

**Fake Terrorism:
Examining terrorist groups' resort to hoaxing
as a mode of attack**

by

Nicole Alexandra Tishler

A thesis submitted to the Faculty of Graduate and Postdoctoral
Affairs in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

in

International Affairs

Norman Paterson School of International Affairs
Carleton University
Ottawa, Ontario

© 2017, Nicole Alexandra Tishler

Abstract

Little academic attention has been accorded to terrorism hoaxers—i.e. those perpetrators who use lies, benign materials and/or empty threats to give the impression that a terrorist act is or has been underway. This dissertation harnesses under-utilised terrorism events data to build a theory of hoaxes in pursuit of a dual aim: to provide a robust substantive answer to the empirical puzzle of why hoaxes are used, but not by all groups, and not all the time; and to evaluate the degree to which existing data can demystify the hoax phenomenon. The starting point is a rationalist framework for terrorist groups' strategic logics, which emphasizes the relative costs and benefits of hoaxes in relation to serious terrorism activity. In the empirical theory-building chapters, probit regression and qualitative comparative analysis (QCA) are used to identify various organizational conditions that differentiate hoaxers from non-hoaxers, thereby indicating which strategic logics are plausibly at play, and in which contexts. A statistical cluster analysis demonstrates that there are five broad classes of hoaxing terrorist groups, which differ from one another along motivational, structural, and campaign contextual lines. While the unit of analysis throughout is the terrorist group, these analyses rely on cross-national terrorism events databases—predominantly ITERATE and the Monterey WMD Terrorism Database—to identify which groups never hoax, and which groups sometimes do. In the dissertation's final section, earlier findings are tested against a new sample of terrorism perpetrators derived from the recently-released Canadian Incident Database (CIDB). Although the Canada-centric data reveals a biased under-reporting of hoax activity in the cross-national datasets, a QCA analysis of its perpetrators reveals roughly similar conditions differentiating hoaxers from non-hoaxers. The CIDB's comprehensive events coverage is further exploited to test whether these organizational indicators and their associated hypothesized mechanisms hold, when campaign activities are evaluated at the event-level. A fine-grained analysis of event sequencing in Canada's most prolific terrorism campaign (that of the *Front de libération du Québec*) corroborates a range of proposed strategic logics. The observational nature of available data is thus limited in its ability to clarify hoaxers' strategic logics, which are both over-determined and equifinal.

Acknowledgements

My most heartfelt thanks go to my supervisor, Jeremy Littlewood, whose teaching and mentorship style I aspire to emulate in my own career. Over the past five years, you have given me free rein to pursue my intellectual whims, but your extensive knowledge and judicious advice have always kept me from straying too far. I am profoundly grateful for your support and guidance.

To my advisor, Dane Rowlands: thank you for encouraging me to expand my methodological horizons, both at NPSIA and beyond. At risk of sounding platitudinous: this dissertation simply would not be what it is, were it not for your instrumental role in my education.

And to my advisor, Stephanie Carvin: I will be forever grateful for your confidence in—and enthusiasm for—my research. You have always urged me to aim high, and I am a stronger scholar for it.

I am indebted to each of you.

My sincere thanks also go to Josh Kilberg: my role model for success in the NPSIA PhD program, and an invaluable source of data and advice in this project's early days.

Additionally, I am grateful for the generous funding support that this research has received from the Social Sciences and Humanities Research Council (SSHRC), the Province of Ontario (Ontario Graduate Scholarship), and the Canadian Network for Research on Terrorism, Security and Society (TSAS).

Finally, I would like to thank my friends and family for their endless love and unwavering support. Mom and Dad: thank you for instilling in me a passion for learning and the tenacity to see this degree through to its end. This is for you.

Table of Contents

Abstract	ii
Acknowledgements	iii
Table of Contents	iv
List of Tables	xii
List of Figures	xiv
List of Appendices	xvi
1. Introduction.....	1
1.1 Why fake terrorism deserves serious attention: Motivation and significance	3
1.1.1 The costs that hoaxes impose on society	5
1.1.2 Hoaxes' implications for theories of terrorism.....	9
1.1.3 Hoaxes' implications for the evaluation of serious terrorist activity	10
1.2 Available data: How we know what we know about hoaxes.....	11
1.2.1 Challenges with hoax data	14
1.3 Scope and methodological overview	19
1.3.1 Unit of analysis.....	20
1.3.2 Dissertation outline	21
PART I: DEFINING HOAXES, CONCEPTUALLY AND EMPIRICALLY	27
2. Defining hoaxes: Key concepts	27
2.1 Categories of hoaxes	27

2.2	What defines a terrorism hoax?	29
2.3	Hoaxes versus related activities	32
2.4	Hoaxes’ relationship with other terrorism tactics	36
2.4.1	Research Note: Hoaxes as a “mode”—not a “tactic”—of terrorism.....	39
2.5	Conclusions on the concept of hoaxing	40
3.	The scope and nature of terrorist hoax activities worldwide	42
3.1	When do hoaxes occur?	43
3.2	Where do hoaxes occur?	45
3.3	What kinds of hoaxes occur?	51
3.3.1	Hoaxes as a mode of terrorism, revisited: Implications for coding.....	54
3.4	Who hoaxes?	57
3.4.1	What else do hoaxers do?	59
3.5	Conclusions on the scope and nature of terrorist hoax activities worldwide	62
4.	Insights into hoax behaviour: Previous empirical research	63
4.1	How terrorists decide among tactics	64
4.1.1	Rational choice.....	65
4.1.1.1	The formal rational choice model and its applications to hoaxing.....	65
4.1.1.2	The strategic logic of terrorist (hoax) behaviour	67
4.1.1.3	The unitary actor assumption and principal-agent relationships.....	71
4.1.2	Ideology and group motivation	73
4.2	The nature of violence	74
4.2.1	Lethality: Hoaxing’s inverse	74
4.2.1.1	Lethality and ideology.....	74
4.2.1.2	Lethality and group structure	76

4.2.1.3	Lethality and group origins	77
4.2.2	Suicide terrorism: A conceptual analog to hoaxing as a mode of terror	77
4.2.3	CBRN terrorism: Parallel data sources and frequently-hoaxed weaponry	78
4.3	Other data sources, methods, and contexts for examining hoaxes.....	81
4.4	Conclusions on existing literature.....	82
5.	Conceptual framework: Proto-theories for hoaxing	84
5.1	Hoaxes as a low-cost alternative to serious activity.....	85
5.1.1	Hoaxes as a low-cost means for exploiting intergroup dynamics	88
5.2	Hoaxes as a complement to serious activity.....	89
5.2.1	Hoaxes as a component of mixed strategies	89
5.2.2	Hoaxes to directly facilitate serious activities	91
5.2.3	Hoaxes to drown out serious activities.....	92
5.3	Hoaxes as a less-severe alternative to serious activity.....	93
5.3.1	Hoaxes as an outlet for incomplete radicalization to violence	93
5.3.2	Hoaxes to signal restraint; or to intentionally disrupt, not destruct.....	96
5.4	Explanations for seemingly irrational hoaxes	97
5.4.1	Improper attribution.....	97
5.4.2	Mistaken failed serious attacks	98
5.4.3	Irrational actors.....	99
5.5	Conclusions on the strategic logics for hoaxing	99
	PART II: INDUCTIVE THEORY-BUILDING	101
6.	Characteristics of likely hoaxers: Insights from regression	101
6.1	Hypotheses.....	101

6.1.1	Hypotheses relating to group motivation.....	102
6.1.2	Hypotheses relating to group structure	106
6.1.3	Hypotheses relating to campaign context.....	108
6.2	The data	110
6.2.1	Unit of analysis and sample: Serious perpetrator groups	110
6.2.2	The dependent variable: Hoax = yes/no?.....	113
6.2.3	Independent variables	115
6.2.3.1	Measures of group motivation	117
6.2.3.2	Measures of group structure	118
6.2.3.3	Measures of campaign context.....	120
6.3	Models and methods	121
6.4	Results and discussion.....	123
6.4.1	Animal rights, environmentalist, and anti-abortionist motivation.....	125
6.4.2	(Non-)hierarchical structure	127
6.4.3	Ethno-nationalist motivation.....	129
6.4.4	Religious motivation	129
6.4.5	Radical splinter origins.....	131
6.4.6	Number of serious events perpetrated	132
6.5	Conclusion: The strengths and limitations of a regression-derived answer.....	134
7.	Characteristics of likely hoaxers: Insights from set theoretic relations	137
7.1	The case for using QCA.....	138
7.2	Hypotheses revisited.....	141
7.2.1	Hypotheses relating to group motivation.....	143
7.2.2	Hypotheses relating to group structure	144

7.2.3	Hypotheses relating to campaign context.....	144
7.3	Required modifications to the data	146
7.4	Models and methods	149
7.5	Results and discussion.....	155
7.5.1	Animal rights, environmentalist, and anti-abortionist motivation.....	157
7.5.2	(Non-)hierarchical structure	160
7.5.3	Ethno-nationalist motivation	162
7.5.4	Religious motivation	164
7.5.5	Radical splinter origins.....	165
7.5.6	Number of serious events perpetrated	167
7.5.7	Overview of QCA results	168
7.6	Putting it all together: Probit, QCA, and the strategic logic of hoax perpetrators..	171
7.6.1	Revisiting hypotheses on group motivation	171
7.6.2	Revisiting hypotheses on group structure	176
7.6.3	Revisiting hypotheses on campaign context	178
7.7	Conclusion: On research methods and the characteristics of likely hoax perpetrators	180
8.	A typology of terrorist hoaxers	183
8.1	Typologies in terrorism studies.....	185
8.1.1	Existing typologies of terrorism perpetrators	187
8.1.1.1	Typologies of why terrorists act.....	188
8.1.1.2	Typologies of how terrorists operate: methods, means, and organizational structure	189
8.1.1.3	Typologies based on groups' origins.....	191

8.1.2	Cluster analysis in terrorism studies	191
8.2	The data	194
8.2.1	The unit of analysis and sample: All hoaxing groups.....	194
8.2.2	Clustering variables.....	196
8.3	Methodology.....	200
8.3.1	Step 1: Choosing a measure of proximity	201
8.3.2	Step 2: Choosing a clustering algorithm	206
8.3.3	Step 3: Choosing the optimal number of clusters	209
8.3.4	Step 4: Interpreting the clusters	210
8.4	Results.....	212
8.4.1	A statistically-derived typology for hoaxers	212
8.4.1.1	Limited serious hoaxers	215
8.4.1.2	Single-issue hoaxers	216
8.4.1.3	Ethno-nationalist hoaxers	217
8.4.1.4	Religious hoaxers	219
8.4.1.5	Economically leftist hoaxers.....	220
8.4.2	Comparing clusters against the non-hoaxer sample	221
8.5	Conclusions from the group-based cluster analysis	224
PART III: THEORY-TESTING AND CONCLUSIONS		227
9.	Hoaxing in Canada: Testing the theory.....	227
9.1	Terrorism hoaxes in the CIDB	229
9.1.1	Hoaxed versus serious terrorism events within Canada	230
9.1.2	The sample: Perpetrator groups responsible for Canadian events.....	233
9.1.2.1	Comparisons with previous samples of terrorism hoax perpetrators.....	234

9.2	Expected findings.....	238
9.3	Methodology.....	239
9.4	Results and discussion: Comparing Canadian hoaxers and non-hoaxers.....	242
9.4.1	Findings relating to campaign context	243
9.4.2	Findings relating to group motivation	245
9.4.3	Findings relating to organizational structure	247
9.5	Case analysis: The strategic use of hoaxes by the FLQ.....	248
9.6	Evaluating the theory of terrorism hoaxes in light of Canada’s hoax history.....	258
10.	Conclusions and steps for future research.....	263
10.1	The limits of hoax data, theory, and empirics.....	264
10.1.1	Hoax data and the study of perpetrator groups.....	264
10.1.2	The (mis-)alignment of theory and empirics	272
10.2	Implications for policy and research	275
10.2.1	Implications for counterterrorism and national security policy.....	275
10.2.2	Implications for the study of hoaxes (and terrorism-related phenomena)	277
10.2.2.1	Large-sample research.....	277
10.2.2.2	Small- and medium-sample research	279
10.3	Conceptual extensions	281
Appendices.....		286
Appendix A	Supplementary materials for Chapter 6	286
Appendix B	Supplementary materials for Chapter 8	292
Appendix C	Supplementary materials for Chapter 9	302
Appendix D	List of acronyms.....	305

References 307

List of Tables

Table 1.1 The frequency of all incidents and hoaxes in ITERATE, WMDDDB, and CIDB.....	18
Table 2.1 Hoax subcategories	28
Table 2.2 Hoax-like activities and the features that distinguish them from hoaxes	32
Table 3.1 Regional rankings for hoax and overall terrorism prevalence, 1968-2012	48
Table 6.1 Cleaning the sample	111
Table 6.2 Description of group-based variables for regression analysis	116
Table 6.3 Average marginal effects for probit models	124
Table 7.1 Guidelines for crisp and fuzzy set scoring of organizational conditions	148
Table 7.2 QCA models and their minimized solutions.....	156
Table 8.1 Description of variables for clustering (n=40).....	199
Table 8.2 Counts of binary outcomes for two objects.....	203
Table 8.3 Statistically-derived typology of terrorism hoaxers.....	213
Table 8.4 Statistically-derived typology of non-hoaxing terrorists	222
Table 9.1 Cleaning the sample of perpetrators operating within Canada	234
Table 9.2 Comparison of perpetrator samples in the present and previous analyses (% = 1)	236
Table 9.3 QCA models and their minimized solutions (CIDB data)	243
Appendix Table A.1 Organizational characteristics of terrorist groups in the sample...	286
Appendix Table A.2 Average marginal effects for probit models testing alternate indicators of organizational structure	290

Appendix Table A.3 Cross-tabulation of single-issue motivation as a sufficient condition for hoaxing	291
Appendix Table A.4 Cross-tabulation of single-issue motivation as a subset of non- hierarchy	291
Appendix Table B.1 List of groups in hoaxer sample and associated ID labels	292
Appendix Table B.2 Basic characteristics of serious terrorist groups who hoax	293
Appendix Table B.3 List of groups in non-hoaxer sample and associated ID labels.....	295
Appendix Table B.4 Basic characteristics of serious terrorist groups who never hoax..	296
Appendix Table B.5 Core cluster results for hoaxing groups, by proportion of terrorist groups within each cluster displaying given characteristics ($\geq 60\%$ highlighted) ...	298
Appendix Table B.6 Sub-cluster results for hoaxing groups, by proportion of terrorist groups within each cluster displaying given characteristics ($\geq 60\%$ highlighted) ...	299
Appendix Table B.7 Core cluster results for non-hoaxing groups, by proportion of terrorist groups within each cluster displaying given characteristics ($\geq 60\%$ highlighted)	301
Appendix Table C.1 List of groups in sample of Canadian perpetrators and associated ID labels	302
Appendix Table C.2 Organizational characteristics of terrorist groups in the sample...	303

List of Figures

Figure 1.1 Hoax frequency in cross-national datasets, 1968-2012	5
Figure 1.2 Frequency of event type for CBRN terrorism incidents, 1968-2012	7
Figure 3.1 Hoax frequency in cross-national datasets, 1968-2012	43
Figure 3.2 Number of CBRN terrorism hoaxes by region, 1968-2012	45
Figure 3.3 Number of transnational terrorism hoaxes by region, 1968-2012	46
Figure 3.4 Prevalence of transnational terrorism hoaxes per region, 1968-2012	49
Figure 3.5 Prevalence of CBRN terrorism hoaxes per region, 1968-2012	50
Figure 3.6 Type of weapon used in CBRN terrorism incidents, 1968-2012	52
Figure 3.7 Delivery mechanism of CBRN terrorism incidents, 1968-2012	53
Figure 3.8 Type and number of ITERATE incidents coded with weapon type = “Hoax, no weapons used,” 1968-2012	56
Figure 3.9 Frequency of motives according to CBRN event type, 1968-2012	58
Figure 3.10 Types of transnational events perpetrated by groups engaging in both hoaxes and serious terrorism, 1968-2012	60
Figure 7.1 Two-by-two illustration of necessary and sufficient conditions	140
Figure 8.1 Visual comparison of linkage methods	208
Figure 8.2 Dendrogram for cluster analysis based on Jaccard similarity measure	214
Figure 9.1 Number of hoaxed and serious terrorism events occurring per year in Canada, 1960-2015	231

Figure 9.2 Percent likelihood of weapons use being hoaxed or serious for terrorism events in Canada, 1960-2015	233
Figure 9.3 Distribution of hoax perpetrators per typology class	238
Appendix Figure B.1 Dendrogram for cluster analysis based on Jaccard similarity measure, sample of non-hoaxers	300

List of Appendices

Appendix A	Supplementary materials for Chapter 6	286
Appendix B	Supplementary materials for Chapter 8	292
Appendix C	Supplementary materials for Chapter 9	302
Appendix D	List of acronyms.....	305

1. Introduction

It is morning rush hour at a busy train station in an Algerian city on the Mediterranean Sea. Three bombs have been placed around the station, in various locations. Two of these bombs will explode, implanting steel shrapnel into the bodies of all those unfortunate enough to be in the vicinity. The third bomb, however, is a fake. It is not designed to explode, but to lure Algerian security forces to the scene. Once these responders are in place, focusing their attention on a benign imitation of a weapon, the perpetrators will trigger the explosion of the real bombs via a cellphone detonator—killing innocent passers-by and security forces alike. With only a minimal additional investment of risk and resources, this fake bomb—a *hoax* bomb—will generate perfect conditions for a deadly attack. Or at least this is how the perpetrators imagined it. Whether by accident or design, a real bomb exploded first, injuring four people. When security forces arrived at the scene, they identified and defused the second real bomb before it could explode; the hoax bomb was only ever discovered after-the-fact.¹

In this instance, the terrorists' hoax activities occurred simultaneously to serious ones.² In others, hoax activities occur independently within the context of a broader

¹ This incident was carried out by the *Groupe salafiste pour la prédication et le combat* (GSPC; Salafist Group for Preaching and Fighting) on May 30, 2006. Event details are chronicled in the *Global Terrorism Database* (GTD; event ID #200605300005).

² While a technical definition of “serious” terrorism is provided in Chapter 6, footnote 56, the term is used in this dissertation to describe terrorism events that involve the use of violence, or a credible threat thereof (i.e. the perpetrators are in possession of the required weapons and materials to carry out a threatened activity). There is no requirement that a terrorist event result in casualties or significant property damage to be classified as serious. “Serious” terrorism is contrasted with “non-serious” activity, which includes hoaxes as well as plots and threats where the perpetrator does not appear to take meaningful steps for follow-through.

terrorism campaign. In the post-9/11 world, such incidents have most frequently manifested themselves in the form of “anthrax” letter hoaxes—in which a benign white powder is posted in the mail with a threatening note to give the impression of an anthrax attack—and, more recently, as airplane bomb scares, in which a supposed radical falsely declares (usually via Twitter) that a bomb is aboard a specified flight. Whatever their form, terrorism hoaxes are those incidents that are believed to be acts of serious terrorism, but do not actually involve any direct risk of harm.

While hoaxes do not directly threaten human life, however, they do impose *indirect* costs on society. For instance, hoaxed anthrax attacks result in building evacuations and investigations by hazardous materials (HAZMAT) teams; and hoax bomb scares result in grounded flights and massive airport delays. The costs of policy response can in turn jeopardize human safety, when response resources are diverted away from real threats, or when the response itself generates dangerous conditions. More insidious, however, are the long- and short-term psychosocial implications of hoaxes that stem from heightened levels of terrorism-related fear in society.

It is thus obvious why some actors would choose to carry out terrorism hoaxes: they are a low-cost, low-risk means of producing sometimes drastic outcomes. While many hoaxes have proven to be jokes made in bad taste or by unorganized individuals who lack both the desire and/or means to carry out serious attacks, the anecdote above illustrates this is not always the case. Although they represent only a minority of hoax perpetrators, there are several terrorist groups who conduct hoaxes in addition to a broad range of serious terrorism offenses. The hoax activities of these groups are particularly

puzzling, since fake terrorism acts signal tactical and operational weakness, and risk undermining future threats of violence. Why do these otherwise capable and motivated terrorist groups hoax?

This dissertation harnesses existing terrorism events data sources to build a theory of terrorism hoaxes in pursuit of a dual aim: to provide a robust substantive answer to the empirical puzzle of why hoaxes are used, but not by all groups, and not all of the time; and to evaluate the degree to which existing data is able to shed light on the hoax phenomenon.

1.1 Why fake terrorism deserves serious attention: Motivation and significance

Little scholarly attention has been dedicated to understanding terrorism hoaxes. To date, hoaxes have never been the primary object of analysis in empirical nor theoretical works; each article referencing hoaxes contributes only small fragments of the total picture. This dissertation thus responds to a wide gap in current scholarly knowledge, bringing coherence to what is presently a sporadic, incomplete, and incoherent literature evaluating terrorist hoax activity. Allocating serious attention to hoaxes is a valuable enterprise, and can provide substantial contributions to broader terrorism studies theory, debates, approaches, and empirics.

Simply put, it is important to understand hoaxes because they are an empirical fact (see Figure 1.1; Table 1.1); terrorists are using them, they are generating costs for society, and their incidence is likely to increase as continued counterterrorism efforts lead terrorists to substitute away from more complex violent alternatives. While most hoaxes

go unattributed (and are likely carried out by individuals who lack the desire and/or means to carry out serious attacks), this is not always the case. Al-Qaeda (post-9/11), Basque Fatherland and Freedom (ETA; pre-2002), Boko Haram, the Irish Republican Army (IRA), and the Liberation Tigers of Tamil Eelam (LTTE) are among the most prolific terrorist groups who use hoaxes in addition to a range of serious tactics.³ The interplay between hoax and serious activities within groups' overarching campaigns is poorly understood and reveals a vulnerability in our ability to counter such groups. Furthermore, since successful counterterrorism has been demonstrated to reduce the sophistication and complexity of terrorists' attacks (Enders and Sandler 2012, 255), it is likely that the continuation of—and improvements in—effective counterterrorism will lead terrorists to rely more significantly on low-complexity tactics. In such a context, the logic of substitution suggests that the frequency of hoaxes might increase, making the ability to distinguish between hoaxes and serious attacks even more important. This imperative can only be pursued by according greater attention to past hoax activity.

³ In fact, of the 25 “top” terrorist groups identified as being responsible for 48% of the world’s wounded terrorism victims between 1970 and 2013 (“A World of Terror” 2017), nine (or 36%) are identified as hoaxers in the datasets considered in the present study: al-Qaeda, Aum Shinrikyo, ETA, Hezbollah, the LTTE, the Manuel Rodriguez Patriotic Front (FPMR), New People’s Army, IRA, and the Taliban. Boko Haram is also on this list of top perpetrators, although its hoax activity (described at length in Mahoney 2017) is not captured in this study’s samples.

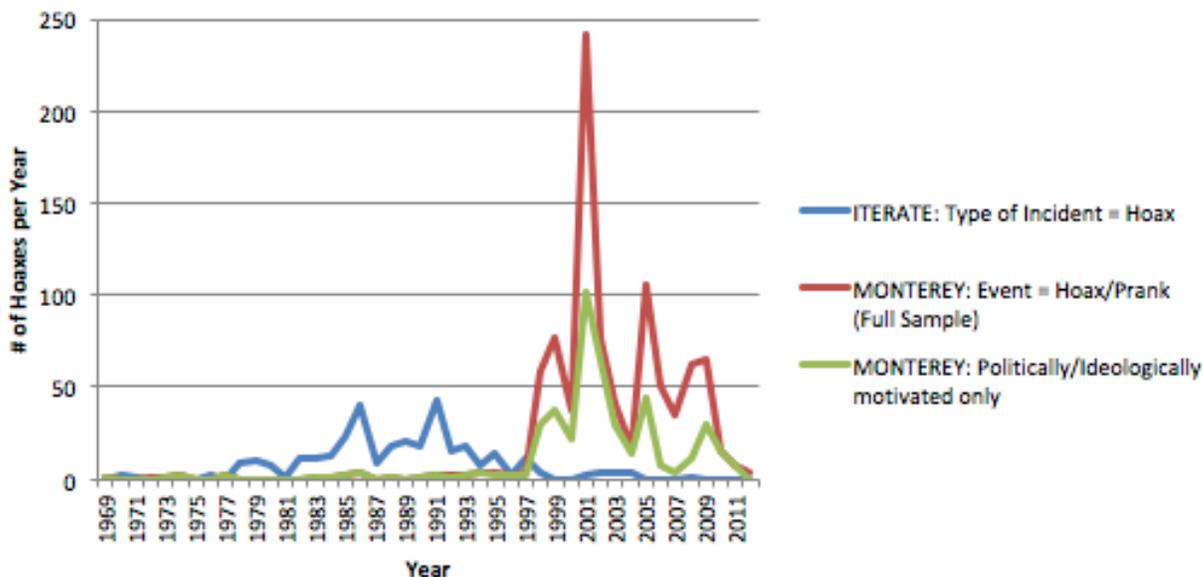


Figure 1.1 Hoax frequency in cross-national datasets, 1968-2012

1.1.1 The costs that hoaxes impose on society

Although terrorism hoaxes do not generate human costs on the same scale as those of a serious terrorism attack, hoaxes affect populations by contributing to an overall climate of terror and disruption in society. At the emotional and psychological level, a credible hoax generates widespread fear, and can cause almost as many socio-psychological outcomes as its serious analog (Dougherty, Green, and Harrington 2001; Giuffrida 1987, 75). Dougherty *et al.* (2001), for instance, examine the psychological impact of two anthrax hoaxes on victims, finding that symptoms of distress and posttraumatic stress were present even after the hoaxes were uncovered as benign. In extreme cases, the *response* to hoaxes can lead to physical human harm. For instance, when a building was

evacuated during a May 1968 hoaxed bomb warning in Ottawa, a young man was killed while attempting to escape a stalled elevator.⁴

The fear generated by hoaxes is particularly high when they purportedly involve chemical, biological, radiological, or nuclear (CBRN) weaponry.⁵ A significant gap has been repeatedly documented between the actual terrorist threat, as identified by experts, and perceptions of the terrorist threat among the lay public (LaFree, Dugan, and Miller 2015, 123; Mueller 2005; Mueller 2006; Mueller and Stewart 2012). This gap is so wide that the terrorist threat need not actually even exist for perceptions of it to generate negative responses (Jenkin 2006; Lee, Lemyre, and Krewski 2010; Slovic 2002).⁶ In presenting themselves as serious terrorism, hoaxes play a role in widening this gap. And whether hoaxed or real, perceptions of terrorism and associated security responses disrupt routine activities.

Because the costs of falsely identifying a serious attack as benign—type II error—may be high, hoaxes must be treated as legitimate terrorist attacks until they are definitively uncovered as deceptions (see Enders and Sandler 2012, 72; Munroe 2009, 299; Ryder 2005, 50–56; 128). They thus place unnecessary strains on emergency responder capabilities (Dishman 2001, 311) and over-extend security service response

⁴ This incident is chronicled in the *Canadian Incident Database* (CIDB event ID #19630523090906001).

⁵ Koblentz (2011) explains the fear induced by CBRN hoaxes and false alarms by the “dreaded risk bias.” “Dreaded risks,” he writes, “are risks that are viewed as being indivisible and involuntary with large-scale, lethal, and long-term effects and indiscriminant, inequitable, and uncontrollable consequences” (509). They thus “provoke a disproportionate psychological response,” and are seen as being more risky and in need of response than other risks that are more statistically probable (Koblentz 2011, 509).

⁶ This gap has been attributed to such behavioural economic phenomena as probability neglect (Sunstein 2003); prospect theory (Kahneman and Tversky 1979); and responding rationally to fear rather than to accurate appraisals of risk (Becker and Rubinstein 2011).

units (Horgan and Morrison 2011, 650; Ryder 2005, 50; 128), taking attention away from serious terrorist activity. The risk of a false negative is particularly dangerous when attacks are believed to involve CBRN weapons. Given that hoaxes are predominantly CBRN in nature (see Figure 1.1) and that hoaxing is the most common CBRN event type (see Figure 1.2), understanding terrorists' hoax behaviour is particularly important. In the United States, responding to anthrax hoaxes alone wasted millions of dollars between October 2001 and March 2002 (Wallace Jr. 2002, 1).

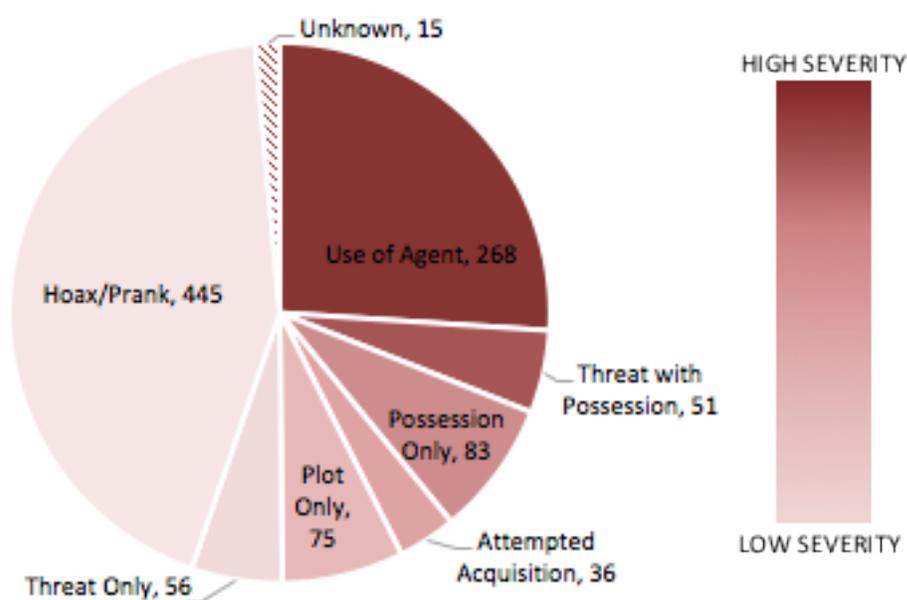


Figure 1.2 Frequency of event type for CBRN terrorism incidents, 1968-2012

But the financial costs of hoaxes are not limited to emergency response. Additional costs are borne by commercial interests when businesses are the victims of attacks—whether or not they are the ultimate target of the campaign. For example, the Animal Liberation Front (ALF)'s product contamination hoax regarding Mars Bars cost the company nine million USD (Monaghan 2000, 258–59; Monaghan 2013, 936); and Direct

Action's anti-Apartheid product contamination hoaxes involving South African wine led to the destruction of 3,000 bottles (Leman-Langlois and Brodeur 2005, 128). Even if companies do not take action to destroy their products, governments may do so, and/or the public may independently choose to boycott the purportedly affected products. In situations where hoaxes could induce tourists to change their travel plans, governments and businesses may be forced to bear additional indirect costs (Enders and Sandler 1991, 53). Indirect costs also include lost productivity from building evacuations during emergency response.

