.

While not ideal, it is not a problem if the HEU used

36 See e.g. Hugh Ward 'Rational Choice' in D. Marsh and G. Stoker (eds) *Theory and Methods in Political Science* 2nd edition (Hampshire:Palgrave MacMillan, 2002) pp.65-89.

37 Amos Tversky and Daniel Kahneman 'Rational Choice and the Framing of Decisions' in Schweers Cook and Levi *The Limits of Rationality* (Chicago: The University of Chicago Press, 1990) pp. 60-89

in a gun consists of material with different degrees of enrichment, so long as the bomb-maker is able to measure the enrichment of each batch so as to calculate the critical mass correctly. If the material comes from a black market broker, it is almost certainly desirable for the terrorist to ensure the material is genuine in any case, and whenever the origin of the material is in any doubt³⁸, measurement of radiation, in particular neutron and alpha radiation³⁹, from the acquired lumps of HEU or Pu should be performed anyway to avoid criticality accidents and failed explosions. Fresh HEU fuel from submarines or research reactors⁴⁰, for example, could be used in a gun design⁴¹, but calculations of critical mass and safety precautions may have to be modified compared to the case of fresh fuel.

Indeed, uranium and plutonium could be used together in the same weapon and while nuclear weapon details are classified, there is no reason to believe this would not be the case in military weapons in states such as the US making military use of both materials. This raises the natural question: why not accept every offer of small batches of both HEU and plutonium until there is enough combined and try to make it explode one way or another? As explained in chapter 3.4 HEU and plutonium pose very different challenges to the bomb maker, and designing a reasonably reliable 'hodgepodge bomb', at least one with a several kilotonne yield, will likely be significantly more difficult than the HEU gun. For one, estimating the critical mass of such a jumble theoretically will be challenging, probably necessitating potentially dangerous experiments.

Furthermore, as Levi points out, the higher the number of purchases the more vulnerable the terrorist becomes to stings 'in which government agents pose as sellers. Repeated transactions would also extend the group's plot over longer timelines, increasing its chances of being detected by law enforcement or

intelligence.⁴² This is just the concept of time failure described above. Sting operations as a disruptive means is further discussed in chapter 6.

It is therefore far preferable, but not necessary, for the terrorist bomb maker to acquire all the fissile material from the same source. A closer examination of this would be an interesting continuation of the work presented herein.

5.7 The gains and weaknesses of gaming

There is an obvious weakness to the whole of the analysis presented, namely that all parameters are as estimated by the terrorist, hence not available to the analyst. This weakness is not primarily one of methodology, however, but a fundamental problem in answering the research question at all as discussed early in this chapter: one simply does not know how terrorists really think.

It is an obvious gain from using a gaming methodology based on an assumption of rationality and intelligence is then that I am able to derive a number of intuitive insights based on only a few assumptions. It may or may not be a good representation of the decision making of a given, real terrorist group, but it is a benchmark. The rational and intelligent assumption was discussed in section 5.5.5 above.

Another obvious gain from formalising the question in terms of a game is that it was a help in conceptualising the question and structuring it in an intuitive way. This point is somewhat subjective; it is always possible that another analyst could have done equally well in finding useful concepts and framework for thinking about the decision in question. Regardless of whether or not this is so, it is certainly the case that the devising and analysis of a game proved a functional way of producing a framework and concepts.

As mentioned in the previous chapter, after laying out all the pros and cons in a matter, reaching a conclusion means finding a way to weigh one against the other. This chapter is another example of this.

Similarly, in the limits listed in section 5.5.4, for example, it is intuitively obvious what the conclusion would be for the terrorist's choice. These could all have been laid out qualitatively, almost trivially. The achievement of the game above and beyond this list,

38 Most probably, such measurement will be deemed necessary even if the material is of certain origin.

39 Alpha radiation can produce neutron radiation indirectly in the vicinity of light elements.

40 Research reactors may also contain lightly irradiated fuel whose radioactivity could be small enough to be useful to terrorists as well.

41 See, for naval fuel: Morten Bremer Maerli 'Timely Options for Increased Transparency and Nonintrusive Verification on Fresh Highly Enriched Uranium Naval Fuel' *Journal of Nuclear Material Management* 31:4 (2003) pp. 18-30. For research reactors: Alexander Glaser 'On the Proliferation Potential of Uranium Fuel for Research Reactors at Various Enrichment Levels' *Science and Global Security* 14 (2006) pp. 1-24

42 Michael Levi *On Nuclear Terrorism* (Cambridge, MA: Harvard University Press, 2007) p.29

however, is twofold. Firstly, the inequality (5.22) is valid also far from these limits when no single concern of the terrorist's dominates all the other. Indeed, the intuitive limits could be seen merely as reality checks of the full result. It is when the system is not close to these limits that the weighing of one strategy against another is difficult and may require a quantitative methodology. Secondly, inspection of inequality (5.22) *immediately* produces such a list, and the analyst can have some confidence (equalling his confidence in the reality of the model itself) that these limits are *all* the limits there are, in which one strategy is sure to be trivially preferred to another. Similar things could be said about other formulae in the chapter.

Moreover, the mathematical form of (5.22) shows that tilting a parameter in the direction of opportunism can be expected to have a greater effect than tilting it towards waiting, because of the order of divergence in the limits presented. Whether it would be possible to foretell this with qualitative arguments is difficult to say, but with the present methodology it is a conclusion which is immediately visible to the trained eye in inequality (5.22).

When making assumptions about how a terrorist may think, as we do in this chapter as well as chapters 6 and 7, one is again susceptible to the information problem as discussed in the previous chapter. Herein I have assumed that virtually nothing is known about the terrorist's preferences, and that they conform with those most experts have arrived at, that HEU has a higher probability of success *cetera paribus* than plutonium. For a government facing a particular terrorist group, however, information might emerge which sheds more light on the real preferences of the terrorist, for example if it is known that the organisation has been on the lookout for plutonium in particular, say. Some such information could be easily incorporated into the current model by simply changing the values of the variables, but some kinds of information might necessitate that a different game be devised.

Likewise, the assumption that the terrorist organisation is a single deciding body may become questionable given more information from inside a particular organisation's decision making circles. The particular simplification of assuming an organisation or government to be representable by a single rational mind was discussed in chapter 2. As discussed there, the arrival of such information could let a government or intelligence agency shift its perspective somewhat,

from one which is completely external to the terrorist group's 'mind' to one where some details of the real decision making becomes visible. Should such be the case, the intelligence agency might do better to try and incorporate the new knowledge into their understanding of the thinking of this particular terrorist group, and perhaps abandon the general model presented herein.

5.8 Conclusion: policy implications

I have found that a terrorist's choice between HEU and Pu as fuel for his or her nuclear weapons project depends primarily on two key parameters, both of which are in a government adversary's power to influence. Given that the terrorist shares our perception that a project to make a true nuclear weapon using plutonium is expected to wreak less havoc than is a project whose end product is a HEU gun-type weapon, at least one of the following criteria must be fulfilled for the terrorist to decide to accept the first offer of fissile material even if it is plutonium:

1. The terrorist is highly impatient (low value of δ in a multiplicative discounting system or a short perceived penetration time θ with the discount function formalism of equation (5.22))
2. The terrorist perceives that there is a much greater chance that the next available material will be plutonium than uranium (that is ρ close to 1)
3. The terrorist believes it will be very difficult to obtain any fissile materials at all, and that the waiting time between each opportunity will accordingly be very long.
4. The terrorist is for some reason indifferent between HEU and Pu.

The exact criterion for a preference of one strategy over another is given in (5.22).

These results are in no way counterintuitive, which is reassuring. However, an equation like (5.9) or (5.22) contains both more and more precise information than the enumerated points above, which could both have been derived by qualitative arguments alone.

Point number 1 above, terrorist impatience, forms an important part of the policy implications of the next chapter, where relative deterrence of terrorists from a nuclear onto a conventional path is discussed. While there is some reason to assume from our previous comparison of the two fissile elements that it would be preferable for a government to seek to

create an environment in which the terrorist could rationally opt for a plutonium solution, making the terrorist wait forever for HEU might be even better. The reader should refer to the following chapter for a more complete discussion of the potential of disruptive means to increase terrorist impatience.

Points 2 and 3 may be addressed through safeguards measures as discussed in previous chapters, especially by 'putting HEU first'⁴³. The policy of giving HEU priority in safeguards and stockpile reduction activities is supported by the analysis of this chapter as well as the previous and following chapters. As in chapter 4 the conclusion can be traced back to a qualitative assumption that HEU is preferred over Pu.

All in all the assumptions employed to reach the above conclusions are few and plausible: apart from the way time and moves are modelled, shown in figure 5.1, there are three main assumptions upon which the model and subsequent analysis rest:

- The terrorist has a preference for HEU over plutonium. If one accepts the arguments laid out in the previous two chapters expressing this view, this assumption is implied by the assumption that the terrorist be an intelligent player. Even if the preference is slight, this assumption implies that the strategy 'wait until plutonium is available' is always inferior to either the strategy 'wait until uranium is available' or 'take whatever is available first'.
- Some assumption has been made about how the terrorist discounts future utility; two options were explored herein, a standard multiplicative model and a model where discounting is linked directly to the fear of failure.
- The necessary quantity of fissile material for a weapon is acquired in its entirety from a single source.

In fact, the analysis has come up with more findings than it set out to achieve. When I tried to take a closer look at the real-life interpretation of discounting, it was discovered that my system, a standard model from the economic literature which would ensure that an attack by either means was always preferable to doing nothing, was flawed. There appears to be no plausible way to model a rational *and* absolutely non-deterrable terrorist, excepting the unlikely case where the terrorist has

absolutely no fear of failure. This finding is primarily of theoretical interest since the deterrence situation discussed (where the alternatives are nuclear attack or no attack) is unrealistic and a real scenario involves many other options. It is, however, a suitable prelude to the next chapter, and supports conclusions arrived at by other authors⁴⁴.

While no obvious conclusion has been drawn as to what is the preferable direction for a government to influence the terrorist's choice of fissile materials, the analysis shows that to the extent that it is in the government's power to change the terrorist's assessment of the probability ρ and the time τ , T's calculus is expected to change as well, in ways expected to adhere qualitatively to the relations derived in this chapter. It is a main conclusion of chapters 4 and 6⁴⁵ that safeguards should be strengthened in general, and in chapter 4, moreover, I argue that measures to improve security for HEU should be given priority over those for plutonium. The findings in this chapter and the next analyse how implementing this can change the terrorist's calculus, to the extent that real efforts by the government result in an updating of the terrorist's estimates of relevant parameters.

43 Ferguson and Potter *The Four Faces ...* p.325.

44 Such as Trager and Zagorcheva 'Deterring Terrorists'

45 And also appendix D.

Deterring terrorists from attempted use of nuclear weapons

It is a common notion that deterrence is ineffective against terrorist adversaries¹. Religious terrorists in particular, amongst whom al Qaeda and its accomplices are notorious, are said to be beyond reasoning with. The US National Security Strategy states that 'Traditional concepts of deterrence will not work against a terrorist enemy'² and Brian Michael Jenkins, a world leading expert on terrorism, concludes that 'The Al Qaeda enterprise cannot easily be deterred'³. In a notable speech to the graduates of the West Point Military Academy, President George W. Bush asserted that in his view '[d]eterrence – the promise of massive retaliation against nations – means nothing against shadowy terrorist networks with no nation or citizens to defend.'⁴

Typically, literature discussing terrorist deterrence explores the potential for persuading terrorists to refrain from attacks on the whole⁵. Our own notion of deterrence is of a different nature. In the model of the previous chapter the only alternative to a nuclear path for the terrorist was doing nothing at all. In a more realistic scenario, however, terrorists contemplating embarking on a nuclear project will not compare it to

giving up but to continuing to do what they know best: executing conventional attacks which have already proven to be effective and for which infrastructure may already be in place. This makes for the possibility of a *relative* deterrence from certain *means* of terrorism rather than *absolute* deterrence from terrorist acts on the whole.

In the bipolar world of the Cold War, deterrence theory became a discipline in its own right and to many the word still bears connotations such as the logic of "Mutually Assured Destruction". An aspect of the classical sense of deterrence was that each side adhered to agreed limits of violence and that consequences for non-compliance were well understood; and the 'carrot' for compliance was a continued and relatively peaceful co-existence⁶. Deterring terrorists is a whole different matter not only because of the strong asymmetry of the situation⁷. There is no acceptable limit to terrorist violence other than zero, and co-existence is no incentive to either side. Consequences of non-compliance, moreover, are unpredictable since terrorists have no single 'return address' at which retaliation may be directed.

A primary concern is to be concise about what exactly is meant by 'deterrence'. I shall be fairly inclusive in my understanding of the term here and use Schelling's famous definition: 'persuading a potential enemy that he should in his own interest avoid certain courses of activity'⁸. Where e.g. Jenkins and Davis distinguish between deterrence, influence, co-optation, inducement, dissuasion and persuasion⁹ I will take a simpler approach. Note that this definition is much broader than the narrow definition employed in the quote by President Bush at the start of the chapter (restricted to threats of retaliation against nations).

The property that distinguishes deterrence, thus, is that deterrence measures affect the terrorist psychologically, actively or passively, at the time of *deciding* on a strategy, as opposed to measures that physically disrupt or block the actual execution of such a strategy, although some government actions

1 e.g. Brian O. Lesser et al. *Countering the New Terrorism* (Santa Monica: RAND, 1999)* p. 36, Bruno S. Frey and Simon Luechinger 'Decentralization as a disincentive for terror' *European Journal of Political Economy* 20 (2004) pp. 509-515.

2 President George W. Bush *The National Security Strategy of the United States of America* (Washington DC, 2002)* p.15

3 Brian Michael Jenkins *Countering Al Qaeda: An Appreciation of the Situation and Suggestions for Strategy* (RAND, 2002)* p.17. Note that Jenkins refers to deterrence in a narrower sense than that adopted in this chapter.

4 George W. Bush, Speech at the United States Military Academy, West Point, New York (June 1, 2002)*

5 The most thorough analysis may be that in Paul K. Davis and Brian Michael Jenkins *Deterrence & Influence in Counterterrorism: A Component in the War on al Qaeda* (RAND, 2002)* p.xi-xii. Another report with similar conclusions is Neil J. Smelser and Faith Mitchell (eds.) *Discouraging Terrorism: Some Implications of 9/11* report from the Panel on Understanding Terrorists in Order to Deter Terrorism, Center for Social and Economic Studies, National Research Council (Washington D.C.: National Academies Press, 2002)

6 Jenkins *Countering Al Qaeda* p.25.

7 e.g. Wyn Q. Bowen 'Deterrence and Asymmetry: Non-State Actors and Mass Casualty Terrorism' *Contemporary Security Policy* 25:1 (2004) p.58; 'Deterring Mass-Casualty Terrorism' *Joint Force Quarterly* 31 (Summer 2002)* pp.25-29

8 Thomas Schelling *The Strategy of Conflict* (Cambridge, MA:Harvard University Press, 1960) p.9.

9 Davis and Jenkins *Deterrence & Influence in Counterterrorism* p.xi-xii

can be both as will be argued later. A further defining feature is that the terrorist realises that refraining from a certain course of action is in her own interest, or to use the nomenclature of rational choice theory: that a conviction is conveyed that a nuclear strategy does not represent a maximisation of the terrorist's utility function.

Two mechanisms of deterrence are typically recognised: deterrence by threat and by denial¹⁰, the former involving a credible promise of retaliation and the latter sending the message that a plan of action is too unlikely to succeed to be worth the money and effort. I will argue that for the purpose of relative deterrence of terrorists (as opposed to their potential sponsors), the latter is by far the more important.

In keeping with Schelling's definition, one furthermore distinguishes between deterrence and coercion, the former seeking to preserve the *status quo* (persuading an actor to continue not doing something), the latter to change it (persuading an actor to stop doing something)¹¹. *Status quo* in this case being the situation where no terrorist nuclear weapon programme has yet left its starting block, deterrence is the relevant term; this could change should it emerge that a project of this sort is underway. A number of coercion strategies could be considered in such a case but I shall not focus on coercion in this chapter.

6.1 Research question and scope

This chapter's research question is

Is it possible and if so under which circumstances and by which means, for a target government to deter a terrorist organisation from attempted acquisition of nuclear weapons?

Our scope is limited by a set of assumptions about terrorist reasoning and available means as will be specified below. Thus the research presented is anything but exhaustive, but adds another perspective to the very scarce work already done by others.

As in previous chapters, I shall consider a nuclear terrorism project in which a terrorist group builds its own device from stolen or illicitly bought fissile material. While the theft or purchase of an intact device from a country's arsenal falls within the definition of nuclear terrorism in section 1.2, I exclude it from the analysis in this chapter, primarily for

10 e.g. Bowen 'Deterrence and Asymmetry'

11 *ibid.*

reasons of manageability. The comparison of two possible strategies is useful even if not all possibilities are taken into account¹².

The research question speaks of a generic terrorist organisation, but I shall use al Qaeda extensively as example, causing no significant loss of generality or topicality in doing so.

6.2 Chapter outline

I start by surveying briefly the literature on terrorist deterrence. First the scarce literature on deterrence of nuclear terrorism specifically is discussed where the emphasis on the possibility of state sponsorship is discussed. Thereafter the literature on deterrence of nuclear terrorism in general is discussed in connection with the research question at hand. The fundamental difference between relative and absolute deterrence of terrorists is discussed.

We thereafter devise a model to analyse the terrorist's choice between conventional and nuclear means of attack. The model is then used for a numerical simulation in which particularly a plot of the utility functions of nuclear and conventional utility functions as a function of terrorist impatience yields important insights into the dynamics of the rational terrorist's calculus when faced with this choice. From this a discussion of possible political means of relative deterrence follows from these, along with the policy conclusions of this chapter.

6.3 The literature on terrorism and deterrence

Surprisingly little has been done in the field of terrorism and deterrence in general, and almost nothing on the special issue of deterrence of particular means of terror by addressing the cost-benefit calculation of the terrorists themselves¹³. Nearly all

12 Of course, there are the possibility of chemical, biological and radiological means as well, not treated here. A comparison of the four using a similar methodology as this chapter might be a useful extension for future study. One could hardly ever justify the claim that all possible strategies are included, however; for example no-one had foreseen, to the author's knowledge, that passenger aircraft would be used as piloted missiles to attack buildings. A useful overview of the threats one has thought about: is Richard L. Garwin 'The Many Threats of Terror' *The New York Review of Books* (November 1st, 2001) and its epilogue (see bibliography)*

13 Exceptions include Bowen 'Deterrence and Asymmetry' and Lewis A. Dunn 'Can al Qaeda Be Deterred from Using Nuclear Weapons?' *Occasional paper #3* (Center for the Study of Weapons of Mass Destruction, National Defense

the literature on deterrence of terrorism has concentrated on deterring states from actively supporting terrorist nuclear missions. I will argue that the discourse on deterrence of terrorism in general poses fundamentally different challenges than relative deterrence as discussed here. Indeed, I find that the two often contradict each other and attempts at a unification of the two (by formal methods or otherwise) could form a fertile field of study for the future.

6.3.1 Deterrence of state sponsorship

There exists a number of papers considering the indirect deterrence by deterring state actors from sponsoring a nuclear terrorism project¹⁴. This idea comes up quite naturally if one accepts the two notions that terrorists cannot be deterred directly and that a nuclear-armed state sponsor is the easiest way for a terrorist organisation to acquire a nuclear weapon. As discussed in chapter 1 and 3, state sponsorship is arguably not necessary for nuclear terrorism; it is far from unthinkable that a terrorist group can obtain all it needs for a small-scale proliferation project without governmental help. Moreover, terrorist leaders have demonstrated definite rational capacity, and as such could be deterred. While the discussion of state sponsorship is one that will not be treated in detail in this thesis, it is such an integrated part of the literature on deterrence of nuclear terrorism that it deserves a brief discussion at this point.

A common notion is that state sponsored nuclear terrorism can be deterred by way of attribution¹⁵. By analysing the debris and fallout following a nuclear explosion, experts can in principle work out where the material came from since a particular history of enrichment, conversion and storage of nuclear material leaves specific traces of isotopic composition and chemical properties. If potential sponsors know

that the finger of blame can be pointed with confidence soon after the attack, the argument goes, and retaliatory attacks (perhaps even nuclear¹⁶) were promised in such an instance, this could deter states from helping terrorists develop nuclear weapons in the first place. To establish the nuclear forensics capacity the preferable route, some argue, is to create an international database of fissile material from every possible source worldwide¹⁷. A suggested way to ensure that all countries supply samples of their fissile materials is through a 'Prove Innocence Treaty'¹⁸. Other commentators are more careful in their optimism, yet do not object to the feasibility of deterrence of state actors in principle¹⁹.

There are at least two potential problems with the strategy of deterrence by attribution. The first is of a practical nature: the attribution needs to be perceived as reliably correct by the parties that are to be deterred. Today this seems to be far from the case, as demonstrated by the erroneous conclusion published in 2005 by US analysts at the Oak Ridge National Laboratory that the uranium hexafluoride surrendered by Libya was of North Korean origin 'with a certainty of 90 percent or better'²⁰. An independent study by the International Atomic Energy Agency in Vienna was inconclusive and after a much publicised row in which it was suggested that the original statement was politically motivated, it was concluded that the material was originally Pakistani²¹.

More importantly perhaps, knowing that the material used in a nuclear weapon came from a specific source does not prove that the state in question actively presented the material to the terrorists. The accused state will likely claim to have donated the material unknowingly and unwillingly, and it is very difficult to prove that this is not so; the material could have been obtained by theft or by bribing low-level employees. Even if documentation

University, 2005)*

- 14 Recent propositions for deterrence of state sponsorship and other support mechanisms are Caitlin Talmadge 'Deterring a Nuclear 9/11' *The Washington Quarterly* 30:2 (2007) pp. 21-34 and Daniel Whitenack 'Deterring Terrorists: Thoughts on a Framework' *The Washington Quarterly* 28:3 (2005). Other optimists are Dunlop and Smith, who argue for the establishment of an international nuclear forensics team. William Dunlop and Harold Smith 'Who Did It? Using International Forensics to Detect and Deter Nuclear Terrorism' *Arms Control Today* (October 2006)* pp.6-10
- 15 e.g. Michael A. Levi 'Deterring Nuclear Terrorism' *Issues in Science and Technology Online* (Spring 2004)*

- 16 As suggested by Robert L. Gallucci 'Averting Nuclear Catastrophe: Contemplating Extreme Responses to U.S. Vulnerability' *The Annals of the AAPSS* 607 (2006) pp.51-58
- 17 e.g. Michael May, Jay Davis and Raymond Jeanloz 'Preparing for the Worst' *Nature* 443 (26 October 2006) pp.907-908.
- 18 Hans Binnendijk and Peter D. Zimmerman 'New nuclear deterrents' *The Washington Times* op. ed. (August 19, 2007)*
- 19 Matthew Phillips 'Uncertain Justice for Nuclear Terror: Deterrence of Anonymous Attacks Through Attribution' *Orbis* 51:3 (2007) pp.429-446 and Michael Miller 'Nuclear Attribution as Deterrence' *Nonproliferation Review* 14:1 (2007) pp.33-60
- 20 Phillips 'Uncertain Justice...' p.434
- 21 *ibid.*

can be found that senior officials were explicitly involved, the state can always claim that the person or persons in question acted without permission.

Should a country wish to donate fissile material or other help to a terrorist organisation, it is not difficult to ensure such layers of deniability, and it seems likely that a regime interested in its own survival will do so. This is clearly demonstrated by the case of A. Q. Khan, the head of a division developing the Pakistani nuclear weapons, who set up his own clandestine proliferation network exporting nuclear technology to such states as Libya, Iran and North Korea. Although there has been much speculation that the Pakistani government must have been aware of these activities, proving such claims turned out to be very difficult²². A tacit handover of some tens of kilograms of nuclear material can be done much less spectacularly than the A. Q. Khan deals, and could easily be made to look like the independent actions of an unfaithful servant, one would think.

Whether a country can avoid a retaliatory attack in this manner might depend on the country in question. If the country has a history of suspect nuclear activities outside the Non-proliferation Treaty (or at least arguably so), the threat of retaliation on grounds of suspicion alone can be credible enough to deter terrorist sponsorship, argues Jenkins²³.

Nonetheless, the US National Strategy for Combating Terrorism includes a deterrence stance based on the pillars of attribution and subsequent retaliation:

A new deterrence calculus combines the need to deter terrorists and supporters from contemplating a WMD attack We require a range of deterrence strategies that are tailored to the situation and the adversary. We will make clear that terrorists and those who aid or sponsor a WMD attack would face the prospect of an overwhelming response to any use of such weapons Finally we will ensure that our capacity to determine the source of any such attack is well-known, and that our determination to respond overwhelmingly to any attack is never in doubt.²⁴

22 See Gordon Corraer *Shopping For Bombs* (London: Hurst & co., 2006)

23 Brian Michael Jenkins *Unconquerable Nation: Knowing our Enemy, Strengthening Ourselves* (Santa Monica, CA: Rand, 2006)* pp. 138-141

24 National Security Council (US) *National Strategy for Combating*

While the Strategy also describes the determination to deny terrorists access to necessary materials through safeguards and disruption, the idea of deterrence by denial is not mentioned. It was recently reported that notions of deterrence by punishment are seriously contemplated amongst senior strategic planners of the Bush administration²⁵. (It is notable that this quote from the National Strategy for Combating Terrorism is somewhat at odds with the quote in the introduction to this chapter, from the National Security Strategy, where it was stated that deterrence is not useful against terrorism.)

Notably, there are other very good reasons to develop a good capability to recognise the source of a nuclear weapon after its detonation. As pointed out by Davis²⁶ as well as Dunlop and Smith²⁷, in the political frenzy that will surely follow a nuclear terrorist explosion, it is of vital importance to establish who *didn't* do it to avoid rushed conclusions and hasty allegations which can have very grave consequences.

I will argue in the following that while deterring state sponsors is perhaps more difficult than some believe and only a partial solution, relative deterrence still has an important rôle to play. Assuming the terrorist is a rational cost-benefit calculator, it is both possible and beneficial to try and deter the terrorist from a nuclear path into staying with the tried and trusted conventional means, since the tools of deterrence are to a great extent tools of denial which will help lowering the probability of success of a nuclear terrorism project even should deterrence fail.

6.3.2 Terrorist deterrence in general and the use of formal methodology

The problem of no return address, mentioned above, is one of three problems of terrorist deterrence countered by Trager and Zagorcheva, who make the case for absolute deterrence of terrorism²⁸. Allegedly, terrorists are moreover held to be 'irrational' hence unresponsive to cost-benefit calculations, and finally

Terrorism (Washington D.C., September 2006)* p.14

25 David E. Sanger and Thom Shanker 'U.S. Debates Deterrence for Nuclear Terrorism' *The New York Times* (May 8 2007)*

26 Jay Davis 'The Attribution of WMD Events' *Journal of Homeland Security* (April 2003)*.

27 Dunlop and Smith 'Who Did It?'

28 Robert F. Trager and Dessimslava P. Zagorcheva 'Deterring Terrorism: It Can Be Done' *International Security* 30:3 (2005) pp.87-123.

the motivation of terrorists willing to die for their cause must be so strong that deterrence is impossible. Pape summarises²⁹:

Although the capture and conviction of Timothy McVeigh³⁰ gave reason for some confidence that others with similar political views might be deterred, the deaths of the September 11 hijackers did not, because Americans would have to expect that future Al Qaeda attackers would be equally willing to die.

There should be little reason to doubt that Pape's point is at least partly valid: the individual suicide bomber can hardly be deterred by threats of violence against his person since he is already on a death mission. He might, however, worry about threats against his family, and the possibility of failure; says Lesser and co-workers: 'the terrorists themselves are often concerned by operational risk - they may be willing to risk or give their lives, but not in futile attacks.'³¹ One must distinguish between the bombers and the leaders who mastermind their missions. While footmen may be expendable tools, even terrorist organisations such as al Qaida have goals and resources that are precious to them and that may be held at risk³². As Davis and Jenkins point out, mission success is of great importance to leaders of terrorist groups³³, making for a possible risk-aversion that could give considerable leverage to policy makers, discussed in some detail below. Furthermore, al Qaida, for example, is not only comprised of its leaders and 'muscle' but also 'lieutenants, financiers, logisticians and other facilitators, ... recruiters, supporting population segments, and religious and otherwise ideological figures.'³⁴ Financiers of terrorism living the good life in a secular world may possess less zeal, and such peripheral players, Davis and Jenkins argue, might be easier to influence.

Deterrence as defined here requires that the deterred party be a cost-benefit calculator³⁵. I

concluded in my introductory chapter that it is reasonable to assume that terrorists may be seen as 'sufficiently rational' for modelling purposes, in that they 'usually have a set of hierarchically ordered goals and choose strategies that best advance them.'³⁶ This is the process modelled through the maximisation of a payoff function, forming the backbone of rational choice theory.

The literature treating formally the strategic countering of conventional terrorism is not large but is steadily growing, with professor Sandler still its centre of gravity³⁷. The threat from nuclear terrorism projects has peculiar properties, however, which the existing debate on general terrorist deterrence cannot cover.

6.3.3 *Relative deterrence of terrorists from non-conventional means*

The challenge of 'relative deterrence' is fundamentally different from that of deterring groups from terrorism tactics altogether. Its goal may be said to be intermediate or preliminary: to persuade attackers to stick with their tried and trusted methods, in themselves potentially very destructive. Some available means, such as disruption of finances, are common to both categories of deterrence, yet while the debate on general terrorist deterrence forms a useful backdrop, it is of limited use to my analysis here.

Given the large scholarly interest in the nexus of terrorism and non-conventional arms³⁸, surprisingly few have dealt with the particular strategic interplay between terrorists and antiterrorists when threats of mass casualty means are involved. This author is only aware of a handful of papers of direct relevance, which will be reviewed briefly, plus one by Sandler and Arce exploring the terrorist's decision to opt for 'spectacular' attacks.

On deterrence of non-state actors in general, a very

29 Robert A. Pape 'The Strategic Logic of Suicide Terrorism' *American Political Science Review* 97:3 (2003) p.347

30 perpetrator of the bombing of the Murrah Federal Building in Oklahoma City, 1995, killing 168.

31 Lesser et al. *Countering the New Terrorism* p. 16

32 Trager and Zagorcheva 'Deterring Terrorism' p.88.

33 Davis and Jenkins *Deterrence & Influence in Counterterrorism* p.xii.

34 *Ibid.* p.xi. See also Dunn 'Can al Qaeda Be Deterred ...'

35 Bowen 'Deterrence and Asymmetry' p.59.

36 Trager and Zagorcheva 'Deterring Terrorism' pp.93-94.

37 Recent works of interest: Todd Sandler and Kevin Siqueira 'Global terrorism: deterrence versus pre-emption' *Canadian Journal of Economics* 39:4 (2006) pp.1370-1387, Daniel G. Arce M. and Todd Sandler 'Counterterrorism: A Game-theoretic Analysis' *Journal of Conflict Resolution* 49:2 (2005) pp.183-200, and Kevin Siqueira and Todd Sandler 'Terrorists versus the Government: Strategic interactions, support and sponsorship' *Journal of Conflict Resolution* 50:6 (2006) pp.878-898.

38 See chapter 1 for an overview of the large body of literature on the subject.

thoughtful overview is given by Bowen³⁹. He lays out the concepts, strengths and limitations of terrorism deterrence, lending much theoretical support to the more applied effort presented here. Key aspects of asymmetry and its impact are discussed, and key questions to be answered in order to go from theory to best policy are presented. Bowen's paper thus forms a fundament upon which the present effort rests, its sole weakness being perchance the lack of a clear distinction between absolute and relative deterrence, which I find to be of importance. His paper provides a framework only, and rather than drawing definite conclusions, Bowen points to future directions of study, some of which I pursue here.

Franck and Melese, in an almost unique application of game theory to the strategic deliberations of terrorist acquisition of non-conventional means, makes a useful contribution to the understanding of terrorist's strategic choice of whether or not to embark on non-conventional projects⁴⁰. The authors assume terrorists are primarily motivated by a will to impress their audience (and not alienate them) from which they arrive at some conclusions which may be described as somewhat vague; that fanatical terrorists pose a greater 'WMD threat' than political ones, but that the latter is also worrisome; that terrorist 'WMD decisions' can be influenced; and that 'while the choice of countermeasures to reduce effectiveness of an attack is important, reducing the probability [of a 'WMD attack'] impressing the audience is also important'⁴¹.

Some of the weaknesses of their paper (concerning rather unrealistic assumptions) the authors themselves recognise⁴². The analysis is based on relatively restrictive assumptions, not in itself problematic, but some of the conclusions (such as those mentioned above) are all but corollaries of the assumptions employed, something the authors do not fully appreciate. For example, given the modelling assumption that disruptive measures are effective against non-conventional but not against conventional terrorist projects⁴³, it is an unsurprising conclusion that a target country should choose a disruptive course of action if it knows the terrorist adversary has

'WMD'⁴⁴. This is a general potential problem with game theory as discussed in chapter 2: A model will never give fundamentally new information beyond the assumptions that go into it, but can be a powerful tool to visualise and analyse the consequences of these assumptions. When the conclusions that result could have been reached equally well without the detour of modelling, however, the model serves no purpose other than maybe to complicate. The graphs and diagrams of the paper are useful, but considering the complexity of the model involved, it is my opinion that the conclusions are disappointingly self-evident from the model itself⁴⁵.

Franck and Melese's modelling effort, however, is a useful starting point and commendable for attempting what, mysteriously, nobody else seems to have tried before to the author's knowledge: to model formally the economy of terrorists' employment of non-conventional means.

A somewhat similar analysis is performed by Sandler and Arce⁴⁶, who, like Franck and Melese analyse the incentives of terrorists to choose 'spectacular' attack tactics. The model, in the form of a classic signalling game, is more useful than that of Franck and Melese in that it involves uncertainty as well as discounting and a model of terrorist financing. Distinguishing between 'large' and 'small' attack instead of the narrower but woolly 'WMD'⁴⁷ only makes the analysis clearer. The focus is on

44 *ibid.* p.366

45 The process of modelling implicitly involves asking the question 'what are the important factors that could influence this choice?'. A qualitative analysis answers this question and makes way for the resulting model. When the conclusions from the modelling in turn are of the form 'these are important factors', essentially equalling the factors that went into the model in the first place, it indicates that a simpler and more transparent methodology could probably have done the job.

46 Daniel G. Arce and Todd Sandler 'Terrorist Signalling and the Value of Intelligence' *British Journal of Political Science* 37:4 pp.573-586 (2007) (a preview was kindly provided by Professor Sandler prior to its publication).

47 Franck and Melese treat 'WMD' collectively (understood to mean nuclear, biological, chemical and radiological weapons in this context, although no definition is provided) without emphasising the differences in scale, challenge and consequences among the various projects that would fall under this very large umbrella. 'WMD' they say are very costly and cause large casualties. This is certainly true for a successful nuclear project, yet historically, non-conventional attacks have proven disappointing in terms of body count: Kate Ivanova and Todd Sandler 'CBRN Incidents: Political Regimes, Perpetrators and Targets' *Terrorism and Political Violence* 18 (2006) pp.423-448.

39 Bowen 'Deterrence and Asymmetry'

40 Raymond E. Franck and Francois Melese 'Exploring the Structure of Terrorists' WMD Decisions: A Game Theory Approach' *Defense & Security Analysis* 20:4 pp.355-372.

41 All *ibid.* p. 370

42 *ibid.* p.369

43 *ibid.* p.362

government response to a terrorist attack from an organisation whose objectives, resources and commitment are hitherto unknown, however, and hence lies somewhat to the side of the research question considered here. Somewhat simplified Sandler and Arce conclude that intelligence has become increasingly important in the new world order where political and militant terrorists co-exist, since the best response to a terror attack depends on the nature of the terrorist group. Although not directly relevant to this chapter, the game found in their paper is valuable for its clarity and is an inspiration for that found below.

A different approach is taken by Melese and Angelis, who propose a classical policy of deterrence by punishment of terrorists considering acquisition of 'WMD'⁴⁸. They suggest a 'brinkmanship strategy'⁴⁹, aided by the UN whose Secretary General would promise to 'allow' unilateral pre-emptive strikes from such states as the United States if non-conventional arms were acquired. Melese and Angelis seem to overlook a few important problems with their plan however, notably (but not only) the obvious 'return address problem' mentioned above which, peculiarly, they fail to mention. I shall not consider their proposition further.

Perhaps worth mention is an oft cited paper by Rapoport in which he argues that terrorists will effectively deter themselves from non-conventional means because they have historically been hard to master and rather ineffective⁵⁰. Interestingly, he seems to imply that terrorists choose their methods rationally using some sort of cost-benefit analysis, yet this assumption is not made explicit. Rapoport spends considerable effort lashing out at what he sees as

alarmism from 'those of the physical sciences'⁵¹, and although his arguments build vaguely on some assumption of terrorist rationality, no formal methodology is employed.

The most convincing treatment of whether al Qaeda may be deterred from the use of nuclear weapons might be that provided by Dunn⁵² whose points relevant to us go hand in hand with our formal analysis. Dunn's question is mainly that of deterrence from *use* of nuclear weapons after acquisition, hence his research question only partly overlaps with ours. Another key paper which forms an important guide through our analysis is Jackson's treatment of terrorist technology acquisition⁵³, drawing on experiences from private companies – a lucid discussion of the process of adopting new technology in general which I draw upon to an equal extent.

In conclusion, the literature on terrorism and deterrence is scarce and the literature on deterring the particular use of nuclear means is almost absent despite the large interest in nuclear terrorism in general. The highly enlightening papers by Bowen and Dunn, the former non-applied, the second purely qualitative, point the way to further enquiry. Franck and Melese begin to fill this space, yet their approach is (intentionally) limited in scope. My own approach is not much less limited in its assumptions, but different choices are made. Unlike Franck and Melese I will assume that the terrorist does not care overly much about her audience but simply wishes to cause as much destruction as possible with limited resources and under pressure. Obviously my analysis will then tell me little or nothing about the effect which the project has on the terrorist's audience – for a formal analysis of this the reader is referred to Franck and Melese – instead its conclusions concern the importance of other parameters which are more easily influenced by a government antiterrorist. It is the author's opinion that this makes for a stronger and more useful treatment.

48 A preliminary study is found in Francois Melese and Diana Angelis 'Deterring Terrorists from Using WMD: A Brinkmanship Strategy for the United Nations' *Defense & Security Analysis* 20:4 (2004) p.337-341, while the full analysis is found in a conference paper: Francois Melese 'A Brinkmanship Game Theory Model of Terrorism' presented at the 3rd Conference on Mathematical Methods of Counterterrorism, Washington D.C. (September 2006) (kindly provided to me by Prof. Melese).

49 Brinkmanship: 'The practice, especially in international politics, of seeking advantage by creating the impression that one is willing and able to push a highly dangerous situation to the limit rather than concede.' *The American Heritage Dictionary of the English Language* (Houghton Mifflin, 2006).

50 David C. Rapoport 'Terrorism and Weapons of the Apocalypse' *National Security Studies Quarterly* 5:3 (1999) pp.49-66. The paper only discusses chemical and biological means, apparently seeing nuclear means as too far fetched to mention.