To the degree that governments may be swayed by the threat of serious terrorism, hoaxes may also coerce governments into policy concessions—particularly for credible hoaxes involving nuclear weaponry (Beres 1998, 70; Editor and Post 1990, 168; Leventhal and Chellaney 1988, 449).⁷ Hoaxes' potential coercive power is a function of the high costs of the associated type II error. In extreme cases, a credible hoax may lead to anticipatory or pre-emptive self-defence actions by states (Joyner 1990). Even when hoaxes do not affect policy decisions directly, they may justifiably factor into threat assessments: "Even if the threatened attack turns out to be a hoax, it still means that the terrorist was thinking seriously enough about chemical or biological weapons to develop credible scenarios for their use" (Tucker 2000b, 254; also referenced in Dishman 2001, 311). Increased attention to hoaxed tactics at the public and governmental level, however, can create a "market

⁷ What would constitute a "credible" nuclear threat in these authors' eyes, however, is unclear. Since their writing, terrorists' pursuit of nuclear weaponry has been fortunately ill-fated. Even Aum Shinrikyo, which dedicated millions of dollars in resources to its nuclear program (and achieved a degree of success in its parallel chemical weapons development), never came close to producing a viable nuclear weapon (Cameron 1999; Nehorayoff, Ash, and Smith 2016; Tishler 2017).

pull” effect toward new, and more destructive, technologies (Jackson 2001, 206). Because terrorists learn from earlier hoaxes that a viable attack would evoke a strong public reaction and corresponding source of leverage, technologically capable terrorists may be more inclined to pursue the hoaxed tactics in their real form (Tucker 2000b, 254). As Enders, Parise and Sandler’s (1992) definitively assert: “Clearly, any threat or hoax can impose significant costs on civilian populations and governmental units” (316).

1.1.2 Hoaxes’ implications for theories of terrorism

At the theory level, understanding hoaxes may play a considerable role in fleshing out alternatives—and alternative routes—to violence in radicalization models. McCauley and Moskalenko’s (2014) examination of hoaxer Clayton Waagner, in an attempt to develop a profile of “lone wolf” terrorists and their radicalization to violence, is a first step in this direction, and should be probed further. That anti-abortionist sentiments are 1) the most common motivation for hoaxed CBRN events, and 2) the only perpetrator motivation for which hoaxes are more likely to occur than the use of an actual agent (see Chapter 3, Figure 3.9) suggests that attitudes toward violence and the loss of life may be a fruitful line of inquiry into radicalization models.

Hoaxes are also relevant in the old-versus-new terrorism debate, although the literature has not yet included them in that context. Proponents of “new terrorism” argue that terrorism has shifted toward more excessive and indiscriminate violence in the post-9/11 world. While a substantial increase in the incidence of CBRN terrorism—as a particularly heinous form of violence—would support proponents’ claims, Figure 1.2

illustrates that CBRN hoaxes are nearly one and a half times more frequent than actual CBRN weapons uses (most of which, incidentally, also involve no casualties). Furthermore, that hoaxes are consistently perpetrated even with very low prospects of negotiation success (based on events coded as weapon type = “hoax, no weapons used” in the ITERATE skyjack file) lends support to the quintessentially old-terrorism claim that terrorists simply want “a lot of people watching, not a lot of people dead” (Jenkins 1987, 8). An incident’s propagandistic elements appear to be more significant, in terms of perpetrator motivation, than any immediate benefits—negotiated or otherwise.

A deeper understanding of hoaxes may also help advance economic approaches to the study of terrorism. Since terrorist hoaxes “do not really use up resources” (Enders, Sandler, and Cauley 1990, 101), they pose a challenge to standard models that rely on budget constraints to explain terrorists’ allocation of effort to various terrorist and non-terrorist activities (as described in Anderton and Carter 2009b). With respect to game theoretical approaches, that some hoax perpetrators are also serious attack perpetrators suggests terrorists may be employing mixed strategies—not just with respect to attacking over not attacking, or of choosing among different tactics, but of instigating hoaxed attacks over serious ones.

1.1.3 Hoaxes’ implications for the evaluation of serious terrorist activity

At a practical level, according appropriate attention to hoaxes will help to refine coding in large-*n* terrorism datasets, since putting the spotlight on hoaxes—as section 3.3.1 illustrates—unveils challenges in the coding structure of existing datasets. Addressing

these coding inconsistencies will directly support research on hoaxes, but it will also clarify statistical understandings of other forms of terrorism by providing a mechanism to accurately filter out hoaxed events while retaining the capacity to assess implications of serious *and* hoaxed terrorism for responders and the public.

Lastly, careful observation of hoax activity can also contribute to our knowledge of serious terrorist activity inasmuch as hoaxes may be perpetrated: while terrorist groups are accumulating resources and planning to conduct serious terrorist acts; to gain tactical intelligence for—and to practice dry-runs of—planned serious attacks; as a function of an incomplete radicalization to violence or of a group’s (ideological) aversion to more severe degrees of violence; or, as a an element of a group’s mixed-strategy-informed attack profile. Whether the study of hoaxes is taken seriously as an end in itself or as the means to a more comprehensive understanding of serious terrorism, it is time to accord terrorism hoaxes the attention they warrant.

1.2 Available data: How we know what we know about hoaxes

Which types of incidents are included in events datasets and how these data are coded have implications for how we understand the phenomena at hand. In their discussion of the “constitutive relationship between terrorism and statistics,” Zulaika and Douglass (2008, 30) describe how, when “threats and hoaxes” were included in the Central Intelligence Agency (CIA)’s annual report on terrorism, the number of terrorist incidents documented between 1968 and publication increased from 3,336 incidents in the 1979 report, to 6,714 in the 1980 report. While many plausible explanations may exist for the

CIA's change in coding, the redesign may have been implemented "to suit the foreign policy objectives of the Reagan Administration" (Ross 1988, 232) and inflate public support for counterterrorism activities. Knowing which terrorism data sources include or exclude hoaxes from their samples, and why, can thus further both our understanding of hoaxes and of serious terrorism.

Not all publicly-available terrorism data sources provide information on hoaxes, which at least partly accounts for hoaxes' marginalization in empirical literature. For instance, the *Global Terrorism Database* (GTD)—often recognized as the world's most comprehensive terrorism dataset—does not include hoaxes as a possible event or weapon type (National Consortium for the Study of Terrorism and Responses to Terrorism [START] 2013a). It explicitly excludes foiled and failed plots, as well as "threats to act where no action is taken" (START 2013b) due to its requirement that coded attacks "must actually be attempted to qualify for inclusion in the database."⁸ Similarly, the methodology of the *Worldwide Incidents Tracking System* (WITS)—active 2005 through 2012 under the National Counterterrorism Centre (NCTC)—required that terrorists "must have initiated and executed an attack [...] result[ing] in some sort of active, kinetic effect—such as an explosion or inflicted injury or damage" for inclusion in the dataset (NCTC 2012, 7). Hoaxes—along with failed and foiled attacks, spontaneous hate crimes, and

⁸ The most recent update to the GTD codebook clarifies its practices with respect to plots, conspiracies, and unsuccessful attacks: "For an event to be included in the GTD, the attackers must be "out the door," on route to execute the attack. Planning, reconnaissance, and acquiring supplies do not meet this threshold" (START 2017, 10). The dataset does, however, "include attacks that were attempted but ultimately unsuccessful," simply marking failed attacks using an indicator variable for success (START 2017, 10). This clarification of inclusion criteria is reflective of a broader trend in terrorism studies toward recognition that serious terrorism events must be contextualized relative to foiled and failed plots in order to appropriately evaluate the full range of terrorist activity (discussed in section 10.2.2.1).

genocide—are explicitly excluded. Additionally, while some datasets conceptually allow for hoaxes’ inclusion, the empirical record may not: the *Energy Infrastructure Attack Database* (EIAD), which documents non-state attacks on energy infrastructure worldwide, does not include any hoax incidents, despite explicitly coding for them (Giroux, Burgherr, and Melkunaite 2013, 120).

The two cross-national terrorism databases that explicitly *do* code for hoaxes are the *Monterey Weapons of Mass Destruction (WMD) Terrorism Database* (herein WMDDDB; until recently, the world’s most comprehensive open-source data set on CBRN terrorism events)⁹ and the *International Terrorism: Attributes of Terrorist Events* dataset (ITERATE; the most frequently-cited dataset on transnational terrorist activities). In the WMDDDB (Monterey Terrorism Research and Education Program 2012), hoaxes are specifically coded for as a “type of event,” applying to those cases where, “although use was threatened, the threat was not credible, or an ‘empty threat’” (Cameron et al. 2000, 158).¹⁰ The WMDDDB codes for events involving CBRN weaponry as early as 1933, but its early coding is particularly sparse: only 10 politically- or ideologically-motivated incidents occurred prior to 1968—none of which were hoaxes. ITERATE (Mickolus et al. 2012) codes for hoaxes as a “type of incident,” or *tactic*, in the common file of all incidents (Mickolus

⁹ The WMDDDB is no longer actively managed. The new *Profiles of Incidents Involving CBRN Use by Non-State Actors* (POICN) database has since been created to address known deficiencies in the WMDDDB’s data, namely with respect to evaluating source reliability (discussed in footnote 157). However, the POICN data is neither publicly available, nor does it code for hoaxes. I thus rely on WMDDDB data here, with awareness of its limitations, which are discussed in detail in subsequent chapters.

¹⁰ The Center for Non-proliferation Studies (CNS) at the Monterey (now Middlebury) Institute of International Studies, which created and maintained the WMDDDB, did not publish its own codebook; Cameron et al. (2000) is the most thorough published source of information on the dataset’s methodology.

et al. 2011). Incidentally, it is the only one of twenty-five listed incident types for which an accompanying description is provided: “e.g., claiming a nonexistent bomb.” Hoaxing is considered a distinct tactic from “threat with no subsequent terrorist action.” While ITERATE is not limited to any particular weapon type, it only includes incidents of a transnational nature: “through the nationality or foreign ties of its perpetrators, its location, the nature of its institutional or human victims, or the mechanics of its resolution, [the event’s] ramifications transcend national boundaries”(Mickolus et al. 2011).¹¹

Within a single country context, the *Canadian Incident Database* (CIDB)—on which I worked as a data coder and validator—contributes data on serious and hoaxed incidents of terrorism and violent extremism relating to Canada between 1960 and 2015 (CIDB 2016). Although the bulk of CIDB incidents occurred within Canada, the database also accounts for incidents with Canadian victims, targets, or perpetrators abroad.

1.2.1 Challenges with hoax data

Terrorism events datasets in general suffer from a number of well-documented flaws (LaFree 2010, 24; LaFree, Dugan, and Miller 2015, 22–24; Mickolus 2002, 151–54;

¹¹ While the language of “transnational” and “international” terrorism is frequently used, there are a range of definitions applied thereto, which limits comparability across studies (Schmid et al. 1988, 41). This difficulty is exacerbated by the frequent unreliability (or lack) of information pertaining to the required dimensions of nationality (LaFree, Dugan, and Miller 2015, 147). By accepting the definitions provided within ITERATE’s data codebook, I ensure conceptual alignment with the data I use, although my findings are limited by the quality of data available to ITERATE coders. ITERATE’s creators differentiate transnational from international terrorism by the perpetrator’s relationship with the state: “International terrorism is such action when carried out by individuals or groups controlled by a sovereign state, whereas transnational terrorism is carried out by basically autonomous non-state actors” (Mickolus et al. 2011).

Sheehan 2012), many of which are accentuated when dealing with hoaxes. These problems are particularly acute when datasets are based on open source news reporting—which most are (Ackerman and Pinson 2016, 622; LaFree 2010, 24; LaFree and Dugan 2007, 188; LaFree, Dugan, and Miller 2015, 15).¹²

First, media accounts are not always reliable. Given the likelihood of incomplete or conflicting reports as events unfold (Ackerman and Pinson 2014, 235; LaFree 2010, 24; LaFree, Dugan, and Miller 2015, 22), even reputable news outlets may circulate errors and inaccuracies (LaFree 2010, 24; Silke 2001, 5–6). Since they are designed to look like something that they are not, hoaxes are particularly vulnerable to being described by conflicting reports: in the absence of official information on the true nature of an incident, news outlets will frequently spread “disinformation and rumors” (see Erickson and Barratt 2004, 3). Even where reports are factually accurate, however, they are subject to description bias in terms of how the events they cover are represented (Ackerman and Pinson 2016, 620; LaFree, Dugan, and Miller 2015, 13; Silke 2001, 6); applying the “terrorism” label is a subjective decision. Bias is accentuated in cases where researchers may lack the required “audience context”—including such elements as cultural norms, jokes, and irony—to appropriately interpret the news coverage they encounter (Silke 2001, 6–7).¹³ Any inaccuracies in the source materials or data coders’ interpretation thereof will be directly translated into distortions of the events data themselves.

¹² Conversely, there are also advantages to media-based research in the terrorism context. Since terrorists are explicitly engaged in seeking media attention for their activities, news outlets may be a particularly apt measure of their occurrence (LaFree 2010, 24; LaFree, Dugan, and Miller 2015, 24).

¹³ Importantly, these challenges are not limited to media-based sources. They also apply to interpretations of official government sources, despite their veneer of authoritativeness and objectivity (Ackerman and Pinson 2016, 622; LaFree, Dugan, and Miller 2015, 22; Silke 2001, 7).

Second, news outlets are only able to report on events that they *know* about, are *free to report* about, and *want* to report about. In autocratic regimes with tight censorship rules, terrorist incidents are unlikely to be publicized. Attacks in these regions of the world will be correspondingly underreported in events databases (LaFree and Dugan 2007, 188; LaFree, Dugan, and Miller 2015, 22). Elsewhere, the media has a tendency to over-represent more “newsworthy” forms of terrorism (Ackerman and Pinson 2016, 619; LaFree and Dugan 2007, 188; LaFree, Dugan, and Miller 2015, 22). This means that unsuccessful attacks, as well as those prevented by authorities at the planning stage, are more likely to be absent from news reporting (LaFree and Dugan 2007, 188; LaFree, Dugan, and Miller 2015, 22).

Hoaxes are particularly likely to be underreported in media-based accounts (Dahl 2011, 625; Giroux, Burgherr, and Melkunaite 2013, 120; Horgan and Morrison 2011, 650; Iglarsh 1989, 65; Ross 1988, 219). By virtue of their interest in newsworthiness, media sources may be uninterested in such non-lethal incidents (Horgan and Morrison 2011, 650), unless they have serious economic consequences (Enders and Sandler 2012, 73). Furthermore, the victims of hoaxes—especially businesses and owners of critical infrastructure—have an incentive to keep confidential any threats to their assets (Giroux, Burgherr, and Melkunaite 2013, 120; Pellegrini 2008, 170). More narrowly, ITERATE may have experienced an underreporting of hoaxes since mid-1996: the point at which the Foreign Broadcast Information Service (FBIS) *Daily Reports*—which had included threats and hoaxes not documented in other sources—became unavailable to ITERATE coders (Enders and Sandler 2012, 73).

In addition to the underreporting of hoax events themselves, missing data presents a major challenge (Iglarsh 1989, 65; LaFree and Dugan 2007, 189; LaFree, Dugan, and Miller 2015, 23). With respect to hoaxes, questions of attribution—and *certainty* of attribution—affect how incidents are coded (Mickolus 2002, 153). In any attack, there may be multiple or false claims of responsibility (LaFree 2010, 24; LaFree, Dugan, and Miller 2015, 76). These claims introduce ambiguity surrounding perpetrator attribution. For instance, if a group falsely claims responsibility for actions it did not cause, the event would presumably be reported—and thereby coded—as a real attack perpetrated by that group until evidence comes to light that the group’s claim was a hoax. Even where perpetrators do not actively introduce misinformation to this end, the media is frequently unable to identify the specific person or organization responsible for an attack (LaFree and Dugan 2007, 188; LaFree, Dugan, and Miller 2015, 75). A comparison of Table 1.1’s columns B and D reveals that perpetrators are identified for less than half of the events coded in large-*n* datasets, with an even weaker identification record for hoaxes than the sample of all incidents. Additionally, even where incidents are appropriately attributed to their perpetrators, they may be wrongly identified as hoaxes. This occurs when serious attacks are so badly bungled that responders assume the incident was intended as a hoax. Based on observation alone, it is difficult to distinguish between an intentional hoax, and a serious plot designed by incompetents.

Table 1.1 The frequency of all incidents and hoaxes in ITERATE, WMDDDB, and CIDB¹⁴

	(A) FULL SAMPLE		(B) TERRORISM ONLY		(C) TERRORISM INCIDENTS IN CANADA		(D) TERRORISM INCIDENTS WITH IDENTIFIED PERPETRATOR	
	All incidents	Hoaxes only	All incidents	Hoaxes only	All Incidents	Hoaxes only	All Incidents	Hoaxes Only
ITERATE (1968-2012) ¹⁵	13,514	322	13,514	322	54	8	6,380	56
WMDDDB (1968-2012) ¹⁶	1,716	924	1,029	445	49	41	333	37
CIDB (1968-2012)	1,411	192	801	115	564	103	383	14
CIDB (1960-2015)	1,844	251	1,133	168	871	153	568	27

There are also limitations with respect to the scope of hoax activity captured in existing datasets. While ITERATE and the WMDDDB do code for hoaxes, neither dataset covers domestic incidents using conventional weapons. This limited scope is problematic, since most terrorist activity is domestic (LaFree 2010, 25; LaFree, Dugan, and Miller 2015, 16, 150; Young 2016, 12), and involves “relatively unsophisticated tactics and readily

¹⁴ All values reflect my cleaning of the data, removing duplicate events and re-classifying incidents wrongly coded as hoaxes. Column B accounts only for those events perpetrated with a political or ideological motive. For the WMDDDB, these values exclude criminally-motivated incidents and those with unknown motivation. For the CIDB, these figures exclude incidences of violent extremism for which the terrorism label is in doubt. Column C accounts only for those events that occurred within Canada. For ITERATE, these values reflect the “location end” variable. For the CIDB, these figures exclude events perpetrated abroad with Canadian victims, targets or perpetrators. Column D accounts only for those events with specific perpetrators identified, whether they be individuals or organizations. It thus excludes ambiguous perpetrator identifiers, such as “criminals,” “no group involved,” “unknown,” as well as various categories of “indeterminate” nationals, “agents,” “rebels,” “sympathizers,” and “students” of a given country.

¹⁵ While ITERATE is current to 2016, I have only accessed data until 2012 due to the costs of data access. Per Figure 1.1, it is unlikely that ITERATE has captured many (if any) hoax events in subsequent years.

¹⁶ While the WMDDDB includes events as early as 1933, its sample is reduced to post-1968 here 1) to mirror ITERATE’s coverage, and 2) to filter out the WMDDDB’s sparse coverage of incidents prior to that year.

available weapons” (LaFree, Dugan, and Miller 2015, 201). Accordingly, in focusing only on subsets of terrorist behaviour, ITERATE and the WMDDDB necessarily underrepresent hoax activity worldwide. Although it speaks only to the Canadian context, the CIDB is particularly useful for highlighting the degree of hoax under-reporting in these cross-national datasets. Table 1.1’s column C reveals that, when periods of coverage are held constant, the cross-national datasets capture only 8-10% of all terrorist activity in Canada, and 7-40% of hoax activity. The CIDB’s success in event coverage is partially due to its inclusion of domestic, non-CBRN activity, but also due to its explicit and intentional inclusion of hoaxes (coded as a “yes”/“no” indicator variable) from the beginning of data collection. The CIDB’s success is further attributed to the wide range of data sources that its coders consulted in the data collection process. In accessing reliable sources beyond media reports, the CIDB circumvents some of the challenges with other hoax data described above.

Despite their limitations, these hoax datasets exist, and are the best data presently available for a comprehensive evaluation of hoax activity worldwide. My methodological choices (discussed below) are made in recognition of these deficiencies, to minimize their effect.

1.3 Scope and methodological overview

At the most basic level, this dissertation brings coherence to a presently sporadic, incomplete, and incoherent literature about hoaxes. It also presents the first focused examination of existing hoax data. More narrowly, the research responds to the

literature's failure to present a coherent response to the empirical puzzle of why hoaxes are widely used, but not by all groups, and not all the time. The dissertation includes a conceptual and empirical theory-building component, as well as a test of this theory as it applies in the Canadian context. Throughout, a multi-method approach leverages triangulation of results based on probit regression analysis; qualitative comparative analysis (QCA); statistical cluster analysis; and process tracing.

1.3.1 Unit of analysis

While both individuals and terrorist organizations may perpetrate hoaxes, the primary unit of analysis throughout this dissertation is the terrorist group. At a practical level, I focus on group-based perpetrators because attribution of hoaxes to particular individuals is difficult. Individual perpetrators are rarely identified in terrorism events databases.¹⁷ In terms of counterterrorism practice, group-based perpetrators are also the most significant. While not accounting for the majority of hoax incidents, hoaxes by established groups are likely to be perceived as more credible than those perpetrated by (previously) unknown individuals. Established terrorist groups are also the most intellectually appealing category of hoaxers, since their choice to hoax is not exclusively mandated by a lack of resources or capability in alternative modes of attack. As an intentional strategy,

¹⁷ In the literature examining the predictors of CBRN terrorist activity, individual perpetrators (as well as those operating in small groups and temporary splinter factions) have generally been excluded from the scope of analysis because of data limitations (Asal, Ackerman, and Rethemeyer 2012, 238; Asal and Rethemeyer 2008, 441; Asal and Rethemeyer 2009). Furthermore, when perpetrators are unknown, accurately classifying incidents as "terrorist" becomes difficult or impossible: "all definitions of terrorism rely to some extent on the putative motives of attackers" (LaFree and Dugan 2007, 188).

hoaxing is interesting because it seemingly contradicts the logic of “costly signalling” in terrorist behaviour outlined by Kydd and Walter (2006).

In the dissertation’s final chapter, the unit of analysis briefly shifts to the terrorist event. Examining individual hoaxed and serious events allows me to trace whether the processes implied at the organizational level of analysis are substantiated within the context of a terrorist campaign.

1.3.2 Dissertation outline

The dissertation proceeds in three main parts. First, I define hoaxes and existing hoax-related knowledge in view of deriving proto-theories for hoax behaviour. Second, I exploit existing terrorism events datasets—spanning all countries, 1968-2012—to test specific hypotheses relating to the likely organizational characteristics of hoax perpetrators, in view of elaborating a more robust, empirically-substantiated, theory of hoax behaviour. Third, I test these findings against an alternate dataset chronicling Canada’s experience with terrorism and terrorism hoaxes, 1960-2015.

Defining hoaxes, conceptually and empirically, is the first step in building a theory of hoaxes as a mode of terrorism. In Chapter 2, I present my conceptual definition for terrorism hoaxes. This chapter differentiates hoaxes from related activities, and clarifies various subcategories of hoax behaviour. In Chapter 3, I define hoaxes empirically, per ITERATE and WMDDDB data, elaborating the “who,” “what,” “where,” and “when” of terrorist hoax activities worldwide. This chapter expands upon section 1.2.1’s discussion of challenges with hoax data to highlight further limitations in the way that existing cross-

national datasets have accounted for the hoax phenomenon. While previous chapters reference existing literature in order to appropriately define and describe terrorism hoaxes, Chapter 4 presents a more traditional review of literature. It contextualizes the limited body of scholarly works addressing hoaxes within the broader study of terrorists' tactical choices. To conclude the chapter's first part, Chapter 5 presents my conceptual framework for hoaxers' strategic logic. It proposes a range of rationalist proto-theories (derived from the literature and induced from Chapter 3's descriptive statistics) to explain why and how hoaxes may be strategically employed by terrorist groups and their members. Importantly, it is theorized that different actors will hoax for different reasons, and that even within individual groups' campaigns, multiple strategic logics may be simultaneously at play. Equifinality and complex causality are likely features of any empirically-substantiated explanation for hoax behaviour.

The dissertation's second part constitutes a theory-building exercise, using conclusions drawn from the preceding chapters' discussion of existing data and literature to guide hypothesis formation. It harnesses cross-national events data and a range of methodologies to provide insight into which proto-theories hold, and under what conditions. In the absence of systematic interview data describing hoaxers' strategic logic in carrying out their hoax activities, various organizational indicators are used to provide insight into the proto-theories' validity. These emphasize group motivation (animal rights, environmentalist, or anti-abortionist ideology; ethno-nationalism; and religion), organizational structure (degrees of hierarchy); and perpetrators' campaign contexts

(how prolific their serious terrorism campaign is; and whether they originated as a splinter offshoot of a more moderate group or movement).

In Chapter 6, I apply probit regression to estimate the net effects of these motivational, structural, and campaign contextual indicators on whether or not a terrorist group ever carries out hoax attacks. All else held equal, having a hierarchical organizational structure and/or having originated as a radical splinter of a more moderate group or movement is negatively correlated with a group's likelihood of ever using hoaxes. Conversely, a group's likelihood of employing hoaxes increases when it carries out serious terrorist acts with greater frequency; or when it follows a purely ethno-nationalist or purely religious ideology. The regression identifies specific single-issue ideologies (animal rights, environmentalist, and anti-abortionist motivations) as perfect predictors of hoax use. Statistically robust though these findings are, they are limited by the correlational logic of probit regression—which, as the subsequent chapter reveals, is not necessarily reflective of the empirical reality.

In Chapter 7, I use QCA as an alternative methodological approach to elaborate upon—and in some cases amend—the probit-derived conclusions described above. Rather than seeking net marginal effects of individual variables, QCA searches for *combinations* of factors that are collectively sufficient to produce the outcome, with an explicit recognition of asymmetry in relationships. Overall, QCA suggest that the incidence of hoaxing is best explained by a lack of hierarchy, when combined with a non-religious ideology (either purely ethno-nationalist or broadly left- or right-wing). The single-issue ideologies (animal rights, environmentalist, and anti-abortionist), on their face, play a

consistently sufficient role in determining whether a group will hoax, but given the sample's lack of hierarchically-structured single-issue groups, non-hierarchical organizational structure tells a more compelling story. While these pathways to hoaxing are relatively consistent across models, they account for only a moderate proportion of groups who hoax. Pathways to non-hoaxing are more complex, but explain a far greater proportion of non-hoaxing groups. Hierarchical groups with either purely ethno-nationalist or broadly right- or left-wing orientations are likely non-hoaxers, as are religious groups with either a combined ethno-nationalist motivation or a radical splinter history.

At a difference to Chapters 6 and 7's identification of factors differentiating hoaxers from non-hoaxers, Chapter 8 uses statistical cluster analysis to look for patterns amongst a sample of groups who all use hoaxes to some degree.¹⁸ The cluster analysis inductively generates a multidimensional typology of terrorist groups engaged in hoax activity, based on their ideological orientations; organizational structures; and other

¹⁸ Although cluster analysis is an inductive approach, frequently used to build theory and hypotheses, it follows the regression and QCA analyses (and their defined hypotheses) in the present study, because it is applied to a different end: whereas the regression and QCA differentiate hoaxers from non-hoaxers, the cluster analysis is used to test whether there exist meaningful differences between sub-types of hoaxers. While existing scholarship in related areas of terrorism studies could be applied to guide hypothesis formation regarding the likely characteristics of hoaxers and non-hoaxers, there is no such literature available to guide expectations regarding differences amongst hoaxing groups. The findings from the regression and QCA analyses, however, could be used to guide variable specification for the clustering. Indeed, the greater-than-three-serious-events threshold used in the clustering as a measure of general campaign activity flows directly from the set analyses of Chapter 7. Furthermore, factors that were shown in Chapters 6 and 7 to be rare amongst hoaxers as compared with non-hoaxers would be particularly valuable in a clustering of only hoaxing groups, since they would indicate clusters of outliers relative to the broader sample of hoaxers. While the primary value added of the cluster analysis in the present context is that it serves as a baseline for evaluating differences in the constitution of groups across different samples (as illustrated in section 9.1.2.1 and Figure 9.3), I anticipate that it may be applied to iteratively guide hypothesis formation in future analyses of terrorist groups' hoax behaviour.

features of their campaign context. Five classes of hoax perpetrators emerge, predominantly defined by ideological orientations' alignment with particular structural and contextual elements: single-issue groups who are non-hierarchical, CBRN-capable, and domestically-oriented; ethno-nationalist groups who are bureaucratically structured and transnationally-oriented; religious groups who are CBRN-capable and did not originate as radical splinters; economically leftist groups who are transnationally oriented; and limited serious groups—not defined by ideology—perpetrating three or fewer serious events, with a domestic orientation, non-hierarchical group structure, and interest only in conventional weapons. This statistically-derived classification of hoaxers differs from an analogous clustering of non-hoaxing groups, who are more effectively distinguished per their organizational structure, with little regard for ideological orientation.

The dissertation's third and final part evaluates Chapter 5's proto-theories and the conclusions drawn from Part II's theory-building empirics in light of more granular data and nuanced approaches. In Chapter 9, I use a new dataset—the Canada-centric CIDB—to evaluate my previously-established findings with respect to the factors affecting terrorists' hoax behaviour. A QCA analysis of CIDB data generally corroborates findings from the cross-national sample used previously. The only substantial difference in results gleaned from the test and theory-building samples pertains to the effect of a group's splinter origins on hoax behaviour. Here, splinter origins can contribute to both the hoax and non-hoax outcome, magnifying the expected effect of the corresponding motivational (religion or ethno-nationalism) or structural (hierarchy or non-hierarchy)

feature—in either direction. A second component of Chapter 9 shifts the unit of analysis from the perpetrator group to the individual terrorist event, within a single group’s campaign context: that of Canada’s most prolific terrorist group, the *Front de libération du Québec* (FLQ). Exploring the FLQ’s event sequencing and associated tactical choices allows me to trace processes that were obscured in the previous cross-sectional analyses, which emphasized the effects of discrete organizational traits. While many strategic logics are consistent with observed activity at the event-level, this analysis highlights the substantial degree of equifinality and multifinality inherent in groups’ use of terrorism hoaxes; even a single group with relatively constant organizational features appears to hoax for different reasons at different times, and the same hoax incidents may be plausibly explained by multiple logics simultaneously.

Chapter 10 concludes with a discussion of the limited extent to which existing, observational data can confirm the strategic logic(s) underpinning terrorist groups’ hoax activity. It also describes how this dissertation’s findings—both substantively, and methodologically—can serve as a launch-pad for future research into hoaxes as a mode of terrorism.

PART I: DEFINING HOAXES, CONCEPTUALLY AND EMPIRICALLY

2. Defining hoaxes: Key concepts

At the most superficial level, terrorism hoaxes are those incidents that are believed to be acts of serious terrorism, but do not actually involve any direct risk of harm. Definitions of what exactly hoaxes comprise, however, are rare in the hoax-related literature, which is itself sparse and fragmented. Where definitions are presented, they are often inconsistent and frequently conflated with related terms. This chapter evaluates the ways in which hoaxes have variously been addressed in the literature, and clarifies my definition of the term.

2.1 Categories of hoaxes

References to terrorism hoaxes generally speak to one of four hoax subcategories. Table 2.1 presents these subcategories in increasing order of the hoaxers' investment of resources and tactical effort, from lowest to highest.

Table 2.1 Hoax subcategories

Hoax Claims of Responsibility	Perpetrators opportunistically claim responsibility for incidents they did not cause. These incidents may be A) attacks truthfully carried out by other groups; ¹⁹ or B) the product of accidents/natural occurrences. ²⁰
Hoax Warnings	Perpetrators A) threaten a future terrorist act, with no intention of following through; ²¹ or B) lie about past actions, falsely indicating that they have already carried out some sort of terrorist activity. ²²
Hoax Devices	Perpetrators use benign materials to give the impression that a terrorist act is, or has been, underway. ²³

Despite clear conceptual distinctions, however, scholars are rarely explicit about which category they are speaking about, and do not tend to contextualize their evaluation of various subcategories (or components of subcategories) into a broader framework of hoax activity. For instance, Kearns, Conlon, and Young (2014) examine “false claims of

¹⁹ Kearns, Conlon, and Young (2014, 430–31) discuss how multiple groups claimed responsibility for an April 1999 bombing campaign against various London targets. Although the real perpetrator—David Copeland—acted alone, four groups claimed responsibility for the act, with varying degrees of credibility. Combat 18 was the most well-known and powerful of the four groups, but because police had infiltrated the group prior to Copeland’s attack, its credit claims were easily diagnosed as hoaxes. The White Wolves actively tried to make their false credit claim appear more credible, by simultaneously sending death threats to a range of London-based minority group leaders. Little is known about the other two groups.

²⁰ The Rajah Solaiman Islamic Movement (RSIM) claimed responsibility for an October 19, 2007, explosion in Makati City, the Philippines, but the police investigation later determined the blast to have resulted from an industrial accident (see Banlaoi 2009). Similarly, “rejectionist Palestinian groups often claimed credit for virtually anything that went wrong in Israel—such as explosions from faulty gas lines” (Mickolus 2002, 153).

²¹ Mahoney (2017) outlines a selection of hoax warnings perpetrated by the Nigeria-based Islamist insurgent organization, Boko Haram. While the group did carry out seventy-eight serious attacks against Nigerian educational institutions since 2009, it issued eleven threats of future action on which it never followed through. These hoax warnings still produced tangible outcomes: in ten of the eleven cases, schools were closed and/or law enforcement augmented its security presence.

²² On March 1st of this year, several Canadian news outlets received emailed bomb threats stating that “artisanal amateur explosive devices” had been placed in various Concordia University locations, and that the explosives would be detonated if the university did not prevent religious activities (particularly those of Muslim students) on campus. In response, classes were cancelled; thousands of students, faculty, and staff were evacuated; the university increased its private security patrols; and Montreal law enforcement and canine units swept the area for bombs. Ultimately, however, this was a hoax warning involving lies about past activity: no bombs had actually been placed, nor were there any meaningful indications of a serious plan to follow through on the threat.

²³ White powder anthrax hoaxes are hoax devices. While the white powder (usually talc, or ground chalk) will not cause anybody direct harm, the benign powder itself is tangible and can be discovered by a victim even without perpetrator warning.

responsibility”—but only those pertaining to actual terrorist attacks committed by other groups; they do not consider that groups might claim responsibility for incidents that were not maliciously-intended or that resulted from natural occurrences. Their examination of false claims is contextualized within a broader framework of “lies” that terrorists tell, which also includes false blaming (i.e. where a group falsely attributes an attack to a rival group; “false flag” terrorism) and lies of omission (i.e. where a group does not claim actions that it *did* cause). Mahoney’s (2017) treatment of terrorist “bluffing” encompasses terrorists’ “empty threats,” but does not account for hoax warnings where groups have lied about their past activity. His analysis of empty threats is contextualized in relation to serious incidents, where terrorists follow through on their threatened action. Horgan and Morrison (2011) treat “hoax devices” as a category of Violent Dissident Republican (VDR) terrorism in Northern Ireland, although their empirical figures conflate this category with “warnings” (650). Whether their use of “warnings” refers to hoaxed warnings or truthful warnings regarding impending serious activity, however, is unclear. Conceptually, their discussion of hoax devices and (hoax) warnings is framed as an example of VDR groups’ desire to “cause routine disturbance” in the population (650).

2.2 What defines a terrorism hoax?

Not all hoaxes are “terrorism” hoaxes, in the same way that not all violent criminal activity is terrorist activity.²⁴ There is no consensus regarding the definition of terrorism; Schmid

²⁴ For instance, Laura Pellegrini’s (2008) doctoral research on *criminal* hoaxes explicitly excludes hoaxes for terrorist ends from her analysis, for their significant differences from other criminal hoaxes in terms of

(2011) chronicles over 250 such definitions used by academics, governments, and intergovernmental organizations. In Canada, “terrorist activity” is formally defined within the *Criminal Code* (R.S.C. [1985], c. C-46, s. 83.01) as the use, attempted use, or threat of any act:

- Motivated by politics, religion, or ideology and intending to intimidate the public or compel individuals, groups, or governments;
- Intentionally causing death or serious bodily harm, endangering a human life, causing health or safety risks to (a portion of) the public, causing substantial property damage, or interfering with the provision of an essential service; and
- Committed outside of armed conflict.

In this dissertation, I accept this Canadian legal definition of terrorism (consistent with the CIDB), while recognizing that different terrorism events datasets may involve varying criteria for inclusion. While the language used in the WMDDDB (see Cameron et al. 2000) and ITERATE (Mickolus et al. 2011) codebooks vary, these datasets share the same core identifiers for terrorism events: 1) they involve a political or otherwise ideological motivation; 2) they intentionally involve violence or the threat of violence; and 3) the perpetrators must be sub-national actors (although they may receive state support). Thus, for a hoax to be a *terrorism* hoax, it must satisfy the conditions of at least one hoax subcategory, as defined in Table 2.1, and satisfy each of these three conditions for terrorism.

likely “international connections” and “different motivators” (2). She focuses instead on four categories of criminal hoax: monetary, attention-seeking, racial, and hate crime.

It is important to note, however, that not all hoax activity by terrorist actors is appropriately classified as terrorism hoax events. For instance, Murray (1984) describes a Provisional Irish Republican Army (Provisional IRA) campaign where terrorists made hoax telephone calls to law enforcement authorities reporting “crimes” as a means of luring police officers to a place where they could then be discriminately killed. In some cases, the Provisional IRA actually staged crimes to coax police officers into response. However, although these terrorists exploited false claims and staged activities (akin to hoax warnings and devices, respectively) as a means of facilitating an act of serious terrorism, the hoaxed component itself was not intended to intimidate or coerce targets, nor did the hoaxed component necessarily involve violence or the threat of violence. These instrumental activities were hoaxes by terrorist actors, but they were not *terrorism* hoaxes.

Still, some hoaxes used by terrorists as tactical facilitators of serious activities do satisfy the requirements of terrorism hoax events. When the instrumental hoax itself is designed to look like a terrorist attack—independent of whatever serious activities the group may be simultaneously carrying out—it counts as a terrorism hoax, and should be classified per its corresponding kind: hoax claim of responsibility, warning, or device. For instance, hoax warnings and devices might be used as decoys, to distract intelligence services and law enforcement from groups’ simultaneous serious activities. Chapter 5 elaborates this (and other) strategic logic(s) for hoaxing.

2.3 Hoaxes versus related activities

There is frequent ambiguity in the literature surrounding hoaxes' distinction from other forms of "non-serious" terrorist-related activities, such as threats, pranks, false alarms, and foiled and failed attacks. These are equally ill-defined in the literature, but can be differentiated in terms of their degree of intentionality—whether the action resulted from some conscious effort of a perpetrator—and scope of intended harm. These are summarized in Table 2.2, in increasing order of associated risk.

Table 2.2 Hoax-like activities and the features that distinguish them from hoaxes

False alarm	A benign object, innocently present in a locale, is perceived as a threat by the public or responding entities There is no consciously-responsible, ill-intended perpetrator
Prank	Perpetrator intends to cause mischief or alarm, not to intimidate/coerce a target audience Perpetrator does not have political or ideological aims
Threat	Perpetrator intends to follow through with the threatened action
Foiled attack	Perpetrator took steps to carry out a serious action, but was prevented from successfully executing the attack by external intervention (law enforcement; security forces)
Failed attack	Perpetrator took steps to carry out a serious action, but the attack was not successfully executed as a function of the perpetrators' own tactical or operational failings In some cases, failed attacks may be indistinguishable (based on observed characteristics) from hoaxes

Most frequent is a discursive conflation of *hoaxes* and *threats* (e.g. Ross 1988, 219; Veness 2001, 409), reflected methodologically in empirical works that collapse threats and hoaxes into a single series and evaluate them collectively (e.g. Enders and Sandler 2012, 72–73). Even where studies accord threats and hoaxes their own respective series, they do not always explain the distinction with definitions, nor retain the distinction in

their presentation of results (e.g. Enders, Sandler, and Cauley 1990). And even where distinguishing definitions are provided, the two are still sometimes evaluated together (as in Enders, Parise, and Sandler 1992, 309).

Threats are most easily confused with hoax warnings, since neither involves observable, kinetic, outcomes (as would be present with a hoaxed claim of responsibility, or a tangible hoax device). Enders, Parise, and Sandler (1992) distinguish the two temporally, based on when the supposed action occurred/occurs in relation to the warning: “A threat occurs when a *future* action is promised, while a hoax occurs when an action that has not taken place is claimed to have occurred” (309). However, since this promise of future action may be made hoaxedly (as an empty threat) *or* honestly (as a pre-attack warning), a clearer conceptual distinction can be drawn with respect to perpetrators’ intent.²⁵

A threat is best understood as a verbal or written promise of future action, or ultimatum where the perpetrator *intends to carry out* the threatened activity. With hoax warnings, threats of future action are empty; the perpetrator has no intention of following through. Intent is demonstrated in cases where the threatened action actually follows the threat’s issuance.²⁶ In other cases, there will be no follow-up serious activity— but not for lack of will or effort on the terrorists’ part.

²⁵ This understanding is in line with Cameron et al. (2000)’s discussion of WMDDDB data: “*Hoax* means that although use was threatened, the threat was not credible, or an “empty threat” (158).

²⁶ See, for instance Brown’s (2015) discussion of cases where terrorists provide detailed, pre-attack warnings as a means of reducing casualties in their serious activities.

This is where hoaxes are meaningfully distinguished from foiled and failed attacks (discussed in Crenshaw and LaFree 2017; Dahl 2011). Plots may fail due to internal problems, such as faulty operational or tactical practices within the group (see Matesan and Berger 2017), or they may be foiled by external intervention. In either case, such incidents are distinct from hoaxes because—even if they were ultimately unsuccessful—the terrorists initially did invest resources and effort into pursuing serious terrorist activity. The perpetrators' planning details or possession of required materials are evidence of a serious plot. For hoax warnings, there will be no such observable indicators. When a hoax device is present, it is more difficult to distinguish serious intent: did the perpetrators aim to successfully execute a hoax? Or did they attempt to execute a serious attack, and fail?

In their measure of terrorist intent in targeting versus attack outcomes, Jackson and Frelinger (2008, 588) give terrorists “the benefit of the doubt,” erring toward assuming that attacks are completed as planned. In this case, giving terrorists the benefit of the doubt would—in the absence of information indicating serious intent—lead such cases to be classified as hoax devices. While these subjective decisions introduce noise into any dataset, they are a necessary outcome of evaluations of groups' intent: “In the absence of extensive information on a terrorist group's internal deliberations and attack planning, statements about the intent of particular operations will invariably be educated guesses rather than certain statements of fact” (Jackson and Frelinger 2008, 588). Capability is not a requirement for distinguishing hoaxes from threats or even failed

serious events, although a known or likely lack of group capability would diminish a hoax event's credibility.

Some studies attempt to distinguish between hoaxes and *pranks*, but they do so inconsistently.²⁷ While both are intended to provoke an audience response, the terrorist hoaxer's intent is to intimidate or coerce a target audience to achieve political or ideological aims, whereas the prankster only intends to cause mischief or alarm. A prankster is not a "terrorist" actor, since it does not pursue a political or ideological goal. The WMDDDB conflates hoaxes and pranks in a single category, but they can be easily disentangled by filtering for motive: terrorism hoaxes are those with a political or ideological motivation; pranks are coded with criminal or unknown motivations. "Swatting," where false 9-1-1 calls are made to send emergency response or law enforcement to another person's address, is an example of such non-terrorist, prank behaviour. While the resource drain on emergency response is comparable to that caused by a terrorism hoax, the perpetrator's motivation is vastly divergent.

Finally, hoaxes are related to—but distinct from—*false alarms*. False alarms are those incidents where a benign object, innocently present in a locale, is perceived as a threat. In example, when two schoolboys stashed their bags at Belfast City Hall while skipping school for the day, it "prompted a major security alert"; and when a nervous

²⁷ Ivanova and Sandler (2006) define hoaxes and pranks as "a claimed past act that either did not occur or involved a fake substance" (427). While this characterization is fair, it does not acknowledge the differences in intended scope of harm between terrorist hoaxers and pranksters, nor does it acknowledge the possibility of hoax warnings. The authors' distinction between hoaxes and pranks in a follow-up study—that hoaxes involve a *fake* substance, whereas pranks involve *no* substance use (Ivanova and Sandler 2006, 276)—appears inconsistent with conventional understandings of perpetrator intent. Their later distinction is more comparable with the distinction presented here between hoax devices and hoax warnings.

bride left behind her wedding dress on a downtown Belfast city center bus, the suspicious package “led to gridlock for over an hour while an [Ammunition Technical Officer] checked it out” (Ryder 2005, 50). False alarms are thus reflective of “the ‘nerves’ of the population” (ibid.), whereas hoaxes are “deliberately planted by the dissident groups to disrupt the normalisation of [...] society” (Horgan and Morrison 2011, 650). Like with failed serious attacks, however, it is sometimes difficult to distinguish hoaxes from false alarms, particularly when hoaxers do not claim responsibility for their actions (see Joyner 1990, 80).

2.4 Hoaxes’ relationship with other terrorism tactics

While there is ambiguity surrounding hoaxes’ definition, the sources that acknowledge hoaxed acts consistently characterize them as a specific tactic or method of terrorism (see Aguirre 2009, 47; Enders, Sandler, and Cauley 1990; Horgan and Morrison 2011; Richards 2014, 222). As a tactic, however, hoaxes have an ambiguous relationship with other methods of terrorism—especially from the perspective of emergency responders: a terrorist bomb hoax *is* a hoax, but until it is uncovered as such, it must be treated as a *real* terrorist bombing (Enders and Sandler 2012, 72; Munroe 2009, 299; Ryder 2005, 50–56; 128). Almost any type of serious terrorist activity can also be hoaxed, from biological attacks involving “white powder” letters, to product contamination incidents, to bombs and explosions, and even assassinations and hostage scenarios.

The most commonly hoaxed tactic discussed in the literature is a biological attack whereby a substance, usually said to be anthrax, is delivered by mail to the targeted victim

(see Donohue and Kayyem 2002, 5; Erickson and Barratt 2004; McCauley and Moskalenko 2014, 75). Product contamination hoaxes, supposedly involving various forms of chemicals and poisons, are also present in the literature.²⁸ Nuclear hoaxes are also discussed in the literature, although to a far lesser degree, and usually as a hypothetical case.²⁹ The potential high costs of viable CBRN devices make these types of hoaxes particularly likely to inflict broad-based fear, and to garner serious attention from responders and policymakers.

Hoax bombs and explosions are the second-most discussed hoaxed tactic in the literature. Bomb hoaxes may comprise hoax devices,³⁰ hoax warnings,³¹ or hoax claims of

²⁸ As examples of product contamination hoaxes in the literature, Monaghan (2000, 258–59; 2013, 936) describes a 1984 Animal Liberation Front hoax involving Mars Bars, as well as other contamination threats like mercury in turkeys and toothpaste, spiked baby oil, and bleach in shampoo; and Leman-Langlois and Brodeur (2005, 128) discuss hoaxes in Canada intended to protest the South African Apartheid government, by claiming to have poisoned South African fruit; and hoaxes perpetrated by Direct Action, who claimed to have poisoned bottles of wine.

²⁹ As examples of nuclear hoaxes in the literature, Jenkins (1983) writes that, since 1970, there had been approximately sixty nuclear threats against American cities, most of which were extortion threats (i.e. hoax warnings); writing only three years later, Beckman (1986, 360) claims that over one hundred nuclear hoaxes had already taken place.

³⁰ As examples of bomb hoax devices in the literature, Monaghan (2000, 260) discusses the Animal Liberation Front's use of hoax letter bombs against individuals suspected of being informers; and Nesser (2012, 66, 72) discusses the hoax bomb device placed by Dutch-British "wannabe" jihadist Nicholas Roddis on a United Kingdom bus in 2007.

³¹ As examples of bomb hoax warnings in the literature, Enders and Sandler (1991) discuss the ETA campaign against foreign tourists in Spain, beginning April 1985; and Breemer (1983, 455) and Kashubsky (2011, 141) discuss the anonymous hoax phone calls received by the Philips Petroleum Company in Yarmouth, England, that delayed-action fuses had been attached to underwater charges of offshore production platforms. The platforms were evacuated, and nearly \$500,000 USD in costs were incurred prior to the incident's identification as a hoax two days later.

responsibility.³² Arson hoaxes are also addressed within the literature as a possibility, but with no confirmed cases.³³

Lentini and Bakashmar (2007) present examples of two less frequently—and more imaginatively—hoaxed tactics: an assassination and a hostage scenario. The assassination hoax involved supposed video evidence from the Badr Brigade showing members disguised as the Iraqi National Guard beheading an Iraqi journalist. The incident was uncovered as a hoax when it was discovered that the video actually depicted a Nepalese citizen’s execution from two years prior (320). The hostage hoax occurred when the Mujahideen Brigades posted a photo online of an African-American soldier, “John Adams,” claiming that they would behead him. The incident was uncovered as a hoax when it was determined that the soldier in the photo was really an action figure (322).

In all of these cases, the hoaxes in question corresponded to a specific tactic of serious terrorism. The immediate responses to the hoax involved the same reactions that actors would have in response to a real incident. For instance, a 1984 Mars Bars product contamination hoax cost the company an estimated nine million USD (Monaghan 2000, 258–59; 2013, 936), and Clayton Waagner’s 2001 anti-abortion hoaxed anthrax letter campaign “seriously disrupted” clinic operations (McCauley and Moskalenko 2014, 75). The hoax label is only applied to these incidents *ex post*, when the act’s incorporeality is

³² As examples of bomb hoax claims of responsibility in the literature, Banlaoi (2009, 64, 73) discusses a Rajah Solaiman Islamic Movement (RSIM) hoax, where the group claimed responsibility for an October 19, 2007 blast in Makati City, Philippines, which was later determined to be the result of an industrial accident; and Kearns, Conlon, and Young (2014, 430–31) describe the hoaxes of four right-wing extremist groups claiming responsibility for an April 1999 bombing campaign in London, truthfully perpetrated by a lone actor.

³³ Joyner (1990) discusses the fire at Libya’s Rabta chemical factory on March 14, 1990, which may have been a false alarm, a hoax, or an incidence of foreign sabotage.

uncovered. For responders, policymakers, and affected citizens, immediate reactions are a function of the hoax's serious terrorism counterpart, not its hoaxed quality.

2.4.1 Research Note: Hoaxes as a “mode” —not a “tactic”—of terrorism

Although the literature consistently characterizes hoaxes as a specific terrorism tactic, doing so obscures the relationship between hoaxes and the serious terrorism act of the corresponding kind. Given the wide range of terrorist acts that can be hoaxed, it is my contention that terrorist hoaxes themselves should be treated as a *mode* rather than a *tactic* of terrorism. In other words, rather than being viewed as an alternative type of terrorist event, the hoax label refers to the “way or manner in which a thing [i.e. a terrorist act] is done” (*Concise Canadian Oxford Dictionary*, 2005 ed., s.v. “mode”). Treating hoaxes as a *mode* rather than a distinct *tactic* of terrorism would always involve qualifying the word “hoax” with a tactic descriptor (e.g. a *bomb* hoax; a *product contamination* hoax; a *biological letter* hoax), when discussing a particular event.

This distinction in vocabulary draws from a grammatical analogy. Terrorism *tactics* may be understood as *verbs* in the terrorist's vocabulary: to bomb, to assassinate, to hijack, and so on. A grammatical mode—also called a “mood”—puts this verb into context, serving to indicate the degree of reality (or unreality) and objectiveness (or subjectivity) of the corresponding verb. The most common verbal mode is the *indicative*, which portrays statements of fact, as well as opinions, and even questions. In the present analogy, verbs in the indicative mode denote those terrorist tactics whose kinetic outcomes (explosions, injuries, deaths, etc.) are readily observable. Contrastingly, verbs

conjugated in the *subjunctive* mode are “used to express a condition, wish, fear, possibility, command, suggestion, uncertainty, or hypothetical situation” (*Concise Canadian Oxford Dictionary*, 2005 ed., s.v. “subjunctive mood”). Hoaxes are essentially subjunctive terrorism: they relate to a given tactic, but are coloured with an element of incorporeality.

“Hoax” can therefore be applied in the same way that “suicide” has been applied to describe certain types of terrorist activity (e.g. a suicide bombing, a suicide skyjack mission, etc.). The specific tactic, for instance, is a bombing, but it can be carried out by a perpetrator whose death is integral to the mission’s completion (i.e. a suicide bombing), or by a perpetrator who uses benign materials to give the impression that a bomb has been placed (i.e. a hoax bombing device).

This distinction is not just one of semantics. Because hoaxes initially generate the same responses as their corresponding serious tactics, it is essential to consider them in relation to—rather than as alternatives to—these serious forms, if one is to fully appreciate the tactic’s societal impact. More importantly (as discussed in section 3.3.1), the distinction between “tactic” and “mode” helps to resolve significant coding challenges in terrorism events datasets.

2.5 Conclusions on the concept of hoaxing

In this chapter, I presented the definition for terrorism hoaxes applied in this dissertation: incidents perpetrated with a political, religious, or ideological motivation, that are designed to look like acts of serious terrorism—but by virtue of involving lies, empty

threats, or benign materials do not actually involve any direct risk of harm. Such incidents may involve hoax claims of responsibility, where perpetrators opportunistically claim responsibility for incidents they did not cause; hoax warnings, where perpetrators threaten a future terrorist attack with no intention of following through, or lie about past actions to falsely indicate that a terrorist attack is in progress; or hoax devices, where perpetrators use benign materials to give the impression that such an attack is, or has been, underway.

In the literature, terrorism hoaxes are often confused with related activities, such as false alarms, pranks, threats, and foiled or failed attacks. Although definitions are frequently blurred and imprecise, there are important conceptual distinctions between these categories of activity that align with perpetrators' goals and associated logic of action. Since data coders have not always been guided by such clear conceptual definitions, existing data sources have sometimes misclassified hoax activities per these alternate forms, and vice versa. In the following chapter, I review the empirical contours of hoax activity worldwide, per the WMDDDB and ITERATE datasets. Where possible, given details of event descriptions, I have revised coding decisions to align more closely with the definition for hoaxes used here.

3. The scope and nature of terrorist hoax activities worldwide

In the previous chapter, I presented a conceptually-driven definition of terrorism hoaxing. This chapter presents descriptive statistics to define hoaxes empirically, as they exist in the two cross-national datasets that presently code for them—ITERATE and the WMDDDB—between 1968 and 2012. Far from eschewing the challenges with hoax data explained in section 1.2.1, this chapter also seeks to highlight inconsistencies and ambiguity in these data sources' treatment of hoaxes.

Throughout, ITERATE and WMDDDB data are evaluated separately, since compiling them would represent an incomplete picture of hoax activity: purely domestic incidents involving conventional weapons are absent from both. The datasets' descriptive statistics thus only apply to their respective subsets of terrorist activity. ITERATE only includes terrorist events with a transnational or international dimension: multiple countries must somehow be connected to the incident, whether via the nationality and/or foreign ties of the perpetrators; the incident's location; the incident's victims and targets; or the entities involved in the incident's resolution (Mickolus et al. 2011). Alternatively, the WMDDDB only includes events that involved CBRN weapons or agents. In this chapter, unless otherwise indicated, the data on international or transnational terrorism drawn from ITERATE reflects the location where the attack occurred, not where it originated (i.e. the "location end" variable); and data on CBRN terrorism drawn from the WMDDDB reflects only those incidents with political or ideological motivation.³⁴

³⁴ The values presented reflect my cleaning of this data, including removal of duplicate events and re-classification of certain categories, per the details provided in event descriptions. My analysis thus differs

3.1 When do hoaxes occur?

Figure 3.1 (replicated from Figure 1.1) shows the number of hoaxes per year captured in both ITERATE and the WMDDDB (including a series for all hoaxes—including criminally-motivated ones—and a series for true terrorism hoaxes, i.e. those with political or ideological motivation). The two WMDDDB series demonstrate that terrorism hoaxes broadly follow the same trends as criminal and other hoaxes. In contrast, the discrepancy between the hoaxes covered by the ITERATE and WMDDDB series is drastic.

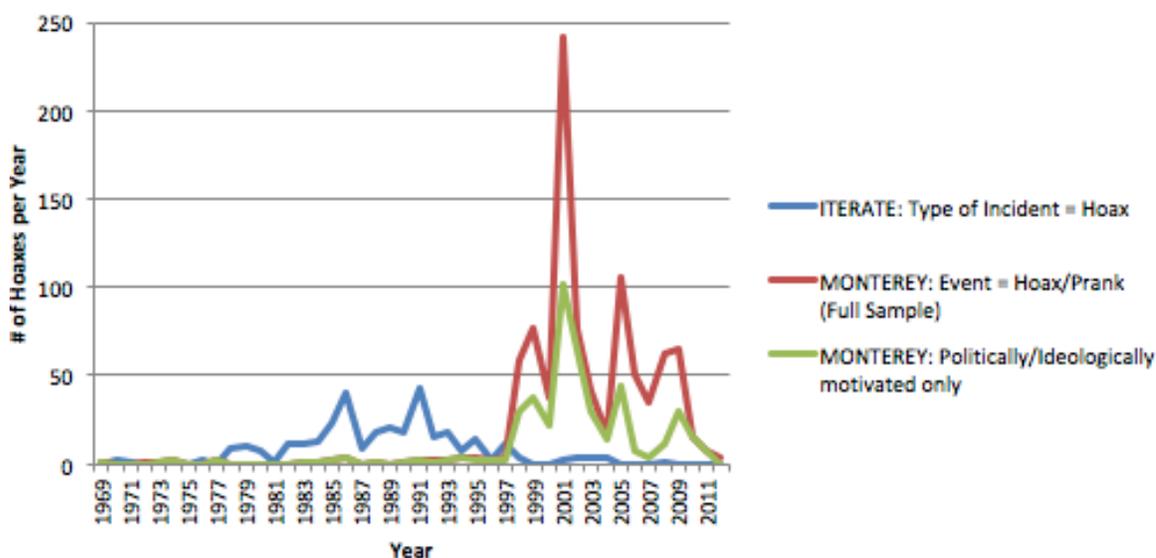


Figure 3.1 Hoax frequency in cross-national datasets, 1968-2012

While ITERATE’s transnational hoaxes occurred predominantly between the late-1970s and mid-1990s, the CBRN hoaxes in the WMDDDB occurred mostly from the late 1990s to the end of the 2000s. The only large-sample empirical studies citing ITERATE’s

in some cases from that presented in my working paper on the same topic (Tishler 2016), which reflected figures from an unverified version of the WMDDDB.

hoax data were published in the 1990s, and so the near extinction of transnational hoaxes that occurred after the 1990s is beyond their scope. The observed drop-off in transnational hoaxes is likely a function of the unavailability of FBIS *Daily Reports* to ITERATE coders after mid-1996 (Enders and Sandler 2012, 73). It may similarly be reflective of a documented worldwide decline in terrorism more generally, beginning in 1992 and “well under way at the end of 1997” (LaFree, Dugan, and Miller 2015, 33); and a similarly documented worldwide decline in the proportion of transnational incidents relative to domestic ones, beginning in the mid-1970s (LaFree, Dugan, and Miller 2015, 171). Neither of these explanations, however, explain the lack of CBRN hoaxes until that point. The initial peak in the WMDDDB series reflects the late-1990s wave of chemical and biological weapons hoaxes and their causes frequently discussed in the literature (Dishman 2001, 311; Stern 2000, 246; Tucker 2000a, 3).³⁵ The spike in 2001 reflects a surge in hoaxed anthrax letter hoaxes, following the real anthrax attacks against American targets in the weeks following 9/11. That the bulk of hoaxes in each dataset diverge so

³⁵ This wave of hoaxes is frequently explained by the “alarmist thinking [that] indirectly fuels criminals and other subversives to threaten civilians and law enforcement authorities with [chemical and biological] weapons” (Dishman 2001, 311): “the fear of WMD terrorism and the publicity associated with each incident within the United States have strengthened the likelihood of more incidents” (Cameron et al. 2000, 159). In the years following 1995, there was “an unprecedented focus on the threat of WMD terrorism by senior US policymakers and elected officials” (Cameron et al. 2000, 159); this timeframe aligns with Aum Shinrikyo’s 1995 demonstrated use of chemical weapons in Tokyo. That Aum had succeeded in deploying chemical weaponry led to an increase in fear and publicity surrounding CBRN terrorism and the threat thereof, which thereby spawned an increase in related hoaxes. The same effect is observed with the increase in anthrax hoaxes following the real incidence of anthrax use in the United States in 2001.

significantly implies that transnational hoaxes are unlikely to involve CBRN weapons, and that CBRN hoaxes are unlikely to be transnational in nature.³⁶

3.2 Where do hoaxes occur?

Figure 3.2 and Figure 3.3 present the geographic distribution of terrorist hoaxes in the WMDDDB and ITERATE, respectively, according to the GTD's regional classifications (START 2013a) of hoaxes' end-points.