51 *ibid.* p.51

52 Dunn 'Can al Qaeda Be Deterred...'

53 Brian A. Jackson 'Technology Acquisition by Terrorist Groups: Threat Assessment Informed by Lessons from Private Sector Technology Adoption' *Studies in Conflict & Terrorism* 24 (2001) pp.183-213. The paper was later followed by an extensive two-volume RAND report: Brian A. Jackson, John C. Baker, Peter Chalk, Kim Cragin, John V. Parachini, Horacio R. Trujillo *Aptitude for Destruction, Volume 1: Organizational Learning in Terrorist Groups and Its Implications for Combating Terrorism and Aptitude for Destruction, Volume 2: Case Studies of Organizational Learning in Five Terrorist Groups* (Santa Monica, CA: RAND, 2005)*

6.4 The model

As in the previous chapter I will assume a rational terrorist motivated by wreaking maximum havoc, presumably the most difficult type of terrorist to deter from nuclear weapons. One could argue that this is simplistic, yet the real strategic process behind a terrorist's choice of nuclear versus conventional means remains a 'crucial unknown'⁵⁴ and some assumption must be made. Precaution makes it reasonable, thus, to assume the 'worst case' for deterrence and be sure not to conclude on the feasibility of deterrence based on unjustified optimism.

6.4.1 Working hypotheses and assumptions

I shall work under the hypothesis that acquisition equals attempted use when it comes to terrorists and nuclear arms. For a state power, arguably the worst number of nuclear weapons to have is one – once that has been used, the targeted country has no incentive not to retaliate with full force. For a terrorist group without a 'return address', the situation is different since retaliation is less well defined. The United States have made it a pillar of their security strategy to 'make no concessions to terrorist demands and strike no deals with them'⁵⁵, meant to deter terrorism in general, but potentially backfiring somewhat for the purposes of relative deterrence, providing an incentive to terrorists to detonate an acquired weapon rather than try to gain political leverage through threats⁵⁶. The common assumption is therefore that al Qaida would surely detonate a nuclear weapon as soon as they got their hands on it, but Dunn concludes that 'it would be ill advised to reject out of hand the possibility that, for Osama bin Laden, nuclear weapons could be too valuable to

detonate'⁵⁷. My assumption is for simplicity and not to be interpreted as an off-hand dismissal of Dunn's conclusions. A detailed discussion of the strategic value of terrorist nuclear weapons is undertaken in chapter 7 of this thesis.

When testing whether deterrence is possible, it seems reasonable to make modelling assumptions that should intuitively make deterrence harder, not easier. One such assumption is that the terrorists are driven solely by the wish to inflict destruction. The author does in no way believe the motivations of the leaders of e.g. al Qaida are so simple. In many contexts, reducing the terrorist adversary to nihilists driven by bloodlust alone is misleading and counterproductive. However, the goal here is not to fully understand how terrorists motivate their outrageous actions – many authors have contributed expertly to such understanding⁵⁸. The question is rather: 'if maximum destruction is (for whatever reason) the preferred outcome for the terrorist, what options still remain?'. A working assumption here, as before, is therefore that terrorist utility from a successful attack is proportional to the damage inflicted.

6.4.2 The game

The game is laid out in illustration 6.1 in the form of a decision theoretical game. Squares denote choice nodes and circles denote chance nodes. Imagine that there is also a shadow antiterrorist player present, in whose power it is to shift some of the parameters of the game, thus taking part indirectly by 'setting the stage'.

In the first round of the game, the terrorist player (player T) can choose to plan and execute a conventional attack. For simplicity, let us assume that these attacks are on a smaller scale than the largest conventional attacks such as September 11, using means that are well known to the terrorist. Such an attack is therefore assumed to take a short time to prepare (planning and execution within one round), has a high probability p_c of success and a relatively small cost C_c to the terrorist. On the down side the

54 Joshua Sinai 'Forecasting Terrorists' Likelihood to Embark on "Conventional" to CBRN Warfare' *International Studies Review* 7:1 (2005) p.151

55 *The National Security Strategy of the United States of America* p.5

56 Holding beyond sway that no concessions will ever be made might deter a politically motivated terrorist who sees murder as a necessary evil, but for the terrorist in possession of terrible weapons, the message would imply that attack is the only option. Conversely, 'allowing' a certain level of small-attack terrorism could make the bloodthirsty terrorist content, dissuading her from seeking more destructive means. This is an interesting dynamic, suitable for further study but not treated in depth here. An investigation of terrorist nuclear blackmail is undertaken in chapter 7.

57 Dunn 'Can al Qaeda Be Deterred ... ' p.25.

58 Rather too much in some commentators' opinion: Avishag Gordon 'Terrorism as an Academic Subject after 9/11: Searching the Internet Reveals a *Stockholm Syndrome* Trend' *Studies in Conflict & Terrorism* 28 (2005) p.45-59. For an analysis of terrorist motivations to employ unconventional weapons, see Bruce Hoffman 'Terrorism and Weapons of Mass Destruction: An Analysis of Trends and Motivations' Rand Paper #8039-1 (1999) *

damage T_c inflicted on the antiterrorist is modest. Alternatively, T can choose to embark on a nuclear project. This project has a high cost C_n , takes a long time to plan and prepare (several rounds) and has a relatively small probability p_n of success. The award for the terrorist upon successful detonation, however, is enormous. The cost of failure, $-\varphi$, is assumed equal for the conventional and nuclear attacks, as an approximation.

The aftermath of the spectacular attacks on September 11 makes it reasonable to believe that such a huge attack as a nuclear detonation would be followed by a period of unrest worldwide. The perpetrators of the attack will be hunted and will most probably need to go into hiding like the al Qaida leadership did after the US invasion of Afghanistan in 2001. I therefore introduce a penalty period following a successful attack during which T is unable to commence another attack plan.

Like before I use discounting to model player T's impatience. As I argued in the previous chapter, there are numerous reasons why T would prefer to see results sooner rather than later. However, this time around I will *not* think of δ as the mathematical representation of the risk of time failure⁵⁹, but in a more abstract sense representing the terrorist's mindset. This is so as not to double count the risk of failure, both through p_n/p_c and a 'time failure' notion, and also because the same discounting is employed for the building period as during the penalty period after an attack, although the price of 'time failing' in the two periods are different. The risk of failure for the nuclear project is surely increased by the fact that it takes several rounds to prepare an attack, but this is assumed to be incorporated into the difference between p_n and p_c . Project failure is not the only thing the terrorist might be concerned about, however: If too much time passes without attack, followers might begin to doubt the credibility of terrorist leaders. Moreover, player T may anticipate that attacking might become more difficult with time, for example because the antiterrorist is taking steps to make nuclear weapon acquisition and use harder.

I will think of the terrorist player as keeping two different accounts: accumulated payoffs and 'real' funds. Per assumption real money has no value to the terrorist beyond the use for attacks, so doing nothing and saving the money for other purposes is not assumed to be an option. While the unit used to

measure payoffs for each round will be dollars, one should bear in mind that neither the utility T extracts from inflicted damage nor her grievance at failing a mission represent gains or losses of real money. Failing a mission could send a message of incompetence and jeopardise the perception amongst followers that 'our side is winning'. The additional failure cost, φ , may be thought of as the sum of money T would have been willing to pay to avoid being in such a situation. Thus, as the game progresses, I keep track of the real money (an attack project of either type may only be commenced if sufficient resources are available, otherwise T can do nothing and must wait for more money to arrive next round - the 'insufficient funds' option is not shown in the figure) but the goal is optimisation of the 'accumulated payoff' account.

Discounting only affects the payoff account whereas all 'real' money, both income and cost, are undiscounted⁶⁰. The preparation time for a conventional attack is short, so one such attack may be performed each round, whereas a nuclear project takes m rounds to prepare. Following a nuclear attack, T must wait another n rounds before another attack may be commenced ($m=3$ and $n=5$ are used in figure 6.1).

As a final note, it must be emphasised that my analysis concerns player T's *decision* whether or not to attempt a nuclear strategy. Thus all parameters are interpreted as perceived by the terrorist at the time of decision, not necessarily representing physical realities. Remember from our discussion of the level of terrorist rationality in chapter 1 that I do not demand for player T to have realistic expectations. Nonetheless, the tool employed here is shown to have great analytical potential independently of whether one agrees with the numbers used, and the reader is of course free to insert his or her own numbers.

For this reason, the exact numerical values of different quantities in the following must not be taken literally: they are all products of the somewhat arbitrary numerical data used in the simulation. The qualitative understanding that comes from considerations of the resulting data, however, is independent of the exact numbers and constitutes the important conclusions of this chapter.

The game is assumed to have an infinite time-horizon, which makes for somewhat different dynamics than if it only lasted for a given number of

59 Which, remember, calls for a discount *function* rather than factor. See chapter 5.

60 thus disregarding such phenomena as interest rates and inflation.

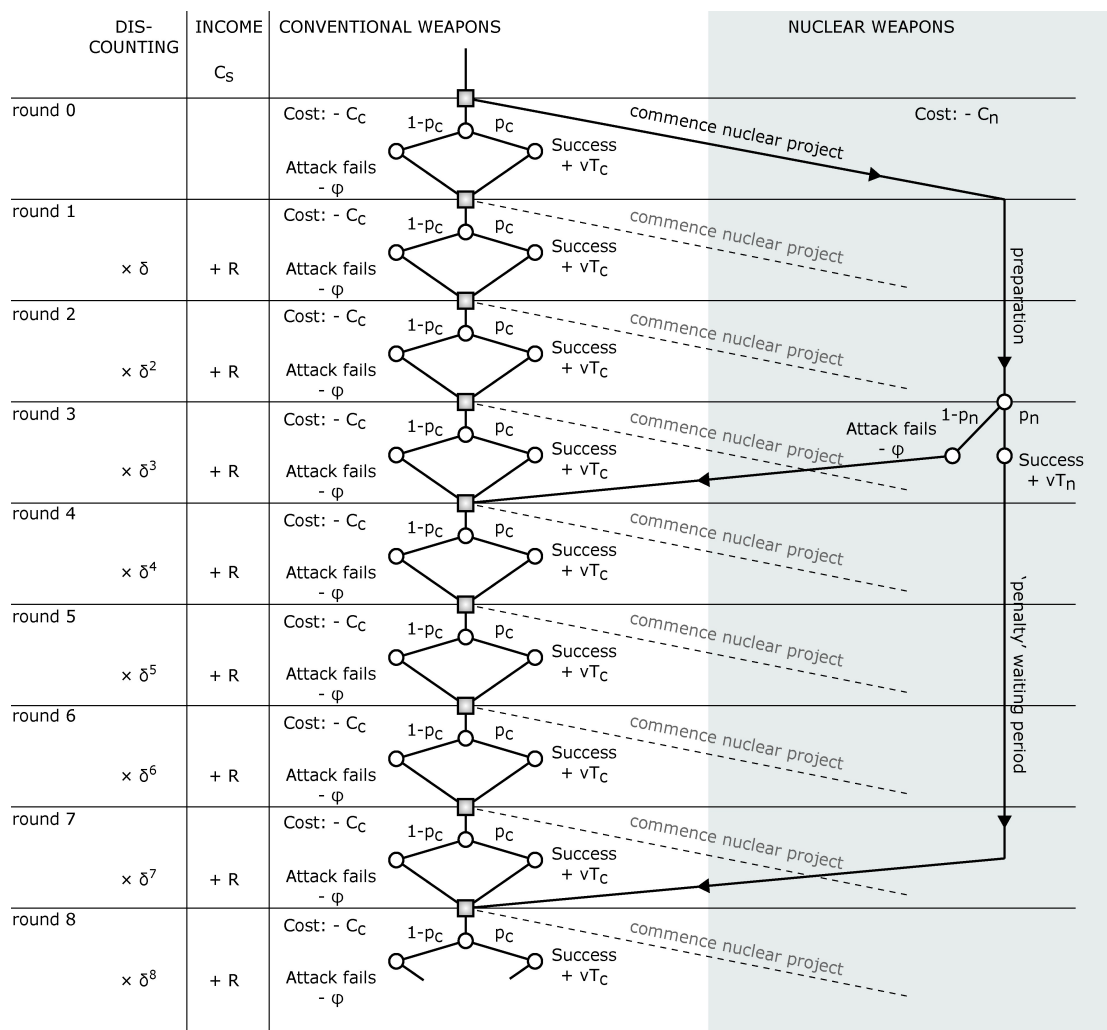


Illustration 6.1.: The game of nuclear vs. conventional strategy.

rounds. If player T knew that there were only m rounds left of the game, say, starting a nuclear project would be more opportune than before, since the game would then end before the penalty period could start. Such tactics are impossible when the time horizon is infinite or unknown. In our actual simulations, the number of rounds is obviously finite, but infinity is mimicked by disallowing tactics that are possible only with a known horizon⁶¹.

6.4.3 Strategies

While complex strategies are thinkable, suffice it for our purposes to compare the two simplest: those in which conventional and nuclear attacks

respectively are attempted as often as sufficient funds are available.

6.5 A discussion of means of deterrence in light of a simulation

The game of figure 6.1 is sufficiently complex that analytical calculations become cumbersome. Using a computer allows us to test the model numerically, however, varying the different parameters on the way. The programme (I have used C++, but any programming language could easily do the job) is detailed in appendix B.

For each run of the programme I choose one parameter to vary while keeping the others constant. The average accumulated payoff per round may then be plotted against the values of parameter which is varied to acquire a graphical representation. For each

⁶¹ The number of rounds in the simulation is so large that payoffs in the last few rounds are for all practical purposes zero due to discounting anyway.

chance node in figure 6.1 I randomise⁶² using predetermined probabilities for success and failure. This means the result can vary greatly - from utter failure to a series of successful attacks. Therefore, each point of a graph represents the average payoff from 10,000 runs of the model in order to remove some of the 'noise' and obtain a graph of the expected payoff which is smooth enough to read⁶³.

Varying δ produces a graph, shown in figure 6.2, which I will interpret in the following. In the figure I have used different probabilities p_n for the nuclear strategy and different levels of damage T_c for the conventional strategy. It is important to note that the absolute values of parameters and payoff are of no importance in themselves - only the qualitative behaviour of the graphs is. The important points to take on board all regard the interplay of the different graphs, independently of the chosen values of numbers themselves⁶⁴.

The numbers used in the numerical exercise deserve some comment. Measures for the damage caused by a conventional attack are deliberately set significantly lower than the approximately \$90 billion of September 11⁶⁵. The 2001 attacks have many properties in common with a nuclear project, particularly in that it was unusually spectacular and involved long preparation time, planning and training. Compared to a nuclear attack, it would be a 'medium size' attack, the option of which I omit from my model for simplicity. Damage costs in the order of a few billion dollars are reasonable for smaller conventional attacks (see section 4.4.1). This also goes with our relatively high estimate of 80% success rate for such attacks. Conversely, the success rate for nuclear strikes is assumed to be low. This assumes that the percentage is as perceived before any preparations have commenced, the necessary expert personnel recruited etc. As the project proceeds the probability of eventual success will increase. Nonetheless, the numbers used (10% and less) are

lower than those suggested by some⁶⁶, but is in agreement with the numbers used in Bunn's mathematical demonstration⁶⁷.

The points of intersection are of particular importance, as they represent border values where the conventional strategy attains a higher expected payoff than the nuclear, or *vice versa*. Between these points lie different regions which all have interpretations and implications.

6.5.1 Interpreting the graphs: points of intersection and noise level

I consider the main graph (the topmost) of figure 6.2⁶⁸ and will interpret what I call the four different areas of the figure, points of interception separating these areas, and the 'noise' of the graphs, going from lower to higher values of δ . The areas are:

- *Impatience area:* (higher conventional payoff)
To the far left at small values of δ , conventional strategies have far higher expected payoff. A small value of the discount factor means high discounting, or in other words, great *impatience*. In this area, player T is in such a hurry to get quick results that waiting the three rounds (in our numerical examples) before a nuclear weapon is ready is simply too long. Note for future reference that the lower the probability of a successful attack, the higher the 'impatience limit' of δ where the graphs intersect.
- *Nuclear first strike area:* (higher nuclear payoff)
Disregarding for a moment the topmost conventional graph, one sees that the nuclear option becomes preferable to conventional means for higher (intermediate) δ . Here, T has

62 The computer equivalent of throwing a die.

63 Instead of randomising, I could simply calculate expected payoff directly, obtaining a smooth graph right away. However, the 'noise' of figure 6.2 visualises risk as will be explained.

64 One should note, however, that all qualitative phenomena discussed do not appear for all sets of parameters. The parameters used are partly chosen so as to demonstrate the full range of important observations.

65 William C. Thompson, Jr., Comptroller, New York City *One Year Later: The Fiscal Impact of 9/11 on New York City* (New York: Office of the Comptroller, September 2002)*

66 Notably, in a testimony before the US Congress in Marc 2007, the prominent physicist and nuclear weapon designer Richard L. Garwin is reported to have estimated the probability of nuclear terrorism against the United States at 20% per year, that is, 90% over a decade! Council of Foreign Relations 'How Likely is a Nuclear Terrorist Attack on the United States' Online Debate between Michael Levi and Graham T. Allison (April 20, 2007)*

67 Matthew Bunn 'A Mathematical Model of the Risk of Nuclear Terrorism' *The Annals of the AAPSS* 607 p.103-120 (2006). Note that Bunn's probabilities are measured per year, since they are estimated as perceived by the government player, who does not know when a nuclear project starts. The terrorist will obviously have this knowledge and is only concerned with the success probability of the single project at hand.

68 Points of intersection are shown with circles. The 'bloodlust' parameter is $v=10^4$.

the time to wait the time it takes for a nuclear device to be finished and reaps great rewards in the case of a successful attack, but δ is still small enough that the penalty time is too far into the future to be a major concern. This area is the primary area in which terrorists would choose to opt for nuclear weapons.

- *Long time bombing area:* (higher conventional payoff) At even higher δ , the tables turn once more as is more clearly seen in the inset to the right of the figure. This area, marked with grey for the lower-damage conventional graph, is where the penalty time following a nuclear detonation becomes a strong argument for the conventional option. In this area, higher expected payoff is harvested from repeated conventional attacks over a long period than can be matched by a nuclear attack, because of the penalty the latter entails. This region is small, the marginal gain from conventional means is small, and it only even occurs for special combinations of parameters against very patient terrorists. It might be seen as representing the power of promised retaliation following an attack.
- *Second nuclear strike area:* (higher nuclear payoff) For values of δ very close to unity, the nuclear graphs rise once more and surpass the conventional option as the inset shows. The in this case extremely patient terrorist is contemplating a second nuclear attack far into the future. This region comes up as a consequence of modelling, but is hardly relevant for our analysis henceforth.

The 'noise' of some of the graphs is notable, and of importance. Bearing in mind that each point of the graph is the average of 10,000 runs of the model, the fact that much noise still remains (and the reader should note that the ordinate axis is logarithmic, masking much of the effect) signifies that although the expectation value of a nuclear plot may be high, the outcome is very uncertain⁶⁹. As one should expect, this is especially the case in the situation where the chances of succeeding with a nuclear project are very small. Here, the probability of failure (modelled to represent a negative payoff of \$500,000) is almost overwhelming, but a success, in the rare cases that it happens, entail an enormous payoff of \$160 million in this case, making for a positive expected payoff, but at

⁶⁹ Statistically: the mean value is well defined, but has a large standard deviation. See below.

an enormous financial risk. Thus, expected utility alone is an insufficient measure of how inviting the nuclear project is to the terrorist in that it does not capture the risk of a strategy⁷⁰; her willingness to place vast resources on a wildcard must be taken into account as well.

6.5.2 *Decreasing the terrorists' chance of success*

Arguably, the region of the graph where the antiterrorist player would like T to be is the impatience region. Quite apart from the fact that making it harder for the terrorist to obtain nuclear weapons provides a very physical and concrete protection for the antiterrorist player (who, I assume, believes he will be the target of all attacks), if player T can be made to *believe* that such a project has a low probability p_n of success (p_n , remember, is player T's belief⁷¹), this could drastically raise the limiting value of δ below which a nuclear project takes too long.

It is thinkable that T's perception of p_n could be shifted by other means than actually trying to block her way to the Bomb, say by a propaganda campaign or deliberate spread of misinformation by the government. With the assumption that the terrorist is rational, this is a possible strategy, but in applying theory to practice, such a plan has drawbacks. First, the effect of such a deterrence strategy is hard to measure. Secondly and more importantly, perhaps, the assumption of terrorist rationality, while highly useful in understanding terrorist reasoning in many cases, ought to be made with some care.

Increased nuclear security would have the very beneficial side effect that it will aid physically blocking a nuclear project, should deterrence fail (whether by a breach of the rationality assumption or not). A lesson to take on board is therefore to make sure that security against nuclear acquisition is not put in place tacitly, but its publication used in a way so as to convey the message that 'it's getting harder'. The graphs, as one would hope and expect, clearly make the case in favour of safeguards efforts, not only for physical threat reduction, but also as part of a

⁷⁰ See chapter 2.

⁷¹ One school of thought amongst statisticians, the Bayesian school, upholds that *all* probabilities must be interpreted as subjective, like p_n is for player T. Hence they discard the standard 'frequency interpretation' of probability as the average rate of success if the event could have been repeated indefinitely. For our purposes I note that p_n and p_e need not be at all the same when perceived by player T as by player A.

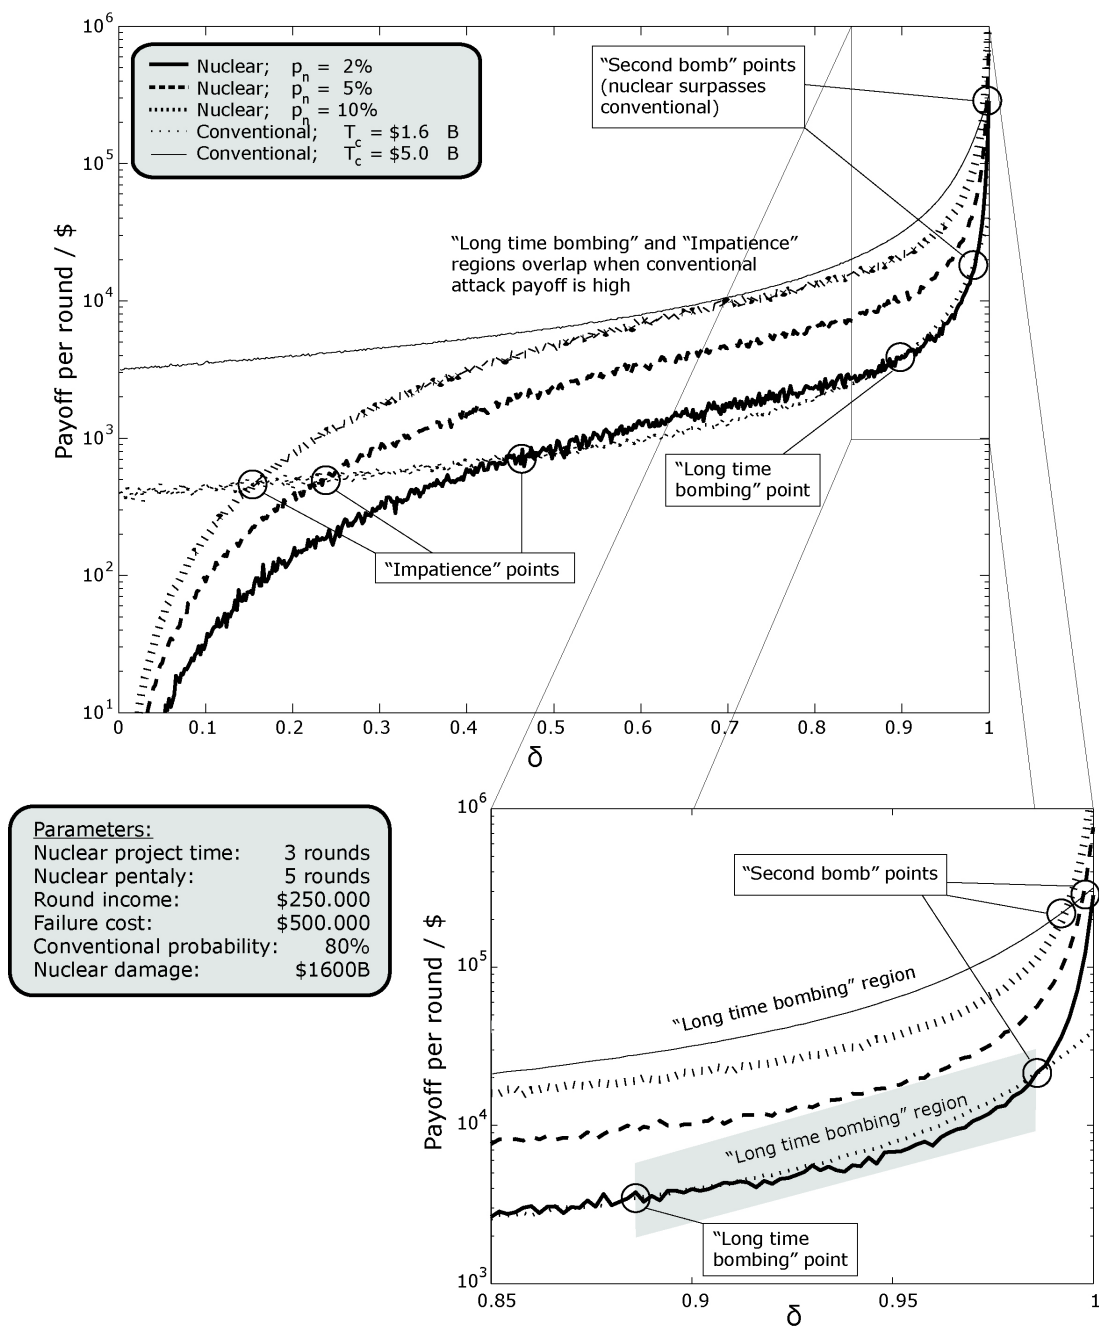


Illustration 6.2.: Average payoff per round for different values of δ .

deterrent. We see quantitatively what Ferguson and Potter recognise by non-formal arguments: 'Prevent access to materials or targets, and the terrorists' decision-making calculus changes Thus, terrorists may be deterred from nuclear terrorism by being denied access to key materials and/or targets'⁷².

Members of the US administration have also noted this point in recent times⁷³. Of course, for deterrence purposes it does not matter exactly how the

72 Charles D. Ferguson and William C. Potter with Amy Sands, Leonard S. Spector and Fred L. Wehling *The Four Faces of Nuclear Terrorism* (New York: Routledge, 2005)* p.32

73 Stephen Aoki, Deputy Undersecretary of Energy for counterterrorism says: 'Barriers to acquisition also provide an important element of deterrence. If terrorists believe that it will be extremely risky, or impossible, to acquire weapons or materials, they may seek other avenues of attack.' testimony (July 2006)*

government makes a nuclear programme more difficult to achieve for the terrorist, and while safeguards measures may be the most potent, 'second layer of defence' measures will pull in the same direction⁷⁴.

6.5.3 Impatience: the rôle of disruption

As our colloquial dubbing 'impatience area' is meant to indicate, the low- δ area where the nuclear option is not preferred, even for a terrorist with no second thoughts about risk, represents great impatience. Anyone about to murder thousands, perhaps hundreds of thousands of people, would worry about word spreading about the project, since it is very safe to assume that a report about an ongoing nuclear terrorism plot with a credible chance of succeeding will spur an intense endeavour by governments who imagine themselves as potential targets to track down and derail the project before a disaster happens. While it takes a cold mind to perform the kind of calculus by which a nuclear attack on civilians is a best course of action, a nuclear project will be a large and costly undertaking and however cynical he or she may be, the person in charge of it will be deeply concerned with its success. This alone should create an atmosphere of some urgency to get to business.

There are means available to a government player to increase the impatience, possibly to the extent where a terrorist feels the need to abandon the thought of acquiring a nuclear weapon or even discontinue an ongoing project⁷⁵.

One such measure is a credible promise that it is getting ever harder to obtain the necessary assets to do the job. Further drastically improving safeguards with a view to create a leak-free system of handling weapons-usable materials once again enters as an effort of great importance. Its value in the particular case of causing impatience is by no means a primary motivation for improved safeguards⁷⁶: the message could well be interpreted as a 'now or never' for the terrorist on the lookout for fissile materials⁷⁷, which

could work either for or against the government player. In keeping with the findings of the previous chapter, the terrorist might feel under pressure to opt for the first available source of what she believes is usable materials, which could turn out to be either plutonium of some grade, usable only in a more complicated weapon design, or uranium of a suboptimal or even useless isotopic composition or in a form that is difficult to handle. On the other hand, the perceived increase in difficulties of accessing fissile materials might be the push that makes the terrorist decide to commence a nuclear programme while there is still time.

Another government strategy is persistent disruption of terrorist plots by various means, a tactic employed extensively by the United States and its allies since the terrorist attacks on September 11 2001. As Jackson reasons, 'a terrorist group under pressure of pursuit will ... have a serious disincentive to seek out or attempt to adopt new technologies'⁷⁸. Excepting the train bombings in Madrid in 2003, attacks since 2001 have been comparatively small, using conventional means and directed at easier targets⁷⁹, including troops in Iraq and Afghanistan⁸⁰. An explanation has been that increased stress has led al Qaida to eschew complex terror operations for simple but deadly bombings closer to home⁸¹.

The banishing of al Qaida from its safe haven in Afghanistan is an example of a disruptive action that may be argued for based on the impatience facet of terrorist choice of action presented here⁸².

6.5.4 Terrorist financing and risk aversion

There are several aspects of risk and aversion to risk in this context, one of which regards the cost of failure, to be explored further below. In this section I shall focus on the risk of undertaking a project whose potentially enormous payoffs come at the cost of a large probability of failure and what a study of al

74 See appendix D for a review of such measures.

75 Note that this option is not possible within the model, yet is an intuitively obvious continuation of it.

76 Strictly speaking, the impatience created by improved safeguards concerns only the nuclear, not the conventional project, and the notion is thus somewhat at odds with our model; a consequence of simplification.

77 Note that I have not considered the option of parameters, e.g. p_{nv} , changing with time.

78 Jackson 'Technology Acquisition ...' p.195.

79 Martin C. Libicki, Peter Chalk, and Melanie Sisson *Exploring Terrorist Targeting Preferences* (RAND, 2007)*, e.g. figure p.48.

80 Often termed 'insurgency' rather than 'terrorism', a distinction whose subtleties I shall not delve into. Notably, violence against exclusively military targets fall outside my definition of terrorism in section 1.2.

81 Todd Sandler and Walter Enders 'September 11 and Its Aftermath' *International Studies Review* 7 (2005) p.167

82 The legality and ethical aspects of the war and its aftermath is a question not to be addressed here.

Qaida's previous *modus operandi* tells us about their attitude towards risk and what rôle their financial prospects play.

Without doubt, the purely *financial* risks involved in undertaking a nuclear project depend upon the fiscal stature of the group. Lack of funds could make the project infeasible altogether, but a mere stemming of the flow of fresh funds will be sufficient to increase risk. When running the same simulations as above, assuming sufficient initial funds for a single nuclear attack but with no new funds coming in, the graph changes little from that shown in figure 6.2⁸³. However, this may not be interpreted to mean that throttling the flow of money to al Qaida is ineffective. Imagine that you were given a 1% chance of winning \$1 billion at the price of everything you own, which for the sake of argument we assume to be \$100,000. The expected payoff of this deal is a neat \$9.9 million, yet none but the craziest gambler would accept the deal because the likely outcome is a very severe loss. If \$100,000 was an easily affordable loss to you, however, the situation would be very different. In the face of a stemming of the flow of finances to al Qaida, thus, it is probable that spending what they have on a new and unknown enterprise becomes too daring a choice.

One way of considering the risk of a strategy is looking at how much the outcome of that strategy varies. A safe plan will produce a similar outcome every time, while one that is risky could bring both big benefits and big losses. I calculate⁸⁴ the standard deviation of the distribution of the payoff over a large sample (of $N=10,000$ runs of the model). Sample standard deviation σ is a measure of how much a sample of measurements (say of accumulated payoff from one run through our model) varies in value. Specifically it is the root mean square of the deviation of the outcome of one run from its mean value, defined as

$$\sigma^2 = \frac{1}{N} \sum_{i=1}^N (f_i - \mu)^2 \quad (6.1)$$

where f_i is the accumulated payoff of run i and the expectation value μ is found as

$$\mu = \frac{1}{N} \sum_{i=1}^N f_i \quad (6.2)$$

I have calculated σ for a large sample of runs of the model following a nuclear and conventional strategy respectively. From plotting σ as a function of δ in figure 6.3 it is clear that for moderately high values of the discount factor, the region in which the 'first nuclear strike' area is typically found (the area in which the terrorist is arguably most likely to wish to opt for nuclear weapons), the standard deviation of the nuclear sample is 3-5 times larger than that of the conventional sample.

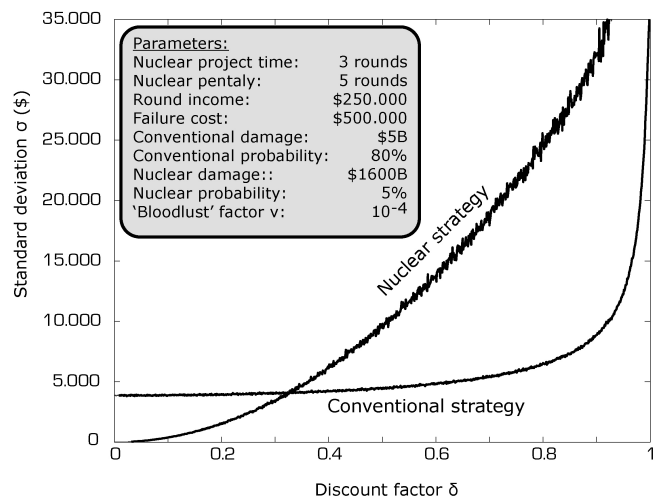


Illustration 6.3.: The standard deviation of accumulated payoff per round.

The reason for the standard deviation for smaller values of δ being smaller than that of a conventional strategy is solely that any payoffs from a nuclear strategy will be small in this area due to the long preparation time: for $\delta=0$, of course, the nuclear payoff is *invariably* zero. A more instructive quantity might be the standard deviation relative to the expected payoff, which we have plotted in figure 6.4. Relative to the expected payoff (approximately equal to μ), the nuclear option has standard deviation four times or more that of the conventional strategy throughout the δ -scale. Indeed for the nuclear strategy the relative standard deviation is more than four times the expected utility itself for most values of δ , and only decreases for very high δ values. This is four times or more the relative standard deviation as for the conventional strategy almost throughout the δ scale. Of course, the exact ratios between these standard deviations and payoffs will depend on the numerical values of the model parameters, but insofar as the

83 The only qualitative change is that there is no 'second nuclear strike' area.

84 Using a slight modification of our C++ programme from before.

data used are not exceedingly far from the truth, these trends are of very general validity.

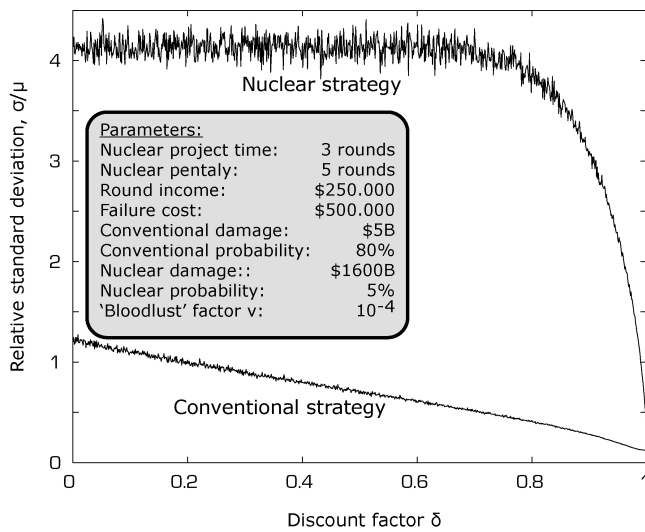


Illustration 6.4.: Standard deviation relative to expected payoff as functions of δ .

I will keep the theoretical analysis of this point on an intuitive level and therefore not draw on economic theory of risk or other relevant theories here. A good way to analyse the problem further would be the extension of rational choice theory called *regret theory*⁸⁵ in which the player's utility does not only depend on the outcome that comes to be, but also that which *could have been* if a different choice had been made, thereby introducing the concepts of regret and rejoicing for having made the wrong/right choice. This theory automatically takes some account of risk aversion in the level of regret associated with a wrong decision.

I turn quickly to the example of the bet once more. Arguably, for the two players (one 'poor' and one 'rich') the most important variable is the relative change of their wealth as a consequence of the gamble. Assume the wealth of the rich player is \$10 million. To him, a loss would imply a reduction of his fortune of only 1%, while on average, taking the bet would double his fortune. To the poor player, the loss would mean a devastating reduction of his fortune by 100%, while accepting the bet would on average increase his fortune 99 times. The present value of the 'fortune' of a terrorist who receives a sum R of money every round is $R/(1-\delta)$ plus whatever she has got to start with. When δ approaches 1 such a regularly

funded terrorist is therefore to be considered very wealthy, provided she has the patience to wait for the money. For this reason, whether the terrorist identifies herself as the 'rich' or the 'poor' player in our analogy will depend on the reliability of the income flow and the patience of the terrorist. Efforts to increase terrorist impatience, therefore, not only push the terrorist towards the impatience region (where she will opt for conventional means) due to lower expected payoff from a nuclear strategy, but also increases risk aversion, adding to the same effect.

Being the group most often seen as a likely aspiring nuclear terrorist, let me turn to the example of al Qaida for a moment. Dunn has studied al Qaida in depth and recognises a distinctive aversion towards risk in their history, a 'persistence in doing what it knows and does well ... Similarly, its choice of targets ... reflects persistence in staying with the tried and true.'⁸⁶ Nonetheless attacks such as September 11 show a definitive ability to think creatively and innovate, even if the novel requirement of those attacks was restricted to learning to fly commercial aircraft.

Religiously zealous terrorists could have additional reasons to fear public failure, argues Jenkins. As he puts it 'jihadists believe that God's will is expressed in success and failure. To succeed is to have God's support. Failure signals God's disapproval. As a consequence, jihadist planners are conservative.'⁸⁷

Another point that could enhance al Qaida's sense of risk is the organisational changes it would probably require, another departure from the tried and trusted. So large are the costs involved (millions of dollars) that a project to build a nuclear device will probably be kept under much closer scrutiny by the leaders than is normally the case with the activities of the many loosely affiliated al Qaida cadres worldwide⁸⁸. A more hierarchical structure could make the organisation more vulnerable to intelligence as well as military action. The successful designing of a working device requires a sound scientific climate where ideas and opinions may be presented freely⁸⁹, the lack of which has been an explanation for Aum Shinrikyo's failure to weaponise biological agents⁹⁰. Finally, it has

85 Graham Loomes and Robert Sugden 'Regret Theory: An Alternative Theory of Rational Choice under Uncertainty' *The Economic Journal* 92 (1982) pp.805-824

86 Dunn 'Can al Qaeda ...' p.16

87 Jenkins *Unconquerable Nation* p.81

88 Sammy Salama and Lydia Hansell 'Does intent Equal Capability?: Al-Qaeda and Weapons of Mass Destruction' *Nonproliferation Review* 12:3 (2006) pp.616-617

89 Ferguson and Potter *The Four Faces...* p.40

90 William Rosenau 'Aum Shinrikyo's Biological Weapons

not been al Qaida's practice to rely on one or a few exceptional individuals for the success of their operations⁹¹; a nuclear project, on the other hand, stands or falls with the scientists and technicians employed⁹².