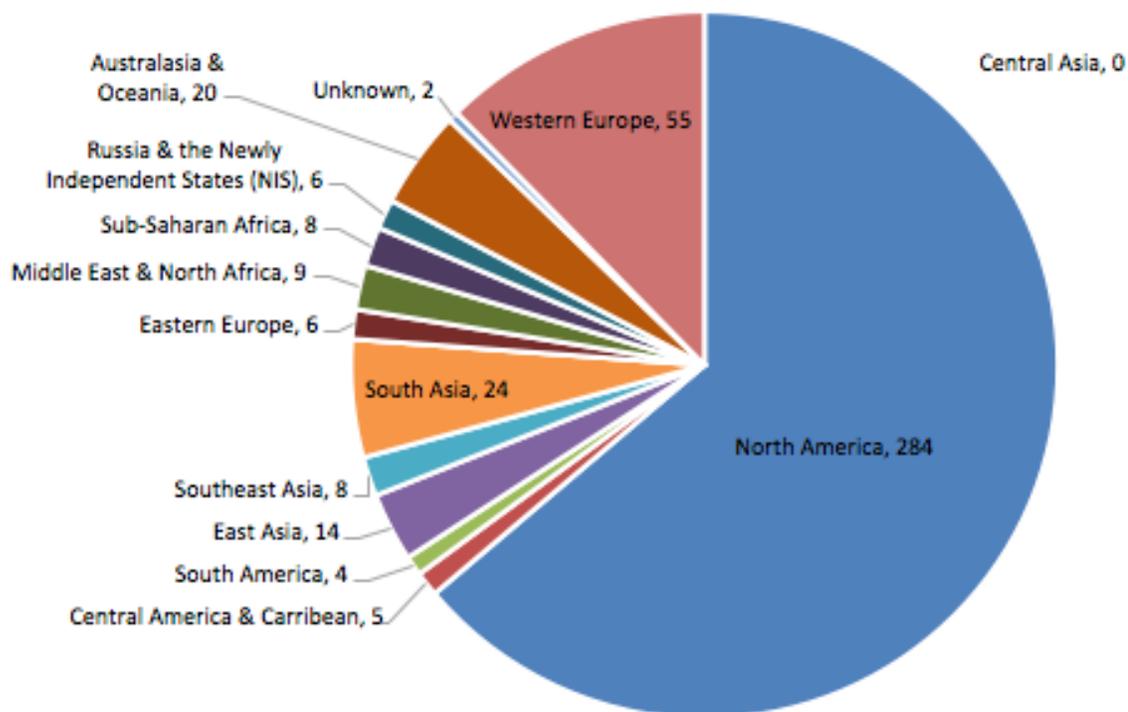


Figure 3.2 Number of CBRN terrorism hoaxes by region, 1968-2012

³⁶ Since the WMDDDB did not rely on the FBIS *Daily Reports* for coding, if transnational CBRN hoaxes *had* been perpetrated in the 2000s, they would have been captured by ITERATE coders using the same open sources as the WMDDDB.

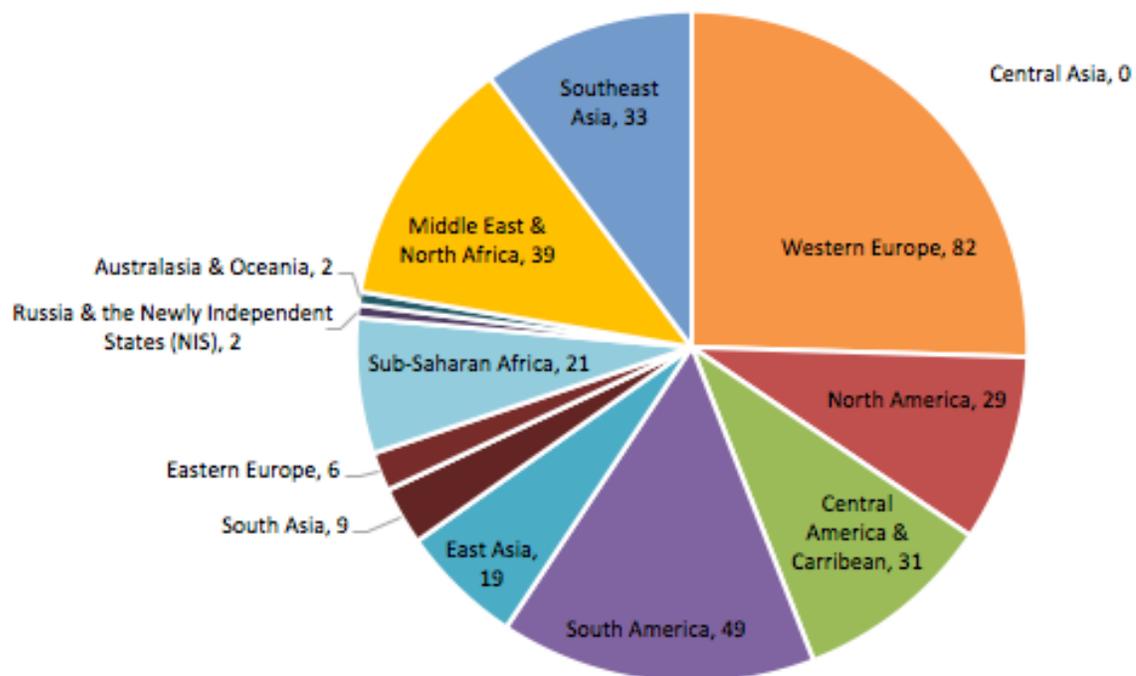


Figure 3.3 Number of transnational terrorism hoaxes by region, 1968-2012

In absolute terms, CBRN hoaxes are predominantly an American phenomenon; of the 284 terrorist hoaxes in North America—which, combined, account for nearly two thirds of the WMDDB’s entire terrorist hoax sample—243 were committed in the United States. That CBRN hoaxes are so strongly geographically concentrated appears to corroborate the finding by Cameron et al. (2000) that “the fear of WMD terrorism and the publicity associated with each incident within the US have strengthened the likelihood of more incidents” (159). The geographic distribution of transnational hoaxes is far less telling. The largest proportion (one quarter) of transnational hoaxes occurred in Western Europe, but the remaining three quarters of incidents are relatively evenly dispersed. Worthy of acknowledgment is the relative infrequency of transnational hoaxes in Russia and the Newly Independent States, as well as Australasia and Oceania.

Neither CBRN nor transnational hoaxes ever occurred in Central Asia (Kazakhstan, Kyrgyzstan, Tajikistan, Turkmenistan, and Uzbekistan). While this absence of hoaxes may be reflective of a true reluctance to use hoaxes in this region, it may also be an artifact of the datasets' reliance on media reports for gathering data (LaFree, Dugan, and Miller 2015, 22); Central Asia is also the lowest-ranked region in both datasets for serious terrorism incidents. In the 2011-2012 *Reporters Without Borders* (2012) "Press Freedom Index," these Central Asian countries all rank in the lowest 40% of countries with regard to "government respect of media freedom." Given governmental incentives to keep information about vulnerabilities quiet, it is likely that media in these un-free countries would be particularly unlikely to receive and/or be able to print in their news outlets details about terrorist events (LaFree, Dugan, and Miller 2015, 22)—hoaxes included.

To account for such systematic biases in the datasets' inclusion of events as well as varying levels of overall terrorist activity across regions, it is instructive to consider the proportion of hoaxes in each region's overall terrorist activity, in addition to the absolute frequency of hoaxes described above. Figure 3.4 compares the absolute numbers of transnational hoaxes (left axis) with the percent of hoaxes among all transnational terrorist activity (right axis) in each region. While Western Europe experienced the highest overall number of hoaxes (n=82), the proportion of hoaxes in its overall transnational terrorist activity (2.1%) was lower than the cross-regional average (3.1%). Conversely, the region experiencing the highest proportion of hoax activity was East Asia (9.3%), despite ranking eighth in absolute number of transnational hoaxes (see Table 3.1 for ranking comparisons), and accounting for less than 6% (n=19) of transnational hoaxes

worldwide. That the regions experiencing higher proportions of hoax activity are *not* for the most part also those regions experiencing higher overall levels of terrorism (North America is the exception, ranking fifth for both overall levels of terrorism as well as the proportion of hoaxes amongst all terrorist activity, at 3.5%) suggests there is some qualitative difference among regions that would lead certain areas to be more attractive hoax targets than others.

Table 3.1 Regional rankings for hoax and overall terrorism prevalence, 1968-2012

	# of CBRN Hoaxes	# Of CBRN Non-Hoaxes	Total # of CBRN Incidents	Hoaxes as a % of Overall CBRN Terrorist Activity	# of Transnational Hoaxes	# of Transnational Non-Hoaxes	Total # of Transnational Incidents	Hoaxes as a % of Overall Transnational Terrorist Activity
Australasia & Oceania	4	10	8	3	12	12	12	7
Central America & Caribbean	11	13	12	2	5	6	6	3
Central Asia	13	12	13	13	13	13	13	13
East Asia	5	4	5	9	8	9	9	1
Eastern Europe	9	11	11	4	10	10	10	4
Middle East & North Africa	6	3	3	12	3	2	2	12
North America	1	1	1	1	6	5	5	5
Russia & the Newly Independent States (NIS)	10	5	6	11	11	11	11	10
South America	12	8	9	10	2	3	3	6
South Asia	3	6	4	7	9	7	8	11
Southeast Asia	7	7	7	8	4	8	7	2
Sub-Saharan Africa	8	9	10	6	7	4	4	9
Western Europe	2	2	2	5	1	1	1	8

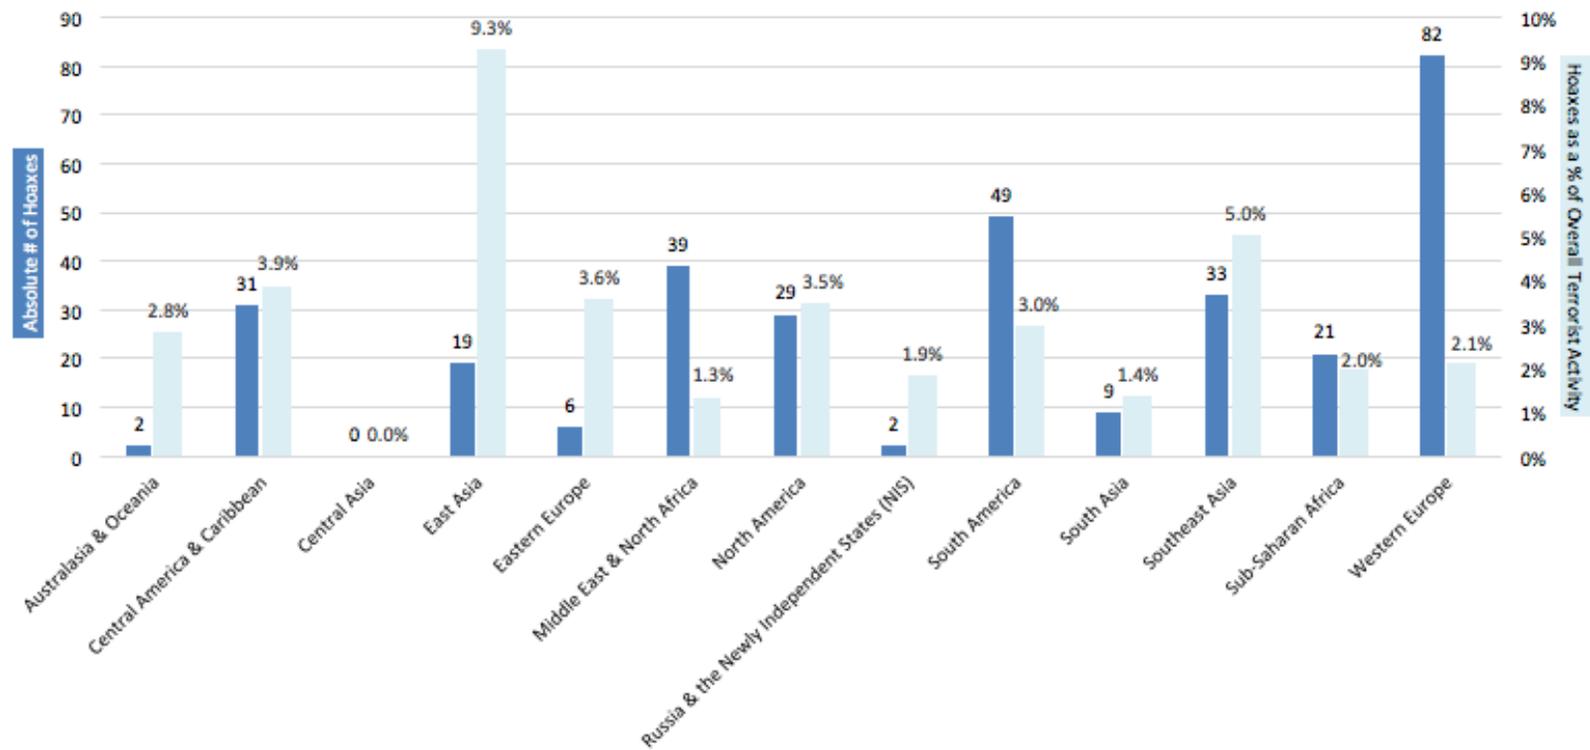


Figure 3.4 Prevalence of transnational terrorism hoaxes per region, 1968-2012

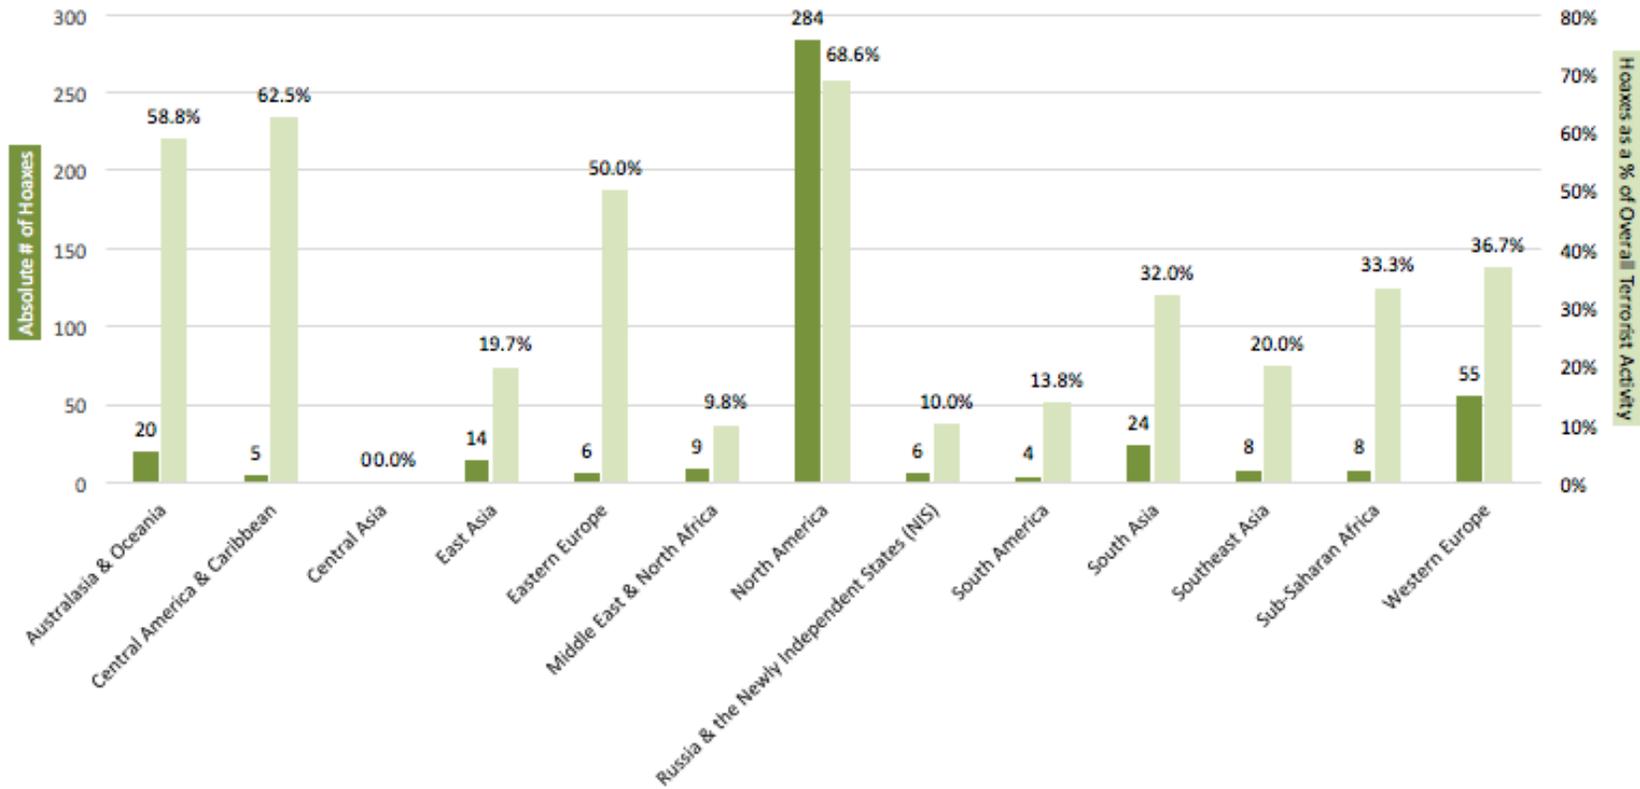


Figure 3.5 Prevalence of CBRN terrorism hoaxes per region, 1968-2012

Similar inconsistencies present themselves for CBRN hoaxes, as illustrated in Figure 3.5. While North America experiences the vast majority of CBRN hoaxes in absolute terms, it also experiences a greater proportion of non-hoax CBRN activity than other regions. While Central America & the Caribbean and Australasia & Oceania experience far lower frequencies of both hoaxed and serious CBRN incidents, as compared with North America, they rank in second and third place respectively for the proportion of hoaxes as a percentage of overall CBRN activity. Eastern Europe (50.0%), Western Europe (36.7%), Sub-Saharan Africa (33.3%), and South Asia (32.0%) also rank above the 31.9% cross-regional average for hoaxes as a percent of overall CBRN terrorist activity.

More broadly, a comparison of Figure 3.4 and Figure 3.5 highlights the more than ten-fold higher prevalence of hoaxes among general CBRN terrorist activity (31.9% cross-regional average) as compared with transnational terrorist activity (3.1% cross-regional average). Evidently, there is something in the nature of CBRN terrorism that makes it particularly prone to hoaxing.

3.3 What kinds of hoaxes occur?

Of the 1,029 CBRN terrorism events recorded in the WMDDDB, hoaxes and pranks account for the greatest proportion (445 incidents, or 43.2%), followed by use of agent (268 incidents, or 26.0%).³⁷ Figure 3.6 compares the CBRN weapon type most frequently used

³⁷ A complete event type breakdown is presented in Chapter 1, Figure 1.2.

in both hoaxed and serious CBRN terrorism events, while Figure 3.7 compares the delivery mechanisms of CBRN agents employed in hoaxes and pranks with actual uses of the agent.

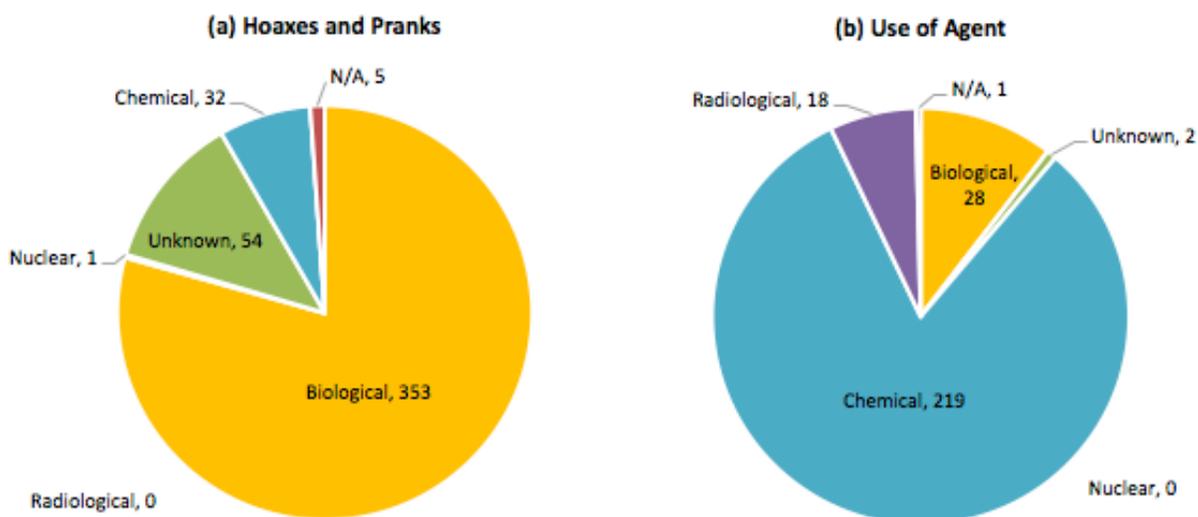


Figure 3.6 Type of weapon used in CBRN terrorism incidents, 1968-2012

Figure 3.6 illustrates that chemical agents dominate in serious attacks ($n=219$, or 81.7% of “use of agent” incidents), but biological weapons are by far the most likely to be used in a hoax or prank ($n=353$, or 79.3%). Whereas nuclear incidents occur exclusively as hoaxes or pranks (although to a very small degree), radiological attacks are never hoaxed or pranked. This discrepancy implies that incidents suggesting the presence of radiological weapons should be taken very seriously.

Figure 3.7 illustrates that CBRN hoaxes and pranks are almost exclusively delivered via mail, letters, or packages. This delivery mechanism is the only known mechanism where a hoax or prank is more likely than actual agent use—and by a large margin. Of the 390 mail/letter/package hoaxes documented, 85.1% ($n=332$) involved biological weapons; all but nine of these biological hoaxes and pranks included a substance that purported to be *Bacillus anthracis* (i.e. anthrax). The remaining mail/letter/package

hoaxes included nine chemical incidents, and 49 incidents where the weapon type was unknown or not applicable. It is interesting to note, by comparison, that of the 52 actual uses of agents delivered by mail, letter, or package, only 13.5% (n=7) involved biological agents (all anthrax). Consistent with the general findings for actual CBRN weapon use, chemical (57.7%; n=30) and radiological (28.8%; n=15) agents delivered by mail, letter, or package were far more common.

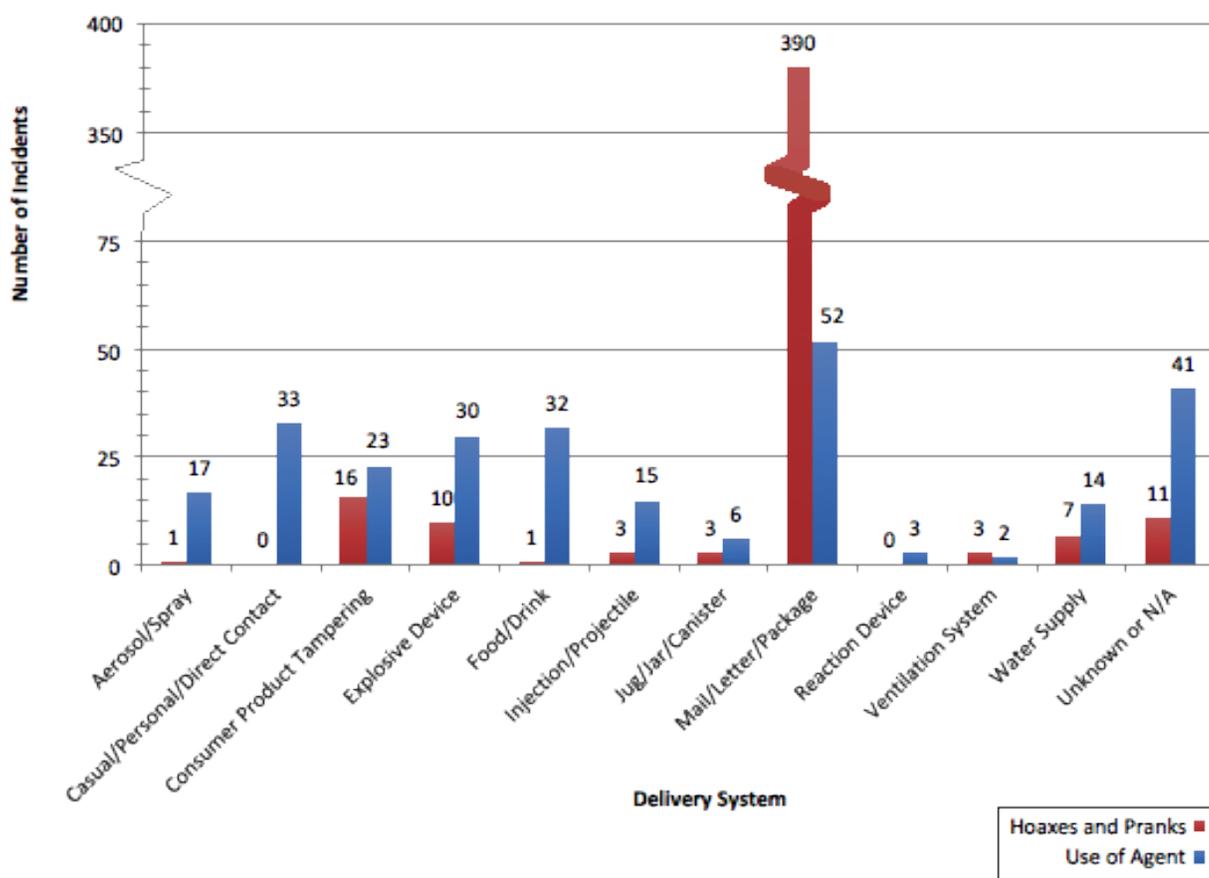


Figure 3.7 Delivery mechanism of CBRN terrorism incidents, 1968-2012

3.3.1 Hoaxes as a mode of terrorism, revisited: Implications for coding

In addition to the practical and theoretical implications of treating hoaxes as a mode versus tactic of terrorism described in section 2.4.1, such classifications have implications for the empirical analysis of terrorist activity. In the WMDDDB, as described above, hoaxes and pranks are together classified as a particular terrorist event category. They are then categorized according to the weapon type and agent purported to be present, as well as the delivery mechanism involved (i.e. an anthrax hoax letter would be coded as involving a biological weapon, with *B. anthracis*—not white powder—as the agent, delivered by mail, letter, or package). This classification is useful, since it provides the same details for a hoax incident that would be included if the attack had occurred in its serious or “indicative” form (which is, of course, how targets and responders would themselves experience the hoax in its immediate aftermath). Classifying events in this way allows for more nuanced analyses of the types of hoaxes carried out, such as those presented in Figure 3.6 and Figure 3.7.

Other datasets, however, are not so clearly and consistently designed. For instance, ITERATE codes for hoaxes as a “type of incident,” alternative to other “types” like assassinations, hijackings, or bombings. In ITERATE, only 2.4% (322) of the total 13,514-event sample are coded as hoaxes. From the variables included in the dataset, there is no mechanism to further subdivide those incidents coded as hoaxes into more detailed weapon-type categories; 100% of those incidents coded as hoaxes for incident type are also coded as “hoax, no weapons used” for weapon type. Therefore, an evaluation of transnational hoax data analogous to that presented in section 3.3 is

impossible: ITERATE treats hoaxes as a unique terrorism tactic, and does not provide any further detail regarding the nature of the hoaxed event.

ITERATE's coding scheme, however, reveals an interesting quirk: a broad range of incident types other than hoaxes are also coded with "hoax, no weapons used" as the type of weapon employed. Figure 3.8 presents the incident type breakdown of all the incidents coded with "hoax, no weapons used." Hoaxes proper account for only 20.4% of this sample. A very wide array of incidents involves hoaxed or no weapons. In fact, the only incident types that are *never* coded as involving hoaxed or no weapons are: sabotages not involving explosives or arson; nuclear-related weapons attacks; sniping at buildings, other facilities; shoot-outs with police; and car and/or suicide bombings.

While it is plausible that certain incident types *could* be carried out without weapon use,³⁸ other incident types cannot. Armed attacks, "exotic pollution," and the various forms of bombings listed as potential incident types, by definition, involve some sort of weapons use. Therefore, the coding of weapon type as "hoax, no weapons used" implies that these incidents are, in effect, hoaxes. ITERATE data thus understates the frequency of hoaxes (when filtering results by incident type), and overstates the occurrence of serious incidents (while concurrently underestimating their lethality), potentially by as much as 12%.

³⁸ Threats with no subsequent terrorist action; conspiracies to commit terrorist action; thefts; arms smuggling; occupations of facilities; barricades; hijackings; and even assassinations could all plausibly be carried out without weapons.

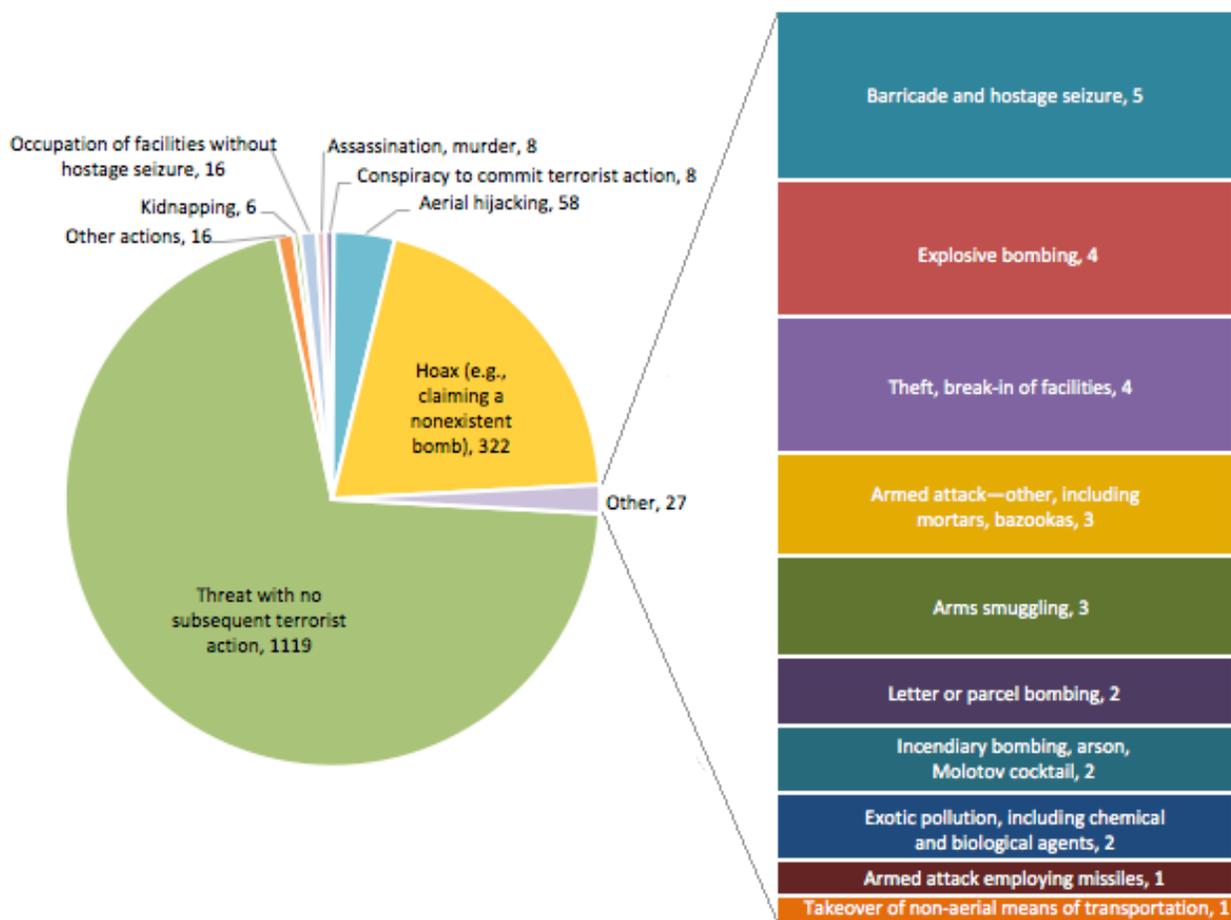


Figure 3.8 Type and number of ITERATE incidents coded with weapon type = “Hoax, no weapons used,” 1968-2012

This confusion could be averted if hoaxing were recognized as an alternative mode—rather than a distinct method or tactic—of terrorism. The ambiguous relationship of hoaxes with their corresponding serious tactic leaves excessive discretionary power to coders. The realities of coding the now-defunct WITS dataset exemplify such challenges: while the WITS methodology required that terrorists “must have initiated and executed an attack [...] result[ing] in some sort of active, kinetic effect—such as an explosion or inflicted injury or damage” (NCTC 2012, 7) to be included in the database, the database allowed analysts to categorize incidents by event type, of which “hoax” was one possible

method (discussed in Richards 2014; Rome 2013, 813). Accordingly, it is likely that WITS vastly underestimated the frequency of hoaxes—since many would have been excluded at the event-identification stage. Like ITERATE, it also failed to provide any coherent information surrounding *which* tactics were hoaxed.