6.5.5 The perceived cost of failure - how important?

In the previous chapter I found that only when the perceived cost of failure is zero is there no theoretical possibility for deterring a terrorist from attacks altogether⁹³. I have performed simulations like those depicted in figure 6.2 in which the cost of failure is varied from zero to several million dollars with other parameters constant. When considering relative deterrence, unlike its absolute counterpart, increasing the cost of failure leads to no qualitatively new features emerging, but plays the rôle of lowering the expected payoff of attacks in general, in particular the nuclear strategy whose probability of failure is large. For the purposes of relative deterrence it is of modest importance, primarily because there is little a government player can do to influence player T's assessment of the value of φ .

6.5.6 The 'bloodlust' parameter: no longer as important

Likewise, the 'bloodlust parameter' ν which played an important rôle in the previous chapter is of less consequence here; the assumption that payoff is proportional to the damage inflicted with some positive proportionality constant, however, is still very important. To understand the rôle of ν in this context, consider the following simplification of the model: Assume just for now that no new funds come in ($R=0$) and that there is no cost of failure ($\varphi = 0$)⁹⁴. Assume furthermore that initial funds suffice for M conventional attacks or one nuclear attack after n rounds of preparation. The utility from a conventional strategy is thus (the first round is not discounted)

$$\frac{1-\delta^M}{1-\delta} p_c \nu T_c \quad (6.3)$$

(see appendix C for mathematical details) while that of a nuclear attack is

$$\delta^{n-1} p_n \nu T_n \quad (6.4)$$

where as in figure 6.1, n is the number of rounds it takes to prepare a nuclear attack. The conventional strategy is preferable in terms of expected payoff given these assumptions iff

$$\frac{1-\delta^M}{(1-\delta)\delta^{n-1}} > \frac{p_n T_n}{p_c T_c} \quad (6.5)$$

The bloodlust factor ν disappears and is of no importance to the choice of strategy! For demonstration, let's put $T_n = \$1.6 \cdot 10^{12}$ and $T_c = \$5.0 \cdot 10^9$. Say furthermore that $p_n/p_c = 0.1$, then the fraction on the left must be larger than approximately 32. If I assume for example that $M=10$ the numerical solution is that conventional strike is preferred for approximately $\delta < 0.2$. The only significance of ν is now determining the sign of the payoff at a given value of φ , of modest interest in the context of relative deterrence.

Clearly, when $\varphi > 0$ and $R > 0$, the bloodlust parameter does not entirely vanish, but a numerical study reveals that its importance is modest.

The level of 'bloodlust' decides whether terrorist action is worthwhile at all, setting the minimum level of damage for the attack to be worth the risk; the more bloodthirsty the terrorist, the smaller the smallest feasible attack. Hence a terrorist with great bloodlust is hard to deter absolutely, but may be satisfied with long term violence on a lower scale, another case where absolute and relative deterrence work against each other. On the whole, however, for a terrorist who extracts utility proportional to damage, the important question is 'how much damage how soon?', regardless of the proportionality factor.

6.5.7 The danger of hardening targets against conventional strikes

The dynamic between absolute and relative deterrence is highlighted when considering the topmost conventional graph in the topmost panel of figure 6.2. When assuming a somewhat higher damage per conventional attack, the 'impatience' and 'long time bombing' regions merge with no 'nuclear

Program: Why Did it Fail?' *Studies in Conflict & Terrorism* 24 (2001) pp.289-301.

91 Dunn points to almost exclusive reliance on personnel trained in al Qaida camps; Dunn 'Can al Qaeda...' p.16

92 e.g. Jackson 'Technology Acquisition by Terrorist Groups' p.201

93 This was rigorously shown only in the context of the model in that chapter, but reasons were given why the conclusion is probably far more general.

94 This is for simplicity only and is hardly a realistic assumption.

first strike' area in between at all. Clearly, if the terrorist is able to inflict more damage (in our example the equivalent of \$5 billion) per conventional attack, the incentive to go nuclear diminishes, the terrorist being content with her current progress in her 'war of a thousand cuts'⁹⁵. Seen in isolation, thus, for the sake of avoiding nuclear terrorism, hardening targets against conventional attacks is counter-effective.

This 'deflection' effect has many faces. Sandler and co-workers have pointed to how one state may deflect attacks onto another by making itself a harder target, initiating a 'deterrence race' in which the target state least able to protect itself gets the most attacks⁹⁶, and how hardening business targets deflects attacks towards tourism⁹⁷. These findings are consistent with previous findings on crime and warfare, quoted by Melese and Angelis: making one kind of crime harder will reduce that crime but increase others, and defending against one kind of military attack may 'drive opponents to new methods of attack and, therefore, generate new ... threats'⁹⁸. For example, the installation of metal detectors in airports in January 1973 led to an immediate change from skyjacking to other terrorist tactics such as hostage taking⁹⁹. As Garwin notes, keeping our planes safe is not a plan for lasting peace, since 'if hijacking passenger aircraft will no longer work, motivated terrorists will doubtless choose something else.'¹⁰⁰ Nuclear weapons are but one option.

The conclusion to be drawn from this is not to 'allow' a certain level of conventional terrorism, but to make sure that for every effort to prevent conventional terrorist attacks, a correspondingly convincing effort should go into preventing nuclear terrorism to avoid the deflection effect. Our analysis shows that a disproportionate emphasis on conventional terrorism, although it may seem like the more immediate threat, could undermine other efforts to achieve relative deterrence from nuclear terrorism.

95 Libicki *et al.* *Exploring Terrorist Targeting Preferences* p.93.

96 Todd Sandler and Kevin Siqueira 'Global terrorism: deterrence versus pre-emption' *Canadian Journal of Economics* 39:4 (2006) pp.1370-1387

97 Todd Sandler and Daniel G. Arce M. 'Terrorism & game theory' *Simulation & Gaming* 34:3 (2003) p.326

98 Melese and Angelis 'Deterring Terrorists from Using WMD' p.338.

99 e.g. Todd Sandler and Walter Enders 'An economic perspective on transnational terrorism' *European Journal of Political Economy* 20 (2004) p.311.

100 Garwin 'The Many Threats of Terror' p.237.

6.5.8 *Smaller payoff from nuclear attack*

I have assumed above that the terrorist extracts payoff proportional to the damage done, which may be something of a worst case terrorist. Some have raised the question whether it is absolutely certain that a terrorist organisation which appears on the surface to want nothing more than to cause death and destruction could not have reservations against killing such a vast number of people in a single blow¹⁰¹.

If this is so it is good news given the previous analysis of deterrence: the only motivation that my modelled terrorist has to go for a nuclear programme in the face of slim chances of success is the enormous payoff she can reap from it. If important strategists have some limitations as to how many they would like to see dead, however, this motivation quickly diminishes. Even if the expected payoff is positive, only a very large payoff is worth a great risk, and even more so if the terrorist is averse to risk in the first place. The policy implications of this are not perhaps great, however, since there is little a government can do to enhance the moral stature of a terrorist adversary. It is a point worth bearing in mind, however, in line with Levi's stance that it may not always be helpful to consider only the worst case¹⁰².

Another point, emphasised in particular by Levi, is that measures can be taken to diminish the effects should a nuclear attack happen¹⁰³. The effects of a successful attack will be dreadful, but if hundreds, perhaps thousands of additional lives could be saved with good emergency response, planning for the worst is likely to be worthwhile based on a simple rationale of self defence. Levi makes the case for a simple sheltering strategy to minimise harm from radioactive fallout¹⁰⁴ for example. Important for our purposes is the possibility that the terrorist could be persuaded that the government is well prepared and will be able to take measures to minimise the damage done. If T_n can be lowered, the effect on relative deterrence is immediate and proportional to the reduction.

101 Dunn 'Can al Qaeda ...' pp. 9-11

102 Michael Levi *On Nuclear Terrorism* (Cambridge, MA: Harvard University Press, 2007) pp. 9-12

103 *ibid.*, pp.61-64

104 *ibid.*, pp.61-64, 161-163

6.6 The advantages of gaming versus a qualitative approach

The advantages and disadvantages of gaming out the terrorist's decision in this chapter are rather similar to what the case was in the previous chapter. The only substantial difference between the two in terms of methodology is that the present chapter performs the gaming numerically whereas the previous chapter was symbolical. Many of the points discussed in section 5.7 such as the limitations of assuming a terrorist organisation be representable as a single, rational actor, also apply to this chapter.

As with the other research chapters of the thesis, the research question in this chapter was chosen partly because it was expected to work well with the chosen methodology. It is another case of a choice between two strategies under uncertainty where the pros and cons of each strategy are easy to lay out in qualitative terms - on the one hand the terrorist would harvest great utility from a successful nuclear acquisition, but on the other it is a difficult, expensive and risky project compared to tried and trusted means - but weighing the pros and cons against each other is more difficult in qualitative terms.

Clearly, after laying out the arguments for and against a nuclear strategy, there could be different ways of going about the weighing, and the gaming approach used here, while a natural choice for someone with a mathematical background, is almost certainly not the only fruitful option. Analysts with a background in psychology might draw on available historical evidence on terrorism to do an analysis of terrorists' risk aversion versus their drive towards spectacular attacks. With information on the way certain terrorist groups are organised, a study might focus on the internal tensions within the group, should different factions dissent on what is the best way forward, in line with Allison's idea of bureaucratic politics¹⁰⁵.

These different approaches should not be seen as competitors as much as complements of each other. The present effort lays out the strictly economical side of the issue in terms of an abstract monetary unit (pleasure in success, aggravation in failure), and thus adds a piece to the jigsaw that is relative deterrence of terrorism. Additional efforts of analysts with different background and training could add to my economical

analysis to paint a more nuanced and complete picture of the effects of efforts towards relative deterrence. The effects of realistic deviations from the purely rational and calculating ideal assumed, for example, could be a valuable addition to the efforts presented in this chapter. It is useful to recognise that one approach alone cannot replace all others, and seek ways of unifying results of different approaches rather than focus on which methodology is best.

Compared to the efforts reported in chapter 5 the formal approach has at least a couple of disadvantages when a numerical procedure is employed rather than the symbolical methods used there. Firstly, one does not end up with simple and powerful criteria such as (5.22) from which a wealth of information may readily be drawn, and which forms a conclusion in itself. The most important information which a model gives away is how a conclusion *varies* when the parameters are varied. An explicit formula permits the variation of every parameter at the same time, but numerically it is normally only feasible to vary one or two parameters while keeping all others constants¹⁰⁶. This is a serious limitation which diminishes the power of the analysis somewhat. The procedure one follows in practice is to experiment with varying different parameters in search of interesting behaviour which can illuminate the problem. I found such behaviour when varying the impatience factor δ , as shown in figure 5.2. It is difficult to ascertain that the model does not hold much *more* information, however, which I did not chance upon.

The second limitation using numerics is that numbers must somehow be chosen for all parameters in order to run the simulation. This introduces an extra layer of application of judgement on the part of the analyst, and fixing numbers to every parameter necessarily reduces the generality of the analysis. Therefore, I must focus on the *qualitative* behaviour of the model when parameters are varied. This was found to be feasible as in figure 5.2 because, as the figure shows, the graphs have the same general shape even if the numbers change its actual values.

Where the analytical treatment in the previous chapter allowed strict criteria, I am here reduced to introducing less rigorous, semi-quantitative concepts for interpreting the graphs: impatience area, first

105 e.g. Graham T. Allison and Morton H. Halperin 'Bureaucratic Politics: A Paradigm and Some Policy Implications' World Politics 24:2 (1972) pp.40-79

106 There exist more sophisticated methods for multivariate systems, but as argued above, such methods could easily obscure the ongoing analysis and introduce another set of *ad hoc* assumptions because they require much information about the function under study.

bomb area and so forth. These are arguably good and intuitively useful concepts, and with reference to Osbourne's view 'if a model enhances our understanding of the world, then it serves its purpose'¹⁰⁷ it is the author's view that the gaming was successful.

However, since these concepts may readily be formulated in a qualitative way, it is less obvious that they could not have been arrived by qualitative arguments alone. The rôle of the gaming and simulation was to formulate the problem in an alternative way to view it from a different angle. When represented graphically as in figure 5.2 the classification of terrorists in terms of their impatience became obvious. With the benefit of hindsight it seems quite possible to arrive at the same taxonomy without the gaming, but whether one could do so in practice is a hypothetical question for the individual analyst to answer.

6.7 Conclusions and policy implications

At a very general level I am in agreement with Melese and Angelis that a relative deterrence project must 'make it relatively more costly for terrorists to acquire' nuclear weapons¹⁰⁸ with respect to conventional means. The means of achieving this, however, differ radically from those suggested by these authors.

Indeed, I show with some clarity that with the assumptions employed about terrorist motivation, the threat of retaliation as suggested by Melese and Angelis has unambiguous effect only under very special circumstances¹⁰⁹. Obviously, one must be careful, because the model used represents only one possibility, and it is thinkable that the value of retaliation is underestimated in our work¹¹⁰. The 'penalty' time for the terrorist organisation represents

the terrorist's belief that retaliation will render her inoperative for a time after a nuclear strike. This causes preferences to change only in the 'long time bombing' area in figure 6.2, a small interval of large values of δ over which the expected payoff of a nuclear strategy is slightly smaller than for conventional means. The strategy only works against the patient terrorist who has high hopes for the long term future but is of little consequence to a terrorist in some rush. While threats may not be altogether impotent, it is shown here that there are better options available to the government intent on dissuading al Qaida or others from nuclear means.

6.7.1 Three effective roads to relative deterrence

I identify three paths by which relative deterrence from nuclear projects may be achieved.

The primary means of relative deterrence should remain continuing to make it harder for terrorists to carry out a nuclear project by protecting nuclear materials from theft because of the beneficial effects of these efforts beyond deterrence. Secondary measures such as border controls and scanning of freight containers, advocated recently by Levi¹¹¹, are probably a useful addition, although no careful cost-benefit analysis of such measures is undertaken in the present thesis for reasons of manageability (some thoughts on this issue are presented in appendix D). As figure 6.2 shows, nothing is as effective in lowering the expected payoff of such an operation as boosting the perceived difficulty of success in the mind of the terrorist.

Improved safety measures furthermore play on terrorist risk aversion and tendency to choose to be safe rather than sorry. Davis and Jenkins say¹¹²

the empirical record shows that even hardened terrorists dislike operational risks and may be deterred by uncertainty and risk. A foot soldier may willingly give his life in a suicide mission, and organizations may be quite willing to sacrifice such pawns, but mission success is very important and leaders are in some ways risk-averse.

I demonstrate that the flow of payoff from nuclear projects is far more volatile than that from staying with the tried and trusted conventional tactics, an

107 Martin J. Osborne *An Introduction to Game Theory* (New York: Oxford University Press, 2004) p.7

108 Melese and Angelis 'Deterring Terrorists from Using WMD' p.338.

109 i.e. very specific values of model parameters.

110 Possible extensions of the current modelling effort which could increase the effect of retaliation would be if the government actor could incur large economic damage to the terrorist. Also, the assumption that payoff depends only on damage inflicted may overlook the often argued position that even the most fanatical terrorists have goals and values they hold dear. However, the penalty time does take account of this to some extent.

111 Levi *On Nuclear Terrorism*

112 Davis and Jenkins *Deterrence and Influence in Counterterrorism* p.xii

effect which becomes more pronounced the smaller the perceived chances of a successful nuclear attack. An explicit modelling of risk aversion could be a fruitful extension of the analysis performed herein.

Defensive measures beyond safeguards, often termed the 'second line of defence' are not discussed in detail in this thesis; the reader may refer to the excellent treatment by Levi¹¹³. Although the effectiveness of secondary means in actually blocking a terrorist organisation is a matter of some dispute, one must not forget the indirect effect which efforts to, say, detect nuclear materials in sea container cargo can have through deterrence. If the message can be communicated that 'there is no safe way to deliver a nuclear attack' this could play effectively on terrorist risk aversion. Terrorists have shown considerable scepticism about recruiting outsiders for specific tasks such as advanced smuggling operations, and any action that can force the terrorist to divert from her known and trusted members and methods will probably increase the perceived risk in the terrorist's mind.

Secondly, measures to disrupt terrorist operations and put strain on the terrorist groups that might be planning a nuclear enterprise (thereby decreasing δ) could, together with decreasing p_n , force the terrorist player into the 'impatience region' of figure 6.2 in which conventional attack is always preferred. The combination of safeguards and disruption is shown in the above analysis to form a powerful tool in persuading terrorists from attempting the nuclear option.

Finally, strangling terrorism funding will further play on terrorist risk aversion. Blocking terrorist access to money is of course a physical way of hindering a project ever taking place, but even creating significant uncertainty about the economic future could well make the terrorist opt for conservative expenditure.

An important additional conclusion is the danger that hardening targets against conventional terrorism could achieve the opposite effect and become an incentive for the terrorist to opt for a nuclear approach. Therefore, spending in deterrence by denial of conventional terrorism (absolute deterrence) must always be accompanied by a proportional spending in anti-nuclear proliferation efforts¹¹⁴. If the right balance is found, our model indicates that defensive measures

can at the same time deter some terrorists altogether and deter others relatively, away from nuclear terrorism.

The United States' *National Strategy to Combat Weapons of Mass Destruction*, interestingly, contains several elements in accordance with our recommendations for relative deterrence, but apparently without fully realising this potential. It reads, for example: 'In addition to our conventional response and defense capabilities, our overall deterrence posture against WMD threats is reinforced by intelligence, surveillance, interdiction, and domestic law enforcement capabilities'¹¹⁵. Intelligence and surveillance concord with our 'disruption' point above, whereas the improved safeguarding of fissile materials, while mentioned briefly as one of many measures¹¹⁶, is not recognised for its value as deterrent in the Strategy, and the impression which is left is that the prevailing ideas of deterrence are still primarily threats of retribution rather than deterrence by denial. One does well to note, of course, that the report in question is primarily concerned with proliferation to states, a different question than that considered herein. Given the amount of attention devoted to nuclear terrorism in later years, however, there seems no particular reason why relative deterrence of nuclear terrorism should not make its way into future deterrence policies.

In conclusion, our analysis strongly indicates that deterring terrorists into opting for conventional means over nuclear ones is not only possible but indeed doable by increasing emphasis on efforts that are already ongoing. While it requires the term 'deterrence' to be freed from its narrow Cold War interpretation, influencing the way terrorists choose their strategies is probably a relevant and important aspect of the struggle to keep the ultimate catastrophe from becoming reality.

113 Levi *On Nuclear Terrorism*

114 The same goes for other forms of non-conventional weapons, although this is not discussed in detail in this thesis.

115 US National Security Council *National Strategy to Combat Weapons of Mass Destruction* (Washington D.C.: NSC, December 2001)* p. 3

116 alongside such measures as recycling of nuclear waste, of much lesser proliferation concern. *Ibid.* pp. 4-5

Nuclear blackmail and other strategic uses of a terrorist nuclear weapon

Assume that an international terrorist organisation has acquired a nuclear weapon. What will it do with it? Up until now we have assumed that the terrorist organisation's plan is always to mount a devastating attack as soon as feasible. Six decades of nuclear-armed states *not* attacking each other despite animosity certifies, however, that there are other possible uses of nuclear weapons besides detonating them at the enemy's doorstep. Deterrence is one: if you can persuade your enemy that in the instance of attack upon you, you will retaliate with nuclear means, he may not dare to attack at all. A more short-term possibility is blackmail: 'concede to these demands or we will blow you up'. There are important differences between states and non-state organisations, however, and part of the aim of this chapter is to explore these.

A few analysts have considered it possible that a terrorist obtaining a nuclear weapon might not necessarily attempt to detonate it right away. Ferguson and Potter comment that '[t]he credible threat created by controlling a nuclear weapon would significantly bolster any political goals of the terrorist group'¹. Levi also devotes four paragraphs of his book² to the question, concluding that nuclear blackmail is difficult to pull off in practice - a point I will return to shortly. Steinhausler mentions nuclear blackmail but focuses on the feasibility of smuggling a weapon into a US harbour while taking the utility of the blackmail strategy itself for granted³. Probably the most careful and thorough analysis of the question is that of Dunn⁴, whose arguments I will consider in

1 Ferguson, Charles D. and William C. Potter with Amy Sands, Leonard S. Spector and Fred L. Wehling *The Four Faces of Nuclear Terrorism* (New York:Routledge 2005)* p.22

2 Michael Levi *On Nuclear Terrorism* (Cambridge, MA: Harvard University Press, 2007) pp.120-121

3 Friedrich Steinhausler 'What It Takes to Become a Nuclear Terrorist' *American Behavioral Scientist* 46:6 (2003) p.791

4 Lewis A. Dunn 'Can al Qaeda Be Deterred from Using Nuclear Weapons?' *Occasional paper #3* (Center for the Study of Weapons of Mass Destruction, National Defense University, 2005)*

more detail below.

On the whole, however, the question of whether there could be alternative terrorist uses of nuclear weapons has remained almost entirely unaddressed despite several statements by the al Qaida leadership indicating that it sees the potential of a nuclear weapon to deter attacks upon itself from western powers or Israel⁵. Of course, such statements should not be accepted uncritically, and the real rationale for making them could be different from what meets the eye. In line with the outset of the thesis, I will assume that the terrorist adversary is rational and will give thorough consideration to the question of what is the best use of a nuclear weapon, should he acquire one. It is therefore worth taking a closer look at the strategic interplay of a few scenarios in order to get a clearer picture of the incentives and threats involved.

7.1 Research question and chapter outline

The question I will address herein is the following is twofold:

Could a rational terrorist plausibly have other strategically beneficial uses of a nuclear weapon than its detonation in an attack, namely extortion or deterrence of attacks upon him/herself?

and furthermore

What stance should a government adopt in the face of attempted extortion by terrorists with the threatened use of a nuclear explosive?

In this chapter I will assume that a terrorist organisation has already acquired a workable nuclear weapon and is considering the strategic options for the use of this asset. The set of assumptions used is somewhat restrictive, and the chapter does not aspire to exhaust this question but merely to present a few of the most plausible scenarios. Of the conceivable alternative uses of nuclear weapons I consider only blackmail and deterrence. A scenario of nuclear terrorist extortion will be gamed out in a classic signalling game, well known from the literature on game theory⁶, in the final sections of the chapter.

5 See references below.

6 See any textbook on Game Theory, for example Robert Gibbons *A Primer in Game Theory* (Hemel Hempstead:

The reader should note how the word 'deterrence' will be used about two in some sense opposite scenarios in this chapter. The first accords with the way the term was employed in chapter 6 and concerns measures by the government player to persuade a terrorist organisation not to commence a nuclear attack, either before a nuclear plot is instigated as in the previous chapter or after the bomb has been acquired. In the second sense of the term the situation is reversed: here the terrorist player mimics a state by having a nuclear weapon at the ready in order to deter attacks upon him from a hostile government or other organisations.

7.2 Example: al Qaida

In a notable paper, Lewis A. Dunn asks whether al Qaida could see other uses of nuclear weapons than simply a clandestinely delivered attack⁷. His analysis is useful as a backdrop for our discussion of the same question.

Dunn considers four aspects of the available knowledge about al Qaida and discusses how well these observed traits concord with the use of nuclear weapons for attack or as a tool either for blackmail or as a deterrent against attacks upon them by, say, the United States. Dunn's approach is well summarised by the four ways in which he approaches his research question of whether al Qaida may choose not to detonate a nuclear weapon they have already acquired:

- *Ground truth*: What does the physical evidence (notably seized documents in Afghanistan) indicate? Dunn concludes that the ground truth proves that al Qaida has attempted acquisition, but gives nothing away about the intended uses of such a weapon post acquisition⁸.
- *Personnel make-up*: Al Qaida, as previously discussed, is made up of members with a range of rôles and responsibilities. While the core members of al Qaida are unlikely to be deterred by the threat of retaliation, some of the more peripheral players such as funders and logisticians may be, Dunn argues⁹. The latter might also have ethical scruples. For the

purposes of the present chapter, which deals with the high-level terrorist strategy, hindering a planned attack by deterrence of peripheral players would be equivalent to any 'second layer of defence' measure¹⁰. It is part of the overall question of al Qaida's capability to mount a nuclear attack with all the different obstacles this entails and is not considered further here.

- *Operational code*: Recognising trends in al Qaida's historical targeting and attack preparations Dunn considers whether a nuclear attack seems a natural extension of the operational code al Qaida has been practising so far. Dunn's conclusions of interest here are that
 - A nuclear attack is consistent with al Qaida's preference for spectacular attacks, 'visually pleasing' destruction and sophistication.¹¹
 - An attack following acquisition is consistent with al Qaida's tendency to want to 'finish the task'¹².
 - The long time spent planning and preparing for attacks (often a year and more) indicates that al Qaida's planners have a long time horizon. This indicates that they might think seriously consider what the best use of a nuclear weapon would be¹³.
- *Consistency with political vision*: al Qaida's proclaimed political vision is the establishment of a Muslim caliphate in an empire stretching from northern Africa to south-east Asia. Perhaps the most important question, Dunn reasons, is whether the detonation of a nuclear weapon will forward this goal or if another usage of the weapon would be preferable¹⁴. Al Qaida's concerns could include alienation of the Muslim populations in case of attack, as well as using the nuclear weapon as a means of deterrence on the way to establishing their kingdom.

Upon extracting its essentials, Dunn's paper tells us some important lessons that will be discussed in the

Harvester-Wheatsheaf, 1992)

7 Dunn 'Can al Qaeda Be Deterred from Using Nuclear Weapons?'

8 *ibid.* pp.5-6

9 *ibid.* p.6-8

10 See appendix D for further details.

11 Dunn 'Can al Qaeda be Deterred...' pp.8-9

12 *ibid.* p.11

13 *ibid.* p.14

14 *ibid.* pp.17-21

following before we engage with the strategic gaming.

7.3 *Principal issues of use or non-use*

Based on the analysis by Dunn, the reasons why a given terrorist organisation might choose not to use an acquired weapon in an attack as soon as possible can be divided into five different categories which will be discussed shortly (Note that these categories are not mutually exclusive):

1. The terrorist group's leadership decides nuclear attack does not serve their purposes or forward their goals. This is the kind of argument that is normally used to assure that a *political*¹⁵ terrorist would not wish to detonate a nuclear weapon¹⁶; a political terrorist group depends on not alienating its followers by excessive violence.
2. The leadership would have preferred to detonate, but decides before delivery the probability of failure is too high.
3. The leadership has been persuaded that a nuclear attack is morally wrong.
4. The leadership decides the nuclear weapon can be more fruitfully used to blackmail a government into concessions.
5. The leadership decides the nuclear weapon is best used as a deterrent against attacks on them by an adversary government or organisation.

Let me consider al Qaida again. Dunn points to the possibility of alienation as a concern which could keep al Qaida from detonating a nuclear device. Bin Laden's stated goal is to rally the Muslims of the world to form his Caliphate, a goal towards which repulsing millions of Muslims with excessive violence would likely be counterproductive. Whether this is his real goal or not is a relevant question which lies to the side of our discussion, but whose answer clearly

matters if one were to answer the question of what use al Qaida specifically would have of a nuclear weapon. In the gaming part of this chapter I consider as usual a more generic terrorist organisation whose motivation is, for generality, part militant and part political.

Worries about alienation conforms to the conventional wisdom as regards political terrorist groups. Until the early 1990s terrorism was typically more or less equated with political terrorism, which made leading analysts such as Brian Michael Jenkins assert that terrorists would have no use of such a massively destructive weapon¹⁷. A political terrorist organisation depends on the sympathy and support of an audience of less radical bystanders, and the killing of innocents in large numbers would almost certainly not serve their purposes.

In a gaming setting this can be translated into an assumption that a political terrorist expects to extract a negative payoff from detonating a nuclear bomb, hence such action is not preferable compared to doing nothing (payoff zero). Presumably, the rational terrorist will in this case probably not have acquired the weapon in the first place.

Point 2 on the list is an example of deterrence by denial as discussed in chapter 6 and is not so important to the analysis undertaken herein. The rational terrorist will have thought through the entire plot from start to finish, including delivery, before attempting to acquire the real weapon, and all else being equal, the project's overall failure probability will be lower after successful acquisition than it was at the outset of the project. However, new information could have become available to the terrorist during the acquisition phase which changes his cost/benefit analysis, but this is not a focus point of this chapter.

There is also some reason to question, argues Dunn, whether bin Laden and his closest associates could still have some moral qualms against killing such a vast number of people (including, most likely, a number of Muslims) in one blow; point 3 above. Such considerations will be little more than informed speculation however, and one is probably ill-advised to base policy on such hopes. In a game, moral concerns would translate to lower expected payoff from a nuclear attack for the terrorist and at least entail that the assumption employed previously, that terrorists extract payoff proportional to the damage done, is not correct. Lower payoffs will serve to enhance self-deterrence as discussed in chapter 6.

15 That is, a group which uses violent means to gain bargaining power as a means to achieve political goals. See discussion in section 5.4.

16 Most famously the series of papers by Brian Michael Jenkins and co-workers: 'Will Terrorists Go Nuclear?' RAND paper (1975)*; 'The potential for nuclear terrorism' RAND paper (1977)*; 'Will Terrorists Go Nuclear?' *Orbis* 29:3 (1985) pp.507-515; 'The Likelihood of Nuclear Terrorism' RAND paper (1985); Peter deLeon, Bruce Hoffman, Konrad Kellen, and Brian Jenkins 'The Threat of Nuclear Terrorism: A Reexamination' RAND paper (1988)*. These papers argue that terrorists want attention and that a large body count is not a goal in itself to such groups.

17 See footnote 16.

Arguably, the two last points in the above list are the most interesting for the present analysis. Blackmail and deterrence are two ways by which a nuclear weapon could conceivably be of use to a nuclear terrorist without its actual detonation.

Osama bin Laden, for example, has been quoted mentioning the possibility of obtaining nuclear weapons for deterrence purposes. In 1998 he commented that 'It would be a sin for Muslims not to try to possess the weapons that would prevent the infidels from inflicting harm on Muslims.'¹⁸ Two months after the 2001 attacks on the US east coast he told a Pakistani journalist 'I wish to declare that if America used chemical or nuclear weapons against us, then we may retort with chemical and nuclear weapons. We have the weapons as deterrent.'¹⁹ To the author's knowledge, no serious analyst has believed bin Laden's claims that he already has nuclear weapons. Nonetheless, these quotes show that al Qaida has done some thinking about what to do with a nuclear weapon should they ever acquire one.

7.4 *The credible threat*

Assume now that a terrorist organisation has a nuclear weapon and wishes to use it to blackmail a government into certain actions. There are, as several other authors have pointed out in the past, a number of obstacles to overcome in order for such a plan to succeed.

Importantly, the terrorist would have to ensure that his nuclear threat is credible. Assume the terrorist has built his bomb himself. No non-state actor has ever been known to obtain a nuclear capability before, so proving that he has succeeded will not be easy. One could imagine the terrorist providing some physical proof which at least makes the successful acquisition of a bomb probable. He could send a sample of the nuclear material for a start. This could be risky, however, since nuclear forensics might be able to find out where the material originated from and hence help intelligence to track down and intercept the plot²⁰.

Nuclear material alone is not enough to have a bomb, however. A considerable technical capability

must also be in place²¹. One could imagine the terrorists making public a videotape showing production equipment and assembly, again a potentially risky enterprise, should anyone be able to recognise the location and localise the group. Indeed, in his book, John McPhee goes so far as to argue that it may be necessary for the terrorist bomb-maker to make two bombs and detonate one as a demonstration in order to be believed²². It will not be necessary to prove the existence of the weapon beyond all doubt, however - the gravity of the consequences of an attack will probably make the government take the threat very seriously unless it can certify that the threat is *not* real.

In addition to proving the bomb really exists and will plausibly work, however, comes the problem of delivering it. Simply having the weapon hidden in some faraway country may not be enough. This is why a nuclear programme in a proliferating state is normally accompanied by a missile development programme: a bomb without a means of delivery is of limited value²³. Thus a credible threat should also provide some proof that the terrorist is able to deliver the bomb to a target of great value to the government.

As has been extensively discussed by Levi²⁴ smuggling a ready-built nuclear weapon into a different country is not a trivial task at the best of times²⁵, even when the target country is not looking for it. After an official threat has been issued, the inspection of incoming goods and traffic is bound to toughen, hence arguably the most effective blackmail or deterrence scenarios seem to be such in which the bomb is already in place at or near to the target before the threat is pronounced (hidden, say, in an American city ready to be detonated) or where the target is easily reached or difficult to defend but may in this case be somewhat less valuable to the government in question (an example could be American forces in Iraq)²⁶.

I will largely limit myself to considering scenarios

18 Rahimullah Yusufzai 'Conversation with Terror' *TIME magazine* (January 11, 1999)*

19 Hamid Mir 'Osama claims he has nukes: If US uses N-arms it will get same response' *Dawn* (November 10, 2001)*

20 *Levi On Nuclear Terrorism* p.121

21 See chapter 3.

22 John McPhee *The Curve of Binding Energy* (New York:Farrar, Straus and Giroux, 1974) pp.144-145

23 Military nuclear weapons can obviously be delivered by plane or some other means of transport as well.

24 *Levi On Nuclear Terrorism*

25 See appendix D for further details about such 'second layers of defence'.

26 For deterrence purposes it is possible that the small probability that a weapon in a faraway country could be used in an attack upon, say, the US might be a risk large enough to deter attacks. But see further discussions in the following.

in which a government receives a threat that a nuclear weapon is on its way to or is hidden somewhere in one of its bigger cities. The weapon may be claimed to be a deterrent or could come with a threat of detonation unless some demand is met. Given some kind of proof which makes the reality of the threat plausible, it is likely that the government will take the threat seriously. One may remember at this point the episode discussed in section 4.4.1 in which a CIA agent codenamed Dragonfire reported that a plot was underway to detonate a 10kT nuclear bomb at the Grand Central Station in central New York. The threat, even though it was an unconfirmed report by a single agent without (presumably) any supporting physical proof, was taken very seriously²⁷. Few politicians would risk taking the blame for dismissing as bluff what turns out to be a full scale nuclear attack.

One aspect that may have led the authorities to taking the 'Dragonfire' incident so seriously was the implication that the purported weapon was already in the US somewhere. Had the threat been 'we have a weapon ready in the mountains of Pakistan and will explode it in the US unless so and so', the threat would have been less imminent, since it would have allowed the government considerable time to try and intercept the delivery somewhere between Pakistan and New York. With enough resources poured into a short term action, the government can significantly boost the probability of detecting such a big metal object weighing perhaps several tonnes and radiating neutrons and gamma-rays, although no such defence can ever be guaranteed to succeed and the staggering number of ways in which such a smuggling operation could possibly embody itself would be on the terrorists' side²⁸. In most plausible scenarios²⁹, a terrorist in his right mind³⁰ would not issue such a threat before smuggling the device, since the smuggling operation would be complicated a great

deal as a consequence of providing the crucial information that it exists.

7.5 Scenarios and incentives

Assume that a terrorist organisation has successfully hidden a nuclear weapon in a western city. What can it do with it? Our list contains two strategic options apart from detonation: blackmail and deterrence of attacks.

7.5.1 The terrorist's deterrent

If *deterrence* is to work, that is, if the planted nuclear weapon is to create a new order in which the targeted government refrains from attacking the terrorist's interests out of fear, the threat must be sustainable over a long period of time, maybe decades. For any government, having a nuclear weapon of unknown yield, safety and predictability sitting in one of its cities for the foreseeable future is an utterly unacceptable situation, and over time it is hard to imagine that the estimated cost of attempted removal of the deterrent will be enough to keep the government from taking action. For this reason the placement of a nuclear bomb in a target city, say in the US, will almost certainly create a situation too unstable to be upheld beyond the short term.

If bin Laden's talk of deterrence is to be taken seriously, what he may have in mind is to keep a nuclear weapon within the area he wishes to defend and threaten to use it tactically against invading forces or against easy-to-reach targets of some value to the attackers.

Perhaps the most likely scenario is one in which the terrorist organisation holds a nearby city hostage, to which it is credible that a nuclear weapon could be delivered. If a large portion of the casualties belong to the population group whose interests the terrorist purports to defend, it is questionable whether carrying out such a threat is in fact the terrorist's best option if deterrence fails and he is in fact attacked, and of course, once detonated the deterrent is gone. Thus the terrorist's challenge is to make this scenario *sufficiently probable* in the mind of the deterred actors. Deterrence is all about psychology, argues Freedman, and as such 'all deterrence is self-deterrence in that it ultimately depends on the calculations of the deterred'³¹. It works, in simple terms, if the terrorist threat can make the government's estimated cost of

27 See e.g. Graham Allison *Nuclear Terrorism: the Ultimate Preventable Catastrophe* (New York: Times Books, 2004) pp.1-3.

28 e.g. Matthew Bunn and Anthony Wier 'Terrorist Nuclear Weapon Construction: How Difficult?' *The Annals of the AASPP* 607 (2006) p. 142 or Matthew Bunn *Guardians at the Gates of Hell: Estimating the Risk of Nuclear Theft and Terrorism - and Identifying the Highest-Priority Risks of Nuclear Theft* PhD dissertation (Massachusetts Institute of Technology, 2007)* p. 157

29 Possible counterexamples might be such in which the terrorist lied about the location of the weapon to divert the government's defensive measures.

30 Some would argue that terrorists are never in their right mind. But see the discussion of terrorist rationality in section 2.9.

31 Lawrence Freedman *Deterrence* (Cambridge: Polity Press, 2004) p. 30

attacking worse than the continued existence of a nuclear armed terrorist organisation. During the Cold War the consequences of deterrence failure could have been all-out nuclear annihilation, a scenario so bleak that it rendered almost any other option preferable, even if there was a significant probability that an attack would not in fact be retaliated despite the threats. Whether a terrorist could achieve the same end with a single weapon or whether, on the contrary, the acquisition of nuclear arms by a terrorist organisation would provoke 'pre-emptive' attacks upon it, is an important and many-faceted question which will only be considered briefly towards the end of this chapter for reasons of manageability.

7.5.2 *The likely candidate: blackmail*

I will henceforth consider threats of attacks directed against the government itself. In this case I argued that the more likely strategy would be blackmail. One can imagine, for example, an al Qaida nuclear weapon planted in a US city with the threat that it will be detonated unless the US pulls out of Afghanistan. The US government would then be faced with difficult questions. Accept withdrawal or risk a potentially devastating attack? And even if the government does as the terrorist demands, can it trust that the bomb will not be set off anyway? It is worth considering such a scenario more closely.

Accepting withdrawal or some other concession will come at a considerable loss of political prestige. The US National Security Strategy, for example, explicitly states 'The United States will make no concessions to terrorist demands and strike no deals with them.'³² Yet if the consequence is half a million people dead³³, many would argue that such principles should no longer be upheld. The President who dares to stand fast in the face of such a threat will also risk bearing some of the blame for the subsequent attack. On the other hand, pulling forces out of a country is not a quick operation and the time it takes will allow the government to try and hunt down the plot and derail it³⁴. The terrorist must be able to uphold the threat throughout the withdrawal process, otherwise the government will have no further reason to carry through with their promise. Thus one would think that demands for concessions which can be given

almost immediately (the release of prisoners, say) would be more likely to succeed than such long term concessions, all else being equal.

Also, if the terrorist really wants concessions he must create a real incentive for the government to meet his demands. This will surely involve handing over the nuclear device when the demands are deemed to be met, since, if he gets to keep the weapon he can issue new demands again and again, an unacceptable situation for the government. Furthermore, having received a threat the government will work very hard to localise it and, at least if the threat was made publicly, evacuate people from the area (if the terrorist threat is publicised, people living in the large cities may have started fleeing to the countryside on their own initiative). Armed with time and the knowledge that the plot exists the government has a fair chance of locating the bomb.

7.5.3 *Alternative: the fake blackmail*

It is important to recognise at this point that the terrorist can have two very different strategic motives for issuing a nuclear blackmail. On the one hand he could genuinely want the concessions involved, for example the withdrawal of forces or release of prisoners. On the other, if the terrorist realised the demands would not realistically be met, the plot could be devised so as to create some justification for the attack and at the same time place some of the blame on the government in question. This option is missed in all of the literature on nuclear terrorism that the author is aware of.

By issuing a well publicised blackmail threat the terrorist shifts some of the blame for the subsequent attack onto the target government, should it not concede to terrorist demands. It is not difficult for the terrorist to ensure that the government does not concede: the threat can be made to appear incredible (e.g. containing poor or faulty information) prompting the government to dismiss it as a hoax, or the demands could be greater than what the government could possibly grant in the face of which the government has little choice but to take military action. If the terrorist announces that the bomb is in 'a major US city', for example, mass hysteria could be achieved. When the detonation subsequently takes place, additional damage will have been caused to the government beyond the attack itself, and the terrorist may have harvested some sympathy for ostensibly giving the government a chance to avoid the terrible outcome.

32 President George W. Bush *The National Security Strategy of the United States of America* (Washington DC, 2002)* p.5

33 See discussions in section 4.4.1.

34 Levi *On Nuclear Terrorism* p.121

This scenario is interesting not only because it has not been considered in the literature previously, but also for having a peculiar property: When talking of attempts to extort concessions by nuclear hoax one thinks of individuals attempting to gain something by blackmailing the government based on an insubstantial threat. The United States has faced a number of such hoaxes over the years and none of them have even been publicly commented on by the government.³⁵ In the scenario just laid out, on the other hand, it is not the nuclear threat that is a hoax, but the blackmail itself, since no concessions are in reality expected. I will henceforth term the former a 'hoax blackmail' and the latter a 'fake blackmail'.

The ploy is somewhat risky on the terrorist's part, and perhaps for that reason unlikely based on the history of terrorist risk aversion discussed in the previous chapter. The plot means waiting in place probably for several days after the threat has been issued before detonating, while the government may be working to track down the weapon.

7.5.4 *Public or tacit blackmailing?*

Arguably, the fake blackmail strategy discussed above depends on the threat being publicly disseminated to be sure to be effective³⁶. If, on the other hand, the terrorist really wants concessions, it is at least arguable that he will stand a better chance if the threat is communicated to the government secretly.

As previously mentioned, a true blackmail must create a real incentive for the government to concede. In short, the government must estimate that the large political cost of negotiating with terrorists and the expected outcome of doing so is in fact preferable to an even less desirable outcome. This presents the terrorist with a very difficult exercise in communication: he must make the government believe that while he is perfectly willing to slaughter

perhaps a hundred thousand citizens in a single blow, he is nonetheless to be trusted to keep his part of the bargain and hand over the weapon if demands are met. Thus it is necessary for the terrorist to make every effort to convey the message that it is 'playing fair'. Adding enormous pressure on the government by creating public panic is hardly useful towards this end.

At the outset the government appears to have every reason to doubt the terrorist. Unless satisfactory proof is provided, the threat will likely be thought to be a hoax. As previously discussed, the terrorist must make a convincing case for the reality of his nuclear capability whether it really exists or not. A question the government should and probably would ask is why they are seeing a blackmail threat instead of a straightforward attack. An obvious candidate answer, at least when the terrorist is of a type believed likely to wish to kill in large numbers, would be that he does not believe the weapon will work and turns to 'plan B' instead. For this reason a hoax blackmail must be made very believable in order to stand any chance of producing concessions, a conclusion supported by the failure of nuclear hoaxes over the years to produce any response from the US government³⁷.

With proof provided, the government still has reason to doubt whether the terrorist will keep his part of the bargain and hand the weapon over. If the terrorist is known to be militant and extremist, like al Qaida, the government must always fear that its intention is to use the weapon to extort all that it can before finally detonating it, in which case the government would have been better off using force to try and stop the plot right away.

If the organisation is known to be of a more political nature, on the other hand, the situation is somewhat less clear, yet as argued by many in the past³⁸ and discussed further below, a politically motivated group will have a very serious disincentive to actually explode the bomb as it will arguably not serve its purposes, for which reason the threat will not be credible even if the weapon itself is real. Most likely a group with a political agenda will not find a nuclear project worthwhile in the first place, however, and the threat is likely to be a hoax. While such logic makes good sense in principle, to what extent a government will dare to take the risk of trusting the terrorist's own cost/benefit analysis to keep it from

35 Jeffrey T. Richelson 'Defusing Nuclear Terror' *Bulletin of the Atomic Scientists* 58:2 (2002) pp. 38-43. Also *Levi On Nuclear terrorism* p. 120

36 One could of course imagine the terrorist organisation warning the government tacitly, setting off the bomb and then making the public announcement that the government had been warned. The government could deny that any threat had been received, however, whereas a public warning would have been indisputable proof. As importantly, the terrorist would not achieve the massive panic which would make every citizen aware of the government's predicament, possibly persuading some that the government is at least partly at fault for the dreadful outcome.

37 Richelson 'Defusing Nuclear Terror'

38 e.g. Jenkins 'Will Terrorists Go Nuclear?' (1975). See footnote 16.

exploding the weapon is an open question. For much the same reason as Jenkins and co-workers, I find the scenario of a believable nuclear threat from a political terrorist organisation unlikely and will not consider it further.

according to the simple model, forceful response is at least as good as dismissal if

$$p \geq \frac{C_R}{b}. \quad (7.1)$$

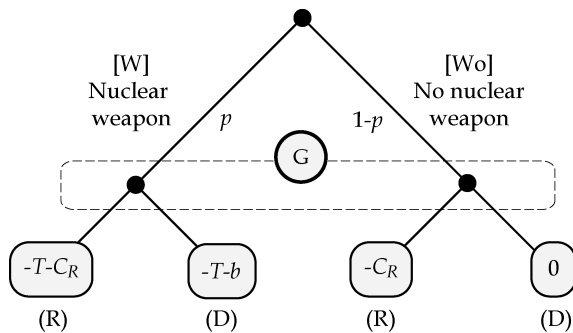


Illustration 7.1.: Simple decision theoretical model of nuclear terrorist blackmail.

While this argument may not encompass the full set of possible scenarios, it strongly suggests that if the terrorist really intends to go into negotiations, it will communicate its threat secretly, in which case the fake blackmail scenario previously analysed is out of the question. If one buys this argument, it leads to an obvious policy implication:

A sufficiently credible nuclear blackmail threat *issued publicly* should be met with force and every effort to try to hinder the threatened attack.

What exactly is meant by 'sufficiently credible' can be formalised by a simple decision theoretical model similar to previous efforts. Consider the situation in figure 7.1 where a threat has been made public by a terrorist who either has a real weapon (state denoted [W], 'with') or does not (denoted [Wo], 'without'). I assume for simplicity that the terrorist with a weapon will always successfully detonate it. The government estimates that [W] has a probability of p and [Wo] of $1-p$. The threat being public, the government concludes (in accordance with the above argument) that concessions are out of the question and chooses between responding with force (denoted (R)) and dismissing the threat (denoted (D)). An attack has a cost of $-T$, with an additional political cost of $-b$ in the case where G dismisses a threat which is real. The cost of forceful response is $-C_R$. The method of analysis is simple and so similar to that of previous chapter that I will skip the details. It is easily verified³⁹ that

Arguably b is likely to be much greater than C_R , in which case dismissal is preferable only when the government is very sure that the threat is a hoax. A refinement of the game in which the government has a finite probability of actually hindering the attack will only strengthen this conclusion.

Thus a straightforward policy conclusion is reached: a publicly announced blackmail whose reality has an estimated probability satisfying (7.1) should be met with a forceful response. Indeed, the indication is that a publicly announced nuclear extortion attempt is in fact either a hoax or a fake blackmail, the latter merely one of the many ways in which a nuclear terrorist attack may embody itself. In the remainder of the chapter I will assume that the threat is issued quietly.

7.6 Strategic interplay of nuclear blackmail: a gaming approach

In order to try and formalise the strategic logic of nuclear blackmail we will now introduce a game theoretic model belonging to the class of games commonly called signalling games. This section will inevitably be tough reading for the reader with no previous experience of this type of game, for which reason I will go through the main concepts of such games as the analysis progresses. Signalling games are not normally difficult mathematically, but can be conceptually challenging, and the reader may wish to refer to an introductory textbook in game theory at this point⁴⁰. I will start the analysis with possibly the simplest non-trivial game which captures some of the essentials of a nuclear blackmail situation and continue by making gradual generalisations of the game to study more detailed aspects of the scenario in keeping with the qualitative discussions above.

As discussed above the game will only deal with the question of nuclear blackmail, and a number of assumptions will be made regarding the situation to be modelled. I assume a nuclear threat is communicated secretly to a government. The scenario

and that for (D) is $U(D) = -pT - pb$, from which (7.1) follows immediately through the inequality $U(R) \geq U(D)$.

³⁹ The expected payoff from opting for (R) is $U(R) = -pT - C_R$,

⁴⁰ e.g. Gibbons *A primer in Game Theory*

one could have in mind is one where a bomb is claimed to be present in or on its way to a major city in a western country. This is the scenario which will be used when making plausibility arguments about the quantitative relations between parameters, but the model is in itself general enough to also capture blackmail against the government's interests outside her own territory (with different values of parameters as appropriate).

We denote the government player G and the terrorist as player T. Since, unlike in previous chapters, the government moves last, player G is assumed female and the terrorist male in this chapter per convention.

After receiving a nuclear threat with a demand for concessions, the government is assumed to have three possible courses of action:

- *Respond forcefully* (R): the government responds forcefully either domestically or abroad in an attempt to hinder the attack.
- *Concede to demands* (C): Concessions are pledged to meet the terrorist's demands in an attempt to hinder the attack.
- *Dismiss terrorist threat* (D): Threat is dismissed as a hoax and no action is taken.

As is standard in game theory in games of incomplete information we assume that there are two *types* of terrorist: the type who really has a nuclear weapon, and a type who does not. The two types have the same sets of possible actions but with different payoff functions.




Illustration 7.2.: The general layout of the two-player nuclear blackmail game.

Later the rounded boxes will contain the payoffs for the two players for each outcome of the game instead of the dots; player G's payoffs above, T's payoff below.

The general game is laid out in figure 7.2 and may be summarised thus: First nature picks the terrorist type from a set of two options: the type *with* nuclear

weapons, [W] and the type *without* [Wo]. [W] is picked with probability p and [Wo] with probability $1-p$. Player T, knowing his own type, then decides whether to blackmail (B) or attack (A). If he has no nuclear weapon [Wo] the attack has no effect (equivalent to doing nothing). If action (B) is chosen, the government can choose from a set of three different actions in response, (R), (C) and (D) as described above.

Each player has a *strategy* which is a set of actions, one for each possible state the game could be in at the time of decision. Since in our game each player only moves once⁴¹ so there is no difference between a strategy and an action. In this simple game there are four different possible sets of strategies for player T: two possible actions for each of the two types. There is only one type of G which has three different possible actions/strategies.

7.6.1 The perfect Bayesian equilibrium: definition

The standard way to analyse signalling games is to look for so-called perfect Bayesian equilibria. A very readable introduction to this type of games is found in Gibbons' book⁴². For the reader who is well versed in signalling games, this section can safely be skipped.

An equilibrium of the Bayesian game⁴³ considered here consists of the following information:

- A set of strategies, one for each player. The terrorist's strategy, remember, consists in turn of two actions, one for each type.
- The belief ρ of player G.

Player G has a *prior* belief at the start of the game, where ρ equals nature's probability p , and a *posterior* belief after a signal is received. Roughly, the primary components of the perfect Bayesian equilibrium (PBE) are that

- It is a Nash equilibrium, that is, for each type of player T, T's strategy is a best response to player G's strategy, and given player G's belief ρ , G's strategy is also a best response to

41 This is not strictly true, as will be seen, since the terrorist [W] is given a choice whether or not to detonate in the case where the government pledges concessions. This choice is trivial to deal with, as will become obvious, and we can play the game as if the terrorist only makes a single choice, between (A) and (B).

42 Robert Gibbons *A Primer in Game Theory* Chapter 4

43 i.e. games of incomplete information, that is, when a player is uncertain about the preferences of another player.

player T's set of strategies. Put in other words: given that all other players and types follow their equilibrium strategies, no player has an incentive to deviate from her/his equilibrium strategy.

- It is Bayesian, that is, player G has a belief about the type of player T, and the belief follows Bayes' rule. Bayes' rule is a mathematical theorem which dictates how a player should update his or her beliefs in a consistent way as a consequence of new information, in this case a *signal*. In the present game the signal which G receives from T is whether (B) or (A) has been chosen⁴⁴. That the equilibrium is Bayesian simply means that the belief of player G must be consistent with the information she has⁴⁵. The information G has here is: what T's strategy is⁴⁶ and what the payoffs are for different outcomes.

One can think of the equilibrium as if the game is played out on paper and the players may change their strategies in response to each other any number of times until both players have no reason to change further. If there is a unique equilibrium, this will have been reached; if there are several equilibria, one of them will have been reached - which one depends on how the gameplay was started, and if there are no equilibria the process will not converge at all. An analogy to the game with a single Nash equilibrium may be a marble set to roll in a bathtub; after a while it will settle in the drain which is the only equilibrium position. The randomness of the roulette wheel, in contrast, is provided by the fact that there are many equilibrium positions in which the ball may finally settle, the thirty-odd 'pockets' of the wheel, and which

one it settles in is determined in a complicated way by exactly how the ball is thrown. Repeated play of a game with no equilibria may behave something like a cat-and-mouse chase; the cat always has an incentive to change its position to get to where the mouse is, for which reason the mouse has an incentive to change its position all the time.

For a more rigorous definition of the PBE, the reader is referred to textbooks in game theory. Several refinements of the equilibrium are available, yet I will as a rule only introduce the minimum of general game theory required for the task at hand.

I will denote T's strategies as for example (A, B), where the first strategy refers to type [W] and the second to type [Wo]. T's strategies (A,A) and (B,B) are called pooling strategies since both types of T choose the same action, and the strategy couples (A,B) and (B,A) are called separating strategies since they dictate different actions for the different types. Note importantly that even though T could well employ the strategy (A,A) which means that G does not get to act at all, G has a strategy nonetheless, which *would have been* used if T's strategy had involved blackmail. This is just what happens in real life: the US is stating that it will never concede to terrorist demands, a policy which exists whether such demands are made or not.

The process of finding an equilibrium goes as follows. First one assumes that a set of strategies is a candidate PBE (in this game there are 12 possible candidates: G has three strategies for each of T's four). Each candidate is then compared to the criteria for a PBE⁴⁷ and candidates not satisfying these are discarded. Remember that a PBE consists of a strategy profile and player G's belief. Specifically the definition of the PBE requires that G cannot believe with finite probability that the state of the game at time of decision is one that could not be reached unless T did not follow his equilibrium strategy. In other words, an equilibrium cannot depend on a player believing it is not an equilibrium.

That the belief is Bayesian in the context of our simple game means only that given a candidate equilibrium strategy profile, G must believe that the state of the game is one which can possibly be reached by all players playing according to this profile. Before the game starts G holds a *prior* belief that the probability of [W] is equal to that with which nature picks between the types, that is $p=p$. Given a signal and a candidate equilibrium, however, G must update her beliefs. If T's strategy is, say (B,A), then if G

44 In the latter case it does not matter what the government chooses to do, but this makes no fundamental difference to the way the game is played. If one pleased, one could equivalently let G choose between (R), (C) and (D) following both (A) or (B), but in the former case let the payoffs be identical independently of G's action.

45 The reader may be confused by looking up the actual Bayes' rule at this point, which in its general form represents methodological overkill in such simple games as these. The concept of consistent beliefs will be treated in a more pragmatic way in the following, suitable for the simple signalling game.

46 This statement is a little simplistic as will be explained. In fact the process of finding equilibria goes 'backwards' so that one starts by assuming what the equilibrium is, in which case all strategies are assumed known. Candidate equilibria which do not satisfy the necessary criteria are thereafter discarded until one is left with the true equilibria.

47 See e.g. chapter 4 of Gibbons *A Primer in Game Theory*

observes a signal (B) (that is, receives a blackmail) it implies that T is of type [W] and thus the posterior belief must be $\rho=1$. Any other value would imply that G believed T might not play according to his equilibrium strategy set (B,A) after all. Likewise for (A,B), a signal (B) implies $\rho=0$. Since G is allowed no action following (A), we need not consider this signal here.

7.6.2 The normative value of an equilibrium

An important question which has been a source of debate for some time is what normative value one can ascribe an equilibrium of a game, should one exist. In the previous cases of decision theory it is relatively unproblematic to think of the solutions as the way the player in question *should* play if he/she is perfectly rational and intelligent. In a multi-player game, however, the interpretation is not so simple, because, as Goeree and Holt put it, 'the best way for one to play a game depends on how others *actually* play the game, not on how some theory dictates that rational people *should* play.'⁴⁸ These authors compare ten experiments where test persons are given one chance to play a given game, and note very significant deviation from the play dictated by the games' Nash equilibria⁴⁹.

This problem is serious, and there may be no way to circumvent it, at least in general⁵⁰. In the present game, however, I will argue that there are certain conclusions one may draw based on plausibility analysis of the dynamics of the game. Even if the game has a unique Nash equilibrium⁵¹ it is not in general true that a player is always better off playing according to this equilibrium⁵², yet in cases where a certain course of action is *dominated* by the other strategies (i.e. it is under no circumstances the preferable one), it would seem on intuitive grounds as a safe normative prescription that this particular course of action should *not* be chosen. In the game

48 Jacob K. Goeree and Charles A. Holt 'Ten Little Treasures of Game Theory and Ten Intuitive Contradictions' *The American Economic Review* 91:5 (2001) p. 1419

49 *ibid.*

50 Nash himself allegedly reached this conclusion. *ibid.* p. 1419.

51 If there are several equilibria, naturally, drawing normative conclusions becomes even more dubious *cetera paribus*. We will find that in the present game there is always only one PBE for a given set of parameter values.

52 This is only so provided all other players play according to the equilibrium.

below we find that one strategy nearly almost dominates the others.

Secondly, upon finding a unique equilibrium one may apply qualitative analysis to assess the 'safety' of playing according to the equilibrium. The relevant question for the government to ask would be 'how grave are the consequences if I play according to equilibrium and it turns out T does not?'. In certain situations G may be fortunate and find that even if the strategy employed by T is off the equilibrium path, the equilibrium response is still the best available, in which case the equilibrium arguably has a normative value. On intuitive grounds it is reasonable that the worse the loss from erroneously assuming T to play rationally, the less the normative value of an equilibrium is. It turns out that I am lucky in this case and while game theory is in general not guaranteed to yield normative prescriptions, the equilibria found in this particular game arguably hold some clear advice for a government facing nuclear blackmail.

7.6.3 The basic game

The first game we will consider is designed to be as simple as possible while still capturing some of the dynamics of the extortion situation. A number of simplifying assumptions are made. The following assumptions will hold throughout the gaming:

- The terrorist is assumed to be bloodthirsty and to extract a positive payoff from detonation.
- The terrorist's cost of acquiring the weapon is not taken into account. This is arguably a reasonable assumption since this sum is already sunken cost for the terrorist.
- The terrorist type with a nuclear weapon can also choose not to issue the threat, but simply detonate without warning (strategy A). The terrorist without a nuclear weapon can also choose to either issue a hoax blackmail (strategy B) or do nothing (also strategy A for simplicity).

The following assumptions will furthermore be made for now and be relaxed later:

- Forceful response by the government will successfully derail the attack. This is obviously unrealistically optimistic as discussed later.
- The real and political cost to the government in responding forcefully is assumed negligible compared to other quantities.

- If the government dismisses the threat and a nuclear attack subsequently occurs, no additional political cost is suffered.
- The terrorist receives no additional payoff for incurring additional terror or political damage beyond the attack itself, such as by a fake blackmail considered in section 7.5.3.

This leads to the game depicted in figure 7.3. The symbols are explained as follows (all quantities are defined positive or zero):

$-T$	G's cost from a nuclear attack
$-C$	G's political cost of conceding to terrorist demands
A	T's payoff for successful detonation
$-\phi$	T's general cost of failure (apprehension, destruction, etc.)
G	T's payoff for harvesting concessions
p	The probability that [W] is picked, equal to G's <i>prior</i> belief that the terrorist has nuclear weapons.
ρ	G's <i>posterior</i> belief that the terrorist has nuclear weapons, given T's strategy.

It is of some importance to bear in mind that it matters *who* estimates the parameters above. Since my goal is to inform policymaking, I assume the parameters are all as estimated by the government. In particular, A and G are the payoff the government *believes* T will extract in these cases, as opposed to T's actual preferences which G cannot know.

Note that the form of the game implies that $A > 0$. If A were negative, the terrorist with a nuclear weapon would rather do nothing (payoff 0) than detonate the weapon, and the game would need to include such an option for [W]. This means the present game implicitly assumes that T is not a purely political terrorist, but has at least a distinct militant streak.

For generality, if G pledges to concede to the terrorist's demands, T is left with the choice whether to detonate the weapon anyway or to give up the weapon and collect the rewards (I assume these are of such a nature as to be impossible for G to take back after the nuclear weapon is given up). In the case where T breaks his promise and detonates anyway, his payoff is A , the same as for a simple attack. This is mainly for simplicity, but could possibly be interpreted thus: while he may have been able to

collect concessions the benefits from this are cancelled by loss of support amongst potential followers due to unchivalrous behaviour (presumably the public will eventually be made aware of the manner in which the drama took place).

One sees from figure 7.3 that in the case where the threat is real, [W], and G decides to concede, (C), T will detonate anyway if $A > G$ and will not if the opposite is true⁵³, and we will consider each of these cases separately. Note, crucially, that whether or not blackmail can succeed depends *not* on what the terrorist's *real* preferences are, but what G *thinks* they are. This will become clearer in a moment.

Assume now that $A > G$, that is, G believes T will detonate his weapon even if concessions are pledged. I consider G's options in the case that a threat *has* been issued, that is T has played strategy (B). G's expected payoffs for the different actions available given her belief ρ are now:

$$\begin{aligned}
 U(R) &= 0 \\
 U(C) &= -\rho T - C \\
 U(D) &= -\rho T.
 \end{aligned}
 \tag{7.2}$$

I now introduce the intuitive concept of dominated strategies, which are strategies that are *never* the preferable one. In (7.2) we see that, excepting the case where $\rho=0$, $U(R)$ is always the highest obtainable utility, that is, whatever T does, G's best response is to respond forcefully. In other words, (C) and (D) are strictly dominated by (R) except when $\rho=0$.

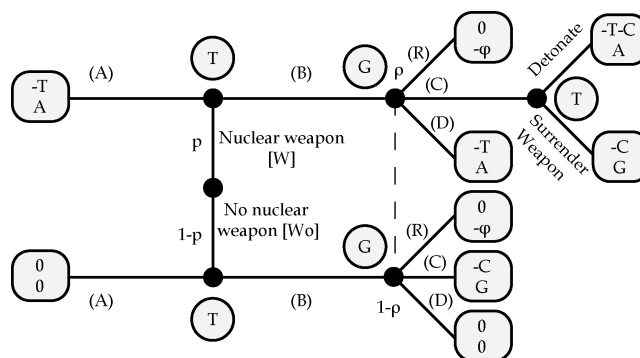


Illustration 7.3.: The game in its simplest embodiment.

This conclusion, reached without any detailed analysis of the game, still holds if we let $A < G$ (since $A > 0$ the terrorist is still believed to wish to kill large numbers of people, but is even more interested in the

53 Since all parameters in the game must be understood as the players' best estimates and therefore somewhat coarse, 'critical' cases where one parameter is exactly equal to some number are irrelevant for practical purposes.

concessions). The payoffs are now

$$\begin{aligned} U(R) &= 0 \\ U(C) &= -C \\ U(D) &= -\rho T. \end{aligned} \quad (7.3)$$

In this case it is no longer true as in (7.2) that (D) also strictly dominates (C), but still $U(R)$ is the best strategy except the case of $\rho=0$ when G is indifferent between (D) and (R).

There are therefore the same two equilibria in both of these games. The first of which is the special case $\rho=0$

$$[(A, B), D, \rho=0]$$

in which the terrorist only blackmails if he does not have a nuclear weapon, whereupon the government dismisses it and nothing happens. This is not an important equilibrium since, in a slight refinement of the model the terrorist will extract a small bonus for the political damage incurred by a government when it dismisses a threat which turns out to be real. In this case the equilibrium no longer holds. It is not a stable equilibrium since it hinges on [Wo] choosing to blackmail even if he is indifferent whether to blackmail or not. The equilibrium, finally, holds little normative value because in practice it requires G to be absolutely certain that T is playing the separating strategy (A, B) - if he is wrong about this, he could in the worst case find himself dismissing a real threat. With reference to the discussion in section 7.6.2 therefore, I can safely discard this equilibrium on qualitative grounds and turn to the more interesting one.

The other equilibrium is

$$[(A, A), R, \rho=p] \quad (7.4)$$

which is the equilibrium whose stability will be tested in the remainder of the gaming section by gradually generalising the game. In this game the terrorist does not issue any blackmail but simply attacks if it has a nuclear weapon. The government on its side has resolved to respond forcefully to any nuclear blackmail threat. Of course, if T follows the strategy (A,A) no such threat will ever be issued in the first place, and so (R) must here be understood as a threat from the government.

Note how this equilibrium chimes well with the declared US policy to not negotiate with terrorists. As such the following analysis tests the robustness of the conclusion that such policy is a best response to

nuclear blackmail threats and thus has very direct implications for the stance of the US and other potential target nations in the face of such threats of extreme violence.

7.6.4 First generalisation: forceful response not guaranteed to work

The equilibrium (7.4) was reached with great ease and appears, within the confines of the very simple game of figure 7.3, to be a very stable equilibrium. However, some of the simplifying assumptions made clearly act to strengthen the assertion that (7.4) represents normatively rational play of the game. In the following I will therefore test the hypothesis that the uniqueness of the equilibrium (7.4) is due to simplifying assumptions, and that more nuanced game-play will emerge upon relaxation of these assumptions.

One such assumption is that forceful response to a nuclear threat is guaranteed to successfully stop the attackers, an assumption which is clearly unrealistically optimistic and makes the government's threat of forceful response more persuasive to the terrorist than is the case in reality. Let me therefore generalise the model so that the efforts by G to stop the attack has a probability P of succeeding.

To avoid excessively lengthy analysis, I also introduce at this stage an additional political cost b for G in the case where he dismisses a threat which turns out to be real, and a bonus β for T when this happens. The resulting game is depicted in figure 7.4.

Consider once more the case of $A > G$. Now the payoffs for G following (B) are

$$\begin{aligned} U(R) &= -\rho(1-P)T \\ U(C) &= -\rho T - C \\ U(D) &= -\rho(T+b). \end{aligned} \quad (7.5)$$

Because of the restrictions on the values the parameters can take the reader will easily verify that

$$\rho(1-P)T < \rho(T+b), \quad \rho T + C$$

and therefore once again (R) is the best response except in the special case $\rho=0$, just like before (which is worse amongst concessions and dismissal is now not certain). Note that neither the introduction of $P < 1$ nor $b > 0$ can change this result. This time around there is no separating equilibrium where T has strategy (A, B)⁵⁴. The only equilibrium is therefore

⁵⁴ If T's strategy is (A, B), G will be indifferent between (R) and

once again (7.4), which holds for all legal values of the gaming parameters. In short, when $G > A$, the equilibrium is arguably even more soundly established than before.

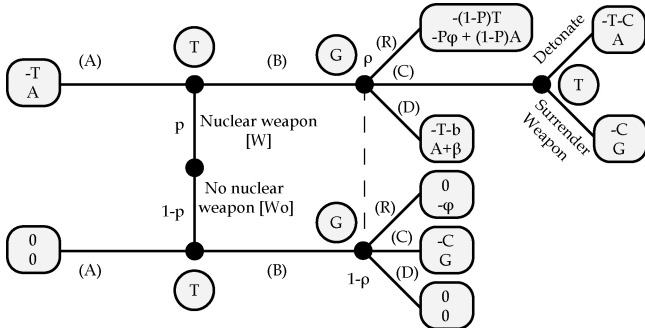


Illustration 7.4.: Generalised game: response could fail, dismissal has cost.

In the opposite case in which $G > A$, payoffs read

$$\begin{aligned} U(R) &= -\rho(1-P)T \\ U(C) &= -C \\ U(D) &= -\rho(T+b). \end{aligned} \quad (7.6)$$

This is a little more interesting since while (R) dominates (D) except for $\rho=0$, (C) can be preferable to (R) depending on the value of ρ . One easily verifies that this is so if

$$\rho \geq \frac{C}{T(1-P)}. \quad (7.7)$$

In this case there is an equilibrium

$$\left[(B, B), C, \rho = p; \quad p \geq \frac{C}{T(1-P)}; G > A \right] \quad (7.8)$$

This equilibrium is interesting and captures some of the dynamics of the game. The situation here is this: the terrorist has succeeded in making the government believe that it truly wants the concessions more than it wants to kill large numbers of innocents, and will in fact give up its weapon in return for government co-operation. The government considers that the terrorist is likely to possess a real nuclear weapon and also that an attempt to try to disarm the blackmailers is quite likely to fail. If she moreover considers making concessions to terrorists a considerably lesser evil than a successful attack, this equilibrium could hold true.

(D) following a signal (B) since she knows there is no real bomb. If he plays (R), however, [Wo] is better off changing strategy, so the only possible equilibrium would be [(A, B), D, $\rho=0$]. However if G is to play (D), [W] is better off changing strategies from (A) to (B), hence this is not an equilibrium.

While depending on a somewhat delicate balance of parameters, this is the situation in which the terrorist could possibly be able to successfully blackmail the government. As argued, however, it may be very difficult to persuade G that $G > A$.

If the inequality (7.7) is reversed, however, the only equilibrium is one like (7.4):

$$\left[(A, A), R, \rho = p; \quad p \leq \frac{C}{T(1-P)}; G > A \right] \quad (7.9)$$

The case of a bloodthirsty terrorist for whom $A > G$, one quickly verifies that (R) once again dominates both (C) and (D) and there is only one equilibrium, namely

$$\left[(A, A), R, \rho = p; \quad A > G \right] \quad (7.10)$$

similar once more to (7.4).

7.6.5 Generalisation: attack without warning has non-unity probability of success

As argued above, T will have great difficulty convincing G that while it is willing to commit such an extreme atrocity as to detonate a nuclear bomb killing thousands, it will nonetheless co-operate if G gives concessions. I turn therefore henceforth to the likely case where $A > G$ (in G's estimation).

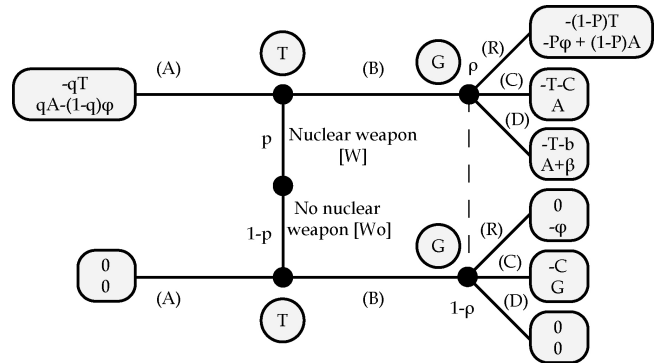


Illustration 7.5.: The game assuming $A > G$ and nonzero direct attack failure rate.

Another assumption employed hitherto is that if the terrorist chooses to attack without deliberate warning he has a 100% chance of success. Since this unrealistic assumption clearly makes the option of simply attacking more lucrative to T than it really is, there is reason to investigate whether generalising this will change the conclusion thus far that if $A > G$ in G's

estimate, the only plausible equilibrium (which also, as argued, has considerable normative value) is one with strategy profile [(A,A), R].

Letting the probability of attack success be q and assuming $A > G$ the game of figure 7.4 is obtained. As before, for G, (R) dominates the alternative strategies whenever $\rho > 0$ in this case, and is arguably the only rational response to a blackmail threat. Knowing G will respond forcefully, however, [Wo] will never be better off blackmailing, but [W] might. Assuming G plays (R), (B) is preferable to (A) for [W] iff

$$qA - (1-q)\varphi \geq A(1-P) - P\varphi,$$

that is, assuming $A > \varphi$,⁵⁵

$$q \geq 1-P.$$

In other words, as long as the probability of a successful attack with no warning is greater than the probability of a successful attack in the face of a forceful response from G, the equilibrium (7.5) still holds. It is reasonable to assume that, except in very special cases⁵⁶, this is always so.

7.6.6 Final generalisation: cost/penalty for failed interdiction

By now it will seem that, at least in the case of the bloodthirsty terrorist who is more interested in attacking than concessions, the only rational play is for the terrorist to not issue a blackmail and for the government to affirm that it will respond forcefully to such a threat should it occur. My final attempt to explore the limits of stability of this equilibrium is to let G have a considerable cost when responding with force. This cost could be the actual price of a large operation and a political cost for not heeding the terrorist threat and thereby placing the population in jeopardy. I denote the cost $-C_B$. The game now looks as depicted in figure 7.6.

With these payoffs one readily finds that following an observation of (B), (R) is preferable to (D) for G if

$$\rho > \frac{C_B}{b+PT}.$$

Since T is a large number, this is probably true except when the probability of the terrorist having a real nuclear weapon is assumed to be very low, C_B is unreasonably large, or when T employs strategy (A,B) so that $\rho = 0$. The latter case is no equilibrium here, however⁵⁷.

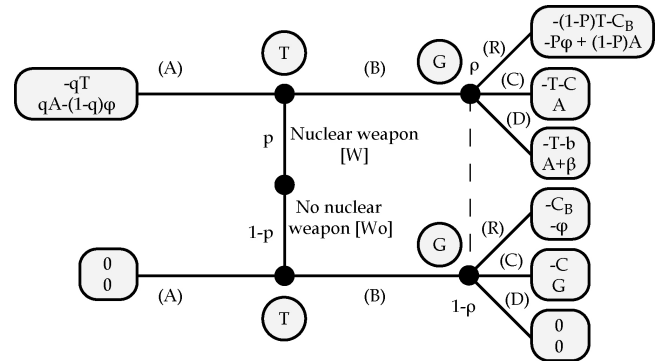


Illustration 7.6.: The game of fig. 7.5 with forceful response penalty.

Similarly, (R) is preferred to (C) iff

$$P\rho > \frac{C_B - C}{T}.$$

This, similarly, is also expected to hold true except for P and ρ very close to 0. Indeed, the more probable estimate will put C greater than C_B (the political cost of agreeing to terrorist demands is arguably expected to be higher than that of responding forcefully to a threat as outrageous as that in question), in which case the fraction is negative and the inequality is true for any $P \geq 0$.

In conclusion, unless unreasonably large, a taxation of the option of forceful response does not render the equilibrium (7.4) notably less stable.

7.7 The merits of modelling

As in previous chapters, I will discuss briefly the gains of applying the chosen methodology, in this case game theory, to the problem at hand. In the present chapter, one does well to notice that much of the chapter is in fact purely qualitative, and many of the new insights stem exclusively from this first part of the chapter. It is only after analysing the different

55 I assume this is true without much discussion. If this is not so, the rational terrorist will not start a nuclear project in the first place, since it will probably have a negative expected payoff and certainly a very large operational risk to the terrorist. See chapter 6 for further discussion.

56 One could imagine the terrorist very cleverly leading governmental efforts on the wrong track thus managing to lower rather than increase the chance of successful attack subsequent to the warning.

57 With a non-zero C_B , G's best response to (A, B) is dismissal, in which case [W] is better off choosing (B).

options in this way that the more specific scenario of blackmail was chosen for gaming.

The blackmail scenario has all the typical traits which point in the direction of a classic signalling game. One player sends a signal which the other player acts upon. This is exactly the kind of scenario which signalling games are designed to deal with, and if one is to study the question of nuclear terrorism blackmail, such a game is a reasonable choice for a first approach.

As I discussed in section 7.6.2, whereas the decision theory exercises in the previous chapters have a clear normative value, such is not guaranteed from a two-player game such as that used in the latter half of this chapter. The fact that the equilibria of the game turned out to hold normative power (according to the pragmatic criteria laid out in section 7.6.2) may be seen as a strike of luck. Had such not been the case, if for example the consequences of erroneously assuming the terrorist to follow an equilibrium path were devastating whereas the equilibrium would not have been, the game would not have proven so useful even if it might still have provided useful insights. It is hard to know beforehand whether this will be so, and the easiest way to find out is perhaps simply to perform the analysis.

An assumption used throughout this chapter is that both the government and the terrorist may be reduced to single rational actors. This assumption and its ramifications were discussed in general terms in chapter 2. In the present setting we found that this simplification may not be so essential for two reasons. Firstly, the analysis is normative and describes a scenario which has not yet happened. Hence one could always argue that, although due to the interaction of various sub-level actors the decision made is not always that of a single rational actor, it *should* be. The recommendation is for the *outcome* of the state's decision making, and how that decision is reached is a separate question. Secondly, one may reasonably argue that, due to the scarce communication between the players, the terrorist and government can be assumed to have little knowledge of the way in which the adversary makes decisions.

Towards the latter point, however, finding the *corrections* to this first approach is clearly an avenue to explore further. For example the terrorist can read up on the target country and find out more about how decisions are actually made, and possibly use this to play sub-state actors against each other. Such scenarios would primarily be expected to add more

detail and nuance to the treatment, but it is also thinkable that new insights emerge which run partially counter to those produced by the simplest method.

It seems clear that the gaming out of nuclear blackmail provided a good deal of insight which would have been hard to obtain by other means. Yet one should notice that a large fraction of the conclusions reached stem from the qualitative analysis of the first half of the chapter.