In large-*n* datasets, the most effective way to treat hoaxes as a mode of terrorism is to use the approach followed in the WMDDDB, specifying weapon type and form of delivery of the indicative tactic in question, and then classifying the event according to the degree to which the incident was carried out (the WMDDDB offers a range of alternatives to “use of agent” and “hoax”). For those datasets like GTD, ITERATE, or WITS, which maintain more traditional classifications of event type, hoax incidents would be most effectively coded per their corresponding serious terrorism tactic, with the hoax mode indicated by a dummy variable. Not only would such a treatment clarify research that addresses hoaxes, but it would refine and clarify statistical understandings of other forms of terrorism by providing a mechanism to accurately filter out hoaxed events while still retaining the capacity to assess the implications of each terrorist “verb”—however conjugated—for responders and the public.

3.4 Who hoaxes?

On their own, the WMDDDB and ITERATE provide little information on hoax perpetrators. Both databases provide perpetrator names only for incidents perpetrated by groups, not individuals, and both have challenges identifying hoax perpetrators: only 8.3% of hoax events in the WMDDDB and 17.4% in ITERATE have an identified perpetrator (see Table

1.1). Still, hoax perpetrator characteristics and their intent may be deduced from characteristics of the incident (Jackson and Frelinger 2008, 588).

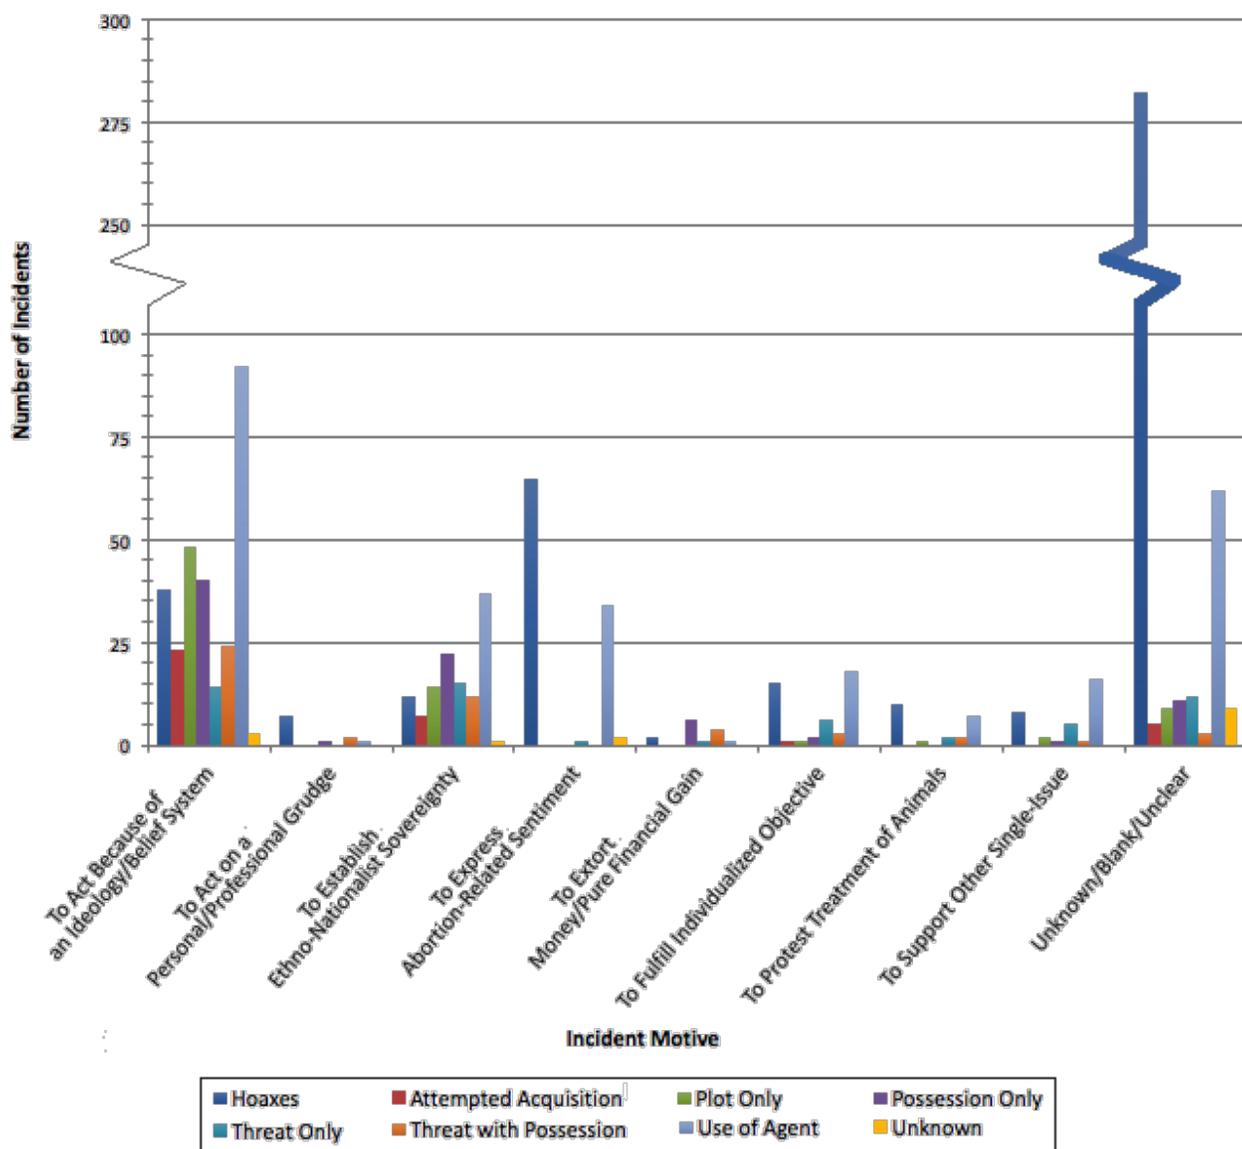


Figure 3.9 Frequency of motives according to CBRN event type, 1968-2012

Figure 3.9's presentation of CBRN event type as a function of motive shows that, while the vast majority of hoaxes are committed with unknown or unclear motives, hoaxes are nearly twice as likely to support single-issue terrorism (including abortion-

related sentiments and animal rights) than all other motives combined. Among these single issues, abortion-related sentiments are the most frequent motive for hoax events.

3.4.1 What else do hoaxers do?

Many terrorist groups who perpetrate hoaxes are not exclusively hoaxers. Of the twenty-six specific groups named as perpetrators of transnational hoaxes in ITERATE, only five committed hoaxes as their sole transnational terrorism tactic. The remaining twenty-one committed a total of 1,372 non-hoax transnational terrorist acts. These hoaxers are credible: on average, they perpetrate serious transnational terrorism nearly twenty-seven times more frequently than they hoax, and they never make threats without subsequent terrorist action.

Figure 3.10 breaks down the types of transnational attacks perpetrated by terrorist groups engaged in both hoaxes and serious terrorism. As per ITERATE's coding methodology, hoaxes are one of twenty-five possible terrorist tactics; they represent a small minority, only 3.6%, of these terrorists' total attack repertoire. The most common event perpetrated by these groups is armed attacks employing missiles (42.0% of incidents), followed by explosive bombings (10.5%), and incendiary bombings, arson, and Molotov cocktails (9.3%).

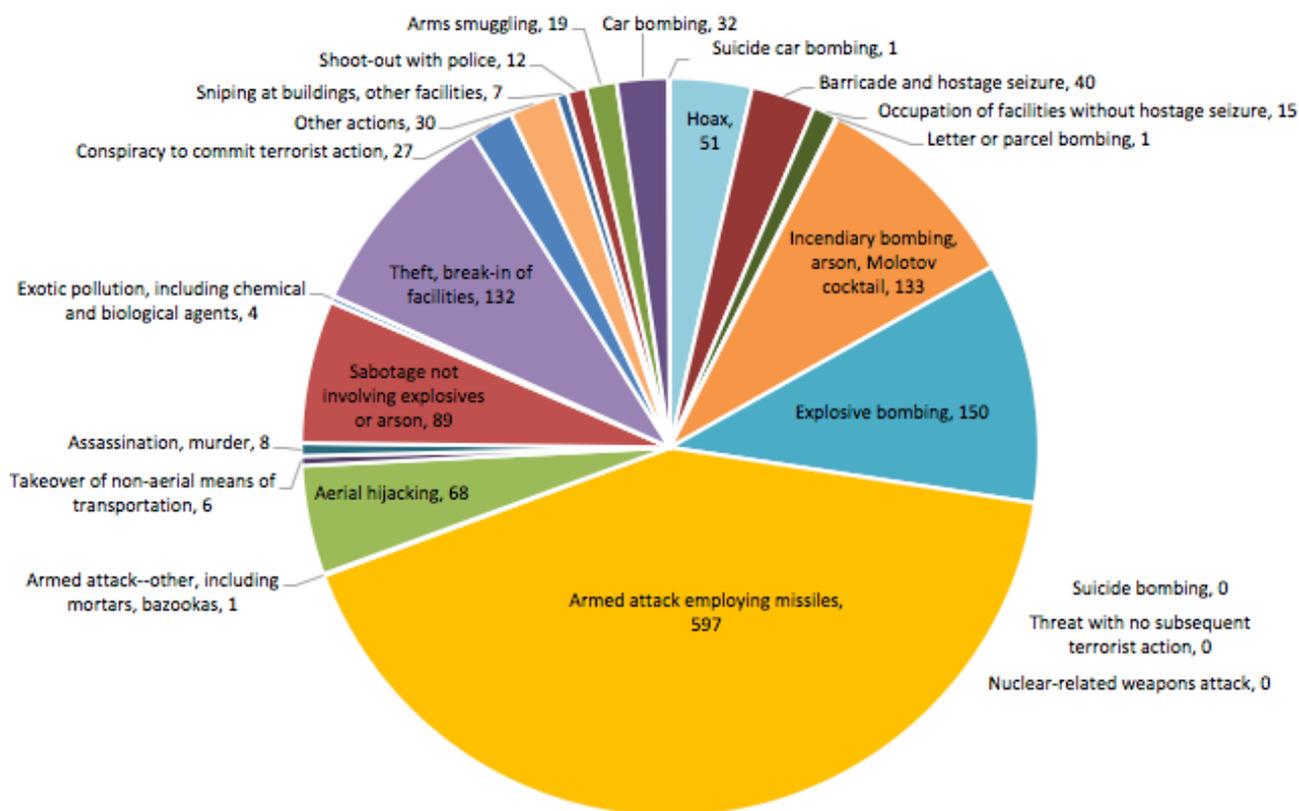


Figure 3.10 *Types of transnational events perpetrated by groups engaging in both hoaxes and serious terrorism, 1968-2012*

In the WMDDDB, there are 29 named perpetrators who committed CBRN hoaxes or pranks. Of these, sixteen groups never carried out serious CBRN activity. The remaining thirteen, however, carried out a total of 95 serious CBRN terrorism incidents. For instance, the Scottish National Liberation Army (SNLA)—the WMDDDB’s most prolific hoaxer, responsible for a collective six biological and chemical hoax incidents—also used chemical agents (sodium hydroxide) on three occasions. However, all but one of the group’s hoax events preceded their serious acts by at least two months; and all but one of those hoaxes involved biological agents (*B. anthracis*) or the toxin ricin—not the sodium hydroxide that the group ultimately deployed. Even when the group conducted a chemical hoax

subsequent to its initial serious chemical weapons uses, it hoaxed with a different agent: lead sulfate. The SNLA's hoaxes thus lacked credibility, and do not appear to be indicative of any attempt to hoax as a dry run of subsequent serious activity. Although the SNLA's hoaxes clearly demonstrated the group's interest in a range of CBRN materials, they never demonstrated their capacity with respect to the hoaxed agents, suggesting that hoaxes were likely used to inflate assessments of their tactical repertoire and to spread fear.

All that said, addressing only the transnational and CBRN subsets of terrorist events vastly underestimates the scope of groups' activities. For instance, when cross-referenced with the GTD (to expand attack profiles beyond the WMDDB's CBRN-only scope), these same 29 named perpetrators are associated with an additional 5,197 serious terrorism events. Whereas a sole focus on these groups' CBRN events suggests that hoaxes account for 38.9% of all activity, including the GTD's events reduces this proportion to a mere 0.7%. In excluding non-CBRN events, two groups that the WMDDB identifies as pure hoaxers are actually responsible for serious terrorism. These include Direct Action, which the GTD codes as responsible for one bombing event; and the Earth Liberation Front (ELF), responsible for 84 serious events—predominantly incendiary attacks on facilities and infrastructure—in the GTD. Even where the WMDDB identifies groups as serious terrorist perpetrators, it vastly underestimates the scope of their serious activity. As extreme examples, the WMDDB attributes only eleven serious events to the LTTE and four to the Taliban, but the GTD contributes an additional 1,607 and 2,467 serious events to these groups, respectively.

While it is true that a focus on only transnational events necessarily “produces some artifacts in the analysis” (Jackson and Frelinger 2008, 596), this is even more-so the case with a focus on only CBRN activity; CBRN events reflect only a small subset of most groups’ overall tactical repertoire. When analysing groups’ tactical choices, then, it is important to consider the full range of their activities and not be bound by the scope of existing datasets.

3.5 Conclusions on the scope and nature of terrorist hoax activities worldwide

In this chapter, I presented a descriptive, empirical overview of terrorist hoax activity worldwide, between 1968 and 2012. The data presented flow from two distinct data sources: the transnationally-oriented ITERATE, and the CBRN-centric WMDDDB. In each case, the datasets’ coding structure determines the amount and type of data available for evaluating the characteristics of these hoaxes. The WMDDDB offers a more detailed picture of what tactics are hoaxed, how they are delivered, and who is responsible for them. ITERATE offers a more complete—albeit still limited—picture with respect to how hoaxes fit within groups’ overall tactical repertoires. Given their respective strengths, both datasets contribute the basis for the clear majority of empirical research addressing terrorism hoaxes to date. In the following chapter, I review this literature, along with literature from related areas of terrorism studies that can be applied to better evaluate the hoax phenomenon.

4. Insights into hoax behaviour: Previous empirical research

There is very little scholarly work dedicated to understanding terrorism hoaxes. In some cases, hoaxes are explicitly culled from analyses of terrorists' tactical choices (see, for instance, Asal, Ackerman, and Rethemeyer 2012; Ivanova and Sandler 2006; Ivanova and Sandler 2007), usually on the grounds that such events are not reflective of "serious" terrorism risk.³⁹ In others (Enders, Parise, and Sandler 1992; Enders, Sandler, and Cauley 1990; Prunckun Jr. and Mohr 1997; Rowlands, Littlewood, and Kilberg 2012; Tishler 2013a; Tishler 2013b), hoaxes are addressed—but always in relation (and usually subsidiary) to other tactical choices; they are never treated as the primary object of analysis.

Although there exist some less-well-known sources of statistical and qualitative data on hoaxes,⁴⁰ most empirical works addressing hoaxes rely on data from ITERATE or the WMDDb, and thus speak to only transnational or CBRN hoaxes, respectively. Accordingly, the literature to date contributes only small fragments of the total hoax-use

³⁹ Asal, Ackerman, and Rethemeyer (2012) exclude "hoaxes, pranks, and groundless threats involving CBRN weapons" (237) from their examination of the factors that influence the terrorist organizational decision to pursue CBRN weapons; they do not explain their data culling decision, but it can be assumed that hoaxes, pranks, and groundless threats are eliminated because they do not reflect the *actual pursuit* of CBRN weapons by the organization. In a study examining the relationship between regime characteristics and the likelihood of CBRN terrorist incidents, Ivanova and Sandler (2006) cull from their sample those instances where there is no evidence of CBRN agent possession; this means they exclude "plot only" incidents as well as threats without agent possession, hoaxes, and pranks (427). They do not provide an explanation for this culling. Ivanova and Sandler (2007) make the same culling decision, but this time present a justification: "Such incidents are removed from our sample, because they do not really represent a CBRN risk even though they may result in inconvenience and costs" (276).

⁴⁰ These include the open-source-based *Violent Dissident Republican (VDR) Project* examining VDR activity in Northern Ireland between 1997 and 2010 (used in Horgan and Morrison 2011); a qualitative, narrative-based dataset of terrorism in Canada between 1973-2003 (used in Leman-Langlois and Brodeur 2005); and the comparative case studies of McCauley and Moskalenko (2014) and Tucker (2000c). See section 4.3 for more details.

picture, and significant gaps remain. Fortunately, insights into why groups might choose to employ hoaxes in their terrorist campaigns can be drawn from non-hoax-specific areas of terrorism studies, as well as more conceptually-oriented frameworks that incorporate various components of hoax subcategories.

The review of literature below frames the study of hoaxes in the context of terrorism studies more broadly, highlighting—where they exist—works that explicitly acknowledge hoaxes,⁴¹ and logical extensions to hoaxes.

4.1 How terrorists decide among tactics

Within the terrorism studies literature is a robust corpus of material addressing why some insurgent or non-state groups resort to the use of terrorism as a tactic, instead of more peaceful means or other forms of violent conflict. These explanations encompass the structural, or “root causes,” debates; the impact of organizational, group-based factors; and individual terrorists’ personal biographies and psychological features—a comprehensive review of which is beyond the scope of this dissertation. A smaller body of literature seeks to explain how terrorist groups select among tactics once they have made the decision to engage in terrorism. Since hoaxes are understood in the literature as a particularly low-severity terrorism tactic (see Prunckun Jr. and Mohr 1997), alternate to other tactics such as bombing or hijacking, this literature can shed light on the decision to use hoaxes within a violent terrorism campaign.

⁴¹ Other studies may exist that subsume hoaxes in their general findings, but if they do not accord specific attention to hoaxes in their description of the data or findings, they are not directly relevant to furthering our knowledge of hoaxes as a distinct mode of terrorism.

4.1.1 Rational choice

The rational choice model is the most prevalent explanation for terrorists' level of activity and target selection. It is sometimes articulated as a formal model, and is sometimes more implicit in terms of describing terrorists' strategic logic.

4.1.1.1 The formal rational choice model and its applications to hoaxing

As articulated in Anderton and Carter (2009a), the formal rational choice model's fundamental parameters include income (and a budget constraint), prices (of terrorist activity and of consuming some alternate, composite good), and preferences (as delineated in indifference curves). The model can be used to show how terrorist organizations will react to changes in their available budgets and "price policies" that affect relative costs. As a function of price policies, terrorists may substitute into other activities (e.g. other weapons technologies and even non-terrorist activities) or toward other targets. The substitution effect reflects the incentive to substitute away from goods that have become more expensive. Income effects result when an increase in a good's price causes real income (that is, the size of the actor's budget) to decrease.

These rational choice dynamics have been used to explain the incidence of transnational hoaxes relative to alternative attack types, based on ITERATE data. Enders, Parise, and Sandler (1992) use a time series spectral analysis to show how the tactics used in terrorism events follow regular trends and cycles. Of all the series they test, "threats and hoaxes" (a combined category) had the shortest cycle (only 3.6 quarters), illustrating

the authors' general finding that a tactic's logistical complexity is positively related with its cycle length; in other words, resource-intensive incidents occur with far lower frequency than events requiring little resource input. Additionally, since 3.6 quarters is roughly equivalent to a one-year cycle, the authors propose that threats and hoaxes may have a seasonal component: "seasonal factors, such as high tourism in the summer months, may account for the cyclical behaviour of the Threats [and Hoaxes] series" (Enders, Parise, and Sandler 1992, 316). The authors suggest that tactical choices are driven by terrorists' resource availability, given strict budget constraints. However, their analysis suffers from a fallacy of division: the cycles exist across the entire sample of ITERATE data—it does not take into account *which* groups are hoaxing or conducting serious attacks and *when*. It may be that the relative frequency of hoaxes is affected by a larger number of hoaxing perpetrators than by specific perpetrators choosing to hoax more frequently than they conduct serious attacks.

Hoaxes are also affected by substitution and complementarity dynamics across tactics. For instance, when the installation of metal detectors in airports (a price policy to raise the costs of skyjacking) resulted in a tactical substitution away from skyjackings and into non-aviation hostage incidents and assassinations (Enders and Sandler 1993), it also produced a complementarity effect for the threats and hoaxes series: "threats and hoaxes became less credible in [skyjacking] situations and, hence, were curtailed" (Enders, Parise, and Sandler 1992, 308). In Enders, Sandler, and Cauley's (1990) analysis of the substitution between different types of transnational terrorist activity in relation to various counterterrorism interventions, the authors identify a substitution into hoaxes

and threats when resource-intensive terrorism incidents declined following the United States' 1986 retaliatory raid against Libya. Prunckun Jr. and Mohr (1997) confirm this shift toward less violent activity following the Libya raid, collapsing the individual ITERATE data series used in Enders, Parise, and Sandler (1992) into three broad categories of low-, medium-, and high-severity terrorism. Hoaxes, per Prunckun Jr. and Mohr's methodology, rank second-lowest in the low-severity classification. Per their framework, hoaxes are used when counterterrorism activities render more severe, logistically complex, and resource-intensive tactical alternatives difficult.

However, since hoaxes "do not really use up resources" (Enders, Sandler, and Cauley 1990, 101) and can accordingly "occur at [terrorists'] will" (Enders and Sandler 2012, 73), they pose a challenge to rational choice models, which rely on budget constraints to explain terrorists' allocation of resources and effort across various activities. In such a framework, we might anticipate an indefinite supply of hoaxes—irrespective of the level of serious terrorism groups are carrying out. Clearly, there are costs to hoax behaviour that are inadequately accounted for in these models.

4.1.1.2 The strategic logic of terrorist (hoax) behaviour

Studies articulating strategic logics for terrorist behaviour emphasize that terrorists carry out violent acts to signal their capacity to inflict harm and, in doing so, to convince opponents to concede to their demands (e.g. Kydd and Walter 2006; Pape 2003; Siegel and Young 2009).

Kydd and Walter's "The Strategies of Terrorism" (2006) has been lauded as "the most comprehensive framework of rationalist explanations for terrorism" (Kearns, Conlon, and Young 2014, 422). In it, the authors identify five enduring terrorist goals—regime change; territorial change; policy change; social control; and status quo maintenance—and delineate five primary signaling strategies by which terrorists attempt to achieve these goals. First, in an attrition strategy, "terrorists seek to persuade the enemy that the terrorists are strong enough to impose considerable costs if the enemy continues a particular policy" (Kydd and Walter 2006, 51). Second, in an intimidation strategy, terrorists "try to convince the population that the terrorists are strong enough to punish disobedience and that the government is too weak to stop them" (Ibid.), Third, in a provocation strategy, terrorists "attempt to induce the enemy to respond to terrorism with indiscriminate violence, which radicalizes the population and moves them to support the terrorists" (Ibid.). Fourth, in a spoiling strategy, terrorists "attack in an effort to persuade the enemy that moderates on the terrorists' side are weak and untrustworthy, thus undermining attempts to reach a peace settlement" (Ibid.). Finally, in an outbidding strategy, violence is used "to convince the public that the terrorists have greater resolve to fight the enemy than rival groups, and therefore are worthy of support" (Ibid.). Terrorists will select a particular strategy given the conditions under which each is likely to succeed. Their strategic choices are neither exclusive nor static; terrorists will frequently use more than one strategy in a campaign, and their choices therein may vary over time.

As articulated by Abrahms (2008), such rationalist explanations for terrorist behavior are underpinned by three core assumptions: 1) terrorists' political goals are relatively stable/consistent, and are explicitly formulated by the terrorist organization; 2) costs and benefits of options are weighted, and terrorism is selected because alternate political avenues are blocked; 3) the option with the highest expected utility is selected, and terrorists rely on experience to believe that terrorism will be effective—they will abandon the strategy if it is observed to be consistently ineffective. The literature on terrorists' strategies explains a range of terrorist behaviours, usually from the starting point that they pose a *prima facie* challenge to these rationalist assumptions.⁴² Many of these seeming contradictions, however, can be explained in rationalist terms by recognizing the diversity of terrorist groups' objectives;⁴³ and/or by recognizing that terrorist organizations are not homogeneous, unitary actors (as elaborated in section 4.1.1.3 below).

⁴² For instance, Abrahms (2008) articulates seven puzzling behaviours of terrorist groups that would appear to violate the rational model's core assumptions. Most significantly—and the topic of his earlier (2006) research in this area—is the logic of targeting civilians writ large. As a coercive strategy, he argues, attacking civilians is ineffective, if not overtly counterproductive, since it usually generates a negative impact on public opinion that would make target audiences particularly unlikely to grant concessions (Abrahms 2006; Abrahms 2008, 82–84). Conversely, Pape's (2003) strategic logic for suicide terrorism is premised on the idea that terrorist groups have learned suicide attacks *are* an effective coercive strategy: past campaigns have taught terrorists the lesson that suicide terrorism can further their political cause and lead to territorial concessions.

⁴³ A range of studies highlight terrorists' objectives beyond the five enduring terrorist goals highlighted in Kydd and Walter's (2006) framework. For instance, Abrahms (2008) argues that terrorist groups are truly motivated by social bonds among members, and not policy concessions that would satisfy the organization's political goals; Jackson and Frelinger (2008) describe a diversity of terrorist intents, even within a single incident: "to injure or kill; kidnap; damage or destroy; hijack; enter or control; or to stage a symbolic strike" (587); and Marsden (2012) presents a comprehensive, multi-level framework for terrorist's goals, ranging from short-term tactical objectives, to mid-range organizational and strategic objectives, to the long-term ultimate political aims described by Kydd and Walter. Appropriately specifying groups' goals is essential to evaluating the efficiency and efficacy of the tactics they choose.

Recent and ongoing research has used this sort of strategic framework to understand components of hoax activity. A recent conference paper by Charles Mahoney (2017) presents the strategic logic for terrorist bluffing, or what I have called the “empty threat” component of hoax warnings. He explicitly traces how terrorist bluffing poses a challenge to existing rationalist theory. Bluffing imposes “audience costs” that would reduce estimates of the terrorists’ capabilities and credibility. Accordingly, if terrorists execute violent acts as a form of costly signalling—to demonstrate their capability and resolve, so as to coerce audiences into policy concessions or behavioural changes to avoid the costs of further violence—then they should rarely or never hoax. However, Mahoney’s framework goes on to explore a range of conditions under which terrorists’ costs from bluffing may not always be greater than their potential benefits (discussed in greater detail in Chapter 5). In particular, Mahoney outlines various operational and strategic conditions under which groups may bluff rationally, in order to further their overarching political aims. While most of his work is conceptual, he presents an empirical test of one such scenario: a case study chronicling how Boko Haram has sequenced bluffs and legitimate warnings in its campaign against educational institutions to “maximize the total leverage of their threats and attacks to a level that is greater than successfully realized attacks alone” (12).

Kearns, Conlon, and Young (2014) present a rationalist explanation for hoax claims of responsibility and other lies that terrorists tell about their violent activities. They broach the topic via the challenge to rational choice posed by non-credible claims: if terrorism’s strategic logic is to communicate the threat of future violence to target

audiences, then lies would obstruct the link between the violent act and expected political outcomes. Of the various terrorist lies they explore, the logic of costly signalling maps most neatly onto false claims of responsibility, since such hoax claimers are trying to convince target audiences that they pose a credible threat. Making explicit reference to Kydd and Walter's strategies of terrorism, they show how hoax claims of responsibility are compatible with both attrition and outbidding strategies. For attrition, "[a] group that falsely claims another's attack could lack the power it is attempting to display, or it could have the capacity to carry out the attack itself but opportunistically claim credit to prevent another group from demonstrating its power" (Kearns, Conlon, and Young 2014, 426). For outbidding, since multiple claims "cause doubt among the population over the effectiveness of the rival" (426), a hoax claim of responsibility can prevent a rival group from receiving additional support, and thus increase the hoaxer's relative advantage.

4.1.1.3 The unitary actor assumption and principal-agent relationships

In addition to the above-described strategic reasons for which a group would intentionally make a hoax claim of responsibility for another group's attack, Kearns, Conlon, and Young (2014, 425) propose an alternative explanation for terrorists' lies: a principal-agent problem. Principal-agent problems arise when individual terrorist group members (the "agents") make decisions and carry out activities on behalf of the central terrorist organization and its leaders (the "principal"). Terrorist groups are not internally homogenous, and so the unitary actor assumption of rational behaviour cannot be expected to hold across the entire group (Abrahms and Conrad 2017, 281; Abrahms and

Potter 2015, 334; Pittel and Rübhelke 2011; Siegel and Young 2009, 768; Weinstein 2007, 44). Principal-agent dynamics—and the divergent incentive structures of leaders and followers—have been used to explain a range of seemingly irrational terrorist activities, such as targeting civilians (Abrahms and Potter 2015); failing to claim credit for attacks (Abrahms and Conrad 2017; Kearns, Conlon, and Young 2014); blaming attacks on other groups (Kearns, Conlon, and Young 2014); and even individuals' participation in suicide missions (Pittel and Rübhelke 2011). In each of these cases, rationalist explanations show how the observed activities were carried out by rational agents who simply were acting in pursuit of individualized incentives which diverged from their broader organizational (political) goals.

Principal-agent problems are exacerbated in more horizontal group structures. The more hierarchical groups are, the clearer their lines of communication and corresponding command and control, but also the more vulnerable they are to infiltration (Abrahms and Potter 2015, 318–19; Enders and Sandler 2012, 239; Heger, Jung, and Wong 2012, 745–47; Siegel and Young 2009, 768; Weinstein 2007, 131). By adopting more horizontal organizational structures that support their need for secrecy and security, terrorist groups thus expose themselves to the costs of principal-agent divergence. Importantly, principal-agent problems emerge whether “leadership deficits” result from a conscious organizational choice to structure the organization horizontally, or whether control is degraded as a result of such exogenous factors as counterterrorist decapitation campaigns (Abrahms and Potter 2015, 319).