7.8 Conclusions and policy implications

I asked at the beginning of the chapter whether there could be other fruitful uses of a terrorist nuclear weapon than immediate detonation. I argue that both of the two specific options, deterrence and blackmail, may be of use to the terrorist under special circumstances, but that there are significant drawbacks with both of these strategies which would be serious disincentives for the nuclear armed terrorist contemplating deterrence or blackmail.

It is argued that for a terrorist, hiding a nuclear weapon on the territory of a target country in order to deter it from taking action against the terrorist's interests is not a viable strategy, since such a situation would be too intolerable to the targeted government for the deterrence situation to be stable. A more likely scenario would be one in which the terrorist held the nuclear weapon in an area it deemed safe and threatened to use it against nearby or easy to reach targets. Also in this case upholding the deterrent over time will be a problem for the terrorist, and the announcement of a nuclear capability could well provoke the pre-emptive attacks it is supposed to deter.

I analyse a few different scenarios of nuclear blackmail, outlining the strategic concerns of the terrorist player and government respectively. I argue with the help of a signalling game that a government can very plausibly make nuclear blackmail a suboptimal strategy for the terrorist both in the case where he has control of a nuclear weapon and in the case where he has not. For nuclear blackmail to work, the terrorist must overcome the great challenge of convincing the government that it will in fact keep its side of a deal to surrender the weapon undetonated in return for concessions; if the government is unconvinced it will have little reason not to respond with force, in which case the terrorist is arguably

worse off than if the attack had been conducted without warning.

For this reason if a nuclear blackmail attempt is issued publicly by the terrorist I argue that it is unlikely that the terrorist truly expects the concessions, especially so if the group in question has a history of militancy. In the (arguably) highly unlikely scenario where a politically motivated group has acquired a nuclear weapon and uses it for extortion, the situation is less clear, and blackmail might be successful in this case, although it appears unlikely that such a group will have gone to the trouble of acquiring the weapon in the first place.

In the face of nuclear blackmail by a group known to be militant, I find that the best response is nearly always for the government to make every effort to hinder the attack with force (despite a terrorist threat to detonate if the government takes such action). This fact is what should most likely deter the nuclear armed terrorist from attempting blackmail and rather opt for a non-advertised attack instead. The corresponding equilibrium situation is one in which no blackmail happens. A test of the restrictive assumptions made reveals that the normative value of this equilibrium appears soundly established: a government should respond with force against a nuclear terrorist blackmail threat except possibly under very special circumstances as described above.

As for defences, the ability to respond in a forceful and directed way to a credible nuclear threat when it occurs is essential. Because there is such a multitude of shapes and forms which such a threat can take, building a static defence to take on all possibilities is probably not as helpful as developing the necessary flexibility and mobility of finances, equipment and manpower to mobilise rapidly if and when necessary, to support the defence mechanisms which are already in place.

All else equal, it may be favourable if the terrorist attempts extortion, since it buys the defending government time to hamstring the attack. However, what appears to be the best response to nuclear blackmail in most cases, forceful defensive action, is also the response which will likely deter blackmail in the first place. Incidentally such a stance will also form part of a defence which has the very desirable side effect of potentially deterring a nuclear terrorism plot before it even gets started, as treated in the previous chapter. Given these considerations, provoking an extortion scenario (e.g. by sending the message that a nuclear threat will not be taken

seriously⁵⁸) would risk damaging deterrence of nuclear terrorism in general and is thus not worthwhile.

Thus the stance taken by the United States, that no deals will ever be struck with terrorists, is probably a sensible one with respect to nuclear terrorist blackmail⁵⁹. While it does create an incentive to the terrorist to detonate the weapon once obtained rather than try to use it for negotiation, the prior declaration of a 'no deals'-policy could allow the government to respond with force to such threats and at the same time appear consistent in its policy.

On the part of a bloodthirsty terrorist, the above analysis indicates that nuclear blackmail and attempts to use an acquired nuclear weapon as a deterrent are probably suboptimal choices compared to simply using the weapon for attack. The likely conclusion is thus that as long as he believes it probable to succeed, the rational terrorist who extracts utility from death and destruction will probably not try to draw on other strategic potentials of possessing a powerful weapon, but try to use it in an attack.

58 This is very far from the political signals sent, particularly by the United States, today.

59 Note carefully that this refers to the response to a direct threat of grand scale terrorist violence. It is *not* to be understood to mean that no negotiations should in general take place with groups that have been defined as terrorists by the United States. Such considerations lie far beyond the scope of this thesis.

Conclusions and outlook

In chapters 3 through 7 I have analysed different questions related to nuclear terrorism as I defined it in chapter 1. I shall here summarise the conclusions reached and briefly discuss the bottom lines from different chapters. In particular, where the conclusions reached are of direct applicability to policy, it is of interest to discuss how different policy prescriptions combine: do they contradict each other or reinforce each other. I find that the overlap between the different chapters of the thesis is limited, but discuss the combined implications for fissile material security.

In the present thesis I have considered a series of separate sub-questions to the overall goal of establishing best policy approaches to the threat of nuclear terrorism, so that the conclusions reached do not form a complete strategy. In particular, an important issue not treated in detail in this thesis is the question of second layers of defence, analysed in detail in Levi's recent book¹ and reviewed in appendix D.

Finally I discuss how the different research efforts in the thesis open new questions and indicate future directions of research on nuclear terrorism using formal methodology.

8.1 Conclusions from research chapters

Here I briefly recapitulate the main conclusions from the research chapters of the thesis.

8.1.1 Safeguards against nuclear terrorism: HEU vs. Pu

In chapter 3 I analysed qualitatively a (presumably well informed) terrorist's choice between using highly enriched uranium (HEU) or plutonium (Pu) for the purpose of building a crude nuclear weapon and concluded, as have a number of analysts in the past, that HEU is a far better choice for a number of reasons. This discussion is not new, but forms a

1 Michael Levi *On Nuclear Terrorism* (Cambridge, MA: Harvard University Press, 2007)

necessary backdrop for the decision theoretical analysis in chapter 4, where it is argued on the basis of cost and benefit, that as concerns safeguards against nuclear terrorism, a heavy spending emphasis on HEU as compared to Pu is justified.

While this conclusion could have been reached by qualitative arguments², the results from the quantitative analysis undertaken are arguably powerful tools for informing policy in that they reduce the question of the division of available funds between safeguards branches into one of estimating numerical values of a small number of parameters.

Indeed, I find there are two key quantities to determine the ideal level of expenditure of the two branches: the threat level (importantly, the probability of successful terrorist attack) and the value for money (more accurately: expected threat reduction per monetary unit spent) of different safeguards efforts. Notably, the amount of money already spent is irrelevant to the question of future spending, as is consistent with economic theory.

A rough numerical estimation of key parameters shows³ that Pu safeguards are possibly overfunded at present, whereas HEU measures are almost certainly underfunded. The former conclusion might seem counter-intuitive since there is every reason to believe that some storage facilities holding quantities of Pu are still inadequately protected. It is important to note, however, that while further spending in plutonium safeguards may not be warranted today, extensive upgrades in HEU security could once again make such measures worthwhile in the future, as is demonstrated in a further numerical analysis in section 4.5.

For policy-makers it will be useful to note that the value for money of safeguards measures is as important as the funding level itself. While this may be intuitively obvious, the clarity with which this emerges from the mathematical results of our gaming is notable. It is essential, therefore, to sustain and improve co-operation with countries of proliferation concern, for example Russia and Pakistan.

8.1.2 HEU or Pu: the terrorist's choice

In chapter 5 I turn the tables and analyse the choice between HEU and Pu as building material for a

2 As do Charles D. Ferguson and William C. Potter with Amy Sands, Leonard S. Spector and Fred L. Wehling *The Four Faces of Nuclear Terrorism* (New York:Routledge, 2005)*

3 See section 4.4.

simple nuclear weapon from the terrorist's point of view, a continuation from the qualitative study in chapter 3. This chapter is of a more theoretical nature than the other research chapters in that policy conclusions do not flow quite so directly from the decision theoretical endeavour as they do in other parts of the thesis.

I establish a formal framework for thinking about this choice in terms of concepts such as the penetration time (the time the terrorist may expect to spend on a given project before some incident is likely to cause it to fail), the deterrence time (the project duration above which the probability of success is so small it is no longer preferable to doing nothing) and the critical period (the average time between two fissile material acquisition opportunities above which it is preferable to accept the first available option rather than wait for the optimal choice of material, HEU). A number of nontrivial relations between these quantities are established, several of which, it is indicated, apply more generally than the specific payoff model employed in the chapter.

The terrorist's choice, I demonstrate, is (at least within the confines of the model employed) governed by two key parameters:

1. Terrorist impatience represented by the discount function. In particular, when the impatience is due to fear of failure, it is well described by the intuitive concept of a penetration time - roughly the time the terrorist believes he can wait before being apprehended.
2. The perceived probability that the next opportunity to obtain material involves HEU.
3. The estimated time between opportunities for obtaining nuclear materials.
4. The terrorist's degree of preference for HEU over plutonium.

All of these parameters (with the possible exception of the terrorist's preference for fissile materials) are, notably, within a government's power to manipulate to some extent. Most notably a terrorist can be forced to accept a suboptimal choice of fissile material for her nuclear project if she is stressed into impatience and the ideal fissile material (HEU in metal form) is harder to obtain than other less favourable options. On the other hand, luring the terrorist into waiting forever for the material of preference might be as good an

option, hence it is not altogether obvious which way a target government might wish to influence the terrorist's decision. I find quantitative criteria for the values of these parameters which will change the terrorist's ideal strategy, and the relations produced from the modelling endeavour can in any case inform a government as to how its actions may affect the terrorist's choice of strategy.

A conclusion of theoretical interest which comes out of the decision theoretical analysis in chapter 5 is a criterion for when *absolute* deterrence of terrorism is possible, namely whenever the terrorist extracts a negative payoff from a failed attack. A terrorist with no fear of failure, in other words, cannot be absolutely deterred, while a terrorist with a fear of failure in principle always can be.

8.1.3 *Relative deterrence of nuclear terrorism*

In chapter 6 I discuss the question of whether and under what circumstances a terrorist can be deterred from embarking on a nuclear weapons project in the first place. A decision theoretical model is devised which is too complicated for algebraic analysis to be useful, but which is readily treated with numerical means. The policy conclusions from this chapter are several and directly applicable to policy.

I recognise three main routes to achieve relative deterrence (i.e., convincing a terrorist that it is in her own interest to stick to her tried and trusted conventional means of attack), in order of importance:

1. Decreasing the perceived probability of success of a nuclear terrorism mission is the most effective means of deterrence. It both lowers the expected payoff of the nuclear strategy and increases the operational risk.
2. Disruptive measures to induce impatience and pressure may furthermore deter the terrorist from long-term ambitious projects into continued reliance on tried and trusted conventional means.
3. Threatening the future of terrorist funding increases investment risk from a nuclear project and forms a tertiary means of relative deterrence.

These conclusions come from analysis of the conditions under which the terrorist's *expected utility* of a nuclear project is lower than for a series of

conventional attacks coupled with qualitative considerations of risk which are not captured by utility theory (which is inherently risk-neutral). The observation that terrorists have historically exhibited a distinctive aversion towards operational risks further increases the potency of these measures towards deterring the terrorists from long term, expensive and ambitious projects in general, nuclear terrorism in particular⁴. I conclude with confidence that relative deterrence of nuclear terrorism is both possible and doable by these means, and may indeed be at work already⁵.

Threats of retaliation against the terrorist's own interests has been suggested by some analysts in the past as a possible means to deter terrorists in the classical Cold War sense of deterrence⁶. I argue in light of the decision theoretical results, however, that the threat of retaliation will probably have but a modest effect on terrorist decision making⁷. Indirectly, however, such threats could deter peripheral terrorist members such as funders and logistical support from taking part in a nuclear project, which could lower the terrorist's perceived chances of successful attack, which will in turn help to shift the terrorist's cost/benefit analysis towards conventional attacks. Threats of retaliation thus primarily count as just another 'second layer of defence' in the context of deterrence and aids relative deterrence as part of point 1 above⁸.

Another policy relevant conclusion is that hardening targets against conventional terrorism can create an incentive for the terrorist to opt for unconventional means. To the extent that a government prefers a 'war of a thousand cuts' to a single nuclear attack, this recommendation should be taken carefully into account. An obvious conclusion to draw is that in order not to undercut the above mentioned measures towards achieving relative deterrence of nuclear terrorism, resources poured into

anti conventional terrorism efforts must be matched by resources allocated to the denial of nuclear terrorism.

Notably, relative deterrence of nuclear terrorism does not depend *directly* on the values of the parameters mentioned in points 1-3 above, but on how the terrorist herself *perceives* these. Thus it is important that the right signals are sent and that any improvement of safeguards, intended disruptive actions and economic sanctions are well publicised in order to affect the terrorist's internal cost/benefit calculus.

8.1.4 Nuclear terrorist blackmail and a terrorist nuclear deterrent

In chapter 7 I discuss the feasibility of other strategic uses of a terrorist nuclear weapon, namely as deterrent or for extortion. I argue that neither of these options are probably preferable options to a terrorist with a nuclear weapon who is militant and extracts utility from massive killings (a terrorist of political nature, who prefers not to detonate a nuclear weapon, I argue, will most likely not have the required incentive to acquire such a capacity in the first place).

Faced with a nuclear blackmail threat from a terrorist group known to be militant, I argue by means of a signalling game that the best response of a targeted government is nearly always to respond forcefully in an effort to hinder a detonation, even though the terrorist threatens to detonate her weapon if such measures be taken. Proclaiming the stance that a nuclear blackmail threat will always be met with force (which is close to what e.g. the US is doing today) will most likely deter the rational and nuclear armed terrorist from attempting blackmail and rather attempt to use the weapon directly in a clandestine attack.

All else equal, provoking a blackmail threat from a nuclear terrorist rather than an immediate attack might be beneficial for a target government in that it allows a little extra time to attempt to stop the attack. However, building the rapid response defensive capacity which can make use of this time slot is the measure which will most likely deter nuclear blackmail from occurring in the first place. Combined with the relative deterrence effect which all defensive measures have (as detailed in chapter 6) the positive aspects of a defence which can respond rapidly to a reported threat (be it a deliberate warning from the terrorist or otherwise) far outweigh the concern that

4 See e.g. Lewis A. Dunn 'Can al Qaeda Be Deterred from Using Nuclear Weapons?' *Occasional paper #3* (Center for the Study of Weapons of Mass Destruction, National Defense University, 2005) p.16

5 This point and several related points are also made by Michael Levi *On Nuclear Terrorism* (Cambridge, MA: Harvard University Press, 2007).

6 Robert F. Trager and Dessislava P. Zagorcheva 'Deterring Terrorism: It Can Be Done' *International Security* 30:3 (2005) pp.87-123

7 The deterrence of state sponsorship of terrorism is a somewhat different question not treated in chapter 6, but see section 6.3.1 for a discussion and references.

8 See appendix D for more details.

should an attack nonetheless occur, no blackmail letter will probably have been received beforehand⁹.

The terrorist in possession of a nuclear weapon might find that keeping it as a deterrent against attacks upon it could be a feasible strategic use of the weapon. Hiding a nuclear bomb on the territory of a potential attacker country is not a feasible deterrence strategy, since such a situation cannot be upheld in the long term, but the threatened use against a target near to the weapon's location or one which is relatively easy to hit could be. In any instance, the declaration of a credible terrorist nuclear deterrent is not sure to create the long-term safety that the terrorist might hope to achieve with it. Quite likely the situation would be unpredictable and unstable and such a declaration could indeed provoke the very attacks it is meant to deter.

8.2 Comparing policy implications of different chapters

The overlap between the different policy conclusions in this thesis is on the whole quite modest. Chapters 3-5 deal exclusively with the acquisition of nuclear materials, whereas chapter 6 is concerned with the terrorist's decision whether or not to attempt a nuclear project, a choice which arguably only depends on the overall probability of success for the entire project, of which materials acquisition is only one part. Chapter 7 deals with a situation in which the terrorist has already acquired a nuclear weapon, and as such is quite separate from the other research chapters. The primary field in which conclusions interact is therefore safeguards of fissile materials.

8.2.1 Safeguards against nuclear terrorism

In chapter 3 I recount an assertion which is agreed upon by nearly all analysts of nuclear terrorism¹⁰, namely that when it comes to hindering a planned nuclear acquisition project, there is nothing more effective than to keep the terrorists from laying hands on the necessary uranium or plutonium. While a nuclear project can take a large number of different

paths beyond this point, no such project can eschew the need to acquire the material, and the probability of overall success will be at least proportional to the probability of acquiring the fissile material¹¹.

Indeed, our quantitative analysis section 4.4 demonstrates that improved HEU safeguards gives excellent value for money if one assumes that a nuclear plot is under way but has not yet acquired the necessary explosive elements. More exactly, by spending money to improve HEU security the reduction in expected cost from nuclear terrorism is significantly greater than the cost of improved security.

Furthermore, as demonstrated in chapter 5, making the ideal material for a terrorist bomb harder to acquire will force the terrorist into attempting use of suboptimal fissile materials with a higher probability of failure. Materials other than metallic HEU will generally be more difficult to handle (more radioactive, sometimes pyrophoric, greater quantities are needed) and the probability of a dysfunctional bomb or accidents (by criticality or radiation) during assembly is increased. The technical challenge of making a workable design is also made greater by having to make do with a secondary choice of fissile weapon fuel¹².

In appendix D of this thesis I review Michael Levi's argument¹³ that the value of good material protection goes beyond the mere blocking of terrorist acquisition attempts. Were the terrorists to succeed in acquiring the material despite protection methods, good procedures for material accountability will increase the chances that the theft is discovered quickly. A quick warning will almost certainly increase the probability that the nuclear terrorist plot be intercepted at a later stage.

Thus the soundness of safeguards measures as a means to counter nuclear terrorism seems very well established, but the fortunate effects of material protection are even more far-reaching. I demonstrate in chapter 6 that decreasing the terrorist's assessed probability of success of a nuclear terrorism project has the added effect of working to deter the terrorist from attempting a nuclear ploy in the first instance in

9 It is of course far from certain that the terrorist would have contemplated blackmail in any case.

10 Excepting, of course, those who argue this threat is not real and thus warrants no particular policy response whatsoever. See chapter 1 for a review.

11 Michael Levi effectively argues (*On Nuclear Terrorism*) that the dependence is more than linear. See appendix D and discussion below.

12 A further elaboration of the technical hurdles in this respect are found in Michael Levi's *On Nuclear Terrorism* pp. 66-97. Note that this discussion is implicitly taken into account in the game of chapter 4.

13 Michael Levi *On Nuclear Terrorism*

favour of staying with her tried and trusted means of conventional attacks. While it is true that what counts for deterrence purposes is not in fact how difficult a nuclear terrorism project is, but how difficult the terrorist *perceives* it to be, the only reliable way to achieve this is probably to keep improving defences against nuclear terrorism and make sure to publicise all progress made.

Note how this contrasts with deterrence in the Cold War sense, by threats of punishments; in a situation of mutual deterrence by threats, steps to increase deterrence, i.e. steps to make the consequences of non-compliance more severe, tended to make the consequences more terrible should deterrence fail. The ultimate consequence of a failure of deterrence during the Cold War might have been all-out nuclear war. When it comes to deterrence by denial, however, the situation is the opposite: the measures which may be taken to persuade an adversary that an attempt at a particular act will fail are the same efforts which will hinder the adversary from succeeding should he try nonetheless.

In summary, all of our considerations point to the importance of fissile materials protection as a means to defend against nuclear terrorism, and none point against it. This fact is a notable trait of the problem of nuclear terrorism.

8.3 The proof of the pudding: has the research approach been effective?

In this thesis a methodology has been used which, it was showed, has hardly been applied to the field of nuclear terrorism at all before. In the chapter where the methodology was discussed I quoted Osborne's pragmatic criterion for a game's justification: 'As always, the proof of the pudding is in the eating: if a model enhances our understanding of the world, then it serves its purpose.'¹⁴ In chapters 4 through 7 I have analysed four different games. Now for the verdict: did they serve their purposes? And more generally, how well was the methodology used suited to the study of nuclear terrorism? I will analyse these questions in this section.

This question has been treated briefly at the close of each gaming chapter, and the reader who has looked at these will already know the conclusion: this author did indeed find the use of rational choice theory to be illuminating and a great help in

enhancing understanding of the questions at hand. There is more to the question than this bottom line, however.

When inspecting Osborne's criterion more closely, one realises that his criterion is not as clear cut as it first appears. Rather, it brings up some new questions which must be answered before it can be put to use in an evaluation. Firstly, what does it mean to 'enhance our understanding'? Whose understanding? Secondly, what is one to compare the merits of a game to; what is the alternative? As I will argue, there is no clear distinction between qualitative methodology and rational choice theory. And thirdly, I have chosen research sub-questions which match the choice of methodology. Will the evaluation then not be lopsided, somewhat like comparing the merits of a hammer and a saw with respect to cutting planks in half?

I shall not here attempt to assay in any detail the achievements of each of the four models I employ, since this was treated in the individual chapters. Rather, I will attempt to provide a more nuanced picture of how to understand and evaluate the merits of rational choice theory in general, and in the context of this thesis.

8.3.1 Understanding - for whom?

If one were to accept Osborne's criterion as a guideline for measuring the success of the rational choice approach, a natural question comes up: what exactly is 'understanding'? This question has been central in epistemological philosophy for a long time, a debate I shall not go into. More specifically, Osborne speaks of 'our' understanding; his use of language seems to imply that there exists some objective and common entity which is 'our understanding', and that researchers keep adding more material to this entity.

I would venture to argue a different view, that 'understanding' is inescapably subjective, at least in part. One and the same piece of research can be highly illuminating to one reader, but provide no new insight to others. Whether a reader will extract new understanding from a research report could depend on a number of aspects both of the report itself and of the reader. Beyond the message which is communicated in the text and what the reader is technically able to grasp there is a question of what it takes to give a reader the sense of having understood something. This will depend on what he already knows and, crucially, the way he is used to thinking about scientific analysis. Contrary to, say, a historian,

¹⁴ Martin J. Osborne *An Introduction to Game Theory* (Oxford: Oxford University Press, 2004) p.7

a reader with a mathematical background might find a formal analysis more illuminating than a qualitative one arguing the same conclusions, just because what he understands as knowledge can be expressed in formulas and associations are easily made to the knowledge he already has. This makes an interdisciplinary effort such as the present thesis challenging: it is hard to present material in a way so that it is understandable and chimes well with scholars from very different fields.

When I argue above that the analyses of the various chapters have indeed enhanced understanding, this must therefore be understood as partly subjective. What I can say with certainty is that I myself found that the analysis in each of the four gaming chapters proved an important help for me to conceptualise, systemise and visualise the questions at hand, and thereby improve my understanding of them. To the extent that I typify a class of researchers, there is reason to believe that others will find these models enlightening as well. If one was not worried about the impact of one's research, such an argument might be sufficient.

Arguably, however, the more readers a piece of research can inform, the better. The use of formal methodology therefore comes at the price of making the research less accessible to some who might otherwise have been interested. Walt is therefore right in warning against 'saying what everyone knows in a way no-one can understand'¹⁵ as discussed in section 2.7. If a conclusion could have been reached in a purely qualitative way, modelling might just be a complicating detour. The sections at the end of chapters 4 through 7 where the merits of the games are evaluated have therefore focussed on whether conclusions reached might have been arrived at solely by qualitative arguments.

Bearing in mind the inescapable element of subjectivity in such an evaluation, the rest of this section is therefore primarily devoted to the question of how much of the insight gained in my research chapters could have been obtained without the help of gaming, and how *that* question is to be understood.

8.3.2 Tools and tasks

I argued in section 2.10 that gaming may be understood as a tool for laying out the implications of a set of assumptions in a systematic way. Evaluating a

tool, however, makes more sense when placed in the context of a specific task. Underneath a broad headline like 'Nuclear Terrorism' there exists a wealth of different sub-questions, and the analyst must choose some of these to focus on. This is well exemplified, for example, by comparing the sub-questions considered in the three monographs by Allison¹⁶, Ferguson and Potter¹⁷, and Levi¹⁸, all considering the same overall topic but only partly overlapping.

A tool and a task must be suited for each other, whether the tool is chosen for the task, or the task is chosen in order to make use of a specific tool. Because the project of the thesis was to apply a certain methodology to the field of nuclear terrorism, the sub-questions chosen for closer examination were chosen and framed in such a way that they were expected *a priori* to be suitable for gaming out. To the analyst who believes in the potential value of gaming in general it is therefore not so surprising that the analysis seems to have worked; after all it was designed to work.

Clearly there are many other sub-questions beneath the nuclear terrorism umbrella for which rational choice theory would *not* be suitable. All the sub-questions considered in this thesis have the trait that each actor has a small set of available strategies with arguments for and against each strategy. These are the kind of questions which may most easily be modelled and expressed in an 'economic' language. But during my discussion and modelling a number of other questions have come up which had to be discussed by qualitative means or had been previously by others.

Thus, while I have argued extensively above that my analyses have indeed provided useful insights, gaming can certainly not be the *only* approach used by researchers of this questions for at least three reasons: (1) a qualitative analysis is always necessary before modelling (discussed further below); (2) there are questions which are not suitable for gaming; and (3) a gaming approach will typically capture one out of several aspects of a question, and always represents a simplified representation of reality. To examine the finer details or the full picture, other methods must be

15 Stephen M. Walt 'Rigor or Rigor Mortis?: Rational Choice and Security Studies' *International Security* 23:4 (1999) pp.5-48

16 Graham Allison *Nuclear Terrorism: The Ultimate Preventable Catastrophe* (New York: Times Books, 2004)

17 Charles D. Ferguson and William C. Potter with Amy Sands, Leonard S. Spector and Fred L. Wehling *The Four Faces of Nuclear Terrorism* (New York: Routledge 2005)*

18 Michael Levi *On Nuclear Terrorism* (Cambridge, MA: Harvard University Press, 2007)

employed in addition.

8.3.3 *The role of qualitative analysis*

It is easy to underestimate the qualitative aspect in all gaming. When designing the game, deciding what should be in it and what shouldn't, and which assumptions to make, qualitative arguments and judgement must be used. And at the end of the gaming, a qualitative understanding of the resulting equations or graphs should be extracted. A game can illuminate a real-life scenario only in conjunction with qualitative considerations.

Each of the gaming chapters in the thesis begin with a qualitative study of the question at hand. Of course, this is a natural way to introduce the problem, review the relevant literature, etc., but it also has the rôle of justifying the model which is thereafter created. Only when a picture of the dilemma is drawn in qualitative terms does it become possible to decide which aspects the model should account for, what the limitations of the model are, and under what circumstances the model must be changed in order to be applicable. The alternative to rational choice theory to which its merits should be compared is therefore not qualitative methods, but qualitative methods *alone*.

The question I have focussed on in the sections towards the end of each research chapter where the merits of the methodology are surveyed has been whether the same conclusions could have been reached without the use of a game. In chapters 5 and 7 my answer to this was that, while it is impossible to prove the impossibility of arriving at the same arguments in verbal form, the production of explicit and formulaic criteria for a strategy over another is something which qualitative analysis alone would not produce. Likewise in chapter 4, while the conclusions about US priorities of HEU versus plutonium have been reached by qualitative means before, I find that the symbolic analysis provides a level of precision (in form of a set of simple inequalities) and depth (as exemplified by the study of how plutonium efforts could possibly be overfunded) which I do not believe I could have achieved otherwise.

The chapter where it is least obvious that qualitative methodology could not have produced the same results is probably chapter 6. The primary reason for this was recognised to be that the model was too complex for analytical treatment to be fruitful, and the extraction of intuitively useful knowledge had to be made from graphs, at an

intuitive and rough (rather than formal and precise) level. Even so, the methodology was found to have its definite merits here as well. The process of gaming is in itself a valuable tool for the analyst for thinking through the problem, and experimenting with simulations using different parameter values was found to be a good way to obtain a certain familiarity with and intuition for how the decision behaved under varying conditions.

The taxonomy of terrorists in terms of their impatience, which resulted from the study of figure 6.2, is an example of an (arguably) illuminating insight which it is reasonable to imagine could have resulted also from a purely qualitative approach. As explained, the exercise of simulation certainly did its job as a tool to enhance this author's understanding of relative deterrence, but may not have been the only tool which could have done so. Different angles of approach will probably highlight different parts of a problem, and re-thinking the question of relative deterrence in a different light, perhaps taking into account real information about decision making in terrorist organisations, will probably be worthwhile and could well be equally illuminating.

Just as it is easy to underestimate the rôle of the qualitative aspect, it is tempting to trust the mathematical results overly, since mathematics is a form of argument which typically connotes rigour and precision. It is certainly true that, as long as the laws of formal logic are followed, the formulas flow from the model with inevitability, but it is necessary to keep in mind that the model itself was arrived at qualitatively. Thus even the mathematical formulas have a qualitative aspect to them. This makes it necessary, before making use of the formulas, to understand the assumptions and arguments behind the model, so as to realise the range within which the formulas are applicable.

8.4 *Where to from here?*

The possibilities for extending the models used or devising different ones are almost infinite. Nonetheless, some directions seem more natural than others given the nature of the problem under scrutiny.

With the exception of chapter 7 I have focussed in this thesis on simple decision theoretical models with no real-time interaction between the players. The major strength of such models is obviously that it makes for simple, transparent and powerful analytical calculations with fairly obvious interpretations. For exactly this reason, it was deemed that such games

were most suitable as a first approach. To look at future directions, one should consider the weaknesses of such very simple games, however, and which extensions might improve on the approach herein with respect to these.

A possible weakness of my games is the lack of communication between the players. In every game I see the situation from one player's perspective while the other player is present merely by determining the value of certain parameters that the player in focus must take into account. I argue that this is adequate as a first approximation, yet there is reason to believe that additional insight might be derived from modelling the communication between the players and its impact.

A natural way to go about this might be to model intelligence gathering. It seems reasonable to prescribe some kind of Bayesian game for this¹⁹. A simple signalling game such as has been used for terrorism gaming before²⁰ and made use of in chapter 7 might form a start, yet as is pointed out, the communication from terrorist to intelligence is sporadic and stochastic in nature, which might call for a more complex model in which time is explicitly incorporated.

It is in principle not difficult to devise complicated models for use with computer simulations, and this might well form the most fruitful path in extension of our simple models. The challenge, then, is perhaps not so much devising the models and writing computer programmes to run them, as it is to make sense of the model's behaviour and determine numerical values for the parameters of the model. This should not deter the clever social scientist, however, and many modelling approaches have been developed on which the researcher may draw²¹.

Another field in which much may be done is that regarding the cooperation between states in the face of a grave threat such as nuclear terrorism. It is most

peculiar to notice that crucial safeguards projects between great powers such as the United States and Russia can be halted for months by small problems such as uncertainty over responsibility in case an accident should happen²², and it might be a problem for researchers in international relations to investigate whether such disagreements do not run deeper than meets the eye. Some authors, moreover, have called for a global coalition against nuclear terrorism²³, a prospect well suited for gaming, along the lines of some of Professor Sandler's work²⁴.

Altogether different paths that seem very fruitful involve real-time role play games, such as that orchestrated by RAND²⁵, in which policy makers and other leading figures in society are faced with scenarios such as nuclear terrorist attacks, intelligence reports of such, threats of detonation *et cetera*.

An advantage of gaming and role play over more technical studies of feasibility of a nuclear project such as that by Zimmerman and Lewis²⁶ is that one is less likely to run into ethical dilemmas where the need to create awareness about the dangers of nuclear terrorism must be weighed against the danger of providing useful information to potential perpetrators. Arguably, enough knowledge about nuclear weapon design is already public and widely published, that a researcher can form a qualified opinion without the need to explore the matter further, hence the potential, although probably not exhausted, seems somewhat limited along a technical line of research.

8.4.1 Modelling gives rise to new questions

One does not have to depart very far from the efforts presented in this thesis in order to extract new knowledge from the use of formal social science

19 Bayesian games are named after the famous statistician Thomas Bayes, 1702-1761, who is credited with formulating the Bayes Theorem or Bayes' Rule, which provides a formal method of updating beliefs following a signal. Cf. e.g. Robert Gibbons *A Primer in Game Theory* (Hemel Hempstead: Harvester Wheatsheaf, 1992) chapters 3 and 4.

20 e.g. Daniel Arce M. and Todd Sandler 'Terrorist Signalling and the Value of Intelligence' *British Journal of Political Science* 37:4 (2007) pp.573-586; and Basuchoudhary and Razzolini: 'Hiding in plain sight: using signals to detect terrorists' *Public Choice* 128 (2006) pp.245-255

21 A good reference to study is Nigel Gilbert and Klaus G. Troitzsch *Simulation for the Social Scientist* 2nd edition (Buckham: Open University Press, 2005)

22 Matthew Bunn and Anthony Wier *Securing the Bomb 2006 Project Managing the Atom* (Belfer Center for Science and International Affairs, Harvard University, 2006)*, p.14-15.

23 e.g. Matthew Bunn *Securing the Bomb 2007* report from Project Managing the Atom (Belfer Center for Science and International Affairs, Harvard University, 2007)*, pp. 97-106.

24 e.g. Daniel G. Arce M. and Todd Sandler 'Counterterrorism: A Game-Theoretic Analysis' *Journal of Conflict Resolution* 49:2 (2005) pp.183-200.

25 Charles Meade and Roger C. Molander *Considering the Effects of a Catastrophic Terrorist Attack* Technical Report (RAND, 2006)*

26 Peter D. Zimmerman and Jeffrey G. Lewis 'The Bomb in the Backyard' *Foreign Policy* (Nov/Dec 2006)

methods applied to nuclear terrorism, however. For each gaming effort presented herein, new questions and ideas for further inquiry have come up, suitable perhaps for future PhD or Masters theses. The games in this thesis have deliberately been designed to be as simple as possible, because the goal was to examine their potential and arrive at easily interpreted results of a conceptual nature. I have concentrated therefore on developing an analytical framework for thinking about these problems, rather than attempting to exhaust the potential they have for examining the situation as it stands in the world today. On a general note, therefore, there is much one could do in tying the models closer to reality, primarily by generalising the assumptions made in order to make the models richer, and undertaking efforts to estimate required numerical parameters as accurately as possible.

In connection with the game of chapter 4 I undertake a coarse numerical study to demonstrate how the equations obtained through modelling may be applied to a real world scenario. An interesting question to pursue is: how accurately can one reasonably estimate such numerical values based on open access sources? Clearly, much relevant information is exempt from the public domain, yet given the large amount of published data it may be possible to form rather confident bounds on several of these parameters.

Another interesting research enterprise might be to compare different countries in models such as that of chapter 4. My numerical study was limited to the case of the United States, a country which regards itself as a target for such acts of terrorism, and for whose non-proliferation programmes relevant information is readily available. Other countries may be uncertain whether they are the intended target or not, which calls for a generalisation of the model to allow the attack to be directed towards a different government, resulting in a much lower, collateral damage in the case of an attack. Presumably, interesting dynamics between different countries would emerge which could tell important lessons about how nations should co-operate and share the burden in nuclear non-proliferation issues.

The model used in chapter 4 has no inherent time and is as such 'instantaneous'. Certain effects which could be expected to become of importance are to do with changing conditions over time, however, which would require a model which introduces a proper timeline. One such example, that of changing stockpiles of HEU and plutonium, is discussed in

section 4.4.3, and a gaming study of this might be a suitable study for a master's thesis, say.

My model in chapter 4 moreover assumes that the terrorist cannot be deterred from a nuclear project. The analysis in chapter 6, however, clearly shows that it is possible to sway the terrorist from a nuclear path back to the tried and trusted means of conventional bombs, and incorporating the possibility of deterrence by denial into the model in chapter 4 would seem an obvious next step. The expected result (as I argue on an intuitive basis) would be an even stronger case for safeguards in that improved fissile materials security has the added effect of deterring the potential nuclear proliferator.

Several times in chapter 5 I come across examples of a conflict between relative deterrence of certain means of terrorism and absolute deterrence of terrorism on the whole. The most important example of this dynamic might be the deflection effect: hardening targets against conventional attacks can force the terrorist into attempting unconventional modes of violence, and conversely, that 'allowing' a certain level of conventional terrorism can remove the terrorist's incentive to employ more destructive means. It seems to the author that the dynamics between these forms of deterrence could be a fruitful area of study for the theoretically inclined.

A somewhat similar topic is the relationship between the possibility for absolute deterrence and the terrorist's fear of failure. It seems from the analysis of chapter 5 that these two concepts are closely related, and from a theoretical point of view it would be interesting to see whether more general statements could be made about the exact relation between them, independently of a specific model.

A number of concepts are developed in chapter 5 and studied in relation to the model in that chapter. These are for example the deterrence time, the penetration time and the critical period. It would be interesting to see if these concepts, which all have intuitive interpretations, could be generalised for a broad class of different models. It would seem to the author that these are very useful and general concepts for use in the study of terrorism deterrence and 'time failure' more generally.

In chapter 6 the only two options available to the terrorist are the nuclear and the conventional. In reality the number of different modes of attack is larger, making for a possible straightforward generalisation of the model in figure 6.1. It is possible that more complex behaviour could emerge upon

addition of a third option, whereas additions beyond this will likely not increase generality much further.

The terrorist as modelled in chapter 6 and elsewhere is only concerned with inflicting damage, which is clearly a simplification. A generalisation of the model in that chapter in which terrorist preferences are made more realistic, incorporating such effects as the fear of alienating an audience, would be interesting and probably not too difficult. Indeed, the preferences could be parameterised so that the consequences of terrorist motivations for relative deterrence purposes could be investigated directly in reply to the question 'how worried must a terrorist be about alienating her audience before she should rationally abandon her nuclear plans?'

The most important improvement to the model in chapter 6 would arguably be to explicitly incorporate terrorist risk aversion. Utility theory is intrinsically risk-neutral (unless it is explicitly modelled into the terrorist's preferences), so the most fruitful approach could be to use another fundamental theory such as regret theory²⁷. This could allow the considerations of risk aversion which are only qualitative in the present analysis, to be incorporated quantitatively in a direct way.

A final question, which arises from our gaming in chapter 7, is what happens if a terrorist organisation were to obtain a nuclear weapon and use it as a deterrent? This question is probably best answered by somebody with a strong background in international relations. It seems likely that some countries will wish to attack the terrorist pre-emptively in order to remove the nuclear threat whereas countries directly threatened by the terrorist bomb might oppose this, since such an attack could jeopardise the lives of hundreds of thousands of their citizens. The situation which occurs would probably be unpredictable, and planning ahead of such an incident might be worthwhile to prevent possible dreadful consequences.

In this thesis I have focussed on nuclear weapons which the terrorists fabricate themselves, so-called 'improvised nuclear devices'. The prospect of terrorist acquisition of an intact military nuclear weapon is, however, also a serious concern. Suitable for gaming similar to that performed in chapter 5 might be the terrorist's choice between these two different 'faces' of

nuclear terrorism as I define it²⁸. Further choices could include that between different classes of military weapons, according to yield, portability, security lock systems and availability. As in the quest for fissile material for a construction project, a choice between opportunism and patience might enter here as well. The information for use as input for modelling would likely be limited by military secrecy for such a project, which could be a problem and which might incline the researcher towards a qualitative approach. On the other hand, a well designed game could be a useful tool for a systematic treatment of the information which *is* available, and laying out the implications of it.