Returning to hoax claims of responsibility, Kearns, Conlon and Young (2014, 428–29) describe how different group structures affect the incidence of credit claiming, due to communications breakdowns and corresponding principal-agent dynamics. While principals may issue false claims of responsibility in the strategic logic of attrition or outbidding, as described in section 4.1.1.2, they may also do so in situations where an agent “tell[s] his superiors that he committed an attack that he did not such that the leadership is unaware that they are falsely claiming responsibility for another group’s attack” (429). The authors explicitly describe how such principal-agent problems occur more frequently in groups that are comprised of horizontal cells, rather than centrally-coordinated hierarchies.

4.1.2 Ideology and group motivation

Terrorists’ tactical choice may alternatively be a function of groups’ motivations for violence (Crenshaw Hutchinson 1972; Drake 1998; Hoffman 2006, 229–40; McAllister and Schmid 2011, 233; Rapoport 2002). Once a group has decided to use violent means, ideology becomes a determinant of the group’s range of perceived “legitimate” targets and means, since it shapes people and objects into “representative symbols” that, in the terrorists’ rationalization, are deserving of the violence with which they are targeted (Drake 1998, 59).

Of the various ideologies inspiring terrorism, religion has received particular attention, due to its perceived recent increase in prevalence. This literature has addressed religious groups’ remarkable lethality when they resort to terrorist violence (discussed

below, in section 4.2.1.1), and their propensity toward targeting civilians in the first place. Asal, Brown, and Schulzke (2015) find that a religious orientation makes violent non-state organizations more than twice as likely to target civilians than groups with other ideological orientations.

4.2 The nature of violence

While the literature discussed above generally applies to one or both of domestic and transnational terrorism (as defined by perpetrators', victims', and targets' nationalities), various clusters of research have been demarcated according to the particular tactics and/or weaponry involved, as well as certain attack outcomes. For the study of hoaxes, the literatures on terrorists' lethality, as well as on suicide and CBRN terrorism, are particularly instructive.

4.2.1 Lethality: Hoaxing's inverse

Since hoaxes are a particularly low-severity mode of terrorism that does not actually generate direct harm to people or property, it may be useful to evaluate the factors contributing to their incidence as the inverse of those which lead to particularly lethal or indiscriminate attacks.

4.2.1.1 Lethality and ideology

Building on the effects of group ideology on terrorists' tactical decision-making discussed in the previous section (4.1.2), there is a body of literature that addresses religious groups'

particular lethality when they resort to terrorist violence (Asal, Ackerman, and Rethemeyer 2012, 237; Asal and Rethemeyer 2008; Cronin 2003, 41–42; Hoffman 2006, 82–130; Kilberg 2011; Kilberg 2012). According to Hoffman (2006), this lethality stems from the transcendental elements of religion, which frees religious groups from the social and moral constraints by which secular groups are bound (88–89; 239). This view is consistent with Juergensmeyer’s (1988) contention that religion inherently deals with a “cosmic struggle” between order and an intrinsically violent disorder. Violence perpetrated in the name of religion is seen as justified, since it occurs in the context of this struggle. This dynamic can explain why religious organizations would be more likely to pursue a CBRN capability, as a particularly heinous form of weaponry (Asal, Ackerman, and Rethemeyer 2012); and frames how religious groups’ reliance on a supernatural audience augments their lethal capacity by generating a strong “othering” effect (Asal and Rethemeyer 2008, 439; Cronin 2003, 41–42).

Treating radical, religious communities as “economic clubs” who provide exclusionary social services to their members, Berman (2009; with Laitin 2008) presents an alternate explanation for religious groups’ lethality: religious groups weed out “shirkers” and reduce outside options for socialization and employment, by requiring members to make severe sacrifices in order to benefit from the club’s social services. Such groups are left with a select membership that is unlikely to defect once the choice to use violence has been made, and they are thus able to carry out more complex and lethal attacks.

Conversely, various single-issue motivations for terrorism—such as animal rights and eco-terrorism—have been described as intentionally avoiding causing death or bodily harm (“If a Tree Falls: A Story of the Earth Liberation Front” 2011; Animal Liberation Front 2015; Monaghan 2000, 257–58).

4.2.1.2 Lethality and group structure

Group structure is similarly gaining currency as a popular explanation for terrorist group lethality. Kilberg (2011; 2012), for instance, examines the effect of terrorist groups’ organizational structure on their attack lethality. He discerns four categories of structure, ranging from “market” groups—lacking leadership, centralized control, and functional differentiation or specialized wings—as the most decentralized form, to “bureaucratic” groups—with leaders who exhibit centralized control over the group, and functional differentiation—as the most highly centralized. Bureaucratic groups, he finds, are the most lethal. Using a proxy for group hierarchy based on groups’ bases of operation and non-violent activities, Heger, Jung, and Wong (2012) similarly find that networked groups are less lethal than hierarchical ones. They present a similar theoretical framework to Kilberg’s, identifying three distinct advantages provided by centralized structures: functional differentiation, which lends itself to specialization; clear command and control structures, which facilitate agenda control by the core leadership; and clear lines of accountability, which correspond to the core’s ability to punish or reward agents for execution of their agenda. This latter point maps on to the authors’ argument relating to group structure and the principal-agent problem described in section 4.1.1.3: hierarchical

organizations have “tightened principal-agent relationships” which make them “particularly effective” (746).

4.2.1.3 *Lethality and group origins*

Finally, Asal, Brown, and Schulzke (2015) examine groups’ splinter history as a predictor of violence. They find organizational splits to be a positive predictor of violence in general, but not of targeting civilians more narrowly.

Whatever the underlying explanation, if religious orientation, hierarchical structure, and origins as a splinter organization make groups more likely to adopt increasingly violent forms, we can infer that these same features might make groups particularly unlikely to pursue hoaxes, since these do not produce any direct harm (indiscriminate or otherwise).

4.2.2 Suicide terrorism: A conceptual analog to hoaxing as a mode of terror

Suicide terrorism is accorded its own subset of the literature on tactical choices, and is treated with a variety of theoretical approaches.⁴⁴ As with the literature addressing hoaxes, Crenshaw’s (2007) canonical review maintains that this literature has produced neither a consensus, nor even clearly demarcated schools of thought. The concept remains imprecise, the facts are poorly substantiated, and neither explanations nor policy

⁴⁴ See, for instance, Gambetta’s (2006) organizational studies approach; Atran’s (2004) emphasis on radicalization and socialization processes; Israeli’s (1997) socio-cultural “Islamikaze” thesis; and utility-based models at the group-level (Berman 2009; Berman and Laitin 2008; Pape 2003) and at the individual-level (Ferraro 2006).

recommendations sufficiently distinguish suicide terrorism from other terrorist methods, nor do they account for variations within the phenomenon. Crenshaw attributes these deficiencies, in part, to a relatively low number of suicide attacks: decisions on which cases to include or exclude significantly shape the conclusions that are drawn. What the literature *does* present is a complex process, contingent on social, psychological, and political interactions. Suicide terrorism is *not* rooted in psychopathology, fanaticism, or individual grievances; it is an adaptable and controllable tactic, with instrumental value for an organization.

As a distinct mode of carrying out terrorism activity, suicide terrorism can be seen as an analog to hoax activity: Hoffman (2006, 132) points out that suicide terrorism has a high likelihood of success, is relatively inexpensive, and more logistically simple than other attack modes. These are all characteristics shared by hoaxes. Hoaxes and suicide terrorism differ substantially in their degree of lethality and in the element of perpetrator sacrifice, but their similarities are conceptually instructive. Importantly, they can both be understood as a *mode* of terrorism: a way it is carried out, rather than an alternative tactic. Various tactics can occur as part of a suicide event or of a hoax (e.g. a suicide/hoax bombing; a suicide/hoax hijacking). Accordingly, the robust literature on suicide terrorism and its relationship with other forms of tactics can inform how hoaxes are understood.

4.2.3 CBRN terrorism: Parallel data sources and frequently-hoaxed weaponry

Another subset of the literature on tactical choices pertains to terrorists' choice to use unconventional—i.e. CBRN—over conventional weapons. This literature is particularly

useful for informing the study of hoaxes because its major data source (the WMDDDB) includes hoaxes as a possible event type, but also due to the high proportion of hoaxes in overall CBRN activity: as described in Chapter 3, there is a nearly ten-fold higher prevalence of hoaxes amongst CBRN terrorist activity as compared with transnational terrorist activity—31.9% and 3.1%, respectively.

The CBRN literature tends to emphasize motivational and organizational variables that would affect groups' propensity to adopt unconventional weaponry. Ivanova and Sandler's (2007) study is a significant contribution to the literature on CBRN terrorism using large-sample research design, and constitutes a substantial improvement over the authors' earlier (2006) work on the topic. This study identifies the likely perpetrators and regime targets of CBRN terrorism activity—although it explicitly excludes hoaxes from the analysis. Regarding attack perpetrators, the authors find that a group's past CBRN use is a significant determinant of its future use; and that religious cults (but not religious fundamentalists) and transnationally-oriented groups are the most likely CBRN users. In Tishler (2013a; 2013b), multinomial logit regression—using chemical, biological, radiological, or nuclear as a categorical dependent variable—is used to test the technological determinism assumption that the more readily CBRN technology, materials, and knowledge are accessible to terrorists, the more likely the terrorists will be to use unconventional weapons of the corresponding kind. The study tests the technological determinism hypothesis on the full sample of WMDDDB incidents, and then compares a stratified sample of serious attacks with that of hoaxes. While it was expected that technological determinism would be stronger for serious attacks—since they actually

require materials, whereas hoaxers are unconstrained by reality—the regression results show the opposite: “the widely held technological deterministic hypothesis, while seemingly false for perpetrators of serious attacks, appears to be internalized by hoaxers and pranksters who may draw on the availability of technology and knowledge to make their claims and acts more believable” (Tishler 2013a, 66; Tishler 2013b, 22). This finding poses a particular challenge to policymakers, who are also inclined to believe the technological determinism hypothesis: if hoaxers and policymakers are both informed by the same logic, then hoaxes conforming to this hypothesis will be viewed by policymakers as more plausible and credible (Tishler 2013a, 67; Tishler 2013b, 23).

Other studies combine WMDDDB data with external datasets encompassing groups’ non-CBRN activity to examine the organizational and contextual factors influencing terrorist groups’ decision to pursue CBRN over conventional weapons. Although they exclude hoaxes from their analysis, Asal, Ackerman, and Rethemeyer (2012) find that organizations with alliance structures—but *not* those with religious motivations—and those based in strongly globalized, authoritarian countries are more likely to pursue CBRN weapons. Rowlands, Littlewood, and Kilberg (2012) retain hoaxes in their analysis, using Kilberg’s (2011; 2012) framework for evaluating group structure. They find that strongly religious terrorist groups and those organized in a “hub-spoke” structure are more likely to perpetrate hoaxes over serious attacks. Although they emphasize these group-based characteristics, the authors recognize that hoaxes are twice as likely to be perpetrated by individuals than by terrorist groups.

4.3 Other data sources, methods, and contexts for examining hoaxes

Beyond the illustrative case studies referenced in Chapter 2, there are a handful of more rigorous empirical studies addressing hoaxes that do not rely on ITERATE or WMDDDB data. Horgan and Morrison (2011) use data from the open-source-based *Violent Dissident Republican Project* at Pennsylvania State University's International Center for the Study of Terrorism. Examining VDR activity in Northern Ireland between 1997 and 2010, they find that hoax devices and warnings increased substantially as of 2009. Hoax incidents were the third most common VDR activity in 2010, only minimally trailing shootings and punishment attacks, and defused bomb incidents. Hoax incidents were marginally more common than detonated bomb incidents, and substantially more frequent than petrol bomb incidents, assaults, violent riots, arson riots, violent robberies, and other types of VDR activity (651).

Leman-Langlois and Brodeur (2005) also include hoaxes in their database, which is a narrative-based qualitative dataset regarding terrorism in Canada between 1973 and 2003. While the authors explicitly intend to consider acts together, rather than as discrete events like incident-based datasets do, they present some count data "for the more quantitatively inclined" (123). Of a total of 400 entries, they identify seventy-three "disruptive hoaxes," along with fifty-nine "threats," and fourteen "plots and failed attacks" (123). While they use these distinct categories, the authors do not define how they differentiated between them in their classification.

Hoaxes have also been addressed using comparative case study analysis. McCauley and Moskalenko (2014) present a case study comparison of lone wolf terrorists

in an attempt to profile them. One of the cases they include is that of Clayton Waagner, who committed a large-scale, anti-abortionist, anthrax letter hoax campaign, but never actually crossed the threshold into violent action. Examining their cases in relation to two pyramidal models of radicalization, one for violent *opinion* and the other for violent *action*, they suggest that Waagner moved from the apex of the opinion pyramid to that of the action pyramid when he began “stalk[ing] abortion providers with the intention of killing them,” but then moved down the action pyramid to continue his fight by harmless means (i.e. hoaxing) when he “could not pull the trigger” (McCauley and Moskalenko 2014, 76). Individuals and groups who perpetrate *only* hoaxes may thus indicate only partial radicalization to violence—or a different radicalization track altogether—than their serious-attack-perpetrating counterparts.

Tucker (2000c) presents an edited compilation of case studies examining various chemical and biological terrorism incidents. While none of the historical case studies deal exclusively with hoaxes, Tucker (2000b) addresses them explicitly in the key observations he draws from the multi-case methodology. He concludes that elaborate and plausible hoaxes “should not be dismissed as irrelevant to [chemical and biological weapons] terrorism,” but that simple and implausible attacks like anthrax letter hoaxes “are more of a nuisance than a serious danger” (254).

4.4 Conclusions on existing literature

Although several scholars mention hoaxes, their insights have never before been evaluated in a comprehensive manner. Furthermore, their insights have insufficient

breadth to systematically assess which terrorist groups hoax, or to effectively explain why they do so. The review of literature presented above highlights that a variety of perspectives and angles can offer valuable insights into understanding groups' tactical decision-making, which can explain certain groups' hoax use. In particular, rational choice dynamics—affected by group ideology and structure—can be used to explain groups' overarching tactical strategies. In the absence of a dedicated body of hoax literature, the study of hoaxes may be meaningfully contextualized relative to the literature on group lethality, which relies upon understandings of group ideology, structure, and also group origins, in addition to the related literatures on suicide and CBRN terrorism. Insights from these approaches are used to guide the following chapter's elaboration of hoaxing groups' strategic logics, and to shape subsequent chapters' more narrow hypothesis formation.

5. Conceptual framework: Proto-theories for hoaxing

Current literature dedicated to understanding the strategic logic of hoax behaviour does not address the full breadth of hoax activity. In examining claims of responsibility for attacks perpetrated by other groups, Kearns, Conlon, and Young (2014) address only one sub-component of hoax claims of responsibility (they exclude claims for accidents or natural occurrences, as well as all other hoax subcategories). Similarly, in examining empty threats or “bluffs” regarding future terrorist activity, Mahoney’s (2017) ongoing work addresses only one subcomponent of hoax warnings: empty threats of future activity. While Horgan and Morrison (2011) address hoax devices and warnings, they only present a conceptual evaluation for hoax devices (despite lumping the two together in their empirical treatment). Until now, no studies have presented a comprehensive conceptual framework which ties all of these behaviours together.

This chapter presents proto-theories for the full spectrum of hoax behaviour, deduced from the descriptive statistics and literature presented in previous chapters. These proto-theories are rationalist. They assume that terrorists act in pursuit of ultimate political aims, with a range of more immediate tactical, organizational, and strategic goals that contribute to the achievement of their overarching goals (see Marsden 2012); that terrorists weight the costs and benefits of different modes of violent and non-violent action; and that terrorists choose to carry out hoaxes because they provide higher utility in the given circumstances than alternate methods. While Chapters 6 and 7 outline specific hypotheses regarding organizational characteristics expected to affect

perpetrators' cost-benefit analysis when deciding whether to incorporate hoaxes into their campaign activities, the proto-theories explain terrorists' overarching strategic logic(s) in doing so. They highlight under which conditions certain factors become increasingly or decreasingly salient in the terrorists' motivational calculus; which objectives the terrorists are privileging in each case; and which hoax subcategories best satisfy each logic.

Importantly, it is expected that different actors will hoax for different reasons. Groups have different needs from one another, and their needs will shift under different conditions. Even within individual groups' campaigns, hoaxes may work simultaneously in service of multiple strategic logics. It is therefore highly unlikely that any one of these proto-theories will serve as a blanket explanation for all hoax activity. In the same vein, the different proto-theories are not mutually exclusive. Equifinality and complex causality are likely features of any accurate characterization of the strategic logics motivating terrorist groups' use of hoaxes.

5.1 Hoaxes as a low-cost alternative to serious activity

The most salient advantage of terrorism hoaxes relative to serious attacks is their low resource requirement for perpetrators (Enders, Sandler, and Cauley 1990, 101). Because of their low cost, "dramatic changes" in hoaxes' frequency can occur at the terrorists' will (Enders and Sandler 2012, 73). Hoaxes also put individual terrorist agents at far lower risk than serious activities: while hoaxes are punishable by criminal law in most countries, their penalties are less severe than for serious acts; and they leave fewer traces for

detection, as compared with serious plots. Furthermore, although hoaxes may indicate what serious tactics terrorist are thinking about (Tucker 2000b, 254; also referenced in Dishman 2001, 311), they do not reveal any information regarding operatives' actual capabilities, thereby conferring another operational advantage over serious attacks. Accordingly, hoaxes allow terrorists to inflate their attack profile and reduce the length of time between incidents, without incurring the same risk or expense as carrying out an equivalent number of serious incidents. They are thus attractive to groups who have insufficient resources or capacity to carry out as much serious activity as they would like; or to operate in the geographic areas that they wish to target. They are also attractive to groups who are conserving resources while they prepare for serious attacks.

Groups will be more willing to accept lowered assessments of their credibility in contexts where they have low physical or human capital: when they are either early in their lifecycle or toward its end, or when they have suffered losses due to counterterrorism strikes against them. Mahoney (2017, 15) shows how bluffs might be strategic for small, young groups, since they have little to lose in terms of audience costs, and hoax warnings can generate exposure and publicity. Similarly, Kearns, Conlon, and Young (2014, 435) show that new or weak groups may make false claims in order to project greater strength, although their claims will not be perceived as credible. Conversely, Mahoney (2017, 20) also shows how bluffs may be carried out as "legacy acts" for declining, formerly-powerful groups, to signal their continued significance in revolutionary movements; once-powerful groups have name recognition and

credibility—“reputational assets”—that persist even once their physical resource base declines. This same logic may be extended across all hoax subcategories.

Hoaxes thus have low costs for perpetrators, and also confer potentially large benefits. Until a credible hoax is uncovered as false, it can be expected to achieve many of the same shorter-term objectives as serious attacks: publicizing the perpetrators’ cause, spreading fear, and wasting government response resources. Accordingly, hoaxes can play into a modified attrition strategy: they may not successfully convince the authorities that the perpetrator is strong, but they will keep sustained pressure on counterterrorist activities, which gradually drains resources.

Hoaxes’ relative benefit thus increases in contexts where the costs they impose on society are large. To spread fear, a target audience must know of the hoax’s existence. This is most likely in countries with free media. Since the hoax does not on its own generate any observable kinetic outcomes (like an explosion, or death), reporting of the hoax is required if it is to generate publicity for the group. To drain emergency response resources, the target country must have emergency response services who deploy proactively at the threat of violence, before individuals are hurt or property is damaged. And finally, if hoaxes are to generate any coercive effect on policy, they must target countries with a low threshold for pain, and a desire to protect citizens at all costs—commonly-described features of democracies (Wilkinson 2011). Hoaxes are thus more likely to be used in countries where governments have conceded to terrorist demands in the past. Even where governments have not made policy concessions, if past hoaxes have

proven to impose any societal costs in a given context, perpetrators will be more likely to continue using them.

While this explanation emphasises how hoaxes as a low-cost alternative to serious terrorism might be strategically desirable for groups, it is possible that some groups' central decision-making bodies will weigh the credibility costs of hoaxing as more important than any possible benefit they could confer—but lower-level foot soldiers might hoax anyway. In this principal-agent problem, the agent chooses to hoax because hoaxes are a low-cost alternative to serious activity, in terms of individual risk accrued. On-the-ground agents have magnified “incentives to defect” given “the level and immediacy of the risks [they] face” (Weinstein 2007, 131). In order to appear more effective, they may report the hoax as a serious incident to their managers, who in turn claim responsibility (as discussed in Kearns, Conlon, and Young 2014, 429), or they may claim the hoaxed activity in the name of the group themselves. In either case, the agent hoaxes to exploit the relative advantage conferred by hoaxes' low requirements for investment of resources and risk.

5.1.1 Hoaxes as a low-cost means for exploiting intergroup dynamics

Inasmuch as hoaxes look like serious attacks, they can be used in attrition or outbidding strategies to gain relative advantage over rival groups. While groups might risk undermining assessments of their own resolve if and when hoaxed warnings or devices are uncovered as false, opportunistically claiming responsibility for another group's attacks does not generate a net negative effect for the false claimer: when a group claims

responsibility for another group's attack, it introduces doubt regarding the rival's effectiveness and undermines the rival's demonstration of its power (Kearns, Conlon, and Young 2014, 426). Therefore, even if the hoaxing group does not anticipate gaining supporters directly, a hoax claim of responsibility will diminish the true perpetrator's relative advantage.

5.2 Hoaxes as a complement to serious activity

As described in section 5.1, hoaxes' low cost requirements make them, in some cases, a strategic substitute for serious terrorist events. However, hoaxes may also be used to complement and even facilitate serious terrorist activities. Hoaxes used for this purpose should be of particular concern to counterterrorists, since they augment groups' actual violent capacity.

5.2.1 Hoaxes as a component of mixed strategies

Hoaxes drain response resources because of the potentially drastic costs of type II error—that is, of providing a false negative, or wrongly concluding that a threat is a hoax when indeed it poses a serious threat. Groups can exploit their targets' sensitivity to type II error by sequencing serious and hoaxed activities in such a way as to keep the counterterrorists guessing. If the counterterrorists cannot be certain that an attack is a hoax, then they *must* respond as if it were serious. When they are sequenced in combination with serious activities, hoaxes thus confer all of the perpetrator benefits described in section 5.1, but with the added benefit of appearing more credible.

In game theory, this type of behaviour could be modelled as a mixed strategy game. In a mixed game, the actor (i.e. the terrorist group) uses two or more strategies (i.e. to hoax or to commit some form of serious terrorism) according to a well-defined probability. When the game is repeated (i.e. the terrorist conducts more than one attack) and the probability of hoaxing is optimally set, mixed strategies can help groups inflate their attack profiles—and achieve the corresponding benefits of a large number of attacks—without committing extra resources. Doing so plays into a strategy of attrition, since the sustained pace of attacks tricks governments into wasting response resources and maintains heightened levels of fear.

It does not appear that mixed strategy games have been formally modelled in the terrorism literature with respect to tactical choices, although this application is a clear analog to the Operation Overlord game (or “Normandy Game”) applied by Dixit and Nalebuff (1991, 195) cited in the field of conflict economics. It is, however, implicit in Mahoney’s (2017, 11–14) discussion of how terrorists sequence bluffs in their ongoing campaigns. His case study of Boko Haram’s campaign against Nigeria’s educational institutions illustrates that the group used empty threats in this way: for ten of the group’s eleven hoaxed warnings, the targets responded with enhanced security and/or shutting down schools. Interestingly, Mahoney finds that Boko Haram used bluffs in a mixed strategy to increase its geographic scope of influence: the bluffs predominantly occurred in areas outside of Borno, where most of its serious incidents occurred. The empty threats thus compensated for the group’s lack of capacity in these areas, and because they were

sequenced amidst real attacks elsewhere, they were taken seriously. This dynamic can be extended to apply for hoax devices, as well as warnings.

Mixed strategies will be most effective for groups that have already carried out serious terrorism acts. It makes little sense to begin playing a mixed strategy without first establishing credibility via demonstrations of capability in a corresponding serious mode. Within mixed strategies, individual hoaxes will have greater effect where hoaxes follow a similar *modus operandi* to groups' serious attacks; and when the hoax-to-serious event ratio is low. When hoaxes are sequenced with similar serious attacks (i.e. with similar targets, weapons types, delivery systems, and agents), they will have greater credibility. The higher the proportion of serious-to-hoaxed attacks, the less vulnerable the perpetrators will be to any signal the hoaxes might send implying the perpetrators' weakness. Additionally, if audiences realize groups are strategically bluffing in this way, "it may not decrease their estimation of groups' credibility" (Mahoney 2017, 12). In other words, if the counterterrorists recognize that the terrorists are playing mixed strategies, they will not downgrade their assessments of the groups' potential to impose serious costs—even when hoaxes are uncovered as false—since the hoax behaviour is a strategic choice, and not one mandated by lack of will or desire to inflict serious pain.

5.2.2 Hoaxes to directly facilitate serious activities

Hoaxes may also be used instrumentally to directly improve the effectiveness of serious activities, without wasting valuable materials. For instance, perpetrators might hoax as a dry run, to practice their own operational activities; or as a tactical intelligence-gathering

mission, to test the adversary's defences and see what reactions to a serious event would be. Such instrumental hoaxes increase the serious attack's chances of success and effectiveness, in terms of reaching the desired target and/or generating maximum harm. Hoaxes following this logic can be identified when perpetrators hoax a new tactic (including weapon type, delivery system, and/or agent) that they subsequently deploy in its serious form.

In their discussion of lone-actor terrorists, Gill, Horgan, and Deckert (2014, 430) find that nearly 30% of violent lone actors carried out dry runs of their activities, and more than half of these followed up with an eventual attack within one year. The authors also found that propensity to engage in dry runs was linked to perpetrator motivation, with single-issue lone actor terrorists being the most likely perpetrators, right-wing lone actors the least likely, and al-Qaeda-inspired lone actors somewhere in the middle (Ibid., 432). These findings may be plausibly extrapolated to group-based actors.

5.2.3 Hoaxes to drown out serious activities

Finally, hoaxes might be employed as decoys, to flood the system and distract from perpetrators' "real" plots and activities. Doing so achieves all the same strategic and tactical benefits described above, and play similarly into an attrition strategy to keep the counterterrorism and intelligence forces of the adversary busy. In addition to diluting the authorities' attention to the perpetrators' serious activities, this sort of hoax might lull the counterterrorists into inaction—a "cry wolf" phenomenon—which would make them even more vulnerable when a real attack strikes. Decoy hoaxes benefit from hoaxes' low-

cost requirements and logistical simplicity. They can be easily carried out in large numbers without detracting from whatever serious activity the group has planned.

5.3 Hoaxes as a less-severe alternative to serious activity

The strategic logics for hoaxing described above are premised on hoaxes' ability to confer benefits by appearing like serious activity. In other cases, however, hoaxes' main asset may precisely be that they do not produce the same severe outcomes (physical harm to people and property) as serious attacks. It may be, as Brian Jenkins (1987) famously observed, that "simply killing a lot of people has seldom been a terrorist objective. Terrorists want a lot of people watching, not a lot of people dead" (8). Hoaxes can thus be used to satisfy organizational and even ultimate aims (described below) to which causing widespread destruction would be antithetical. In this context, hoax perpetrators get the benefit of a low-cost, logistically-simple activity, without the audience costs faced by terrorists who want to be perceived as coercively powerful.

5.3.1 Hoaxes as an outlet for incomplete radicalization to violence

Since they do not generate direct physical harm, hoaxes may be the output of actors who have not fully radicalized to violence. In the language of McCauley and Moskaleiko's (2014) pyramidal radicalization model, such hoaxers would be at the apex of the opinion radicalization pyramid, feeling a personal moral obligation to act (71), but they would fall short of the rank of "terrorist," at the apex of the action radicalization pyramid, since they do not "target lethal violence at civilians" (73). McCauley and Moskaleiko would instead

classify such hoaxers as “radicals, engaged in illegal political action that may include violence” (73). While both of these categories may be classified as terrorists, per legal definitions of the term,⁴⁵ their differentiation does reflect a qualitative difference in terms of degree of radicalization between those most hardened terrorist actors and those who are similarly motivated, but less inclined to use violent force.

At the group level, this tendency toward less violent forms may be indicated by certain ideologies and their corresponding norms against the use of physical violence. Asal, Brown, and Schulzke (2015, 602) explicitly test the mediating role of self-professed democratic norms on groups’ perpetration of generalized violence, as well as violence targeted toward civilians. They find that groups without a democratic ideology are about one-third more likely to use violence, and about twice as likely to target civilians, than the baseline.⁴⁶ Some ideologies thus have a mediating role on the use of violence. Certain animal rights (Monaghan 2000; Monaghan 2013) and environmentalist (“If a Tree Falls: A Story of the Earth Liberation Front” 2011) groups have been described as explicitly avoiding causing harm to people, precisely because their ideological orientation does not allow for violence to be perpetrated against sentient beings. By extension, one might expect anti-abortionist groups—concerned, by definition, with the preservation of life—to have a similar ideological predilection toward less-severe methods. That anti-abortionist sentiments are the most frequent motivation for CBRN hoax activities (section

⁴⁵ The Canadian legal definition for terrorism provided in Chapter 2 recognizes the threat of violence—whether or not it is directed toward civilians—as sufficient to classify as a terrorist act.

⁴⁶ They explain this finding by two main mechanisms. First, democratically-oriented groups are opposed to the “absolutism and intolerance” that leads other ideologies, such as religious ones, toward more extreme forms of violence (Asal, Brown, and Schulzke 2015, 594). Second, the democratic peace hypothesis for inter-state war can be applied to non-state actors (Ibid., 595).

3.4; Figure 3.9) supports this extrapolation. Serious acts by groups with anti-abortionist, animal rights, and environmentalist orientations tend to emphasize property damage, rather than death or personal injury.⁴⁷ Hoax activity provides an additional morally-consistent means for these groups to publicise their cause, spread fear, and drain responder resources.

In groups that do carry out attacks against civilians, hoax activity might still follow this logic of an incomplete radicalization to violence under conditions that aggravate principal-agent dynamics. Any person may claim membership in market structured (or “leaderless resistance”) groups. This means that an individual who might never independently carry out an act of physical violence could carry out a hoax in the name of a group that sometimes does. Such an individual may believe their hoax actions are contributing to the cause, even if they contradict the group’s core *modus operandi*. Compounding this effect, groups with ideologies that are consistent with norms against violence are particularly likely to attract such members.