I may finally mention the ideas for future research which came out of the modest efforts I made on the question of second layers of defence, appended in appendix D. This topic was not pursued in this thesis, but as detailed in section D.5, opens up a range of questions which are probably very well suited for analysis of the kind that I have undertaken in chapters 4-7 of the thesis. I conclude in chapter 4 that the important quantity to establish when setting ideal cost levels for various elements of a defence is the security achieved per monetary unit spent. Establishing this for the complex and many-faceted second layers of defence, however, is not trivial, and developing methods for estimating the value of various elements of the second layers taking into account the way the different elements interact and strengthen each other, is both important and complicated. I propose that future researchers of this issue might wish to draw on optimisation theory from the engineering sciences in combination with decision theory games such as that of chapter 4. Such a task, I believe, could form the backbone of a future PhD thesis.

As a general conclusion of the thesis, it is clear that the employment of rational choice theory has proven a fruitful means of research in the field of nuclear terrorism as demonstrated by the number of conclusions reached by relatively simple calculations and simulations, and the number of new questions which these efforts have given rise to. In a field of research whose results may be of great importance for the safety of our future, there is plenty left to do for

27 e.g. Graham Loomes and Robert Sugden 'Regret Theory: An Alternative Theory of Rational Choice under Uncertainty' *The Economic Journal* 92 (1982) pp.805-824

28 I refer here to the title of Ferguson and Potter's *The Four Faces of Nuclear Terrorism*, *op.cit.* See e.g. chapter 3 of that reference for an overview of nuclear terrorism with intact weaponry. Note from chapter 1 that I define nuclear terrorism to include only nuclear fission, hence only two of the four faces fit my definition: improvised nuclear devices and assembled military devices.

the somewhat mathematically inclined analyst.

- A -

List of seizures of attempted smuggling of fissile materials

The list is adapted from that provided by the Nuclear Threat Initiative (NTI) Research Library¹

CASE NAME & DATE OF DIVERSION	MATERIAL DIVERTED	ORIGIN OF MATERIAL	RECOVERY OF MATERIAL
Podolsk 5/92-9/92	1.5 kg of 90% HEU	Luch Scientific Production Association, Podolsk, Russia	10/9/92: Russian police operation intercepted the smuggler, an employee of Luch facility, in the Podolsk train station.
Vilnius, Lithuania early 1992	About 100 g of 50% HEU	Institute of Physics and Power Engineering, Obninsk, Russia	5/93: Discovered in Vilnius bank vault embedded in portions of a shipment of beryllium.
Andreyeva Guba 7/29/93	1.8 kg of 36% HEU	Naval base storage facility, Andreeva Guba, Russia	7/29/93: Russian security forces arrested the thieves before they could smuggle the material out of Russia.
Tengen Unknown	6.15 g of Plutonium-239	Unconfirmed; possibly Arzamas-16, Russia	5/10/94: Police in suspect's apartment for another reason, stumbled upon the cache of plutonium.
Landshut Unknown	800 mg of 87.7% HEU	Unconfirmed; likely Obninsk	6/13/94: Undercover German police acted as potential customers in a sting operation.
Sevmorput 11/27/93	4.5 kg of 20% HEU	Naval shipyard, Sevmorput, Russia	6/94: The brother of a suspect asked a co-worker for help finding a customer. The co-worker notified authorities.
Munich Unknown	560 g MOX fuel; 363 g of Plutonium-239	Unconfirmed; likely Obninsk	8/10/94: Undercover German police acted as potential customers in a sting operation.
Prague Unknown	2.7 kg of 87.7% HEU	Unconfirmed; likely Obninsk	12/14/94: Anonymous tip to police giving the material's location (a parked car). In two instances in June 1995, Czech authorities recovered small additional amounts of HEU believed to be from the same source.
St. Petersburg [2] Unknown	3.05 kg of 90% HEU	Unconfirmed; likely Machine Building Plant, Elektrostal, Russia	6/8/94: Russian Federal Security Service agents arrested three suspects attempting to sell the material. Russian officials have confirmed the incident.
Moscow May 1994	1.7 kg of HEU	Elektrostal	6/8/95: In a sting operation, Russian Federal Security Service agents arrested three suspects trying to sell HEU, one of whom was an employee of Elektrostal.
Sukhumi Unknown	Approximately 2 kg of 90% HEU	I.N. Vekua Physics and Technology Institute, Sukhumi, Georgia	12/97: Russian inspection team visited facility, which had been closed by 1992 Abkhazian-Georgian conflict, and found facility abandoned, and material included in 1992 inventory missing. Material has not

¹ Part of the NTI webpage; online: http://www.nti.org/e_research/e3_special_nuctrafficking.html

			been recovered.
Chelyabinsk Oblast, Russia Unknown	18.5 kg of HEU (enrichment level unspecified)	Unknown, possibly Mayak Production Association, Chelyabinsk-70, or Zlatoust-36	12/17/98: Russian Federal Security service reports that it thwarted an attempt by workers at a nuclear facility in Chelyabinsk Oblast to steal 18.5 kg of nuclear material.
Dunav Most, Bulgaria Unknown	10 g of 76% HEU	Unknown	5/29/99: Bulgarian customs officers discovered HEU hidden in the trunk of a car crossing from Bulgaria into Romania. Driver said he had obtained material in Moldova.
Batumi, Georgia Unknown	920 g 30% HEU	Unknown	4/19/00: Georgian police arrested four suspects and seized HEU.
Elektrostal, Russia Unknown	3.7 kg of 21% HEU	Unconfirmed., possibly Elektrostal, Bochvar Institute (VNIINM), or Politekh Enterprise, Russia	5/2000: A resident of Elektrostal was detained during an attempt to sell 3.7 kg of uranium enriched to 21 percent U-235. Incident was reported by Gosatomnadzor.
Tbilisi, Georgia Unknown	0.4 g of plutonium powder	Unknown	5/2000: An individual was arrested for illegal possession of a small quantity of mixed powder containing about 0.4 g of plutonium and 0.8 g of low-enriched uranium.
Paris, France	~5 g of 70-80% HEU	Unknown, Russian/NIS origin suspected	7/16/2001: French police arrested three men and confiscated approximately 5 g of HEU.
Sadahlo, Georgia 6/26/03	170 g of nearly 90% HEU	Unknown, Russian origin suspected	6/26/03: Georgian border guards arrested a man trying to transport the material across the Georgian-Armenian border.
Tbilisi, Georgia 2/01/06	79.5g of 89% HEU	Unknown, Russian origin suspected	2/01/06: Georgian security services in a string operation arrested a Russian national in Tbilisi attempting to sell 79.5 grams of HEU.

- B -

Numerical programme used in chapter 6

Below is presented diagrammatically the structure of the simple programme used for simulations in chapter 6. For simplicity, all parameters in the programme are global and may be manipulated by any subroutine. The routines 'vary [parameter]', thus simply vary the global parameter [variable] and calls up the procedure which runs the actual simulation (same for all the 'vary' routines).

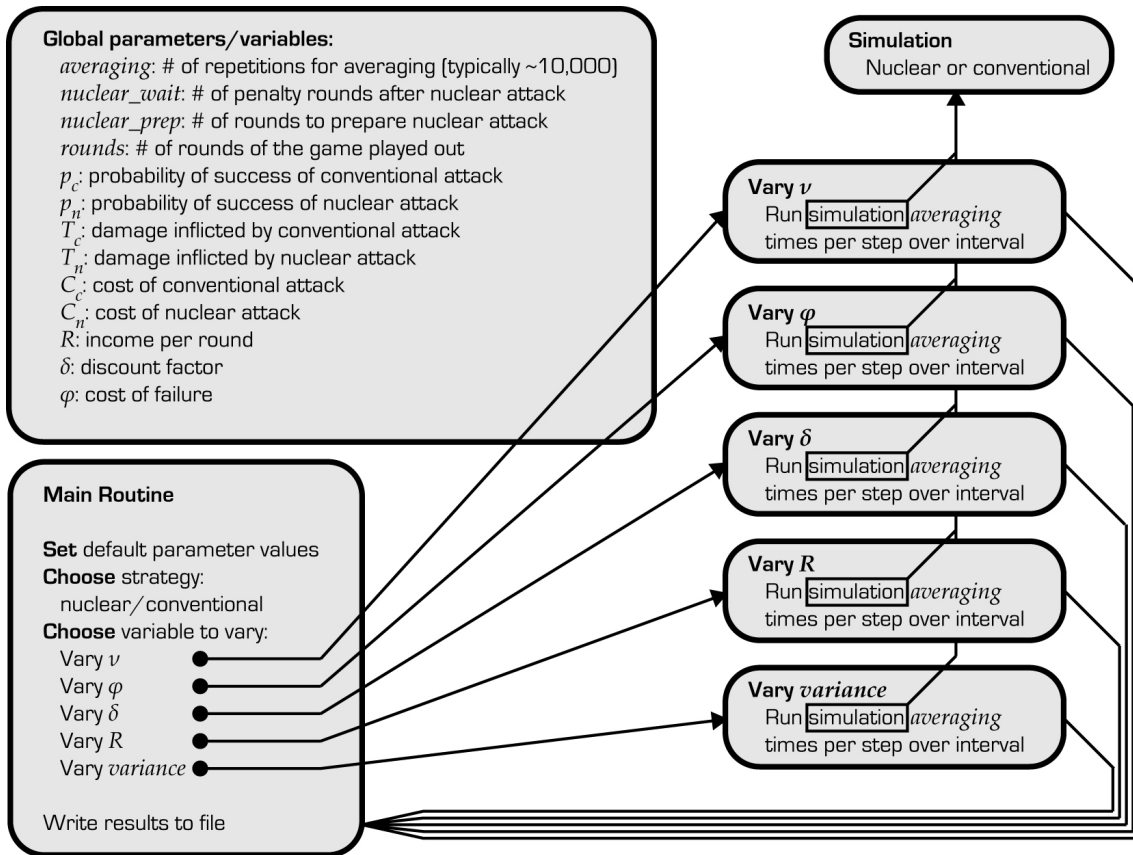


Illustration B.1.: Structure of programme of chapter 6

Only pseudo-code and structural overview is provided, yet the reader with some programming experience will nonetheless be able to translate to actual programme code. The language used for the

actual application was C++, but almost any programming language could have done the job with ease.

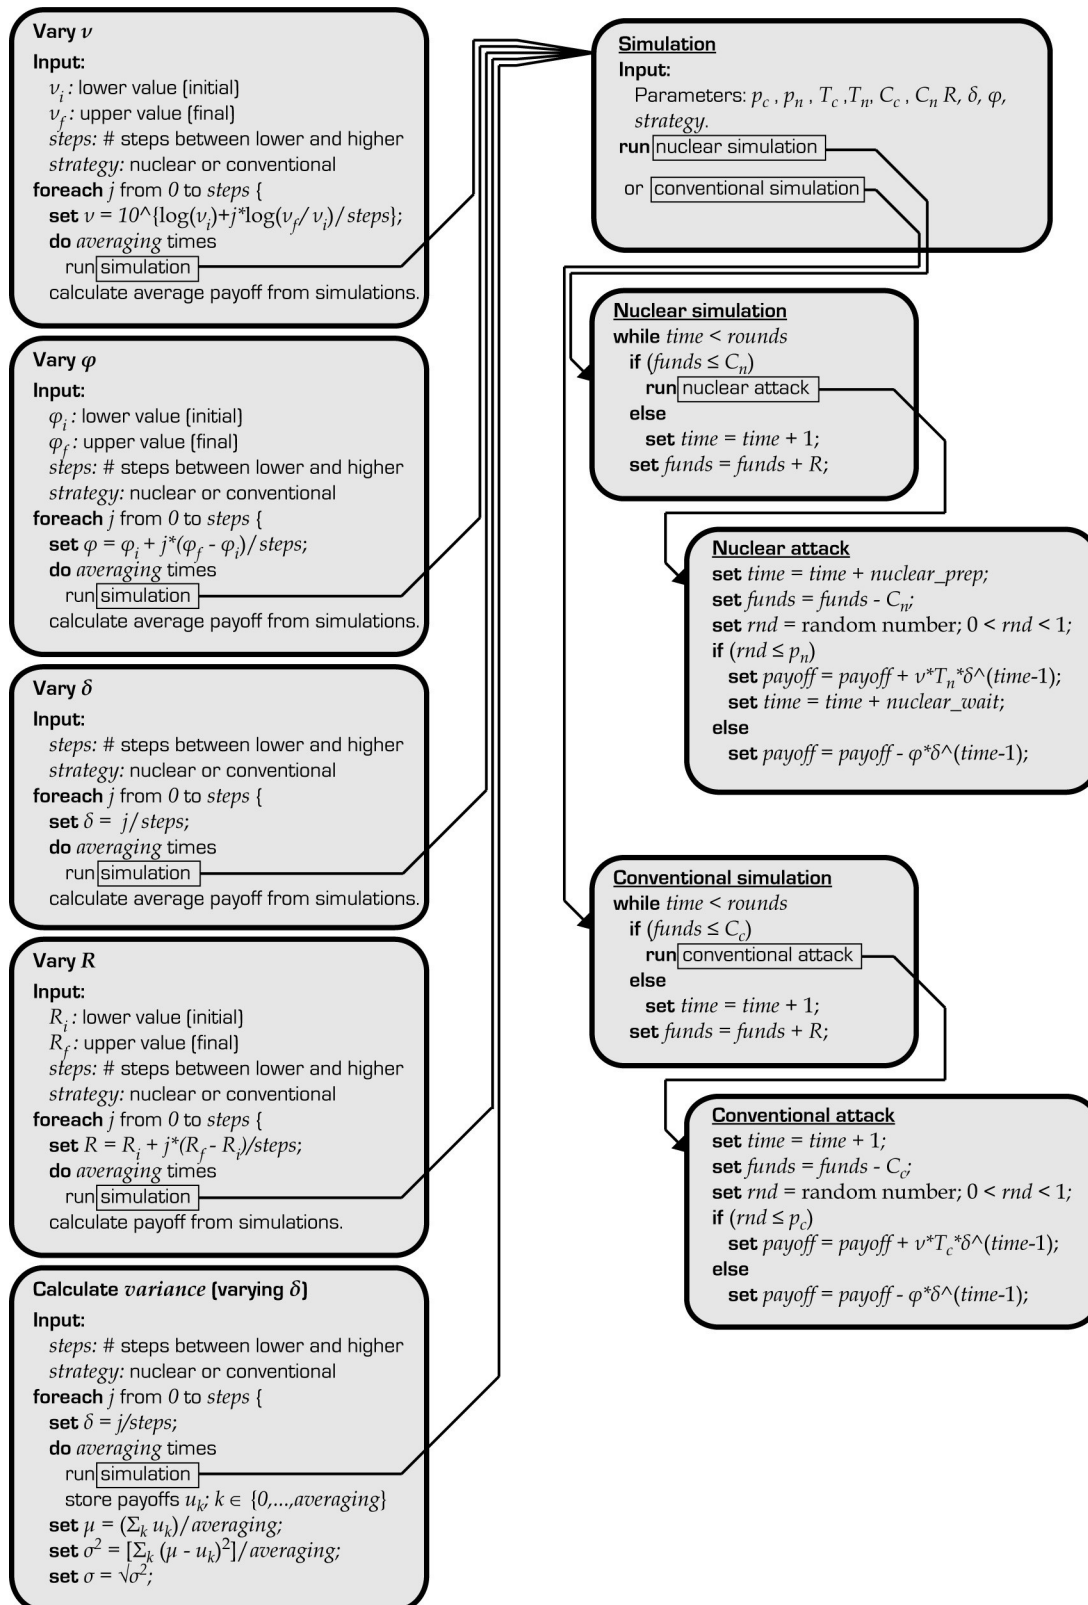


Illustration B.2.: The programme structure in more detail (with pseudocode).

- C -

Some mathematics

This chapter contains some mathematical support to our chapters. It is assumed throughout chapters 4 to 7 that the reader has a basic command of algebra and calculus, yet we provide below some support for reference to the less mathematically inclined reader. Our goal can neither be completeness nor rigour, hence for a deeper understanding the reader should refer to one of the hundreds of undergraduate level calculus textbooks available, close to all of which will cover all contents of this appendix in much greater depth.

We will focus on an intuitive understanding of the techniques employed rather than the algebraic implementation of these techniques themselves. For example, a graphical explanation of differentiation for finding local extrema is given, whereas the techniques for working out derivatives in practice are not provided. The below, thus, is not intended to provide the reader with the adequate mathematical background to *perform* the calculations presented (this would be too lofty an ambition for a mere appendix) but provide explanations which will hopefully suffice to aid the reader's understanding of the logic behind the treatment of the models.

C.1 Functions

In our treatment of all models, we speak of the *utility function*. Let us therefore take an intuitive look at what exactly constitutes a 'function' ('utility' or otherwise). The 'function' in mathematics is aptly named, for it is a close analogy to what we call a function in everyday life. A man's function is the action he does: the blacksmith transforms a lump of metal into a horse's shoe or a sword, a painter transforms paint and canvas into an artwork. Likewise, a mathematical function transforms one mathematical object (or a set of objects) into another object. It is often likened with a black box: you put something in and something else comes out the other end. Importantly, what comes out depends only on what you put in and on how the machine transforms the input. The function concept is extremely general, but here we will only consider simple functions that transform one or two input numbers, or *scalars*, into

one output number.

Let's look at a couple of simple examples. The café chain Costa is running its own charity providing education to children in third world countries. In May 2007, their money collecting boxes, found in all their cafés, read something like 'Costa will double all gifts collected in May'. Very commendable. This is a simple example of a function: the costumer gives £1, Costa *transforms* it into £2. The costumer gives £11.50, Costa transforms it into £23. More generally: the costumer gives x and Costa transforms it into $2x$. If we call Costa's function f , we thus have $f(x) = 2x$. So the notation $f(x)$ means the function f takes the variable x as input and gives you back the number $f(x)$. In this case both x and f have *dimension* money (measured in *units* of £). Generally, f and x will have different dimension.

A function can take in more than one number, too. An example is the body mass index, widely used in medicine to determine if a person's weight is healthy. The body mass index is defined as the weight, measured in kg, divided by the height, measured in metres, squared. Let the mass be m and the height h . Let the body mass index be B . Then, as the reader should verify,

$$B(m, h) = \frac{m}{h^2} . \quad (C.1)$$

As dimensions go, m has dimension mass, h has dimension length and B has dimension mass per length squared.

C.1.1 A little note about dimensions

We see that all variables need not be of the same dimension, nor does the function need to have the same dimension as any of its variables. But dimensions *must* be used consistently throughout all calculations. In our actual calculations in chapters 4 to 7, symbolised quantities are of only two different dimensions: money (such as costs, symbolised by e.g. C , T and φ with different subindices) and quantities with *no* dimension (simply numbers, or of 'dimension 1', if you will – examples are probabilities p , q and ρ , which are just numbers between 0 and 1).

In fact, dimensions can be thought of and treated much in the same way as mathematical symbols themselves. We will demonstrate. Say the variable M is of dimension mass. Then a numerical value of M is not merely a number: a unit of mass is also required. Any mass unit will do (kg, lb, tonnes, you name it)

but the *number* in M will change according to what unit is used. Say we use kg, being the SI unit. Then $M = [\text{a number}]$ multiplied by $[\text{a unit}]$. For example: $M = 5 \text{ kg}$: the *number* five multiplied by the *mass* 1kg. We normally omit the '.' between number and unit, but it is there nonetheless. Now let $m = 8 \text{ kg}$. Then by the basic laws of algebra:

$$m \cdot M = (8 \cdot \text{kg}) \cdot (5 \cdot \text{kg}) = 40 \cdot \text{kg}^2$$

and

$$m + M = (8 + 5) \cdot \text{kg} = 13 \cdot \text{kg}$$

and

$$\frac{m}{M} = \frac{8 \cdot \cancel{\text{kg}}}{5 \cdot \cancel{\text{kg}}} = \frac{8}{5}.$$

Note how we treat 'kg' like any other algebraic factor. In the left expression, we change the order of multiplication (dimensions are *commutative*), in the middle expression, 'kg' is a common factor, so we can move it outside the parentheses (dimensions are *distributive*). Finally in the rightmost expression, 'kg' is deleted above and below the fraction, just like we would with another common factor.

So what happens if we try to add together two numbers of different dimensions? Dimensions are then not a common factor and cannot be placed outside parentheses like in the middle above. We must conclude that we can *never* add (or, equivalently, subtract) quantities of different dimensions - it makes no sense to do so!

C.2 Graphs and plots

Once a function has been defined, such as $f(x)=2x$ or $B(m,h)$ in (C.1) above, we may plot the function (provided it takes no more than two variables as input) in an axes system. Figure C.1 shows graphs of Costa's function (left) and $B(m,h)$ (right).

When a function takes two variables in, its graph must be made three-dimensional, like a landscape. In fact, a map of a terrain can be thought of as a function of two variables, latitude and longitude for example, and whose output is height above sea level. This can be a good way of thinking about functions of two variables.

C.3 The derivative: the slope of the function

The next concept we introduce is the *derivative* of a function. For functions of a single variable, the

derivative has a very simple interpretation: it is the slope of the function when it's drawn as a graph. Imagine we travel along the graph from left to right along the abscissa (that's the horizontal axis). Then the derivative is just how much uphill or downhill we travel; how many units we must go vertically for every unit we travel horizontally.

We will not go into how to work out the derivative of a function – the reader can find out from any calculus textbook. It is the intuitive interpretation we're interested in. The derivative of our Costa function, $f(x) = 2x$, is

$$\frac{df}{dx} = 2.$$

That df/dx equals 2 is obvious from the graph to the left of C.1 and our explanation of the derivative: for every pound we move along the x -axis, we move 2 pounds upwards on the f -axis¹. Since our function f when plotted with respect to x is a straight line, its slope is constant - the number 2.

Note the notation df/dx , as it is used throughout our text; however, when the function depends on a single variable, a shorthand notation is often used: $df/dx \equiv f'$.

Now regard the function drawn in figure C.2. We will call the function $g(x)$. Never mind its algebraic form for now – we regard it purely geometrically. The slope of the graph in any one point is found by drawing a tangent to the graph. If the tangent points uphill (when going from left to right) the derivative is positive, if it points downhill, it is negative. From this it follows that when the tangent is perfectly flat, parallel with the abscissa, the derivative of the graph is zero.

We have marked the two places in the figure where the slope is zero. From this it is obvious why we will call these points *local extrema*: the points where $dg/dx=0$ are either local minima or maxima of the graph².

-
- 1 In the literature, one will often see used the terms x -axis and y -axis for the horizontal and vertical axes respectively. This is a potentially misleading convention when the variables do not happen to be dubbed x and y (in our case they are x and f). A better set of names is *abscissa* for the horizontal axis and *ordinate* for the vertical axis.
 - 2 In fact one can have zero-slope points that are neither minima nor maxima (for example, $f(x)=x^3$ has zero slope at $x=0$, but has no local minima or maxima), but we keep things simple here.




Illustration C.1.: Examples of functions: 'Costa function' (left) and BMI (right).

So assume we have found a point at which a function has zero slope. How do we know if it is a minimum or a maximum? For a function of a single variable the rule is simple (we will not show it here):

if $\frac{d^2 f}{dx^2} > 0$ at the extremum then the extremum is a minimum

if $\frac{d^2 f}{dx^2} < 0$ at the extremum then the extremum is a maximum.

where the notation d^2/dx^2 means that f is differentiated twice with respect to x . The shorthand notation when f depends on a single variable, is f'' . In more graphic terms: if $d^2f/dx^2 > 0$, it means that $f(x)$ is concave up; if $d^2f/dx^2 < 0$, it means that $f(x)$ is concave down. From figure C.2 it is easy to verify that wherever the graph is concave up, a point of zero slope must be a local minimum, and the opposite if it is concave down. In the special case where $d^2f/dx^2 = 0$, the test is inconclusive.

C.4 Functions of two variables and the partial derivative

When the function depends on more than one variable, the simple interpretation of slope when moving from left to right is not quite enough. Regard the BMI graph to the right of figure C.1, for example:

how 'steep' the hill is for someone walking at some point on the graph depends which direction he is walking.

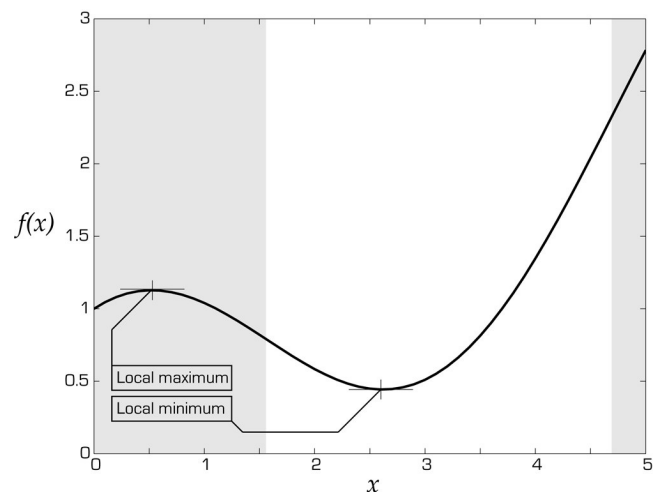


Illustration C.2.: A graph of a function of one variable.

This is why we need to generalise the derivative a little and introduce the *partial* derivative. A deep understanding of the difference between different kinds of derivatives is non-trivial and beyond our scope here. But it suffices for our simple purposes to explain in the following manner: When a function depends on more than one variable, the partial derivative is the derivative with respect to *one* variable, holding all others constant.

Regard the graph of figure C.3, for example. It is, as we can see, a 'hill' with a single maximum (the 'summit'). As is illustrated, the top of the hill may be found by finding the point on the hill where a *plane*

that runs tangent to the point (the *tangent plane*) is perfectly flat. That is: whichever direction one walks on this plane, one walks in zero slope. On figure C.3 there is only one such point, but we can easily imagine different cases where there are several. The tangent plane is the 3D equivalent of the tangent line for measuring slope on a graph like that in figure C.2.




Illustration C.3.: A 3D graph with a single maximum.

So how do we find such a plane? Obvious: the top of the hill is the only place where, by walking in two different directions 'on the map', we find zero slope simultaneously. In fact, as long as the two directions are not parallel or antiparallel³ to each other, any two directions will do the job. Nonetheless, two directions stand out in our graph as the natural choices lie along the the x -axis and y -axis.

Suppose now we start at the point where $y = 0$ and $x = 0.5$ (right above the word 'maximum' in the caption), and walk parallel to the y -axis, that is the direction 'into the paper' more or less, following the landscape. First the slope is uphill, we come to the top of this path (not the top of the hill, however), and then it's downhill again on the other side. Now we claim: at the top of this path the following is satisfied:

$$\frac{\partial g}{\partial y}(x=0.5)=0.$$

That is: the partial derivative of g with respect to y , taken at $x = 0.5$ is zero at the highest point of the path that comes from keeping x constant at 0.5 while varying y . We can do the same thing walking parallel to the x -axis at some constant value of y , and find another point which is the highest of *that* path. And so

3 i.e. point in opposite directions

on and so forth.

In order to find the highest point, however, we have a better method than such trial and failure. We will walk first along the y -axis (say) at some general and unspecified value of x , and find the highest point. This gives an equation involving x and y . Then we do the walk along the x -axis at some general value of y and do the same thing. We now have two equations in two variables and may solve them together to find the top. In figure C.3 it is found to be at $x = y = 1$.

This is exactly the technique used in chapter 4, where our sport is to find the maximum of the antiterrorist's utility function (which, when plotted, might look a little like figure C.3). However, our application is not simply to find the summit: in our particular application, what exactly is the ideal spending level is not so interesting, since the model is so simple and our numbers so rough. Much more powerfully, however, we can determine whether the planned spendings are too low or too high. We'll continue to use figure C.3 as example.

Let us assume that we are trying to find the top of the hill of the figure, but that we are lost somehow and don't know where exactly we are (that is, our xy -coordinate) nor where the summit is. What do we do? Of course: we check the terrain and make sure that whatever we do, we move uphill! Then, at least, we are sure to be getting closer to the top, not further from it.

This is just analogous to what we do in chapter 4. In chapter 4, we make rough estimates of some numbers. This allows us to use the equations we have developed to answer the following question: if we move along the x -axis (analogously: if we increase the value c_p) are we then moving uphill or downhill (analogously: are we then increasing or decreasing the expected utility)? Likewise we determine which way is uphill along the y -axis. Thus we may determine very roughly which way to go and which way not to go in order to get closer to the top. This is called the *method of steepest ascent*.

For the advanced reader, we may mention that the *optimal* direction⁴ to be moving in the xy -plane in order to reach the top is given by the gradient of $g(x,y)$

4 This is not quite true. If one has knowledge of all derivatives of the function and it is analytical one can find the direction which will take one to the summit in a straight line. If the available information is the function value and its first derivatives, however, the gradient is the best direction one can find, and a first order approximation of the real *optimal* direction.

- a vector in the xy -plane given by

$$\nabla g(x,y) \equiv \left[\frac{\partial g}{\partial x}, \frac{\partial g}{\partial y} \right].$$

The gradient points in the direction in which g increases fastest. This technique is particularly useful in the case where a total expenditure level for both branches (HEU and plutonium for example, as in chapter 4) is constant, fixed at a time before the division is performed, and totals significantly less than the required sum that can take us from the *status quo* point (somewhere down the slope) to the summit itself. If thus a small spending C is given each year, say, the government might do well to each year estimate the derivatives involved and spend according to the ideal spending vector.

C.5 Summation of power series

In this section the level of mathematical complexity is raised a little bit. It is a well known result that for a number δ so that $0 < |\delta| < 1$,

$$\sum_{i=0}^{\infty} \delta^i = \frac{1}{1-\delta}. \quad (\text{C.2})$$

We use this in chapters 5 and 6 for calculating the sum of payoffs from infinitely many rounds, multiplicatively discounted. The formula is easily shown. We start by multiplying the sum by the factor $(1-\delta)$. We then get a difference between two infinite sums in such a manner that all terms are cancelled save the first, and thus:

$$\begin{aligned} (1-\delta) \sum_{i=0}^{\infty} \delta^i &= (1+\delta+\delta^2+\delta^3+\delta^4+\dots) \\ &\quad - (\delta+\delta^2+\delta^3+\delta^4+\dots) \\ &= (1+\delta+\delta^2+\delta^3+\delta^4+\dots) \\ &\quad - (\delta+\delta^2+\delta^3+\delta^4+\dots) = 1, \end{aligned}$$

from which (C.2) follows easily.

From (C.2), furthermore, it follows that if we sum instead from $i = 1$,

$$\sum_{i=1}^{\infty} \delta^i = \left(\sum_{i=0}^{\infty} \delta^i \right) - 1 = \frac{1}{1-\delta} - 1 = \frac{\delta}{1-\delta},$$

a result we also make use of.

Finally, here's a trick I use in chapter 5. To understand how it works, the reader must be familiar with the basic rule of differentiation saying that

$$\frac{d}{dx} x^n = n x^{n-1}.$$

Using this, we can evaluate sums where the summation index is a factor in the summand, in some variation of the following example:

$$\sum_{n=1}^{\infty} n \cdot \delta^{n-1} = \sum_{n=1}^{\infty} \frac{d}{d\delta} \delta^n = \frac{d}{d\delta} \sum_{n=1}^{\infty} \delta^n = \frac{d}{d\delta} \left(\frac{\delta}{1-\delta} \right) = \frac{1}{(1-\delta)^2}.$$

Moving the differentiation operator outside the sum, notably, is not *always* allowed for infinite sums (the sum's convergence must be *uniform*, to use the correct language – for explanation, see some university level textbook on calculus). But this is ok for power series like that above (so long as $-1 < \delta < 1$ and $\delta \neq 0$).

- D -

Some thoughts on a 'Second Line of Defence'

In the preceding chapters we have focussed on different parts of a terrorist nuclear project. In chapter 6 we discussed the planning phase where the terrorist decides whether to try for a nuclear option or not. In chapters 3 to 5 the challenge of acquiring fissile material for a weapon was the focus point and in chapter 7 I discussed the strategic considerations of a terrorist who has already obtained a finished weapon. However, there are intermediate stages of every project to build and detonate a nuclear weapon which are not considered in the research chapters of the thesis, namely the smuggling of material and the finished weapon, and the technical and logistical challenges of assembling the weapon itself.

The tools available to a government to hamstring a nuclear project in progress at a time later than the acquisition of fissile materials are collectively dubbed the 'second line of defence'. The term spans a wide variety of different measures, typical examples of which could be intelligence and police work, radiation detection at border crossings and in harbours and airports, economic tools such as freezing of terrorists' assets¹, export control on nuclear and dual use equipment² or even some mechanisms which are not directed towards terrorism specifically but could still get in the way of a successful execution of a terrorist plot.

The issue of second layers of defence is not a central topic in this thesis and the analysis herein is but a brief discussion of how such measures may be thought of in light of the work of the preceding chapters. The present appendix was originally intended to become a research chapter, but with the publishing of Michael Levi's book *On Nuclear Terrorism*³ which I review below, whatever research

effort I had undertaken on the question of second layers of defence could make only incremental additions to Levi's work, and it was deemed that while the work already done on the issue might be interesting reading for some, the level of novelty of the work was no longer sufficient for a research chapter. Hopefully the reader can nonetheless extract some interesting thoughts and ideas from the modest contributions in this appendix.

The diversity of measures which form the second layer makes this part of the defence against nuclear terrorism in some ways more complex than the first layer of defence which focuses on hindering the spread of fissile materials. The complexity reflects the number of different paths a nuclear terrorism plot can take once the fissile material is secured (while no nuclear project can eschew the acquisition phase; as Allison puts it: 'it's a basic matter of physics: without fissile material, you can't have a nuclear bomb. No nuclear bomb, no nuclear terrorism'⁴). Possibilities are many at this stage. The weapon can be assembled in the target country or somewhere else. The possible smuggling routes into any given country are many and diverse, and each of the large number of minor challenges which must be negotiated on the way to a workable weapon presents the terrorist with a set of options. The paths to the bomb is formalised diagrammatically below in figure D.2.

D.1 Second layers in literature and policy: a brief review

The literature on nuclear terrorism has, at least until very recently, had a strong emphasis on the first layer of defence, and many experts have held that acquiring the necessary material is by far the most important obstacle to the nuclear terrorist. Ferguson and Potter, for example, suggest *only* safeguards-related measures in the face of terrorism with nuclear explosives, emphasising (in keeping with our conclusions from chapter 4) the importance of 'putting HEU first' in all priorities.⁵ Allison's comment above indicates a similar view, although he admits that 'even the best efforts to secure weapons and fissile material may not achieve 100 percent success, and that some nuclear material may already be loose, we cannot rely

1 e.g. Sidney Weintraub 'Disrupting the Financing of Terrorism' *The Washington Quarterly* 25:1 (2002) pp. 53-60.

2 James J. Wirtz 'Counter-terrorism via Counter-proliferation' *Terrorism and Political Violence* 14:3 (2002) pp.129-140.

3 Michael Levi *On Nuclear Terrorism* (Cambridge, MA: Harvard University Press, 2007)

4 Graham Allison 'How to Stop Nuclear Terror' *Foreign Affairs* 83 (2004) p.64.

5 Charles D. Ferguson and William C. Potter *The Four Faces of Nuclear Terrorism* (New York:Routledge, 2005)* Chapter 7.

exclusively on any single line of defense'.⁶ Maerli holds that the building of a gun design weapon⁷ is relatively 'easy' compared to acquisition of fissile material in sufficient quantity 'which is a prerequisite and probably the most formidable obstacle to the production of nuclear weapons'⁸.

As for interdicting a nuclear smuggling operation, Bunn, Wier and Holdren are amongst those who paint a bleak picture. Acquiring the material is the primary difficulty, they maintain, and 'once terrorists get or make a nuclear bomb, there is little to stop them delivering it to a U.S. city', leading to the immediate conclusion that fissile materials safeguards must be given the highest possible priority⁹. Several other analyses reach the same conclusion¹⁰.

Policy measures (by the US in particular) have apparently reflected a more positive view of the prospect of second layers. I will briefly look at some examples of US policy on a second line of defence. At the latest Non-Proliferation Treaty (NPT) Conference, in 2006, Thomas D. Lehrman of the US Department of State (DoS) presented his vision in which 'The United States must work together with partner nations and international organizations to develop a global layered defense against this threat.'¹¹

Lehrman mentions the Proliferation Security Initiative (PSI), a loose co-operation between the US and some 75 countries providing the legal framework to allow interdiction of transport of proliferation concern as a primary example of efforts this 'layered defence' should encompass¹². International naval law does not initially permit hindering the free passage of foreign ships on the high seas, one problem the PSI is

meant to salvage by a network of bilateral agreements¹³. Lehrman's vision goes beyond interdiction, however:¹⁴

Important as interdiction is, a comprehensive approach to combating WMD terrorism extends beyond interdiction capabilities. It involves developing and deploying capabilities to prevent and deter the full range of linkages - transport, travel, communications, and financial - between terrorists seeking WMD and their facilitators.

Notably, safeguarding nuclear materials is nowhere mentioned as part of the 'multilayered defense', possibly because it largely concerns nuclear weaponry, only one class of what Lehrman calls 'WMD'.

The US has sought to aid the detection of nuclear smuggling abroad and domestically. Under its 'Second Line of Defense' programme, the US Department of Energy's (DoE) National Nuclear Security Administration (NNSA) had by 2006 provided training and radiation detection equipment to 98 out of a planned 450 non-US border crossings it recognises as most important¹⁵. Caravelli, a former officer of the DoE nonproliferation efforts, grants the programme some success, but holds that in recent years DoE's efforts have been fumbled and mismanaged¹⁶, officers have been more interested in protecting their jobs than doing their work, Caravelli claims, and even accusations of deliberately misinforming Congress about the progress of the programme are made. Reporting from the ground on the border between Russia and Georgia where the NNSA has funded technologically sophisticated

6 Graham Allison *Nuclear Terrorism: The Ultimate Preventable Catastrophe* (New York: Times Books, 2004) p.199.

7 See chapter 3.

8 Morten Bremer Maerli 'Relearning the ABCs: Terrorists and "Weapons of Mass Destruction"' *The Nonproliferation Review* (Summer 2000) p. 113

9 M. Bunn, A. Wier and J.P. Holdren *Controlling Nuclear Warheads and Materials* report of the Project Managing the Atom (Belfer Center for Science and International Affairs, Harvard University, 2003)* p.15.

10 See references in the following.

11 Thomas D. Lehrman, Office of Strategic Planning and Outreach, DoS, 'Building a Layered Defense to Combat Weapons of Mass Destruction (WMD) Terrorism', remarks to the NPT Conference, Washington College of Law (February 9, 2006)*.