⁴⁷ Although some of these groups use CBRN weapons, which are generally considered a particularly extreme form of weaponry, there is a broad range of CBRN capability amongst these groups (see footnote 127) and a wide range in their intended scope of harm. For instance, Fodeman (2015) describes anti-abortionist groups’ use of butyric acid attacks as only “indirect violence”—an entirely different category from more direct event types such as murders, bombings, or even arson attacks (174). That he classifies butyric acid—a CBRN agent—alongside such other extremist event types as “home invasion, vandalism, trespassing, [...] and burglary” (Ibid.) highlights the particularly low- to non-lethal end of the CBRN spectrum at which such groups operate. This mild use of CBRN conforms with the expectation that such groups are disinterested in causing death or personal injury.

5.3.2 Hoaxes to signal restraint; or to intentionally disrupt, not destruct

Even where there is no ideological norm against violence amongst group members, groups may still find it advantageous to pursue less-severe modes of activity under certain conditions. Importantly, when groups want to establish legitimacy amongst constituent populations, they will be inclined toward less-extreme forms. This is particularly the case amongst ethno-nationalist groups, who aim to present themselves as legitimate political alternatives to the existing powers. In example, Hamas's abandoning of suicide terrorism in the late 1990s to early 2000s reflected their desire to regain local support and build a reputation for fairness, as compared with the Palestinian Authority (Cragin 2006, 199). Similarly, dissident groups in Northern Ireland have privileged hoax devices and warnings against transportation infrastructure as a means of "disrupt[ing] the normalisation of Northern Irish society" (Horgan and Morrison 2011, 650), not causing widespread destruction. On the flip side of the coin, religious groups—which are unconcerned with earthly constituents (see section 4.2.1.1)—will be less inclined to use hoaxes, since hoaxes' lower severity does not confer this same kind of strategic advantage for them.

By intentionally disrupting infrastructure, rather than destroying it, groups concerned with constituent support can impose costs on society (albeit to a lesser degree than their serious alternatives), without a corresponding loss of legitimacy amongst constituents. When a group has a history of serious violent activity, hoaxes may be carried out as an intentional signal of restraint, to build legitimacy.

5.4 Explanations for seemingly irrational hoaxes

While this dissertation privileges rationalist explanations for hoax behaviour, there are explanations for hoaxes' occurrence that this framework cannot account for: flaws inherent in observational data, without access to details on groups' inner workings (see Jackson and Frelinger 2008); and the possibility of actors' irrationality.

5.4.1 Improper attribution

Hoaxes may be wrongly attributed when external actors claim hoax activity in a group's name. Pranksters are particularly likely to contribute to improper attribution of hoaxes when they claim hoaxes in a group's name to free-ride on its reputation for serious violent activity, and receive correspondingly greater attention. For instance, a 14-year old girl from Rotterdam was arrested when she tweeted "hello my name's Ibrahim and I'm from Afghanistan. I'm part of Al Qaida [sic] and on June 1st I'm gonna do something really big bye" at American Airlines, in an apparent joke (ITV Report 2014). In this case, it was easy to trace the Tweet's origins, and recognize that the hoax bomb warning was not perpetrated by a real al-Qaeda member. In other cases, however, it may not be so self-evident.

Terrorist perpetrators may also blame their hoax activity on rival groups, for many of the same reasons that they may carry out "false flag" serious terrorism (Kearns, Conlon, and Young 2014). While hoaxes are unlikely to provoke infiltration, stigmatization, or destabilization, in the same way that false flag serious attacks might, they can play into a

strategy of deniability.⁴⁸ Blaming a hoax on another group for deniability is particularly likely when that hoax takes on an instrumental role (as in sections 5.2.2 and 5.2.3)—but the group does not wish to risk downgraded assessments of its capability. Carrying out such a hoax and blaming it on a rival group would confer the same practical benefits for facilitating the group’s serious activity, without the same audience costs. It would even have the added benefit of improving the false flag hoaxer’s relative advantage, by reducing evaluations of the rival’s efficacy.

Whatever the reason, unless deep investigation uncovers an event’s true perpetrator, it is likely that datasets will wrongly attribute observed hoaxes to the perpetrator whose name in which they were claimed.

5.4.2 Mistaken failed serious attacks

Failed serious attacks may also be wrongly labelled as hoaxes. In particular, a hoax designed to intentionally signal restraint may look a lot like a failed serious attack. Such a hoaxer might include real explosives—to trigger bomb-detection responses and to demonstrate that they possess the required materials to carry out a serious attack—but intentionally leave out a detonator, to ensure that no direct harm results. Conversely, a particularly ill-equipped actor intending to cause serious harm might not know that a detonator was required to trigger an explosion. This form of improper identification is

⁴⁸ Jenkins (1988) presents four strategies by which intelligence agencies or rival terrorist groups may use “false colours”: to promote their rivals’ infiltration; to increase their own deniability; to stigmatize the group, when carrying out particularly repugnant acts; and to destabilize the group, by provoking some sort of general crisis. Kearns, Conlon, and Young (2014) reprise these same explanations in their discussion of “false flag terrorism” as a form of terrorists’ lies.

vulnerable to post-hoc perpetrator rationalizations: whatever the reputational costs of a strategic hoax, the costs of a terribly bungled attempt at a serious incident—in terms of perceptions of credibility and capability—are likely worse. While both may be sub-optimal, “weak” is preferable to “incompetent.”

5.4.3 Irrational actors

Finally, hoax perpetrators may simply be irrational actors, acting with no strategic logic, and without weighing costs and benefits across the spectrum of available terrorist activities. While irrationality almost certainly explains some cases (there are always outliers), it offers little guidance to counterterrorists and scholarly inquiry alike. Accordingly, before hoax perpetrators on the whole can be diagnosed as irrational, we must first rule out the menu of strategic logics for hoaxing described above.

5.5 Conclusions on the strategic logics for hoaxing

In summary, existing literature and trends in the available hoax data suggest that a range of non-exclusive strategic logics may plausibly explain groups’ use of hoaxes. First, terrorist actors may use hoaxes to exploit their low costs, relative to serious activities—a logic that could be applied to particularly great effect in contexts where groups are seeking a low-cost means of exploiting intergroup dynamics. Second, hoaxes may be used as a complement to serious activity. Such hoaxes may be incorporated as a component of groups’ mixed strategies, designed to waste counterterrorism resources, and broadly keep responding agencies on their toes; as dry-runs or intelligence-gathering facilitators,

to instrumentally increase future serious attacks' chances of success; or as "noise" to drown out simultaneous serious activities, and thereby decrease the likelihood of appropriate counterterrorist intervention. Third, hoaxes may be used by actors who prefer less-severe alternatives to serious activity, whether as an outlet for incomplete radicalization to violence, or as an intentional signal of restraint guided by normative imperatives and concern for audience costs. In each case, hoaxes may result from deliberate group strategy—as defined by some central executive—or from principal-agent breakdowns and associated actions of fairly autonomous individual group members. When hoaxes do not appear to conform with one of these logics, irrationality may be to blame. So too, however, may be flaws inherent to observational data: improperly attributed events, and mistaken failed serious attacks.

Most of these strategic logics have not been subject to prior empirical evaluation. In the dissertation's next part, I test these logics indirectly, by evaluating whether their corresponding motivational, structural, and contextual indicators are associated with a higher or lower likelihood of groups' use of hoaxes.

PART II: INDUCTIVE THEORY-BUILDING

6. Characteristics of likely hoaxers: Insights from regression

In the last chapter, I presented a range of proto-theories, or strategic logics, for explaining why terrorist groups use hoaxes. In the present chapter, I begin my empirical evaluation of terrorist groups' hoax behaviour. In this first step, I use probit regression to identify organizational characteristics that differentiate terrorist groups who use hoaxes from those who do not. I begin my empirical analysis with this statistical approach, since it is the traditional methodology used in empirical terrorism studies for the comparative analysis of many terrorist organizations. Whatever statistical relationships I uncover herein will be meaningful guides for subsequent statistical analyses incorporating terrorist hoax activity. They will also provide preliminary insight into which strategic logics appear to be substantiated in empirical cases.

6.1 Hypotheses

This dissertation's hypotheses and corresponding indicators are divided into three categories, relating to group motivation, organizational structure, and campaign context. These hypotheses (and associated variables) were selected to encompass a range of plausible explanations, given the constraints of data availability.⁴⁹ Each hypothesis is designed to represent one or more of the rationalist mechanisms proposed in Chapter 5.

⁴⁹ This approach is consistent with other studies in comparable areas. For instance, in Asal, Ackerman, and Rethemeyer's (2012) analysis of the factors affecting groups' decision to use or pursue CBRN weaponry,

6.1.1 Hypotheses relating to group motivation

As described in section 4.1.2, groups' overarching motivations for terrorism affect the quality and nature of the violence they carry out (Crenshaw Hutchinson 1972; Drake 1998; Hoffman 2006; Rapoport 2002), by determining the range of perceived "legitimate" targets and means (Drake 1998, 59). Accordingly, certain motivations should predict the use of lower-severity forms of violence like hoaxes (as elaborated in section 5.3), whereas others might predict a preference for severely lethal alternatives.

While the literature has not systematically evaluated the targeting logics of all motivations (Drake, for instance, only details the targeting patterns among separatist, extreme right, and communist groups—despite identifying "single-issues" as a relevant category of terrorist ideology), it is logical and morally-consistent that groups with anti-abortionist, animal rights, and/or environmentalist motivations will be more inclined to hoax than groups with more universalist motivations, due to their correspondingly narrow objectives and expected predilection for the protection of life. As described in section 5.3, carrying out more severe forms of violence would be antithetical to the achievement of such groups' aims. The ELF, for instance, is explicit in linking its calls for arson attacks that do not harm any life to the group's stated goal: "to protect all life on this planet" ("If a Tree Falls: A Story of the Earth Liberation Front" 2011). Similarly, the ALF and other violent animal rights groups "are characterized by a deeply held conviction

"[t]he decision regarding variables to be selected for investigation was driven by two concerns—data availability and a desire to examine several variables of different types" (232).

that animals have rights parallel and equal to those of human beings [... including] the right to life and freedom from pain and suffering” (Monaghan 2000, 257). In pursuit of protecting animal rights, these groups avoid causing physical harm to any living being—humans included. By logical extension, the same norm against causing harm to human life should apply to anti-abortionist groups. That anti-abortionist sentiments are the most frequent motivation for CBRN hoaxes (section 3.4; Figure 3.9) suggests that this extrapolation may be empirically founded.

For these groups, then, hoaxes are a morally consistent means of publicizing their cause, spreading fear, and draining responder resources. Hoaxes’ low-severity character confers strategic benefits (moral consistency with overarching ideology and norms) that would be more difficult to achieve using serious forms of violence (which would also risk endangering human life, even where property damage alone is the desired goal). For this reason, Hypothesis 1 predicts that groups motivated by any of these three single-issue motivations will be more likely to hoax than groups with other motivations.

Section 5.2.2 offers a complementary explanation for which single-issue⁵⁰ motivations might lead to an increased propensity to hoax. Gill, Horgan, and Deckert’s (2014) finding that lone actors motivated by single issues are the most likely perpetrators

⁵⁰ Note that these three ideologies do not reflect the full spectrum of groups who may be appropriately classified as “single-issue” terrorists. Even within this study’s sample, there are two groups with motivations that can only be classified as “single-issue,” but do not align with one of these three ideologies: the Justice Guerillas, created to protest the treatment of mentally ill patients in Austrian hospitals; and R.I.S.E., created to kill all residents of Midwestern American cities and to build a new master race made up of group members. The single-issue label, then, is conceptually problematic, since there is little reason to assume that such vastly divergent motivations might lead groups to behave in similar ways. Where the label is applied in this dissertation, it is used for simplicity, and intended only to refer to groups with animal rights, environmentalist, or anti-abortionist orientations.

to engage in dry runs of their activities (432) is plausibly extrapolated to group-level motivations. There is a likely consistency between the ideological/normative mechanism described above and the logic of dry runs, since dry runs are a practice that would increase the likelihood that an attack will be carried out as planned. If a group is casualty-averse (and thus desires more discriminate targeting, as groups with the above-described single-issue motivations are), practice runs will increase the likelihood that their violent activities conform with their desired targeting logic.

*H₁: If a terrorist group is motivated by **anti-abortionist, animal rights, or environmentalist** sentiments, then that group is more likely to hoax.*

Even if groups themselves do not hold ideologically-rooted norms opposing the use of violence against humans, they may prefer to avoid modes of action that generate casualties. As described in section 5.3.2, ethno-nationalist groups are particularly concerned with presenting themselves as legitimate political alternatives to existing powers. In the pursuit of legitimacy, then, ethno-nationalist groups should be inclined to privilege less-severe forms of violence that do not impose high costs on their constituents. Hoaxes offer them a low-cost means of publicizing their cause without accruing high audience costs (in terms of lost legitimacy), since they *disrupt* infrastructure, rather than *destroy* it.

*H₂: If a terrorist group is motivated by **ethno-nationalism**, then that group is more likely to hoax.*

On the opposite side of the spectrum, groups motivated by religious ideologies may be less inclined to use hoax activity. As discussed in section 4.2.1.1, there is a body of literature that evaluates religious groups' particular lethality when they resort to terrorist violence. This literature suggests two primary explanations for religious groups' lethality. First, their transcendental, absolutist worldviews free them from moral constraints against the use of severe violence (Asal, Ackerman, and Rethemeyer 2012; Asal, Brown, and Schulzke 2015, 594; Asal and Rethemeyer 2008, 439; Cronin 2003, 41–42; Hoffman 2006, 88–89; Juergensmeyer 1988). Since they find theological justification for extreme violence and are unconcerned with (earthly) constituents, they are unlikely to object to serious violence on ethical or moral grounds, and are similarly unlikely to view hoaxes' inability to cause death or bodily harm as an advantage worth pursuing. Second—because religious groups essentially act as economic clubs, requiring members to make substantial sacrifices in order to benefit from the group's social service provision—religious groups maintain a select membership base that is unlikely to defect (Berman 2009; Berman and Laitin 2008). This reliable membership makes religious groups more capable of carrying out complex and lethal attacks, so that hoaxes' ease of perpetration does not confer a relative advantage over serious tactics. It also makes religious groups less vulnerable to principal-agent problems that might result in hoax activity. In either

case, hoaxes will confer religious groups little relative advantage, and so they are unlikely to be used.

*H₃: If a terrorist group is motivated by **religion**, then that group is less likely to hoax.*

Hypothesis 3, however, is at odds with Rowlands, Littlewood, and Kilberg's (2012) empirical finding that religious groups hoax CBRN incidents twice as frequently as left-wing and nationalist groups (21).⁵¹ If this finding is confirmed in the present study, it might suggest that religious groups are employing mixed strategies (section 5.2.1)—since they are otherwise so lethal, they can afford to hoax with greater frequency while still retaining credibility—or that they are hoaxing to improve the effectiveness of their serious activities (sections 5.2.2 and 5.2.3).

6.1.2 Hypotheses relating to group structure

As described in section 4.1.1.3, the more hierarchical groups are in their structure, the clearer their lines of communication, command, and control, but also the more vulnerable they are to infiltration (Abrahms and Potter 2015, 318–19; Enders and Sandler 2012, 239; Heger, Jung, and Wong 2012, 745–47; Siegel and Young 2009, 768; Weinstein 2007, 131). As groups become increasingly decentralized, they become more resistant to penetration

⁵¹ In Rowlands, Littlewood, and Kilberg's (2012) study, the authors evaluate only those three motivations (left-wing, nationalist, and religious).

by the authorities, but this improved operational security comes at a cost: severed lines of communication between cells and units reduce their ability to carry out large-scale and sophisticated attacks. Not only are the activities of hierarchical groups more sophisticated than horizontally-structured ones, they are also often more lethal (Heger, Jung, and Wong 2012; Kilberg 2011; as elaborated in section 4.2.1.2). Accordingly, Hypothesis 4 expects that groups lacking hierarchy will be more likely to hoax than groups exhibiting any degree of hierarchy, since such groups are less capable: hoaxes' relative simplicity and low costs will be seen as a major benefit over difficult-to-execute serious attacks.

H₄: If a terrorist group displays **no hierarchy**, it is more likely to hoax.

Furthermore, groups lacking centralized structures will be more vulnerable to principal-agent problems (Heger, Jung, and Wong 2012, 746; Kearns, Conlon, and Young 2014, 428–29) that in turn might generate hoax activity. When groups are entirely decentralized, as is the case when they lack leadership in addition to centralized command and control structures, principal-agent problems are likely to be particularly acute. “Market groups”—lacking leadership, centralized control, and functional differentiation—are the most decentralized type in Kilberg’s (2011) framework of terrorist organizational structure, and reflect what has elsewhere been described as “leaderless resistance” (Beam 1992; Gray 2013). While any group lacking command and control would be vulnerable to agents defecting from the group’s strategy to pursue hoaxes’ low-costs and low-risks (section 5.1), market or leaderless resistance structures

may further encourage hoax activities by encouraging activities to be carried out in the group's name by a wide range of loosely-connected individuals. In such cases, individuals who have not themselves fully radicalized to violence might hoax on behalf of the group, even if it is counter to the group's *modus operandi* (section 5.3.1).

6.1.3 Hypotheses relating to campaign context

A group's origins may determine its subsequent propensity to hoax. When a political group or movement spawns a violent splinter organization, that splinter group is likely to be particularly extreme in its use of violence, since there already exists an established outlet for non-violent expressions of its will.⁵² In desiring to maintain their organizational distinction from the "mother" group, splinter group members will be unwilling to accept the audience costs of reduced perceptions of their credibility. Such groups are thus unlikely to use low-severity tactics like hoaxes, since they offer a lower relative advantage (these groups will not be inclined to use hoaxes as a signal of restraint, as discussed in section 5.3.2).⁵³

H₅: If a terrorist group emerged as a radical splinter or offshoot of a less

violent group or movement, that group is less likely to hoax.

⁵² In practice, however, this may not be the case. While Asal, Brown, and Schulzke (2015) find that splinter history predicts groups' use of violence in general, it is not a predictor of their targeting civilians more narrowly. Tactical intensity (i.e. using violence), however, may be a contributing factor in the incidence of group schism to begin with (Asal, Brown, and Dalton 2012).

⁵³ See section 6.2.3.3 for a discussion of what precisely constitutes a "splinter or offshoot" organization in the context of this dissertation.

Finally, a group's overall tactical repertoire may also determine its propensity to hoax, particularly if the group is attempting to play a mixed strategy (section 5.2.1). According to the principles of game theory, a terrorist group in a mixed game would play two or more strategies—that is, to hoax or to commit some sort of serious terrorism (and/or no terrorist activity at all)—according to a well-defined probability. It is only in cases where the game is repeated (i.e. the group attacks more than once) that it would be strategic for groups to start randomizing between hoaxes and non-hoaxes. Furthermore, the goal of the terrorist group in applying a mixed strategy is to keep the counterterrorists guessing. If the counterterrorists cannot be certain that the attack is a hoax, then they *must* respond as if it were a serious attack. Therefore, the more serious events a group carries out, the higher the probability that any single incident is a serious one, and the more likely a hoax will be to impose costs on society (by virtue of appearing more credible). Moreover, perpetrating a large number of serious incidents signals a group's capability to such a degree that hoax activity would be unlikely to signal weakness and undermine the group's credibility in the same way that it might for a less prolific group.

H₆: The more serious terrorism acts a terrorist group commits, the more likely the group is to also hoax.

6.2 The data

6.2.1 Unit of analysis and sample: Serious perpetrator groups

In line with this dissertation's broader goal of uncovering the strategic logic for hoaxing within serious terrorism campaigns, this chapter's unit of analysis is the terrorist *group*, not individual terrorist actors nor the hoax incidents themselves. Since there are many terrorist groups (ITERATE, for instance, names over 1,200 discreet perpetrators) and coding details of group structure and campaign context is an onerous task, the sample of groups is restricted for feasibility to groups listed as event perpetrators in the WMDDDB. This sample of groups was selected for its manageable size, but also because groups committing CBRN offenses represent a particularly hard case for hoaxing: in having crossed or attempted to cross the unconventional weapons threshold, terrorists with an interest in CBRN weaponry are the most extreme in their tactical choices. Their use of hoaxes, at the complete opposite end of the tactical severity spectrum, is particularly puzzling.

While the WMDDDB lists 131 distinct perpetrator names, not all of these are included in the present sample. Table 6.1 outlines the data culling process, which was designed to generate a sample that is true to the concepts of both "terrorism" and "groups," but also to address some of the challenges discussed in Miller (2016) of transforming events-based databases into group-based ones.⁵⁴

⁵⁴ See LaFree, Dugan, and Miller (2015, 72–74) for a discussion of further "organizational linking problems" that emerge when attributing individual terrorist events to their organizational perpetrators.

Table 6.1 Cleaning the sample

Total # of distinct perpetrators in the WMDDDB =	131
# of perpetrators culled as duplicates (spelling/aliases) =	5
# of perpetrators culled for lacking political- or ideological motivation =	4
# of perpetrators culled for referring to unspecific categories =	4
# of "group names" that do not appear to reflect a real organization =	13
# of perpetrators culled for referring to "lone wolves" =	5
# of perpetrators culled for never having carried out "serious" acts =	25
# of groups remaining in sample =	75
# of groups with dynamic structure =	6
Total # of groups in sample =	81

First, some perpetrators were excluded due to redundancies in the data. These include groups with duplicate names (reflecting alternate spellings or group aliases), as well as appropriate aggregations of incidents attributed to individual members or cells within broader organizations. Second, only perpetrators committing politically or ideologically motivated events are retained. Those committing events with purely criminal or unknown motivations are excluded, since they are not reflective of *terrorist* activity. Third, only defined, organized groups are retained. Perpetrator names that refer to individuals unaffiliated with larger entities, as well as ambiguous group labels (e.g. "Chechen militants") or names for which there is no credible evidence of a corresponding group (i.e. the name was applied as a front for an individual or unorganized group's actions) are excluded, since they do not reflect *terrorist groups* with a coherent organizational structure.⁵⁵ Fourth, only "serious" terrorist groups are retained, since these are the groups who fit the puzzle of hoaxing despite also having the capacity and

⁵⁵ Because this chapter is concerned with group-based perpetrators, this criterion helps to address the variation in perpetrators' "entitativity, or the degree to which a collection of individuals is an entity" discussed in Miller (2016, 7). It requires a degree of organization in pursuit of a common aim across multiple individuals as the minimum threshold for consideration as a group.

motivation to carry out serious attacks. In other words, groups must have committed at least one serious event in the WMDDDB, ITERATE, or the GTD between 1968 and 2012.⁵⁶

Groups therefore do not need to be serious CBRN users (their inclusion in the WMDDDB may be purely for plots, threats, or hoaxes), but they must represent a serious terrorist threat as per any of the three cross-national events databases. The three databases are used to evaluate the groups' serious activities, since none on their own reflect a complete picture of groups' campaigns: the WMDDDB is restricted to only events involving CBRN agents; ITERATE is restricted to only events with a transnational dimension; and the GTD—while encompassing both domestic and transnational events, and events with both conventional and unconventional weaponry—does not explicitly include hoax activity. The result is a sample of 75 distinct groups.⁵⁷ Six of these groups

⁵⁶ The 1968 to 2012 period is selected, since these are the years for which events data are consistently available in ITERATE and the WMDDDB. While the WMDDDB does code for events as early as 1933, its coding is particularly sparse until the 1960s, with only 10 politically- or ideologically-motivated incidents occurring prior to 1968—none of which are hoaxes. ITERATE and GTD (which begins only in 1970) extend to 2016 and 2015, respectively, but their end-dates are restricted to 2012 to remain consistent with the WMDDDB's end.

In the WMDDDB, "serious" events are identified—per Ivanova and Sandler's (2006) culling of cases "for which there was no confirmation of a realistic threat of a CBRN substance" (426-427)—by excluding "plot only" incidents, as well as "hoax or prank" and "threat only" incidents that are not accompanied by the possession of the corresponding CBRN agent. In ITERATE, "serious" events include all incident types except for "threat," "conspiracy," and "hoax"; and incidents with degrees of logistical success other than "aborted by terrorists before initiation" and "incident stopped by authorities at planning stage," since these effectively correspond to plots only. In some cases, ITERATE events are further culled when weapon type is coded as "hoax, no weapon used" if it appears the event truly is a hoaxed one, and not a serious event committed without weaponry. While the GTD explicitly excludes foiled and failed plots as well as "threats to act where no action is taken" due to its requirement that coded attacks "must actually be attempted to qualify for inclusion in the database" (START 2013b), it does sometimes include non-serious, hoax incidents. These are identified (and excluded from serious event counts) using the author's subjective evaluation of event descriptions for incidents coded as involving "fake weapons."

⁵⁷ When cleaning aggregations of terrorism events to their organizational perpetrators, a significant decline in the number of perpetrators is common. For instance, when the GTD initially took over the Pinkerton Global Intelligence Services (PGIS) data that forms its basis, coders "ultimately collapsed the more than 5,000 names into 1,599 unique entries" (LaFree, Dugan, and Miller 2015, 72): a more than three-times reduction that is far more substantial than the sample culling performed here.

have dynamic structures, and so they are split to reflect their changed organization over multiple time periods, resulting in a total n of 81 groups. For five groups, I was unable to acquire sufficient details for all independent variables. Accordingly, the bulk of the study is conducted on a sample of $n = 76$; and sometimes on a further-reduced sample of $n = 71$, for reasons elaborated in section 6.3.

6.2.2 The dependent variable: Hoax = yes/no?

The dependent variable is a binary one, identifying groups that never use hoaxes (*hoaxer* = 0), and those that have hoaxed at least once (*hoaxer* = 1). A binary variable is preferred to a continuous variable counting the number of hoaxes or the proportion of hoaxes to serious events in the present analysis because of inconsistency in the coding of coordinated campaigns, and because—before it can be established why some groups hoax more than others—it is important to establish why some groups hoax at all.

Events data from the WMDDDB, ITERATE, and GTD are used to determine whether a group hoaxes.⁵⁸ All three sources of events data are used to determine the hoaxer variable for the same reason that all three are consulted to confirm groups' serious activity: each on its own reflects only certain subcomponents of the group's overall campaign. For events in the WMDDDB and ITERATE, the designation is straightforward: the WMDDDB codes hoaxes and pranks as a possible event type, and ITERATE codes hoaxes as

⁵⁸ This coding of a binary independent variable is analogous to that used by Asal, Ackerman, and Rethemeyer (2012, 238) to distinguish groups who have ever used or attempted to develop a CBRN capability in their period of study. The authors married a sample of groups derived from the TKB to CBRN data from the WMDDDB.

a possible incident type (Mickolus et al. 2012). However, ITERATE includes additional incidents which, by the author's subjective evaluation, were recoded as hoaxes for determining the dependent variable and tactical repertoire variable discussed below. ITERATE codes a broad range of incident types other than hoaxes with "hoax, no weapon used" as the type of weapon. Hoaxes-as-incident-type account for only 20% of these events. While it is plausible that certain incident types (i.e. the non-serious incident types described above, as well as thefts; arms smuggling; occupations of facilities; barricades; hijackings; and even assassinations) *could* be carried out without weapon use, other incident types cannot. Armed attacks, "exotic pollution," and the various forms of listed bombings by definition involve some sort of weapons use; therefore, the coding of weapon type as "hoax, no weapons used" implies that these incidents are really hoaxes. Similarly, while the GTD excludes hoaxes from its possible event types, it does sometimes include them indirectly, through the "fake weapons" weapon type. These incidents are identified as hoaxes using the author's subjective evaluation of the weapon/event type relationship, and available details in the event summary. In some cases, the GTD's event summaries refer to hoax activity that is not otherwise accounted for in any of the datasets; the hoaxer variable accounts for these hoaxes, despite their lack of inclusion as explicit events.

Accordingly, the dependent variable measures whether there is any evidence that a group in the sample of those groups with a demonstrated interest in CBRN weaponry has ever used a terrorism hoax—CBRN-related, or otherwise. Of the 81 groups in the sample, 23 (or 28.4%) are coded as groups that hoax (*hoaxer* = 1).

6.2.3 Independent variables

Unless otherwise noted, the study's explanatory variables have been coded independently by the author, based on open-source data. For all variables pertaining to group characteristics (motivation, structure, and history/origins), the primary sources of data are group descriptions from Terrorist Organization Profiles (TOPs)—developed under the Memorial Institute for the Prevention of Terrorism (MIPT)'s Terrorism Knowledge Base (TKB)—and Jane's World Insurgency and Terrorism—developed by IHS Markit—as well as any information on the perpetrators provided in their corresponding event descriptions (available for all WMDDB events and some in GTD). Where this information is insufficient, secondary sources were consulted, privileging in the following order: government reports, peer-reviewed academic resources, and news archives. For consistency with existing scholarship, coding decisions (where samples overlap) were checked against those assigned in both the Big Allied and Dangerous Database I (BAAD-1) Lethality Data (Asal, Rethemeyer, and Anderson 2011) and Jones and Libicki's (2008) organizational data.⁵⁹ Descriptions of the variables are provided in Table 6.2, with coding decisions presented in Appendix Table A.1.

⁵⁹ In section 10.1.1, I discuss how results are frequently contingent on the sometimes-subjective coding decisions that scholars make. While more nuanced case analyses may incorporate more fine-grained characteristics, the coding decisions here privilege consistency with existing data sources, since a primary goal of this study is to test the limits of existing data.