12 *ibid.*

13 Andreas Persbo and Ian Davis *Sailing Into Uncharted Waters? The Proliferation Security Initiative and the Law of the Sea* (British American Security Information Council, 2004)*.

14 Lehrman 'Building a Layered Defense...'

15 Matthew Bunn *Securing the Bomb 2007* report from Project Managing the Atom (Belfer Center for Science and International Affairs, Harvard University, 2007)* p.133

16 Jack Caravelli *Nuclear Insecurity* (Westport, CT: Praeger Security International, 2008) pp.45-58. Similar criticism has also come from the Government Accountability Office *Nuclear Nonproliferation: U.S. Efforts to Help Other Countries Combat Nuclear Smuggling Need Strengthened Coordination and Planning* GAO-02-426 (May 2002)*

radiation detection equipment at a number of border crossings, Langewiesche all but ridicules the US efforts to secure foreign borders¹⁷. Langewiesche relates how the Americans have built extravagantly expensive border protection instalments (whose potential is far from fully utilised) near the main points of entry while neglecting the rest of the large border almost entirely.

Under the DoE's 'Megaports' initiative and the US Customs and Border Protection's (CBP) 'Container Security Initiative' (CSI)¹⁸, the US has made an attempt to 'push its borders out' by equipping foreign harbours with radiation scanning equipment¹⁹. The equipment installed focuses on identifying 'high risk' containers and scanning them before they are loaded onto ships bound for US harbours. A recent review concluded that these measures have much room for improvement both in identifying the right containers to scan and in covering the routes from the countries of greatest concern²⁰. In the mock smuggling episode staged by ABC News described in section 3.4.3, a 6.8kg piece of depleted uranium was not detected although the container it was transported in was one of the small fraction which was scanned upon arrival in the US²¹.

17 William Langewiesche *The Atomic Bazaar* (New York: Farrar, Straus and Giroux, 2007) pp. 55-60.

18 A similar programme which primarily targets transportation over land rather than sea containers is the 'Customs Trade Partnership Against Terrorism' (C-TPAT), a voluntary programme under which exporters to the US are given benefits in exchange for adhering to a set of guidelines. Jon D. Haveman, Howard J. Shatz, and Ernesto A. Vilchis 'U.S. Port Security Policy after 9/11: Overview and Evaluation' *Journal of Homeland Security and Emergency Management* 2:4 (2005) Article 1.

19 Government Accountability Office (GAO) *Preventing Nuclear Smuggling: DOE Has Made Limited Progress in Installing Radiation Detection Equipment at Highest Priority Foreign Seaports* GAO-05-375 (2005)*, see also Bunn and Wier *Securing the Bomb 2006* report from Project Managing the Atom (Belfer Center for Science and International Affairs, Harvard University, 2006)* p.82.

20 For an overview, see Jon D. Haveman, Ethan M. Jennings, Howard J. Shatz, and Greg C. Wright 'The Container Security Initiative and Ocean Container Threats' *Journal of Homeland Security and Emergency Management* 4:1 (2007) Article 1.

21 Depleted uranium (basically what is 'left' after the ²³⁵U isotope has been extracted in enrichment; almost entirely ²³⁸U) is somewhat less radioactive than HEU, but physicists at the National Research Defense

Domestically, DoE plans to install some 3,000 radiation monitors at US points of entry by 2009 at a cost of \$1.3B²². This programme was recently criticised for being based on unrealistic expectations of detector performance²³. The fissile materials of interest to us - HEU and plutonium - are not particularly radioactive and are relatively easily shielded²⁴. These problems were summarised succinctly long ago by Robert Oppenheimer; when asked by Congress in 1946 how a nuclear weapon on its way to New York smuggled in a crate might be detected, he replied 'With a screwdriver'.²⁵ A recent task force report indicates that little has changed, saying 'Today, it would be easy for adversaries to introduce and detonate a nuclear explosive clandestinely in the United States'.²⁶

The prospect of scanning every container bound e.g. for the US is a daunting prospect. There are some

Council, who provided the uranium, argued that HEU could easily be shielded to give a comparably detectable signal. Charles D. Ferguson and William C. Potter *The Four Faces of Nuclear Terrorism* (New York:Routledge, 2005)* p.141.

22 Gene Aloise, Government Accountability Office (US) 'Combating Nuclear Smuggling: Challenges Facing U.S. Efforts to Deploy Radiation Detection Equipment in Other Countries and in the United States' testimony before the Permanent Subcommittee on Investigations and Governmental Affairs, US Senate. GAO-06-558T (March 28, 2006) p.4

23 Gene Aloise, GAO 'Combating Nuclear Smuggling: DHS's Decision to Procure and Deploy the Next Generation of Radiation Detection Equipment Is Not Supported By Its Cost-Benefit Analysis' Testimony, GAO-07-581T (2007)*

24 As previously noted, HEU has low activity of gamma radiation of low energy (186keV) which is weakly penetrating and easily shielded. Plutonium's gamma signature is somewhat more penetrating (640keV). Both emit neutrons from spontaneous fission but these are slow and harder to detect. The alpha radiation from ²⁴⁰Pu (a primary reason for plutonium's toxicity) will not even penetrate a sheet of paper. See Office of Nonproliferation Research and Engineering *Technology R&D for Arms Control* David Spears (ed.) (2001)*.

25 Richard L. Garwin 'The Technology of Megaterror' *Technology Review* 105:7 (September 2002)* p.64. Oppenheimer led the nuclear weapons research at Los Alamos during World War 2.

26 *Report of The Defense Science Board Task Force on Preventing and Defending Against Clandestine Nuclear Attack* (Washington D.C.:Office of the Under Secretary of Defense for Acquisition, Technology, and Logistics, 2004)* p.1.

6,000 shipping ports in the world, some 700 of which ship to the US directly²⁷. Expenses due to delays in the shipping current that would follow from scanning every freight container would be staggering. Even the most recently developed detection system under development at Los Alamos and Idaho National Laboratories requires at least 2 minutes of scanning time per container²⁸, and the systems actually deployed probably much more. The world's ten largest container ports had a collective throughput of some 162 million containers in 2007; Singapore alone shipped 28 million container units²⁹. The cost of delaying a significant fraction of these containers just 2 minutes is staggering.

The problem of interdicting smuggling is even larger, however. 'The number of pathways for smuggling a nuclear bomb or its ingredients into the United States is immense,' says Bunn and Wier, 'and intelligent adversaries will choose whichever route remains undefended'.³⁰ The US for example (excluding Alaska and islands) has some 10,000km of border, 8,000 km of coastline³¹ as well as almost countless small ports and marinas for small boats. The quantity of HEU required for a gun design, some 50-60 kg, is easily transported in small private boats, cars and aeroplanes that may be rented almost anywhere. Concentrating on container ships seems analogous to the parable of the drunk who looks for his lost keys under the streetlight instead of where he knows he dropped them: in a dark area.

Despite these pessimistic reports, Michael Levi differs from his expert peers and emphasises in a

recent book the importance of secondary defensive measures³². Levi's argument is twofold³³. Firstly, probabilities multiply. Levi agrees that each individual element of the large operation we dub the 'second layer' may have a slim chance of successfully derailing the nuclear proliferation project when acting in isolation. A radiation detector at a border crossing may not detect the nuclear material, guards can be bribed and intelligence agents can be avoided by a tight-knit group. However, if a number of such measures must be passed one after the other, the probability that *one* of them will catch the terrorist is not so small. As Levi states it³⁴

This perspective turns a cliché about terrorism on its head. It has often been said that defense against terrorism must succeed every time, but the terrorist must succeed only once. This is true from plot to plot, but within each plot, the logic is reversed. Terrorists must succeed at every stage, but the defense needs to succeed only once.

Say each element of a 'serial' defence has a 10% probability of success. Putting 10 of these elements together, the probability that the terrorist will be stopped in one of them is 65%, a much more agreeable figure which could well deter the rational terrorist from trying, as discussed in chapter 6.

Levi's second argument extends this last point. For every part of the defence which the terrorist dodges, he may have to take some detour from the most direct 'route' which could make him vulnerable to some other part of a complete set of defence mechanisms acting together. To illustrate this point Levi uses an analogy from baseball. If the value of a baseball team were to be assessed by evaluating each player in isolation one would come to the conclusion that the whole team was useless and stood no chance. What would be the point of a right-fielder, for example, when the batsman could simply play to the left field? Such an assessment clearly makes no sense: to understand the value of each player he must be seen in relation to the rest of his team³⁵. Transferring the analogy to nuclear smuggling, for example, the

27 Bunn and Wier *Securing the Bomb* 2006 p.82.

28 James L. Jones *et al* 'Detection of shielded nuclear material in a cargo container' *Nuclear Instruments and Methods in Physics Research A* 562 (2006) p.1087. The system is *active*, irradiating the container with high energy photons (10 MeV) and counting gamma and neutron emission. Passive methods (which only measure whatever particles radiate from an object without irradiation) are less potent and require longer scanning times. The standard reference on passive detection methods is Doug Reilly *et al.* (eds) *Passive Nondestructive Assay of Nuclear Materials* (US Nuclear Regulatory Commission, 1991) (kindly lent to the author by Dr. Morten Bremer Maerli)

29 Port of Rotterdam *Port Statistics* 2007 (2008)* p.14.

30 Matthew Bunn and Anthony Wier 'Terrorist Nuclear Weapon Construction: How Difficult?' *The Annals of AAPSS* 607 (2006) p.142

31 Janice Ch. Beaver *U.S. International Borders: Brief Facts* (Congressional Research Service, 2006).

32 Levi *On Nuclear Terrorism*

33 For a more detailed review of Levi's book see Simen A. Ellingsen 'Rethinking the Second Line of Defense' *The Nonproliferation Review* (Summer 2008) To be published.

34 Levi *On Nuclear Terrorism* p. 7

35 *ibid.* p. 6

terrorist who wishes to avoid radiation detection in US harbours by smuggling the material over the land border to Mexico could end up being apprehended by the immigration authorities, a measure not even intended primarily to stem terrorism.

Levi in no way disagrees that a heavy emphasis on securing proliferation attractive fissile materials is justified; on the contrary he argues that amongst the most important premises of a successful secondary defence is warning provided by the security and accountability measures in place where the fissile material is stored. Quite simply: the governments stand a very much better chance of regaining control over stolen material if they know it has been stolen, preferably shortly after the theft took place.

D.2 A diagrammatic outline of the paths to nuclear terrorism

When speaking of different 'layers of defence', it is useful to provide some visual exemplification of what is meant. A suitable procedure is to look at the necessary steps a terrorist must take one way or another in order to acquire a nuclear capacity. Upon doing so, one finds that while options are many, there is a finite number of principle paths leading to this goal. I have outlined this in figure D.2 including for completeness also the possibility for direct acquisition of a ready weapon from a national stockpile³⁶. The diagram is a generalisation and extension of that provided by Bunn, Wier and Holdren a few years ago³⁷.

Figure D.2 is largely self-explanatory. Perhaps its most fundamentally important trait is the limited number of ways via which a terrorist may acquire a nuclear bomb, and that each path involves a certain succession of steps to be taken. Some variation from that shown in the figure is certainly possible especially as regards the ordering of the steps (some tasks can also be performed in parallel). Importantly, however, along a given path *no step may be skipped*³⁸, meaning that if the terrorists are unable to manage even one of the tasks, they will have no bomb, in line with what Levi concluded.

36 This option is included in our definition of nuclear terrorism in section 1.2, but is generally disregarded in this thesis for reasons of manageability of scope.

37 Bunn, Wier and Holdren *Controlling Nuclear Warheads and Materials* p.21.

38 Note that there are thinkable ways the plot could unfold so that not all the points are relevant.

The first task, forming a capable group with extreme motives, may only be hindered in the long run by addressing the root causes of terrorism itself, an interesting and much researched field which lies outside the scope of this thesis. The next point, deciding to escalate to nuclear level of violence, may be addressed through measures towards deterrence, as understood and treated in chapter 6. Beyond this point, the terrorists must be *physically* stopped from completing their mission. The 'first line of defence' or 'safeguards' denotes all efforts to physically block the terrorist from acquiring military nuclear weapons or fissile materials, such measures as are treated in great detail by Bunn and co-workers³⁹. Any efforts to stop the terrorist project other than hindering their acquisition of weapons or fissile materials are commonly dubbed 'second layers'.

D.3 Value for money in second layers: the challenge

In the following I will use Levi's understanding of second layers of defence as a premise and seek to outline how this frame of thinking could be coupled with the decision theoretical method employed in the research chapters of this thesis. A brief comparison with the efforts in chapter 4 shows that the case is somewhat similar in that an evaluation of second layers of defence involves finding a way to prioritise between different efforts which all cost money but improve security by some amount which can in principle be measured in dollar equivalents as reduced nuclear terrorism threat. A central question must then be how to evaluate the value for money of a particular branch of the second line of defence.

In chapter 4 this was rather straightforward because there was a very limited amount of overlap between the security measures concerning HEU and that targeting plutonium. As for second layers, as we have seen, the situation is far more complex. Remembering the baseball analogy it becomes clear that measuring the success rate of each piece of the defence in isolation makes little sense and will certainly lead to an underestimation of the value of each part of the defence.

In principle, the methodology explained in chapter 3 and exemplified in chapter 4 is applicable to second line of defence measures as well, but the challenge of arriving analytically at a reasonably realistic model, evaluate the corresponding utility function U and its

39 e.g. Bunn *Securing the Bomb* 2007

dependency on the potentially large number of parameters at the government player's disposal, is formidable.

The remainder of this chapter is a modest attempt to analyse very crudely the interplay between first and second lines of defence and outline possible approaches through which a thorough evaluation of second layers of defence may be undertaken. We start by devising and analysing a simple model of nuclear terrorism defences, followed by a short presentation of future directions and some early conclusions.

D.4 A simple model of nuclear terrorism countermeasures

To formalise these thoughts a little I'll introduce a very simple game similar to that used in chapter 4. For a further discussion of the various aspects of the analysis, see that chapter. The government first sets the spending levels for first and second layer defences, equal to C_1 and C_2 respectively, whereupon the terrorist player attempts to acquire necessary nuclear material for a simple device (either using HEU or plutonium - I will not distinguish explicitly between them here) and second to smuggle the material, build the device, smuggle it to its target and detonate it, following the steps on the right hand side of figure D.2. The terrorist (player T) is assumed to successfully acquire the material with probability p and to go successfully through all the secondary hurdles (building and transporting) with probability q . Unlike before, I cannot assume that these probabilities are uncorrelated, because there is an important correlation between the quality of safeguards measures and the likelihood that a theft is detected quickly, even if it is not stopped⁴⁰. Thus I assume $p=p(C_1)$ whereas $q=q(C_2, p)$, that is q depends on C_1 as well through the probability p ⁴¹. The game is shown in figure D.1. As before, circles denote either a choice node or end node, and the 'angle' symbols denote a choice from a continuum. I will use the notation

$$\frac{\partial}{\partial C_1} \equiv \partial_1; \quad \frac{\partial}{\partial C_2} \equiv \partial_2.$$

Similar to chapter 4 (and for the very same reasons as therein) we assume

40 See Levi *On Nuclear Terrorism* pp. 98-123

41 Note that this assumption differs from that made in chapter .

$$\frac{dp}{dC_1} \equiv p' \leq 0; \quad \partial_2 q \leq 0.$$

Furthermore, the better the first layer of defence, the better the second, so

$$\partial_1 q = \frac{\partial q}{\partial p} p' \leq 0,$$

which implies that

$$\frac{\partial q}{\partial p} \geq 0.$$

The payoffs of the game are $-C-T$ if the terrorist can pass both the first and the second layer of defence, otherwise it is $-C$, so the probability we need to establish is the that of the terrorist succeeding at both stages. Here $C = C_1 + C_2$ denotes the total cost of all defensive measures spent by the government player (G) and T is the estimated cost of a terrorist (player T) nuclear attack. Although p and q are now assumed to be correlated, the probability of a successful nuclear attack is still merely the product of these two quantities⁴².

From figure D.1 we then readily get player G's expected payoff as

$$U = -C - pqT \quad . \quad (D.1)$$

I will go through the now standard procedure of locating the maximum of the utility function with respect to C_1 and C_2 which is found where the derivatives of U with respect to both variables are zero⁴³. This gives us two equations:

$$(-p') \left(q + p \frac{\partial q}{\partial p} \right) = \frac{1}{T} \quad (D.2)$$

42 Let A be the event that the terrorist successfully acquires sufficient fissile material and B be the probability that she is able to mount an attack. With basic statistics this gives the probability that both events happen as

$P(A \wedge B) = P(B|A)P(A) = P(B)P(A) = qp$ as before because by the law of total probability

$P(B|A) = P(B) - P(B|\bar{A}) = P(B)$. Here a bar means 'not'. That $P(B|\bar{A}) = 0$ is another way of saying as Allison does: 'No fissile material, no nuclear weapon' (see footnote 4).

43 Strictly, this must be checked. I assume as usual that $p'' \geq 0$ and $q'' \geq 0$ so the sufficient criterion is that $\partial_2^2 U / \partial C_1^2 \leq 0$ and $\partial_2^2 U / \partial C_2^2 \leq 0$ are fulfilled everywhere, hence a critical point is a maximum.

$$p(-\partial_2 q) = \frac{1}{T} \quad (D.3)$$

Although the analysis is extremely simple, (D.2) and (D.3) could be useful in their simplicity, because their significance is easily interpreted: if player G's best estimates indicate that $(-p')(q+p \cdot \partial q / \partial p) > 1/T$, it means that safeguards efforts are underfunded while $(-p')(q+p \cdot \partial q / \partial p) < 1/T$ means the opposite: safeguards receive more money than can be defended in terms of threat reduction. Likewise for second layer efforts: '>' in place of '=' in (D.3) means too little money, '<' means too much.




Illustration D.1.: A simple model of second layers.

The quantity $(-\partial_2 q)$ may be said to be the 'value for money' of all second layer efforts while it is not entirely obvious exactly what the value for money for safeguards measures is here; setting it at $(-p')$ will somewhat underestimate the value since the positive effects for the second layer are then not taken into account. A better measure of 'value for money' (which I denote $-P'$) is the left hand side of (D.2) divided by q :

$$(-P') \equiv (-p') \left(1 + \frac{p}{q} \frac{\partial q}{\partial p} \right) \quad (D.4)$$

~ Value for money spent on safeguards.

so that (D.2) reads

$$q(-P') = \frac{1}{T} \quad (D.5)$$

In order to make practical use of these formulae, one must be able to establish estimated values of the

symbols involved. This is a difficult task even when all secondary measures are seen as a whole, yet more troubling perhaps is the fact that rather than regarding the second layer of defence as a single entity, a government will need a tool with which to evaluate each element of this highly complex defensive effort. Making a single gaming representation for each element of the defence will underestimate the value as Levi argues.

D.4.1 Tentative numerical example

My ambitions are modest here, but I find that even a very tentative numerical exercise can add valuable insight into the above. In section 4.4 I used the highly approximate measure $(-p') \approx 3\% / \$1B$ for the US efforts to secure nuclear materials abroad, based roughly on a progress report by Bunn and Wier⁴⁴. Here, let the 'value for money' quantity (D.4) take that tentative value. Furthermore, let us – again very tentatively – use a damage estimate of $T = \$1,000,000,000,000$ (one trillion dollars). I found in section 4.4 that a less conservative estimate indicates a true cost of an attack with a $\sim 10kT$ device in a major city might rather be ~ 3 trillion dollars, but one must bear in mind that the estimate used here encompasses the probability that the bomb may have a yield dramatically smaller than this (a 'fizzle'⁴⁵).

With these numbers one finds that the value of q corresponding with an ideal level of spending on safeguards measures would be $q_{ideal} = [(-P')T]^{-1} \approx 3.3\%$, as seen from (D.5). This is probably lower than the real probability even in light of Levi's optimistic analysis (clearly such an estimate must necessarily have a large uncertainty), although some experts (see review above) will almost certainly argue that this is on the low side⁴⁶. Note how a larger value of T (which is possibly underestimated here) will decrease q_{ideal} , while a smaller P' will increase it. If, after a more elaborate estimation process than I have gone through, finds that q_{ideal} is indeed significantly lower than the actual value of q as estimated by the government player, it means that more money should be spent on safeguards. I believe it more than

44 Bunn and Wier *Securing the Bomb* 2006.

45 See chapter 3 for details.

46 Remember I used a 2-10% estimate for the product of p and q in chapter 7, in accordance with Matthew Bunn 'A Mathematical Model of the Risk of Nuclear Terrorism' *The Annals of the AAPSS* 607 pp.103-120 (2006)

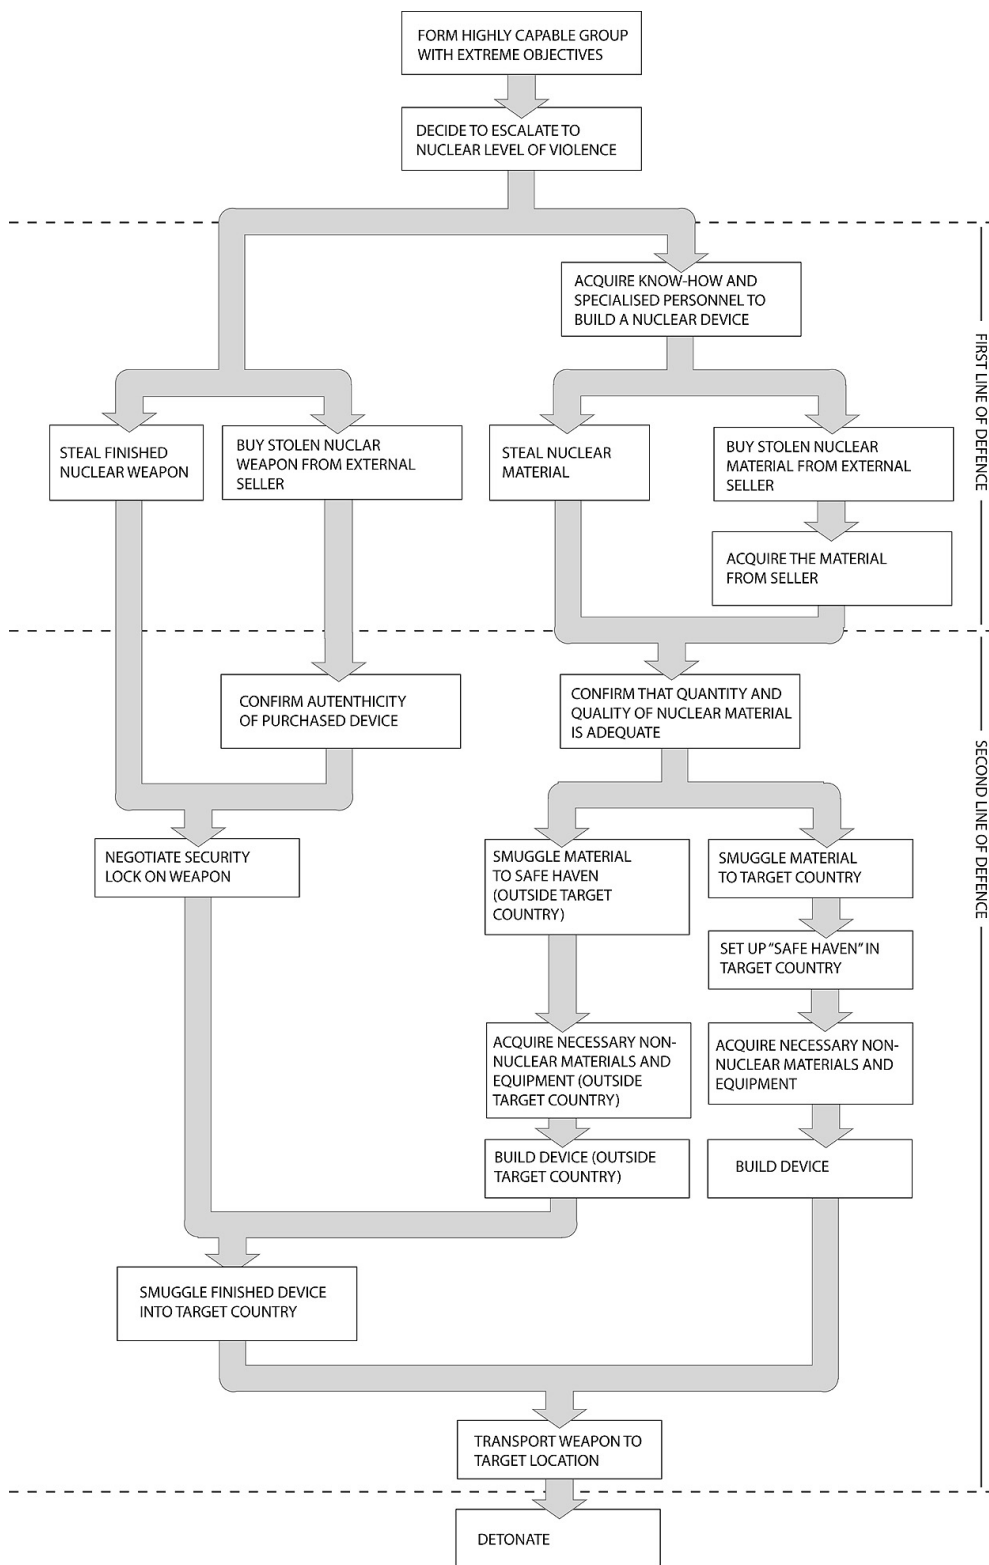


Illustration D.2.: Different paths to a nuclear weapon.

arguable that this conclusion would be correct based on the above figures. Notice that the value of the second term between parentheses in (D.4) which represents the positive effect of improved safeguards on secondary measures, could be significant and tends to further strengthen this conclusion, in line with what Levi emphasises⁴⁷.

Mathematically, a similar analysis could be made for secondary measures as a whole, yet the complex nature of these defences mean that such an effort may not be very helpful. Estimating a 'value for money' for secondary measures as a whole is of limited interest, since this value will depend strongly on exactly how the money is spent⁴⁸ and provides no help with the equally important question of how best to organise and prioritise the different elements of the 'second line' optimally.

Nonetheless the guiding principle for evaluating the second layer must in principle be the same as for safeguards measures, that the money that go into 'improvement' should always be spent so as to maximise threat reduction per monetary unit, that is, in such a way that each dollar (say) makes the greatest difference possible. This should be the guiding principle behind any effort to analyse the value of second layers of defence quantitatively. Devising a good way to ensure this in a useful way, however (as opposed to evaluating the feasibility of second layers as a whole), is a large and complex task which lies beyond the scope of this thesis (indeed, it could well form a thesis worth of research in itself). I will restrict myself to pointing out a direction in which future efforts could be directed.

D.5 Outlook: Optimum funding and organisation of of the 'second line of defence'

In this section I will only provide a few paragraphs of outline of an enterprise which could form a suitable task for one or several PhD theses, namely the devising and analysing of a model to aid optimising the funding and operation of second layers of defence against nuclear terrorism. The methodology will be

47 Levi *On Nuclear Terrorism* Chapter 5.

48 The same is true for safeguards measures as well. For example, it matters whether one safeguards uranium or plutonium as concluded in chapter 4. However, the various measures which sort under 'safeguards' as the term is used herein are far less diverse than are those of the 'second line of defence'.

principally identical to that made use of in chapters 5, 6 and in particular 4 of this thesis, which consists in principle of 3 steps:

1. Make a model of the situation which takes account of the (most) relevant interactions which
 - a. Minimises the number of free parameters to be determined for manageability, but...
 - b. ...is still complex enough to capture the essence of the situation modelled.
2. Work out, either analytically or (more likely in case of a complex project) numerically the utility function of the government player as a function of the modelling parameters.
3. Find a suitable method to
 - a. optimise the utility system with respect to the free parameters and ...
 - b. ... analyse the dependency on the various parameters so as to gain a better and more general understanding.

I will discuss these points briefly one by one.

In the relatively simple cases analysed in chapters 4, 5 and 6, devising a model was fairly straightforward, and yet a fair amount of judgement had to be applied in the process. There are no general prescriptions available as to how to determine what the essential interactions are, so the modelling must be based on a qualitative discussion, and the approach may differ from case to case. Doubtless, some will find that important aspects of the situation studied in those chapters are missing or could have been represented better.

The task of modelling is much more complex when considering secondary layers of defence, however. As Levi argues, in order to evaluate one element of the defence, the interactions of that element with the rest of the second line (which may be defined so as to include measures not directed at nuclear terrorism or even conventional terrorism) must be taken into account. A fruitful approach could be the tools available in the field of network modelling⁴⁹. The second layer of defence could thus be modelled somewhat crudely as a physical network of defensive mechanisms and with the paths to nuclear terrorism having to pass nodes of this network.

49 See e.g. Nigel Gilbert and Klaus G. Troitzsch *Simulation for the Social Scientist* 2nd ed. (New York: Open University Press, 2005) Chapter 3.

The utility function of such a network model is in principle possible to find analytically, but will probably be much too complicated for symbolic analysis such as that undertaken in the present thesis to be fruitful. There are however a number of numerical approaches available from the literature on *optimisation theory*⁵⁰, commonly used in engineering subjects. One such method has already been used in this thesis, namely the 'method of steepest ascent/descent' reviewed in section 4.2.1, where the (multidimensional) gradient of the utility function is calculated in a given state of the world which points in the direction in which U increases most rapidly. Optimisation is then performed by taking incremental steps in the direction of the gradient and re-evaluating the gradient (numerically) for each step. More sophisticated methods are also available to improve convergence⁵¹, but a complete survey is beyond the scope of this section. In simpler models the combination of an explicit (albeit complicated) utility function and optimisation theory could form a potent approach in attacking this complex task.

Perhaps a more realistic approach of purely numerical nature is the use of multi-agent models⁵². In this approach each element of the second layer is simulated as one or more *agents*, that is, it is given 'a mind of its own'. Rules are specified dictating how each agent interacts with other agents. The terrorist player is made an agent as well with explicitly ordered preferences and a goal to successfully execute a nuclear plot from materials acquisition to detonation. A degree of randomness is introduced governing the chances of 'defence agents' stopping the terrorist. Such an approach will effectively perform step 2 and 3 above together and the analysis of parameters will be entirely implicit and numerical rather than explicit and symbolical as was the case in our simple models in previous chapters.

D.6 Some early conclusions

Second layers of defence against nuclear terrorism have been treated to some extent in a large number of publications on this issue, yet has rarely been the

focus of attention. The conventional wisdom has been that once the fissile materials are in terrorist hands, a target government's means to derail the terrorist attempt to acquire nuclear explosives are severely limited, and that for this reason emphasis should be on securing the fissile materials in the first place. In a recent book, however, Michael Levi disputed this pessimism, and while agreeing with the importance of safeguards, introduced a framework within which one can argue that second layers of defence are not as futile as formerly portrayed.

A further inquiry into the consequences of Levi's framework for practical policy prioritisation appears warranted, and would seem to provide fertile ground for future research on nuclear terrorism defences. Herein I propose a very general scheme by which the decision theoretical (economical) tools employed in this thesis may be united with Levi's concept to arrive at more detailed recommendations for optimal government of second layer efforts.

I conclude, similar to previous chapters, that the deciding quantities when evaluating the balance of second layers versus safeguards, as well as between different elements of both layers of defence, are the threat level (most prominently: probability of terrorist success) and value for money (probability reduction per monetary unit).

Two simple formulae are derived giving, provided numerical estimates of the quantities involved, a guideline for the relative emphasis between first and second layer measures, taking into account the positive synergies of good material accounting system for the success of secondary measures. This is, however, merely the beginning of what should be a thorough investigation of this important topic, possibly suitable for a future PhD.

50 There exists a number of textbooks in this field. One good introductory textbook is Hubertus Th. Jongen, Klaus Meer, and Eberhard Triesch *Optimization Theory* (New York: Kluwer Academic Publishers, 2004)

51 See e.g. Jongen et.al *Optimization Theory*.

52 See e.g. Chapter 8 of Gilbert and Troitzsch *Simulation for the Social Scientist* and references therein.

Bibliography

The bibliography contains only items that are explicitly referred to in the footnotes of chapters and appendices. Literature beyond that listed below has been drawn upon for further background. Governmental documents without specified authors are listed by the issuing office. Items are sorted by author(s), then chronologically where several consecutive items by the same author(s) occur. Where URLs are given, the reference has been marked with an asterisk (*) when cited in the main text.

All URLs were checked in September 2008 and refer files at the official website of the webpage of the publishing institution except where explicitly indicated.

Dissertations

Bunn, Matthew *Guardians at the Gates of Hell: Estimating the Risk of Nuclear Theft and Terrorism - and Identifying the Highest-Priority Risks of Nuclear Theft* PhD dissertation (Massachusetts Institute of Technology, 2007). Online:

http://esd.mit.edu/people/dissertations/bunn_matthew.pdf

Maerli, Morten Bremer *Crude Nukes on the Loose?: Preventing Nuclear Terrorism by Means of Optimum Nuclear Husbandry, Transparency, and Non-Intrusive Fissile Material Verification* PhD dissertation (University of Oslo, 2004). Online from the Norwegian Institute of International Affairs (NUPI):

http://www.nupi.no/content/download/1474/39531/version/2/file/MBM_dissertation20042.pdf

Books and Monographs

Allison, Graham T. *Essence of Decision: Explaining the Cuban Missile Crisis* (Boston: Little, Brown, 1971)

Allison, Graham *Nuclear Terrorism: The Ultimate Preventable Catastrophe* (New York: Times Books, 2004)

Alvarez, Luis *The Adventures of a Physicist* (New York: Basic Books 1988)

Barry, Brian *Sociologists, Economists and Democracy: Themes and Issues in Modern Sociology* (Collier-Mac, 1970)

Berger, James O. *Statistical Decision Theory* (Berlin: Springer, 1985)

Chernoff, Herman and Lincoln E. Moses *Elementary Decision Theory* (New York: Courier Dover Publications, 1988)

Corera, Gordon *Shopping for Bombs: Nuclear Proliferation, Global Insecurity, and the Rise and Fall of the A.Q. Khan Network* (London: Hurst & co., 2006)

Coster-Mullen, John *Atom Bombs: The Top Secret Inside Story of Little Boy and Fat Man* (Self-published, 2004) (Available for purchase directly from author)

Dixon, Lloyd and Rachel Kaganoff Stern *Compensation for Losses from the 9/11 attacks* (RAND, 2004). Online: http://www.rand.org/pubs/monographs/2004/RAND_MG264.pdf

Falkenrath, Richard A., Robert D. Newman, and Bradley A. Thayer *America's Achilles' Heel: Nuclear, Biological, and Chemical Terrorism and Covert Attack* (Cambridge MA: MIT Press, 1998)

Ferguson, Charles D. and William C. Potter with Amy Sands, Leonard S. Spector and Fred L. Wehling *The Four Faces of Nuclear Terrorism* (New York: Routledge 2005). Online: http://www.nti.org/c_press/analysis_4faces.pdf

Robert Gibbons *A Primer in Game Theory* (Hemel Hempstead: Harvester Wheatsheaf, 1992)

Gilbert, Nigel and Klaus G. Troitzsch *Simulation for the Social Scientist* 2nd edition (Buckham: Open University Press, 2005)

Glasstone, Samuel and Philip J. Dolan (eds) *The Effects of Nuclear Weapons* 3rd edition (Washington D.C.: GPO, 1977). Online from the IPFM Library: http://www.fissilematerials.org/ipfm/site_down/gla77.pdf

Hoddeson, Lillian, Paul W. Henriksen, Roger A. Meade and Catherine Westfall *Critical Assembly: A Technical History of Los Alamos during the Oppenheimer Years 1943-1945* (Cambridge: Cambridge University Press, 1993)

Jackson, Brian A., John C. Baker, Peter Chalk, Kim Cragin, John V. Parachini, Horacio R. Trujillo *Aptitude for Destruction, Volume 1: Organizational Learning in Terrorist Groups and Its Implications for Combating Terrorism* (Santa Monica, CA: RAND, 2005). Online: http://www.rand.org/pubs/monographs/2005/RAND_MG331.pdf

Jackson, Brian A., John C. Baker, Peter Chalk, Kim Cragin, John V. Parachini, Horacio R. Trujillo *Aptitude for Destruction, Volume 2: Case Studies of Organizational Learning in Five Terrorist Groups* (Santa Monica, CA: RAND, 2005). Online: http://www.rand.org/pubs/monographs/2005/RAND_MG332.pdf

Jenkins, Brian Michael *Unconquerable Nation: Knowing our Enemy Strengthening Ourselves* (Santa Monica: RAND, 2006). Online: http://www.rand.org/pubs/monographs/2006/RAND_MG454.pdf

- Jongen, Hubertus Th., Klaus Meer, and Eberhard Triesch *Optimization Theory* (New York: Kluwer, 2004)
- Langewiesche, William *The Atomic Bazaar* (New York: Farrar, Strauss and Giroux, 2007)
- Lee III, Rensselaer W. *Smuggling Armageddon: The Nuclear Black Market in the Former Soviet Union and Europe* (New York: St. Martins Press, 1998)
- Levi, Michael *On Nuclear Terrorism* (Cambridge, MA: Harvard University Press, 2007)
- Libicki, Martin C., Peter Chalk, and Melanie Sisson *Exploring Terrorist Targeting Preferences* (Santa Monica, CA: RAND, 2007). Online: http://www.rand.org/pubs/monographs/2007/RAND_MG483.pdf
- McPhee, John *The Curve of Binding Energy* (New York: Farrar, Straus & Giroux, 1974)
- Osborne, Martin J. *An Introduction to Game Theory* (New York: Oxford University Press, 2004)
- Reilly, Doug, Norbert Ensslin, Hastings Smith, jr., and Sarah Kreiner (eds) *Passive Nondestructive Assay of Nuclear Materials* (US Nuclear Regulatory Commission, 1991)
- Rhodes, Richard *The Making of the Atomic Bomb* (New York: Simon & Schuster, 1986)
- Serber, Robert *The Los Alamos Primer: The First Lectures on How To Build An Atomic Bomb* (Los Angeles: University of California Press, 1992)
- Schelling, Thomas *The Strategy of Conflict* (Cambridge, MA: Harvard University Press, 1960)
- Trunin, R. F. *Shock Compression of Condensed Materials* (Cambridge: Cambridge University Press, 1998)
- Immediate Action* report from Project Managing the Atom (Belfer Center for Science and International Affairs, Harvard University, 2002). Online: http://www.nti.org/e_research/securing_nuclear_weapons_and_materials_May2002.pdf
- Bunn, Matthew and Anthony Wier *Securing the Bomb 2006* report from Project Managing the Atom (Belfer Center for Science and International Affairs, Harvard University, 2006). Online: http://www.nti.org/e_research/stb06webfull.pdf
- Bunn, Matthew, Anthony Wier and John P. Holdren *Controlling Nuclear Warheads and Materials* report of the Project Managing the Atom (Belfer Center for Science and International Affairs, Harvard University, 2003) Online: http://www.nti.org/e_research/cnwm/cnwm.pdf
- Carter, Ashton B., Michael M. May, and William J. Perry *The Day After: Action in the 24 Hours Following a Nuclear Blast in an American City* Workshop Report (The Preventive Defense Project, Harvard and Stanford Universities, 2007). Online: http://belfercenter.ksg.harvard.edu/files/dayafterworkshopreport_may2007.pdf
- The Commission on the Intelligence Capabilities of the United States Regarding Weapons of Mass Destruction *Report to the President of the United States* (March 2005) Online: http://www.wmd.gov/report/wmd_report.pdf
- Davis, Paul K. and Brian Michael Jenkins *Deterrence & Influence in Counterterrorism: A Component in the War on al Qaeda* (RAND, 2002). Online: http://www.rand.org/pubs/monograph_reports/2005/MR1619.pdf
- Defense Science Board Task Force on Preventing and Defending Against Clandestine Nuclear Attack *Report of the Task Force* (Office of the Under Secretary of Defense for Acquisition, Technology, and Logistics, 2004). Online: http://www.acq.osd.mil/dsb/reports/2004-06-Clandestine_Nuclear_Attack.pdf

Reports (governmental and non-governmental)

- Albright, David and Kimberly Kramer *Global Stocks of Nuclear Explosive Materials* (The Institute for Science and International Security, September 2005). Online: http://www.isis-online.org/global_stocks/end2003/tableofcontents.html
- Beaver, Janice Cheryl U.S. *International Borders: Brief Facts Report RS21729* (Congressional Research Service, 2006) Online: <http://www.opencrs.com/document/RS21729/>
- Bunn, Matthew *Securing the Bomb 2007* report from Project Managing the Atom (Belfer Center for Science and International Affairs, Harvard University, 2007). Online: http://www.nti.org/e_research/securingthebomb07.pdf
- Bunn, Matthew, John P. Holdren and Anthony Wier *Securing Nuclear Weapons and Materials: Seven Steps for*
- Etzioni, Amitai *Pre-empting Nuclear Terrorism in a New Global Order* (The Foreign Policy Center, 2004). Online: <http://fpc.org.uk/fsblob/314.pdf>
- Finlay, Brian and Andrew Grotto *The Race to Secure Russia's Loose Nukes: Progress since 9/11* (The Henry L. Stimson Center, 2005). Online: <http://www.stimson.org/ctr/pdf/LooseNukes.pdf>
- Finlay, Brian D. and Elizabeth Turpen *25 Steps to Prevent Nuclear Terror: A Guide for Policymakers* (The Henry L. Stimson Center, 2007). Online: http://www.stimson.org/cnp/pdf/25_Steps_complete.pdf
- Frank, W.J. (ed.) *Summary Report of the Nth Country Experiment* (Lawrence Radiation Laboratory, 1967) partly unclassified (excised). Online: <http://www.gwu.edu/~nsarchiv/news/20030701/nth->

country.pdf from The National Security Archive.

- Government Accountability Office (US) *Nuclear Nonproliferation: U.S. Efforts to Help Other Countries Combat Nuclear Smuggling Need Strengthened Coordination and Planning* GAO-02-426 (May 2002). Online: <http://www.gao.gov/new.items/d02426.pdf>
- Government Accountability Office (US) *Weapons of Mass Destruction: Nonproliferation Programs Need Better Integration* GAO-05-157 (January 2005). Online: <http://www.gao.gov/new.items/d05157.pdf>
- Government Accountability Office (US) *Preventing Nuclear Smuggling: DOE Has Made Limited Progress in Installing Radiation Detection Equipment at Highest Priority Foreign Seaports* GAO-05-375 (March 2005). Online: <http://www.gao.gov/new.items/d05375.pdf>
- International Panel of Fissile Materials *Global Fissile Material Report 2007* (2007). Online: http://www.fissilematerials.org/ipfm/site_down/gfmr07.pdf
- Jenkins, Brian Michael *Countering Al Qaeda: An Appreciation of the Situation and Suggestions for Strategy* (RAND, 2002). Online: http://www.rand.org/pubs/monograph_reports/2005/MR1620.pdf.
- Lesser, Brian O. , Bruce Hoffman, John Arquilla, David Ronfeldt, Michele Zanini, and Brian Michael Jenkins *Countering the New Terrorism* (RAND, 1999). Online: http://rand.org/pubs/monograph_reports/MR989/
- Meade, Charles and Roger C. Molander *Considering the Effects of a Catastrophic Terrorist Attack* Technical Report (RAND, 2006). Online: http://www.rand.org/pubs/technical_reports/2006/RAND_T_R391.pdf
- National Commission on the Terrorist Attacks Upon the United States *The 9/11 Report* St. Martin's Paperback Edition (New York: St. Martins, 2004) (Official Government Edition, under the title *The 9/11 Commission Report*, 2002 Available Online: <http://www.9-11commission.gov/report/911Report.pdf>)
[Note: page numbers differ between the two editions; page number references refer to the paperback edition]
- Office of Nonproliferation Research and Engineering *Technology R&D for Arms Control* edited by David Spears (Washington D.C.: US Department of Energy, 2001). Online: http://www.partnershipforglobalsecurity.org/Documents/acnt_war.pdf from the Partnership for Global Security Library
- Office of Technology Assessment, US Congress *Nuclear Proliferation and Safeguards* (OTA, June 1977). Online: <http://www.wws.princeton.edu/cgi-bin/byteserv.prl/~ota/disk3/1977/7705/7705.pdf>
- Office of Technology Assessment, US Congress *Technologies Underlying Weapons of Mass Destruction* OTA-BP-ISC-115 (December 1993). Online: <http://www.wws.princeton.edu/cgi-bin/byteserv.prl/~ota/disk1/1993/9344/9344.pdf>
- Persbo, Andreas and Ian Davis *Sailing Into Uncharted Waters? The Proliferation Security Initiative and the Law of the Sea* (British American Security Information Council, 2004). Online: <http://www.basicint.org/pubs/Research/04PSI.htm>
- Project on Government Oversight *U.S. Nuclear Weapons Complex: Security at Risk* (2001). Online: <http://www.pogo.org/p/environment/eo-011003-nuclear.html>
- Smelser, Neil J., and Faith Mitchell (eds) *Discouraging Terrorism: Some Implications of 9/11* report from the Panel on Understanding Terrorists in Order to Deter Terrorism, Center for Social and Economic Studies, National Research Council (Washington D.C.: National Academies Press, 2002)
- Thompson, Jr., William C. (Comptroller, New York City) *One Year Later: The Fiscal Impact of 9/11 on New York City* (New York: Office of the Comptroller, September 2002) Online: <http://www.comptroller.nyc.gov/bureaus/bud/reports/impact-9-11-year-later.pdf>
- Tikhonov, Valentin *Russia's Nuclear and Missile Complex: The Human Factor in Proliferation* (Carnegie Endowment for International Peace, 2001). Online: <http://www.carnegieendowment.org/publications/index.cfm?fa=view&id=656>
- Wier, Anthony and Matthew Bunn *Funding for U.S. Efforts to Improve Controls Over Nuclear Weapons, Materials and Expertise Overseas: Recent Developments and Trends* report from Project Managing the Atom (Belfer Center for Science and International Affairs, Harvard University, 2007). Online: http://www.nti.org/e_research/cnwm/wier_bunn_fy08budget.pdf

Papers, articles, reviews and chapters in edited volumes

- Abrahms, Max 'Are Terrorists Really Rational? The Palestinian Example' *Orbis* (Summer 2004) pp.533-549
- Ackerman, Gary 'WMD Terrorism Research: Where to from Here?' *International Studies Review* 7 (2005) pp.140-143
- Albright, David 'Al Qaeda's Nuclear Program: Through the Window of Seized Documents' *Policy Forum Online* The Nautilus Institute, *Special Forum* 47 (November 2002). Online: http://www.nautilus.org/archives/fora/Special-Policy-Forum/47_Albright.html

- Albright, David, Kathryn Buehler and Holly Higgins 'Bin Laden and the bomb' *The Bulletin of Atomic Scientists* 58:1 (2002) pp.23-24
- Albright, David and Holly Higgins 'A Bomb for the Ummah' *The Bulletin of Atomic Scientists* 59:2 (2003) pp.49-55
- Allison, Graham T. 'Conceptual Models and the Cuban Missile Crisis: Rational Policy, Organization Process, and Beurocratic Politics' RAND paper #3919 (1968) Online: <http://www.rand.org/pubs/papers/2008/P3919.pdf>
- Allison, Graham T. 'Conceptual Models and the Cuban Missile Crisis' *The American Political Science Review* 63:3 (1969) pp.689-718
- Allison, Graham 'How to Stop Nuclear Terror' *Foreign Affairs* 83:1 (2004) pp.64-74
- Allison, Graham 'The ongoing failure of imagination' *Bulletin of the Atomic Scientists* 62:5 (2006) pp.36-41
- Allison, Graham T. and Morton H. Halperin 'Bureaucratic Politics: A Paradigm and Some Policy Implications' *World Politics* 24:2 (1972) pp.40-79
- Allison, Graham and Andrei Kokoshin 'The New Containment: An Alliance against Nuclear Terrorism' in Juliette N. Kayyem and Robyn L. Pangi (eds) *First to Arrive: State and Local Responses to Terrorism* (Cambridge, MA: MIT Press, 2004) pp.9-20.
- Arce M., Daniel G. and Todd Sandler 'Counterterrorism: A Game-theoretic Analysis' *Journal of Conflict Resolution* 49:2 (2005) pp.183-200
- Arce, Daniel G. and Todd Sandler 'Terrorist Signalling and the Value of Intelligence' *British Journal of Political Science* 37:4 (2007) pp.573-586
- Badey, Thomas J. 'Defining International Terrorism: A Pragmatic Approach' *Terrorism and Political Violence* 10:1 (Spring 1998) pp.90-107
- Baker, Richard D., Siegfried S. Hecker, and Delbert R. Harbur 'Plutonium: A Wartime Nightmare but a Metallurgist's Dream' *Los Alamos Science* (Winter/Spring 1983) pp.142-151. Online: <http://www.fas.org/sgp/othergov/doe/lanl/pubs/00416629.pdf>
- Basuchoudhary, Atin and Laura Razzolini 'Hiding in plain sight: using signals to detect terrorists' *Public Choice* 128 (2006) pp.245-255
- Bennett, Peter G. 'Modelling Decisions in International Relations: Game Theory and Beyond' *Mershon International Studies Review* 39 (1995) pp.19-52
- Berkhout, Frans, Anatoli Diakov, Harold Feiveson, Helen Hunt, Edwin Lyman, Marvin Miller, and Frank von Hippel 'Disposition of Separated Plutonium' *Science and Global Security* 3:3 (1993) pp. 161-213
- Bowen, Wyn Q. 'Deterring Mass-Casualty Terrorism' *Joint Force Quarterly* 31 (Summer 2002) pp.25-29. Online: http://www.dtic.mil/doctrine/jel/jfq_pubs/0731.pdf
- Bowen, Wyn Q. 'Deterrence and Asymmetry: Non-State Actors and Mass Casualty Terrorism' *Contemporary Security Policy* 25:1 (2004) pp.54-70
- Broome, John 'Should a Rational Agent Maximize Expected Utility?' in Karen Schweers Cook and Margaret Levi (eds) *The Limits of Rationality* (Chicago: The University of Chicago Press, 1990) pp. 132-145
- Bunn, George, Fritz Steinhausler and Lyudmilla Zaitseva 'Strengthening Nuclear Security Against Terrorists and Thieves Through Better Training' *The Nonproliferation Review* (Fall-Winter 2001) pp.1-13
- Bunn, Matthew 'Preventing a Nuclear 9/11' *Issues in Science and Technology* 21:2 (Winter 2005) pp.55-62 Online: <http://www.issues.org/21.2/bunn.html>
- Bunn, Matthew 'A Mathematical Model of the Risk of Nuclear Terrorism' *The Annals of the American Academy of Political and Social Science* 607 (2006) pp.103-120
- Bunn, Matthew and Anthony Wier 'The Seven Myths of Nuclear Terrorism' *Current History* (April 2005) pp.153-161
- Bunn, Matthew and Anthony Wier 'Terrorist Nuclear Weapon Construction: How Difficult?' *The Annals of the American Academy for Political and Social Science* 607 (2006) 133-149
- Burgess, Mark 'Terrorism: The Problems of Definition' Center for Defense Information (2002) Main issue: 'Explaining Terrorism'. Online: http://www.cdi.org/program/document.cfm?DocumentID=1564&from_page=../index.cfm
- Cameron, Gavin 'WMD Terrorism in the United States: The Threat and Possible Countermeasures' *The Nonproliferation Review* (Spring 2000) pp.162-179
- Caplan, Bryan 'Terrorism: The relevance of the rational choice model' *Public Choice* 128 (2006) pp.91-107
- Carter, Ashton B., Michael M. May, and William J. Perry 'The Day After: Action Following a Nuclear Blast in an American City' *The Washington Quarterly* 30:4 (2007)
- Clarke, R. H., John Dunster, Jean-Claude Nenot, Hylton Smith, and George Voeltz 'The environmental safety and health implications of plutonium' *Journal of Radiological Protection* 16:2 (1996) pp.91-105
- Cohen, Bernard L., 'Hazards from Plutonium toxicity' *Health Physics* 32:5 (1977) pp. 359-379
- Davis, Jay 'The Attribution of WMD Events' *Journal of Homeland Security* (April 2003). Online: <http://www.homelandsecurity.org/journal/Articles/Davis.html>
- deLeon, Peter, Bruce Hoffman, Konrad Kellen, and Brian

- Jenkins 'The Threat of Nuclear Terrorism: A Reexamination' RAND note #2706 (1988). Online: <http://www.rand.org/pubs/notes/2007/N2706.pdf>
- Dolnik, Adam 'Die and Let Die: Exploring Links between Suicide Terrorism and Terrorist Use of Chemical, Biological, Radiological, and Nuclear Weapons' *Studies in Conflict & Terrorism* 26 (2003) p.17-35
- Dunlop, William and Harold Smith 'Who Did It? Using International Forensics to Detect and Deter Nuclear Terrorism' *Arms Control Today* (October 2006) pp.6-10
Online: http://www.armscontrol.org/act/2006_10/CVRForensics.asp
- Dunn, Lewis A. 'Can al Qaeda Be Deterred from Using Nuclear Weapons?' *Occasional paper* 3, July 2005, Center for the Study of Weapons of Mass Destruction, National Defense University. Online: http://www.ndu.edu/WMDCenter/docUploaded/206-186_CSWMD_OCP3WEB.pdf
- Ellingsen, Simen A. 'Strengthening the Second layer of Defense' (Review of Michael Levi's *On Nuclear Terrorism*) *The Nonproliferation Review* 15:2 (2008) pp.399-402
- Ellingsen, Simen A. 'Safeguards Against Nuclear Terrorism: HEU vs Plutonium' *Defense and Security Analysis* 24:2 (2008) pp. 129-146
- Elster, Jon 'When Rationality Fails' in Karen Schweers Cook and Margaret Levi (eds.) *The Limits of Rationality* (Chicago: University of Chicago Press, 1990) pp.19-46.
- Elster, Jon 'Emotions and Political Economy' *Journal of Economic Literature* 36 (1998) pp.47-74
- Falkenrath, Richard A. 'Confronting Nuclear, Biological and Chemical Terrorism' *Survival* 40:3 (1998) pp.43-65
- Ferguson, Charles D., and William C. Potter 'Improvised Nuclear Devices and Nuclear Terrorism' Paper #2, Weapons of Mass Destruction Commission (June 2004)
Online: www.wmdcommission.org/files/no2.pdf
- Fetter, Steve, Valery A. Frolov, Marvin Miller, Robert Mozley, Oleg F. Prilutsky, Stanislav N. Rodionov, and Roald Z. Sagdeev 'Detecting Nuclear Warheads' *Science & Global Security* 1 (1990) pp.225-302
- Franck, Raymond E. and Francois Melese 'Exploring the Structure of Terrorists' WMD Decisions: A Game Theory Approach' *Defense & Security Analysis* 20:4 (2004) pp.355-372
- Frey, Bruno S. and Simon Luechinger 'Decentralization as a disincentive for terror' *European Journal of Political Economy* 20 (2004) pp. 509-515
- Frost, Robin M. 'Nuclear Terrorism Post-9/11: Assessing the Risks' *Global Security* 18:4 (2004) p.397-422
- Frost, Robin M. 'Nuclear Terrorism After 9/11' *Adelphi Paper* 45:378 (2005)
- Gallucci, Robert L. 'Averting Nuclear Catastrophe: Contemplating Extreme Responses to U.S. Vulnerability' *The Annals of the American Academy of Political and Social Science* 607 (2006) pp.51-58
- Garwin, Richard L. 'The Many Threats of Terror' *The New York Review of Books* (November 1st, 2001)
Online with epilogue: <http://www.fas.org/rlg/020300-strikingterror.pdf> from the Federation of American Scientists
- Garwin, Richard L. 'The Technology of Megaterror' *Technology Review* 105:7 (September 2002) p. 64-70
Online: <http://www.technologyreview.com/Energy/12925/>
- Glaser, Alexander 'On the Proliferation Potential of Uranium Fuel for Research Reactors at Various Enrichment Levels' *Science and Global Security* 14 (2006) pp. 1-24
- Goeree, Jacob K. and Charles A. Holt 'Ten Little Treasures of Game Theory and Ten Intuitive Contradictions' *The American Economic Review* 91:5 (2001) pp. 1402-1422
- Gordon, Avishag 'Terrorism as an Academic Subject after 9/11: Searching the Internet Reveals a Stockholm Syndrome Trend' *Studies in Conflict & Terrorism* 28 (2005) pp.45-59
- Gottemoeller, Rose with Rebecca Longworth 'Enhancing Nuclear Security in the Counter-Terrorism Struggle: India and Pakistan as a New Region for Cooperation' Carnegie Endowment Working Paper #29 (August 2002). Online: <http://www.carnegieendowment.org/files/wp29.pdf>
- Gressang IV, Daniel S. 'Audience and Message: Assessing Terrorist WMD Potential' *Terrorism and Political Violence* 13:3 (2001) pp.83-106
- Harney, Robert, Gerald Brown, Matthew Carlyle, Eric Skroch, and Kevin Wood 'Anatomy of a Project to Produce a First Nuclear Weapon' *Science and Global Security* 14 (2006) pp.163-182
- Haveman, Jon D. , Ethan M. Jennings, Howard J. Shatz, and Greg C. Wright 'The Container Security Initiative and Ocean Container Threats' *Journal of Homeland Security and Emergency Management* 4:1 (2007) Article 1.
- Haveman, Jon D. , Howard J. Shatz, and Ernesto A. Vilchis 'U.S. Port Security Policy after 9/11: Overview and Evaluation' *Journal of Homeland Security and Emergency Management* 2:4 (2005) Article 1.
- Hecker, Sigfried 'Toward a Comprehensive Safeguards System: Keeping Fissile Materials Out of Terrorists' Hands' *The Annals of the American Academy of Political and Social Science* 607 (2006) pp.121-132
- Hoffman, Bruce 'Terrorism and Weapons of Mass Destruction: An Analysis of Trends and Motivations'

- RAND Paper #8039-1 (1999). Online:
<http://www.rand.org/pubs/papers/2007/P8039-1.pdf>
- Ivanova, Kate and Todd Sandler 'CBRN Incidents: Political Regimes, Perpetrators, and Targets' *Terrorism and Political Violence* 18 (2006) pp.423-448
- Jackson, Brian A. 'Technology Acquisition by Terrorist Groups: Threat Assessment Informed by Lessons from Private Sector Technology Adoption' *Studies in Conflict & Terrorism* 24 (2001) pp.183-213.
- Jenkins, Brian Michael 'Will Terrorists Go Nuclear?' RAND paper #5541 (1975). Online:
<http://www.rand.org/pubs/papers/2006/P5541.pdf>
- Jenkins, Brian Michael 'The Potential for Nuclear Terrorism' RAND paper #5876 (1977). Online:
<http://www.rand.org/pubs/papers/2006/P5876.pdf>
- Jenkins, Brian Michael 'Will Terrorists Go Nuclear?' *Orbis* 29:3 (1985) pp.507-515
- Jenkins, Brian Michael 'The Likelihood of Nuclear Terrorism' RAND paper #7119 (1985)
- Jenkins, Brian Michael 'The New Age of Terrorism' in David G. Kamien (ed.) *The McGraw-Hill Homeland Security Handbook* (New York: McGraw Hill, 2006) pp. 117-130
Online from RAND:
https://www.rand.org/pubs/reprints/2006/RAND_RP1215.pdf
- Jones, James L. *et al.* 'Detection of shielded nuclear material in a cargo container' *Nuclear Instruments and Methods in Physics Research A* 562 (2006) pp.1085-1088
- Kaplan, Stanley and B. John Garrick 'On the Qualitative Definition of Risk' *Risk Analysis* 1:1 (1981) p.11-27
- Kamp, Karl-Heinz 'An Overrated Nightmare' *The Bulletin of the Atomic Scientists* 52:4 (1996) pp.30-34.
- Kamp, Karl-Heinz, Joseph F. Pilat, Jessica Stern and Richard A. Falkenrath 'WMD Terrorism: An Exchange' *Survival* 40:4 (1998) pp.168-183
- Lapan, Harvey E. and Todd Sandler 'To Bargain or Not to Bargain: That is the question' *American Economic Review* 78:2 (1988) pp.16-20.
- Lee, Rensselaer 'Nuclear Smuggling: Patterns and Responses' *Parameters* 33:1 (Spring 2003) pp.95-111
- Levi, Michael A. 'Deterring Nuclear Terrorism' *Issues in Science and Technology Online* 20:3 (Spring 2004)
Online: <http://www.issues.org/20.3/levi.html>
- Loomes, Graham and Robert Sugden 'Regret Theory: An Alternative Theory of Rational Choice under Uncertainty' *The Economic Journal* 92 (1982) pp.805-824
- Luongo, Kenneth N. and Naeem Salik 'Building Confidence in Pakistan's Nuclear Security' *Arms Control Today* (December 2007). Online:
http://www.armscontrol.org/act/2007_12/Luongo.asp
- Maerli, Morten Bremer 'Relearning the ABCs: Terrorists and «Weapons of Mass Destruction»' *The Nonproliferation Review* (Summer 2000) pp.108-119
- Maerli, Morten Bremer 'Timely Options for Increased Transparency and Nonintrusive Verification on Fresh Highly Enriched Uranium Naval Fuel' *Journal of Nuclear Material Management* 31:4 (2003) pp. 18-30
- Mark, J. Carson 'Reactor-Grade Plutonium's Explosive Properties' Nuclear Control Institute (1990). Online:
<http://www.nci.org/NEW/NT/rgpu-mark-90.pdf>
- Mark, Carson, Theodore Taylor, Eugene Eyster, William Maraman, and Jacob Wechsler 'Can Terrorists Build Nuclear Weapons?' in Leventhal and Alexander (eds) *Preventing Nuclear Terrorism* (Lexington: Lexington Books, 1987) pp.57-66. Online from the Nuclear Control Institute: <http://www.nci.org/k-m/makeab.htm>
- May, Michael 'September 11 and the Need for International Nuclear Agreements', in *Nuclear Issues in the Post-September 11 Era; Recherches & Documents No 30* (Foundation pour la Recherche Stratégique, 2003) pp. 113-124. Online:
http://www.frstrategie.org/barreFRS/publications/rd/Telechargements/rechdoc30_eng.doc
- May, Michael, Jay Davis, and Raymond Jeanloz 'Preparing for the Worst' *Nature* 443 (26 October 2006) pp.907-908.
- Melese, Francois 'A Brinkmanship Game Theory Model of Terrorism' presented at the 3rd Conference on Mathematical Methods of Counterterrorism, Washington D.C. (September 2006) (kindly supplied by the author).
- Melese, Francois and Diana Angelis 'Deterring Terrorists from Using WMD: A Brinkmanship Strategy for the United Nations' *Defense & Security Analysis* 20:4 (2004) p.337-341
- Miller, Michael 'Nuclear Attribution as Deterrence' *Nonproliferation Review* 14:1 (2007) pp.33-60
- Mueller, John 'Simplicity and Spook: Terrorism and the Dynamics of Threat Exaggeration' *International Studies Perspectives* 6:2 (2005) pp. 208-234
- Mueller, John and Karl Mueller 'Sanctions of Mass Destruction' *Foreign Affairs* 78:3 (1999) pp.43-53
- Pape, Robert A. 'The Strategic Logic of Suicide Terrorism' *American Political Science Review* 97:3 (2003) pp.343-361
- Perkovich, George 'Deconflating "WMD"' Paper #17, Weapons of Mass Destruction Commission (October 2004)
Online: www.wmdcommission.org/files/no17.pdf
- Persky, Joseph 'The Ethology of Homo Economicus' *Journal of Economic Perspectives* 9:2 (1995) pp.221-231

- Phillips, Matthew 'Uncertain Justice for Nuclear Terror: Deterrence of Anonymous Attacks Through Attribution' *Orbis* 51:3 (2007) pp.429-446
- Pittel, Karen and Dirk T. Rübbecke 'What Directs a Terrorist?' *Defence and Peace Economics* 17:4 (2006) pp.311-328
- Pluta, Anna M. and Peter D. Zimmerman 'Nuclear Terrorism: A Disheartening Dissent' *Survival* 48:2 (2006) pp.55-70
- Potter, William C. 'Nuclear Terrorism and the Global Politics of Civilian HEU Elimination' *The Nonproliferation Review* 15:2 (2008) pp. 135-158
- Potter, William C., Charles D. Ferguson, and Leonard S. Spector 'The Four Faces of Nuclear Terrorism and the Need for a Prioritized Response' *Foreign Affairs* 83:3 (2004) pp.130-132
- Rabin, Matthew 'Psychology and Economics' *Journal of Economic Literature* 36 (1998) pp.11-46
- Rapoport, David C. 'Terrorism and Weapons of the Apocalypse' *National Security Studies Quarterly* 5:3 (1999) pp.49-66
- Richelson, Jeffrey T. 'Defusing Nuclear Terror' *Bulletin of the Atomic Scientists* 58:2 (2002) pp. 38-43
- Riedel, Bruce 'The Return of the Knights: al-Qaeda and the Fruits of Middle-East Disorder' *Survival* 49:3 (September 2007) pp.107-120
- Rosenau, William 'Aum Shinrikyo's Biological Weapons Program: Why Did it Fail?' *Studies in Conflict & Terrorism* 24 (2001) pp.289-301
- Salama, Sammy and Lydia Hansell 'Does Intent Equal Capability?: Al-Qaeda and Weapons of Mass Destruction' *Nonproliferation Review* 12:3 (2005) pp.615-653
- Sandler, Todd 'Collective Action and Transnational Terrorism' *The World Economy* 26:6 (2003) pp.779-802
- Sandler, Todd and Daniel G. Arce M. 'Terrorism & Game Theory' *Simulation and Gaming* 34:3 (2003) pp.319-337
- Sandler, Todd and Walter Enders 'An economic perspective on transnational terrorism' *European Journal of Political Economy* 20 (2004) pp.301-316
- Sandler, Todd and Walter Enders 'September 11 and Its Aftermath' *International Studies Review* 7 (2005) p.165-168
- Sandler, Todd and Kevin Siqueira 'Global terrorism: deterrence versus pre-emption' *Canadian Journal of Economics* 39:4 (2006) pp.1370-1387
- Sandler, Todd, John T. Tschirhart, and Jon Cauley 'A Theoretical Analysis of Transnational Terrorism' *The American Political Science Review* 77:1 (1983) pp.36-54;
- Simon, Herbert 'From Substantive to Procedural Rationality' in Spiro J. Latsis (ed.) *Methods and Appraisal in Economics* (Cambridge: Cambridge University Press, 1976) pp.130-131
- Sinai, Joshua 'Forecasting Terrorists' Likelihood to Embark on "Conventional" to CBRN Warfare' *International Studies Review* 7:1 (2005) pp.151-153
- Siqueira, Kevin and Todd Sandler 'Terrorists versus the Government: Strategic interactions, support and sponsorship' *Journal of Conflict Resolution* 50:6 (2006) pp.878-898
- Spector, Leonard S. 'The New Landscape of Nuclear Terrorism' in Michael Barletta (ed.) *After 9/11: Preventing Mass-Destruction Terrorism and Weapons Proliferation Occasional Paper #8* (Center for Nonproliferation Studies, Monetary Institute of International Studies, 2002). Online: <http://cns.miis.edu/pubs/opapers/op8/op8.pdf>
- Steinhausler, Friedrich 'What It Takes to Become a Nuclear Terrorist' *American Behavioral Scientist* 46:6 (2003) pp.782-795
- Stober, Dan 'No Experience Necessary' *The Bulletin of the Atomic Scientists* 58:2 (2003) pp.56-63
- Talmadge, Caitlin 'Deterring a Nuclear 9/11' *The Washington Quarterly* 30:2 (2007) pp.21-34
- Trager, Robert F. and Dessislava P. Zagorcheva 'Deterring Terrorism: It Can Be Done' *International Security* 30:3 (2005) pp.87-123
- von Hippel, Frank 'Recommendations for Preventing Nuclear Terrorism' *Journal of the Federation of American Scientists* 54:6 (2001) pp. 1-10. Online: <http://www.fas.org/faspir/2001/v54n6/v54n6.pdf>
- Walt, Stephen M. 'Rigor or Rigor Mortis?: Rational Choice and Security Studies' *International Security* 23:4 (1999) pp.5-48
- Ward, Hugh 'Rational Choice' in D. Marsh and G. Stoker (eds) *Theory and Methods in Political Science* 2nd edition (Hampshire:Palgrave MacMillan, 2002) pp. 65-89
- Weintraub, Sidney 'Disrupting the Financing of Terrorism' *The Washington Quarterly* 25:1 (2002) pp. 53-60
- Whiteneck, Daniel 'Deterring Terrorists: Thoughts on a Framework' *The Washington Quarterly* 28:3 (2005) pp.187-199
- Wik, Mette 'Individual Decision Theory under Risk: Deficiencies and alternatives to expected utility theory' Department of Economics and Resource Management, Norwegian Agricultural University, Discussion Paper #D-23 (1996). Online: <http://www.umb.no/ior/publikasjoner/D9623.pdf>
- Willis, Henry H. 'Guiding Resource Allocations Based on

- Terrorism Risk' *Risk Analysis* 27:3 (2007) pp.597-606
- Wirtz, James J. 'Counter-terrorism via Counter-proliferation' *Terrorism and Political Violence* 14:3 (2002) pp.129-140
- Wolfsthal, Jon B. and Tom Z. Collina 'Nuclear Terrorism and Warhead Control in Russia' *Survival* 44:2 (2002) pp. 71-82
- Zimmerman, Peter D. 'Proliferation: Bronze Medal Technology Is Enough' *Orbis* (Winter 1994) pp.67-82
- Zimmerman, Peter D. and Jeffrey G. Lewis 'The Bomb in the Backyard' *Foreign Policy* (November/December 2006) pp.32-39

Other documents

- Air Transport Association 'Airlines Outline War Impact: Report Predicts Massive Economic Damage' News Release (ATA, March 11 2003). Online: http://www.airlines.org/news/releases/2003/news_3-11-03.htm
- Bush, President George W. *The National Security Strategy of the United States of America* (Washington DC, 2002) Online: <http://www.whitehouse.gov/nsc/nss.pdf>
- Department of Defense (US) *Militarily Critical Technology List Part II: Weapons of Mass Destruction Technologies* (Washington D.C., 1998). Official US web pages (<http://www.dtic.mil/mctl>) only contain a limited version of the list, but the full 1998 document is available from the Federation of American Scientists: Online: <http://www.fas.org/irp/threat/mctl98-2/index.html>
- International Atomic Energy Agency *Convention on the Physical Protection of Nuclear Materials With amendment* (2005). Online: <http://www.iaea.org/Publications/Documents/Conventions/cppnm.html>
- National Security Council (US) *National Strategy to Combat Weapons of Mass Destruction* (Washington D.C.: NSC, December 2001). Online: <http://www.whitehouse.gov/news/releases/2002/12/WMDStrategy.pdf>
- National Security Council (US) *National Strategy for Combating Terrorism* (Washington D.C.: NSC, September 2006). Online: <http://www.whitehouse.gov/nsc/nsct/2006/nsct2006.pdf>
- Nuclear Transfer and Supplier Policy Division, Department of Energy (US) *Annex 3 of the Handbook for Notification of Experts to Iraq (UNSC Resolution 1051)* (DoE, 1998) prepared by Sandia, Argonne and Los Alamos National Laboratories and the Oak Ridge National Security Program Office. Online from Iraq Watch: <http://www.iraqwatch.org/government/US/DOE/DOE-Annex3.htm>

Office of Management and Budget, *The White House Budget of the United States Government Fiscal Year 2008* (The White House, 2007). Online: <http://www.whitehouse.gov/omb/budget/fy2008/>

Port of Rotterdam *Port Statistics 2007* (Port of Rotterdam, 2008). Online: http://www.portofrotterdam.com/mmfiles/port_statistics_2007_tcm26-48060.pdf

Speeches, testimonies and transcripts

- Aloise, Gene, Government Accountability Office (US) 'Combating Nuclear Smuggling: Challenges Facing U.S. Efforts to Deploy Radiation Detection Equipment in Other Countries and in the United States' testimony before the Permanent Subcommittee on Investigations and Governmental Affairs, US Senate. GAO-06-558T (March 28, 2006). Online: <http://www.gao.gov/new.items/d06558t.pdf>
- Aloise, Gene, Government Accountability Office (US) 'Combating Nuclear Smuggling: DHS's Decision to Procure and Deploy the Next Generation of Radiation Detection Equipment Is Not Supported By Its Cost-Benefit Analysis' Testimony before the Subcommittee on Emerging Threats, Cybersecurity, and Science and Technology; Committee on Homeland Security; House of Representatives. GAO-07-581T (March 14, 2007) Online: <http://www.gao.gov/new.items/d07581t.pdf>
- Annan, Kofi, Secretary-General of the United Nations, 'A Global Strategy for Fighting Terrorism' Keynote Address to the Closing Plenary of the International Summit on Democracy, Terrorism and Security (Madrid, 10 March 2005). Online: http://www.unfoundation.org/files/pdf/2005/A_Global_Strategy_for_Fighting_Terrorism.pdf
BBC report from address: <http://news.bbc.co.uk/2/hi/europe/4336713.stm>
- Aoki, Stephen, Deputy Undersecretary of Energy for Counterterrorism 'Countering Terrorist Nuclear Weapons Threats to the Homeland' Testimony before the Senate Judiciary Committee, Subcommittee on Terrorism, Technology, and Homeland Security (July 27, 2006). Online: <http://nnsa.energy.gov/news/1236.htm>
- Bush, George W., President of the United States of America, Speech to graduates at the United States Military Academy, West Point, New York (June 1, 2002). Online: <http://www.whitehouse.gov/news/releases/2002/06/20020601-3.html>
- Kosiak, Stephen M., Center for Strategic and Budgetary Assessments, 'The Global War on Terror: Costs, Cost Growth and Estimating Funding Requirements' Testimony before the Senate Budget Committee (June 6, 2007). Online:

http://www.csbaonline.org/4Publications/PubLibrary/T.20070206.The_Global_War_on_/T.20070206.The_Global_War_on_.pdf

Lehrman, Thomas D. , Office of Strategic Planning and Outreach, US Department of State (DoS) 'Building a Layered Defense to Combat Weapons of Mass Destruction (WMD) Terrorism' remarks to the NPT Conference, Washington College of Law (February 9, 2006). Online:
<http://usinfo.state.gov/is/Archive/2006/Feb/27-135117.html>

Paine, Christopher, The Natural Resources Defense Council 'Preventing Nuclear Terrorism' Testimony before the National Security, Veteran Affairs, and International Relations Subcommittee of the House Government Reform Committee (September 24, 2002). Online from Iraq Watch:
<http://www.iraqwatch.org/government/US/HearingsPreparedstatements/hgrc-paine-092402.htm>

Press reports and op.ed.s

BBC 'The Destructive Forces Unleashed' (September 18 2001). Online:
<http://news.bbc.co.uk/2/hi/science/nature/1550326.stm>

Binnendijk, Hans and Peter D. Zimmerman 'New nuclear deterrents' op.ed. *The Washington Times* (August 19, 2007). Online:
<http://www.washtimes.com/news/2007/aug/19/new-nuclear-deterrents/>

Calabresi, Massimo and Romesh Ratnesar 'Can We Stop the Next Attack?' *TIME Magazine* (March 11, 2002). Online:
<http://www.time.com/time/magazine/article/0,9171,1001961,00.html>

Follath, E. and G.M. Mascolo 'Interview with Mohamed ElBaradei: Al Qaida also wants the Bomb' *Der Spiegel* (February 8 2005). Online:
<http://www.spiegel.de/international/spiegel/0,1518,343030,00.html>

Goldberg, Jeffrey 'On Nov. 4, Remember 9/11' op.ed. *The New York Times* (September 9, 2008). Online:
<http://www.nytimes.com/2008/09/09/opinion/09goldberg.html>

International Herald Tribune 'Report: North Korea provided technical assistance to Syria to build nuclear reactor' (February 18, 2008). Online:
<http://www.ihf.com/articles/ap/2008/02/18/asia/AS-GEN-NKorea-Syria-Nuclear.php>

Kessler, Glenn 'N. Korea, Syria may be at work on nuclear facility' *Washington Post* (September 13, 2007) p. A12. Online:
<http://www.washingtonpost.com/wp-dyn/content/article/2007/09/12/AR2007091202430.html>

Mir, Hamid 'Osama claims he has nukes: If US uses N-arms it will get same response' *Dawn* (November 10, 2001) Online: <http://www.dawn.com/2001/11/10/top1.htm>

Sanger, David E. and Thom Shanker 'U.S. Debates Deterrence for Nuclear Terrorism' *The New York Times* (May 8 2007). Online:
<http://www.nytimes.com/2007/05/08/washington/08nuke.html?th&emc=th>

Yusufzai, Rahimullah 'Conversation with Terror' *TIME Magazine* (January 11, 1999). Online:
<http://www.time.com/time/magazine/article/0,9171,989958-1,00.html>

Webpages and software

Council of Foreign Relations 'How Likely is a Nuclear Terrorist Attack on the United States' Online Debate between Michael Levi and Graham T. Allison (April 20, 2007). Online:
http://www.cfr.org/publication/13097/how_likely_is_a_nuclear_terrorist_attack_on_the_united_states.html

Google Earth (aerial photography). Online download:
<http://earth.google.com>

International Atomic Energy Agency *Illicit Trafficking Database* Online: <http://www-ns.iaea.org/security/itdb.htm> (Available to members only)

Nuclear Threat Initiative 'NTI Website Resources on Nuclear Trafficking' Part of NTI *Reserch Library*. Online:
http://www.nti.org/e_research/e3_special_nucl Trafficking.html

Royland, Lucas 'Nuclear Weapon Effects Calculator', online from the Federation of American Scientists. Based on Glasstone and Dolan *The Effects of Nuclear Weapons op.cit.* Online:
<http://www.fas.org/main/content.jsp?formAction=297&contentId=367>

Sublette, Carey 'The Nuclear Weapon Archive', online resource about nuclear weapons. Online:
<http://www.nuclearweaponarchive.org>