Table 6.2 Description of group-based variables for regression analysis

VARIABLE	DEFINITION	DESCRIPTION (% = 1)
DEPENDENT VARIABLE		
<i>hoaxer</i>	= 1 if group has ever conducted a hoax; 0 otherwise	28.4%
MOTIVATION		
<i>single-issue</i>	= 1 if the group has an anti-abortionist, animal rights, or environmentalist motivation; 0 otherwise	6.2%
<i>purely religious</i>	= 1 if the group is solely motivated by religion; 0 if the group has any other motivation (even if there is a religious component)	27.2%
<i>contains religion</i>	= 1 if the group's ideology contains a religious component; 0 otherwise	49.4%
<i>purely ethno-nationalist</i>	= 1 if the group is solely motivated by ethno-nationalism; 0 if the group has any other motivation (even if there is an ethno-nationalist component)	11.1%
<i>contains ethno-nationalism</i>	= 1 if the group's ideology contains an ethno-nationalist component; 0 otherwise	39.5%
<i>religious ethno-nationalist</i>	= 1 if the group's ideology is both religious and ethno-nationalist; 0 otherwise	18.5%
<i>economically leftist</i>	= 1 if the group's ideology contains an economically left-wing component; 0 otherwise and/or if also <i>single-issue</i>	19.8%
<i>broadly left</i>	= 1 if the group's ideology contains a broadly left-wing orientation, including animal rights and environmentalism; 0 otherwise	27.2%
<i>economically rightist</i>	= 1 if the group's ideology contains an economically right-wing component; 0 otherwise and/or if also <i>single-issue</i>	6.2%
<i>broadly right</i>	= 1 if the group's ideology contains a broadly right-wing orientation, including anti-abortionist sentiments and racial supremacy; 0 otherwise	12.4%
STRUCTURE		
<i>leadership</i>	= 1 if the group has a clear leader/leadership structure, including a leadership council; 0 otherwise	80.3%
<i>functional differentiation</i>	= 1 if there is evidence that the group involves at least one clearly-defined cell or department dedicated to a specialized function (e.g. political wing, bomb centre, communications department)	46.1%
<i>command and control</i>	= 1 if there is evidence that the group's actions are coordinated from a central executive; 0 otherwise	29.0%
<i>hierarchy</i>	= 1 if the group displays any form of hierarchy or centralized command and control (<i>leadership</i> and/or <i>command</i> = 1); 0 otherwise	81.6%
CAMPAIGN CONTEXT		
<i>splinter</i>	= 1 if the group originated as a radical splinter/offshoot of a less-violent group or movement; 0 otherwise	24.7%
<i>number of serious events, logged</i>	Natural logarithm of the number of serious events perpetrated across the WMDDb, ITERATE, and GTD	Range: 0 – 7.82 Mean: 3.38 Std. Dev.: 2.25

6.2.3.1 Measures of group motivation

A range of explanatory and control variables are used to test Hypotheses 1, 2, and 3, pertaining to group motivation. The *single-issue* explanatory variable indicates whether a group has either an anti-abortionist, animal rights, or environmentalist orientation. The “single-issue” label is used for ease of reference to encompass these three motivations, although it should be noted that it does not encompass all single-issue or “special interest” terrorist activity (e.g. the Justice Guerrillas are not coded as *single-issue* = 1, even though their motive is given by the WMDDDB as “other single-issue: to protest the treatment of mentally ill patients in Austrian hospitals”). In other datasets, these perpetrators are usually classified as right-wing (anti-abortionist) or left-wing (animal rights or environmentalist). Motive variables *broadly left* and *broadly right* are included as controls to capture these traditional left- and right-wing motivational classifications, encompassing groups motivated by economic ideology as well as anarchist/environmentalist/animal rights and supremacist/anti-abortionist events respectively. The *economically leftist* and *economically rightist* controls capture groups’ economically left- and right-wing orientations, respectively (single-issue groups are excluded). This coding is consistent with the CIDB’s coding of motivation. In the CIDB, leftist organizations are those who advocate “economically leftist policies such as redistribution of wealth by the government and nationalization of industry,” whereas rightist organizations advocate such approaches as “extreme laissez-faire policies, small government, or policies that favour status quo economic elites” (TSAS 2015, 3–4).

Two dummy variables, *contains religion* and *purely religious*, are used to measure whether groups are motivated by religion, and to what degree. Following the coding norms outlined in the BAAD-1 dataset, *contains religion* is used when a group's ideology contains—but is not necessarily restricted to—a religious dimension, and *purely religious* is used when the group's ideology is exclusively religion-oriented. Similarly, *contains ethno-nationalism* and *purely ethno-nationalist* are used to indicate whether a group's ideology contains an ethno-nationalist element or is exclusively ethno-nationalist in orientation. The *religious ethno-nationalist* control is used to indicate groups that include both religious and ethno-nationalist ideological components.

Except for the variables indicating purely religious or ethno-nationalist motivations, this study's motivational variables are non-exclusive; groups may be coded with multiple motivations, where applicable, to reduce sensitivity to coder judgment in determining a so-called "primary" motivation. The total set of variables is also non-exhaustive, in that groups with other single-issue motivations ("Justice Guerrillas" and "R.I.S.E.") receive zeros across all defined motive categories.

6.2.3.2 Measures of group structure

Group structure, used to test Hypothesis 4, is measured per the methodology elaborated in Kilberg's (2011) dissertation on the effects of organizational structure on terrorist group behaviour, and expanded to the most prolific CBRN users among the WMDDB

sample in Rowlands, Littlewood, and Kilberg (2012).⁶⁰ Structure is coded along three binary dimensions: *leadership*, which indicates whether the group has a clearly defined leader or leadership structure, including a leadership council; *functional differentiation*, which indicates whether the group maintains “at least one clearly separate department or cell that is dedicated to a specialized function” (Kilberg 2011, 32); and centralized *command and control*, which indicates whether there is evidence that the group “coordinates its actions (violence, communication and/or politics) from a central executive” (Kilberg 2011, 32).⁶¹ While all three of these structural features are tested as controls (and as a subsidiary test of the degree of hierarchy on groups’ propensity to hoax), the key explanatory variable, *hierarchy*, indicates whether the group exhibits any form of leadership or centralized control.

While the coarse hierarchy indicator necessarily removes a significant degree of nuance from the analysis, I follow Heger, Jung, and Wong’s (2012) logic in doing so: “for analytical clarity, and also to highlight the tradeoff between extreme organizational forms” (746-47). Furthermore, a focus on this coarse structural distinction helps to hedge against skepticism that functional forms can be easily distinguished. While it is true that some groups may exhibit different structural aspects simultaneously, or vacillate between them, the presence (or absence) of any leadership structure appears to be a relative

⁶⁰ Of the 81 groups in the present study’s sample, 41 retain original structure coding from Rowlands, Littlewood, and Kilberg (2012).

⁶¹ In Kilberg’s (2011; 2012) framework, these three features combine to reflect three organizational types: market structure (*leadership* = 0; *functional differentiation* = 0; *command and control* = 0); all-channel (*leadership* = 1; *functional differentiation* = 0; *command and control* = 0); hub and spoke network (*leadership* = 1; *functional differentiation* = 1; *command and control* = 0); and bureaucratic structure (*leadership* = 1; *functional differentiation* = 1; *command and control* = 1).

constant. Furthermore, as Heger, Jung, and Wong (2012) note: “understanding extreme forms is the first step to understanding how organizational structure affects the differences in lethality [hoaxes, in this case] between violent groups” (747).

6.2.3.3 Measures of campaign context

To test Hypothesis 5, regarding groups’ historical origins, the *splinter* variable indicates whether the group originated as a radical splinter or offshoot of a less-violent group or movement. While some literature on splinter organizations addresses only bifurcations within existing terrorist groups (e.g. Bueno de Mesquita 2008), here I accept a broader definition of the term: splinter groups may originate from other terrorist organizations, political parties, or even broader social movements.⁶² Since all groups have ideological origins somewhere, however, I restrict the notion of splintering from social movements to those cases where there is some degree of organization in the “mother” movement,⁶³ and a traceable continuity of membership participation between the broader movement and the terrorist group. In those cases where the source organization is another terrorist group, a group only qualifies as a *radical* splinter if the original organization appears to be

⁶² See, for instance, Weinberg, Pedahzur, and Perliger’s (2009)’s discussion of the relationships between political parties and terrorist groups. While the authors enumerate seven relationship types, three are relevant here. First, a political party may deliberately support the creation of a “subsidiary terrorist group in order to carry out the party’s goals by violent means” (Weinberg, Pedahzur, and Perliger 2009, 22). Second, a political party may experience factional division, driven by individual leaders’ personal qualities, regional differences, disagreements over party doctrine, attitudes toward compromise, and even “generational rebellion” (Weinberg, Pedahzur, and Perliger 2009, 24). Third, a social movement may simultaneously generate both political parties and terrorist groups (Weinberg, Pedahzur, and Perliger 2009, 27).

⁶³ This requirement for organization is similar to coding decisions taken elsewhere, where a “split” is understood as “when new organizations emerge out of the old ones” (Asal, Brown, and Dalton 2012, 106).

in a period of low or diminishing serious terrorist activity at the time of the new organization's emergence; otherwise, the hypothesized mechanism should not apply.

The explanatory variable measuring groups' tactical repertoire, to test Hypothesis 6, is the natural logarithm of the *number of serious events* carried out by each group, cross-referenced across all three of the WMDDDB, ITERATE, and GTD to remove duplicates.⁶⁴ I use the logged (rather than linear) functional form for this variable because the basic count variable is widely dispersed (range: 1 – 2,492), and unlikely to yield a constant marginal effect. From a theoretical perspective, for instance, the effect of increasing from 2 to 12 serious attacks should have a far greater influence on a group's likelihood of hoaxing than increasing from 2,000 to 2,010.

6.3 Models and methods

I use probit regression to estimate the effect of each hypothesized predictor on the dependent variable of whether or not a terrorist group ever uses hoaxes. The variable indicating the single-issue motivations, however, is a perfect predictor of the hoax outcome, and its associated observations are dropped from this basic model's estimation. As the discussion below will illustrate, these single-issue groups also suffer from another confounding characteristic: they are all non-hierarchical. It is unclear without further analysis whether these groups are theoretically distinct from the other groups in the sample (and thus *should* be excluded from the estimation), or whether all terrorist groups

⁶⁴ To account for variation in coding procedures across datasets, duplicates are identified when groups' events involve similar weapon types and location, and occur within two days of one another.

are governed by the same basic theory, but single-issue motivation just happens to pose an empirical challenge in the estimation. Accordingly, I specify the most parsimonious model for each sample of terrorist perpetrators (with or without the single-issue groups), and then test each model against the alternate sample, to distinguish how much of the result is linked to the observations in the sample versus the predictive value of the independent variables themselves. Model 1A reflects the most parsimonious model for the sample of 71 groups (estimated for the sample of 76 in Model 1B), and Model 2A reflects the most parsimonious model for the sample of 76 groups (estimated for the sample of 71 in Model 2B). In excluding single-issue groups from the sample, Models 1A and 2B effectively treat single-issue groups as theoretically distinct terrorist perpetrators. Alternatively, Models 1B and 2A reveal the important factors associated with hoaxing across the entire sample of perpetrators (including those with single-issue motivations), with the constraint that we cannot observe the hypothesized effect of being a single-issue group.

Models 1A and 1B include five indicators, one per each of Hypotheses 2-6: purely ethno-nationalist motivation; purely religious motivation; hierarchical structure; radical splinter origins; and the logged number of serious events perpetrated per group. As described above, Hypothesis 1 regarding single-issue motivation cannot be tested directly in the regression. Because the single-issue indicator cannot be estimated, Models 2A and 2B use the broadly left-wing motive control in place of the single-issue indicator, to test whether it is the hypothesized normative dimension of the specified single-issue groups that affects the hoax outcome, or the broader left-wing orientation to which these single-

issue groups are traditionally assigned. I included the broadly right-wing motive control in an alternate specification of this model, but it never yielded a significant or approaching-significant estimate; it is thus excluded from the parsimonious Models 2A and 2B.

In Table 6.3, regression coefficients are reported as average marginal effects to facilitate the interpretation of magnitude.⁶⁵ Marginal effects, expressed as a percent, reflect the change in the probability that a group will use hoaxes ($y=1$), given a one-unit change in the independent variable. For binary independent variables, this means that the marginal effect reflects the increase (or decrease) in the probability that $y=1$ when the independent variable changes from $x=0$ to $x=1$. To measure the models' goodness of fit, the percent of correctly predicted values as well as McFadden's pseudo R-squared (based on the probit coefficients, prior to the models' transformation to marginal effects) are reported. Average marginal effects for models testing alternate indicators for group structure are reported in Appendix Table A.2.

This analysis was conducted using the Stata 14 statistical software package.

6.4 Results and discussion

In the subsections that follow, I present and discuss the findings for each hypothesis and corresponding independent variable.

⁶⁵ Average marginal effects, which take the average of the individual marginal effects, are preferred over marginal effects at the means, because of the preponderance of binary variables being tested.

Table 6.3 Average marginal effects for probit models

	Model 1A	Model 1B	Model 2A	Model 2B
<i>purely ethno-nationalist</i>	0.426*** (1.334)	0.359*** (0.140)	0.521*** (0.168)	0.481*** (0.151)
<i>purely religious</i>	0.263*** (0.074)	0.233*** (0.085)	0.393*** (0.117)	0.317*** (0.102)
<i>broadly left</i>	--	--	0.231* (0.123)	0.100 (0.124)
<i>hierarchy</i>	-0.344*** (0.112)	-0.516*** (0.074)	-0.554*** (0.087)	-0.366*** (0.120)
<i>splinter</i>	-0.350*** (0.134)	-0.199* (0.104)	-0.242** (0.111)	-0.346*** (0.133)
<i>number of serious events, logged</i>	0.050*** (0.018)	0.058*** (0.018)	0.052*** (0.018)	0.047*** (0.018)
Observations (<i>n</i>)	71	76	76	71
Pseudo R-squared	0.3530	0.3289	0.3688	0.3615
Percent correctly predicted values	83.10	78.95	81.58	83.10

Note: Coefficients and standard errors (reported in parentheses) have been transformed to reflect average marginal effects. The constant is not shown. The dependent variable is the dummy *hoaxer*, indicating whether (=1) or not (=0) a group ever uses hoaxes.

***, **, * indicate statistical significance at the 0.01, 0.05, and 0.10 levels, respectively, for a one-tailed test.

As described above, Model 1A is the most parsimonious model for the sample excluding single-issue groups, but it does not perform well when tested against the full sample of 76 perpetrators. Clearly, there is an effect that the remaining five indicators fail to capture. When Model 2B adds broadly left-wing orientation to the estimation, it does not have a statistically significant coefficient estimate; the model's slightly larger pseudo R-squared is thus likely attributed only to the presence of an additional regressor (indeed, the percent of correctly predicted values remains constant, at 83.10%). However, on the sample of 76 perpetrators, the broadly-left indicator in Model 2A leads to a substantially improved goodness-of-fit over Model 1B, although it still performs worse than either model estimated on the sample of 71. With the exception of the indicator for groups' radical splinter origins, which only approaches significance in Model 1B, the explanatory

variables are robust across both sample specifications, and so conclusions are extrapolable to all terrorist perpetrators—including those with a single-issue motivation.

6.4.1 Animal rights, environmentalist, and anti-abortionist motivation

Hypothesis 1 posits that the single-issue groups motivated by animal rights, environmentalist, and/or anti-abortionist interests will be more likely than others to hoax. The probit results support this hypothesis. First, the single-issue indicator perfectly predicts the hoaxing outcome. This perfect prediction occurs because of a relationship of sufficiency between the specified single-issue motivations and hoax behaviour (see Appendix Table A.3): for the groups included in the present sample, having environmentalist, animal rights, or anti-abortionist motivation is sufficient for predicting that a group will use hoaxes. In other words, in all cases where one of those three single-issue motivations is present, so too is the decision to hoax. Since the regression will not run with a perfect predictor, the single-issue variable and the five associated observations are dropped from the model's estimation.⁶⁶

Second, economically-oriented left- and right-wing ideologies do not produce a statistically significant effect on the hoax outcome, suggesting that any effect amongst the single-issue groups is attributed to their more narrow animal rights, environmentalist, or anti-abortionist orientation. The broadly right-wing indicator never approaches significance, and so it is not presented in Model 2. The broadly left-wing indicator

⁶⁶ As an additional test, to include all 76 observations in the estimation, I ran Model 1 using an ordinary least squares (OLS) regression. The coefficient for *single-issue* motivation is 0.733, at a significance level of $p = 0.001$, indicating a robust, positive relationship.

approaches significance at the $p < 0.1$ level in Model 2A (where four of the five single-issue groups are classified as broadly left), but loses all significance when the single-issue groups are excluded from the sample. Economically leftist groups on their own do not demonstrate a net increased likelihood of hoaxing over single-issue leftists (animal rights or environmentalist actors). Accordingly, the distinction between general Models 1 and 2 supports the norm-based mechanism underpinning the hypothesized positive relationship between single-issue motivation and hoaxing, and suggests that models including the single-issue groups but not controlling for their motivation are misspecified.

Complicating the acceptance of Hypothesis 1 (which posits a positive relationship between the single-issue motivations and hoaxing), however, is the relationship between the single-issue variable and the indicator for hierarchy: all of the single-issue groups also happen to lack any form of hierarchy (see Appendix Table A.4). It is because of this confounding relationship that the negative marginal effect of hierarchy increases so substantially in models where the single-issue groups are included. Since the sample contains no examples of single-issue groups that are hierarchically structured, it is impossible to determine using the regression whether their hoaxing results from some feature of their ideology and corresponding norms against violence (although the improved fit of Model 1A over 1B implies that ideology does play a role), or if it is primarily driven by these groups' market structure and resultant principal-agent dynamics. The effect of single-issue motivation is over-determined.

In the following sub-section, I address the effect of hierarchy on hoaxing to better explore this dynamic. Subsequent sub-sections address the other hypotheses in their respective order.

6.4.2 (Non-)hierarchical structure

Hypothesis 4 posits that groups lacking hierarchical structure will be more likely than others to hoax. This hypothesized negative relationship between hierarchy and hoaxing holds consistently across all models, lending strong support to Hypothesis 4. When leadership, functional differentiation, and centralized command and control are tested as alternate indicators of groups' organizational structure (Appendix Table A.2),⁶⁷ the coarse hierarchical/non-hierarchical indicator is associated with a larger marginal effect than any of its counterparts: being hierarchically structured decreases a group's likelihood of being a hoaxer by a minimum of 34.4%. Leadership offers the second-largest marginal effect, decreasing the likelihood of a group ever hoaxing by 27.7%; followed by functional differentiation, which decreases the likelihood of a group ever hoaxing by 26.1%. Centralized command and control has the lowest marginal effect (reducing the likelihood of hoaxing by 17.7%), and is only statistically significant at the $p < 0.1$ level. These tests demonstrate that it is the presence of a minimal level of hierarchy that matters most for

⁶⁷ Note that of the six groups coded with dynamic structures (Abu Sayyaf Group [ASG], al-Qaeda, Aum Shinrikyo, ETA, Khmer Rouge, and Students Islamic Movement of India [SIMI]), only two (ASG and SIMI) varied at the basic level of hierarchy (in their second periods, they both lost leadership and attained fully decentralized market structures). For the other groups, structural changes are reflected in functional differentiation and/or command, meaning that they do not actually vary with respect to the hierarchy indicator tested in Models 1-4. They are retained as distinct groups in the sample anyway, since their structural changes coincide with distinct moments of discontinuity (sometimes accompanied by name changes and/or rebranding) which justifies their treatment as separate units.

predicting hoax behaviour. Particular structural features (i.e. leadership, specialized functions and/or centralized control) do not make groups, on average, any less likely to hoax than their minimally hierarchical counterparts characterized either by identifiable leaders *or* centralized command and control.

Because all of the sample's single-issue groups are also non-hierarchical, leaving those groups in the sample but not controlling for their ideological orientation (as in Models 1B and 2A) leads to a substantial increase in the negative marginal effect of hierarchy. In these models, hierarchy reduces the likelihood of hoaxing by between 51.6 and 55.4%. As discussed above, however, this effect is likely exaggerated, since it appears that single-issue motivation does play a role (of an indeterminate size) in determining groups' hoax use.

The negative relationship between hierarchy and hoaxing supports a range of strategic logics for hoax behaviour. On the one hand, hierarchical groups are more capable of sophisticated attacks (because they have greater financial and/or human resources). In this scenario, hoaxes' low-cost requirements engender a lower relative advantage (section 5.1), thus discouraging their use over alternative tactics. On the other hand, hierarchical groups are the least vulnerable to principal-agent problems, and so their members are less likely to defect from the group's overarching strategic logic to independently employ hoaxes in pursuit of their lower relative risk (section 5.1) or for their lower severity (as an outlet for incomplete radicalization to violence; section 5.3.1). From a counterterrorism perspective, then, activities of hierarchical groups should be

taken particularly seriously, since these groups are 26.2 to 46.5% less likely than others to use hoaxes.

6.4.3 Ethno-nationalist motivation

Hypothesis 2 posits a positive relationship between ethno-nationalist motivation and a groups' likelihood of hoaxing. This hypothesis is tested using the indicator measuring purely ethno-nationalist ideology, and I find consistent and robust support for my hypothesis: being purely motivated by ethno-nationalism increases a group's likelihood of being a hoaxer by between 32.2 and 47.0%. Importantly, it is the *pure* ethno-nationalist ideology that matters: other versions of the variable that indicate ethno-nationalism as just one component in a mix of motivations do not have a statistically significant effect. The theorized mechanism underpinning Hypothesis 2—that ethno-nationalist groups hoax in order to signal restraint and/or to intentionally cause disruption (not destruction) in their pursuit of legitimacy amongst constituents (section 5.3.2)—holds only with respect to ethno-nationalists that are dealing with an exclusively terrestrial audience (not in addition to a spiritual or “other-worldly” audience, as with religious groups), and not pursuing an economic agenda.

6.4.4 Religious motivation

Hypothesis 3 posits a negative relationship between religious ideology and a group's likelihood of hoaxing. While the marginal effect of having a purely religious ideology on groups' hoax behaviour is consistently strong, its direction is opposite to that

hypothesized: the probit estimations indicate a *positive* effect of purely religious motivation on a group's propensity to hoax, increasing the probability of hoaxing by between 23.3 and 39.3%. As with ethno-nationalist motivation, the relationship holds only for purely religious groups. When religion is combined with ethno-nationalism and/or left- or right-wing economic policies, any statistically significant result disappears; the presence of a secondary motive moderates religion's positive effect on hoaxing.

The counter-intuitive positive marginal effect of religion may be spurious, attributable to coders' disproportionate attention paid to religious group activity in the so-called age of "new terrorism" and thus a tendency to code more of the events associated with them. It could also be the case that unaffiliated individuals carry out false flag hoaxes in the name of religious groups, to exploit their lethal reputations. It may also, however, be reflective of a different strategic logic for hoaxing. When groups are otherwise highly capable and/or lethal, as the literature suggests to be true of religious groups (see section 4.2.1.1), they can afford to hoax with little reputational damage. Playing mixed strategies (section 5.2.1), groups can achieve high payoffs precisely *because* they are otherwise so capable, while expending very little in terms of human or financial resources. In this sense, and in line with the literature, religious orientation may be interpreted as a proxy for group capability and effectiveness (the indicator for the number of serious events perpetrated per group speaks to perpetrators' ability to successfully conclude serious acts, but says nothing of their effectiveness in doing so).⁶⁸ Conversely,

⁶⁸ It does not appear that the religion indicator is in any way capturing a non-linear effect of the number of serious events indicator. First, the variables have very low correlations, irrespective of functional form for the number of serious events variable: with a linear count, the correlation with the purely religious

religious groups may also be hoaxing to practice serious tactics they wish to implement in the future (section 5.2.2) or to burden counterterrorism efforts by providing more and noisy signals of activity (section 5.2.3). Whatever the explanation, being purely motivated by religion increases a group's likelihood of hoaxing, and vice versa.

6.4.5 Radical splinter origins

Hypothesis 5 posits that groups who emerged as radical splinters or offshoots of less violent groups or movements will be less likely to hoax. Across both models and samples, the variable indicating groups' radical splinter origins generates a negative marginal effect on hoaxing, although to varying degrees of significance. It has the largest marginal effect (decreasing the likelihood of hoaxing by between 34.6 and 35.0%) in the estimations that exclude single-issue groups from the analysis. When the single-issue groups are included in the estimation, the effect of splinter histories is smaller, and not so robust. In Model 1B, the marginal effect of a group's splinter origins only approaches significance at the $p < 0.1$ level. In Model 2A (when the single-issue motivations are partially controlled for by the broadly left indicator), however, it regains significance at the $p < 0.05$ level, with a negative marginal effect of 24.2%. This change in significance is explained by the radical splinter histories of the single-issue groups, and their corresponding broadly left- or right-wing affiliation: three of the five single-issue groups are radical splinter organizations, and

ideology indicator is only -0.025; for the quadratic number of events squared, the correlation is only 0.012; and for the logarithmic form, the correlation is only -0.044. Additionally, when I test an explicit interaction term (the binary purely religious ideology indicator times the natural log of the number of serious events perpetrated), the coefficient estimate is insignificant.

all three of these are also broadly leftist. When the motive of these groups is not otherwise controlled for, their presence in the estimation thus undermines the negative effect of splinter histories. In other words, splinter origins appear to make groups less likely to hoax, but only robustly-so when groups are not also motivated by animal rights, environmentalist, or anti-abortionist ideologies.

From a counterterrorism perspective, this means that activities of radical splinter groups should be perceived as credible. From a theoretical perspective, splinter groups are unlikely to hoax because they are not willing to accept the audience costs of reduced perceptions of their credibility. There already exists a non- (or lesser-) violent outlet for furthering their cause, and so members that choose to partake in the splinter's activities are both unlikely to defect from their new organization (and hoax via principal-agent divergences) and unlikely to find hoaxes' lower severity (section 5.3) attractive.

6.4.6 Number of serious events perpetrated

Hypothesis 6 posits that the more serious terrorism acts a terrorist group commits, the more likely that group will be to hoax. Consistent with this hypothesis, the coefficient estimate for the logged number of serious events a group perpetrates is strongly significant, and in the hypothesized direction: the more (fewer) attacks a group carries out, the more (less) likely it is to hoax. Because the natural logarithm of each group's

serious event count is estimated, this means that doubling the number of serious events perpetrated increases the likelihood of that group hoaxing by between 4.7 and 5.8%.⁶⁹

However robust these marginal effects are, they are substantively minor when compared with the magnitude of other variables' marginal effects. They may also be spurious, due to the media-based nature of terrorism datasets. As described in section 5.4.1, attribution of terrorist events to their perpetrators is difficult. Incidents may be claimed in the name of a group and coded accordingly, but this is no guarantee that the group's core decision-making body sanctioned the activity. If groups are prolific, a prankster wishing to spread fear or disturb the peace may choose to freeride on the group's serious reputation and exploit it to achieve maximum payoff from a not-really-affiliated hoax. Conversely, the more prolific a group is, the more likely news sources are to follow its activities and capture any hoaxes it commits. Hoax activities of a group that is poorly-known may not be on coders' radars. If the result is reflective of a true correlation between the number of serious incidents perpetrated per group and the group's decision to hoax (or not), however, the direction and significance of this coefficient estimate lends support to the strategic logic of hoaxing to support mixed strategies.

⁶⁹ I also tested alternate functional forms for the number of serious events perpetrated variable: a linear version (a simple count of the number of serious events perpetrated), and a quadratic version (the number of serious events squared). Although results were similar, the logged version resulted in consistently larger pseudo R-squareds across all models, and yielded a more parsimonious and easily interpretable marginal effect. Furthermore, as elaborated in section 6.2.3.3, the logarithmic form has greater theoretical resonance.

6.5 Conclusion: The strengths and limitations of a regression-derived answer

In this chapter, I presented results from a probit regression that estimated the effect of various motivational, structural, and campaign contextual indicators on whether or not a terrorist group ever carries out hoax activities. The regression results indicate two conditions that are, on average, associated with a lower likelihood of a group ever employing hoaxes: all else held equal, when a group is hierarchically structured, it is 34.4 to 55.4% less likely to ever hoax; and when a group originated as a radical splinter of a less-violent group or movement, it is 19.9 to 35.0% less likely to ever hoax. Conversely, a range of factors indicate an increased likelihood of a group ever employing hoaxes. The predictor with the largest effect is whether the group is motivated by either animal rights, environmentalist, or anti-abortionist “single issues,” although its effect is confounded by the effect of non-hierarchy. Without being able to disaggregate the effects of a single-issue motivation and non-hierarchy, the single-issue indicator perfectly predicts the hoax outcome: when a group espouses one of these single-issue motivations, there is a 100% likelihood that it will at some point employ hoaxes. To a lesser degree, having a purely ethno-nationalist orientation makes groups 35.9 to 52.1% more likely to employ hoaxes at some point throughout their campaign; and having a purely religious ideology makes groups 23.3 – 39.3% more likely to employ hoaxes at some point. When ethno-nationalist or religious ideology are combined with secondary motivations (including a joint religious ethno-nationalist ideology), they have no statistically significant effect on groups’ hoax behaviour. Additionally, the probit results suggest that a 100% increase in a group’s total number of serious events perpetrated increases its likelihood of also using hoaxes by 4.7

to 5.8%. In all but the case of religion, these statistically significant relationships are in the hypothesized direction.

While the regression results are robust, their practical significance is limited by deficiencies in the sample. First, the sample size is relatively small, and so certain relationships (particularly the perfect prediction of the single-issue indicator) may be over-stated. Second, the sample is limited only to those groups with a demonstrated interest in CBRN weaponry. Accordingly, the findings can only be said to hold amongst this small subset of terrorism perpetrators; hoaxers who rely exclusively on conventional weaponry may exhibit different characteristics. Third, the dependent variable is coded based on events captured in media-based terrorism events datasets. There is thus a likely bias toward capturing hoax events perpetrated by particularly prolific and/or otherwise fearsome groups—a bias which may contribute to the significant, positive coefficient estimates on the number of serious events perpetrated and purely religious ideology variables.

The findings presented in this chapter are also vulnerable to methodological limitations of probit regression, and the correlational logic that underpins it. First, coefficient estimates reflect the influence of each variable, all else held equal, across all cases. While this “on-average” approach may be valuable in the counterterrorism context (we frequently do not have a comprehensive group profile at the moment of attack, and so individual identifiers may be useful), it is also limiting: there may be contexts in which certain variables combine to affect a group’s likelihood of using hoaxes, contradicting the *average* marginal effects uncovered here. Secondly (and perhaps more importantly), the

regression assumes symmetric correlations between each predictor and the dependent variable, which are not always substantiated. As an example, while single-issue motivation perfectly predicts hoaxing, the absence of a single-issue motivation does not come close to predicting non-hoaxing (see Appendix Table A.3). Chapter 7's re-interpretation of the hypotheses (section 7.2) illustrates how many of the relationships tested in the present analysis are in fact asymmetrical. Average marginal effect estimations, by virtue of assuming symmetry, thus over- or under-estimate the influence of certain variables.

In the next chapter, I use QCA to illustrate how alternative methodological approaches can provide insight into the nuanced effects of terrorist group organizational characteristics on hoax behaviour; and how conventional statistics may be deficient in tackling such complex phenomena. In section 7.6, I elaborate upon—and in some cases amend—the probit-derived conclusions described above.

7. Characteristics of likely hoaxers: Insights from set theoretic relations

In Chapter 6, I conducted a probit regression analysis on a sample of hoaxers and non-hoaxers to identify organizational characteristics that differentiate terrorist groups who use hoaxes from those who do not. In this chapter, I use the same binary dependent variable (whether a group ever uses hoaxes in its campaign activities, or not) and analogous data on terrorist groups' organizational characteristics, but use an alternative method of analysis: crisp and fuzzy set qualitative comparative analysis (csQCA and fsQCA, respectively). Whereas the regression revealed the net marginal effect of each independent variable with all else held equal, QCA is an exploratory technique that searches for causal combinations of factors (i.e. pairings or groups of factors that are collectively associated with the outcome). As a function of their differing underlying logics, regression uncovers correlational relationships, whereas QCA identifies necessary and sufficient conditions. Whether these conditions include combinations of factors or conditions in isolation, they are always built on the logic of asymmetric causality. QCA is thus used to supplement the regression results by elucidating factors that are significant in some cases and under certain conditions, if not for the entire sample on average. The results it generates will provide more nuanced insight into which strategic logics for hoaxing hold, and in which contexts.

7.1 The case for using QCA

While probit and/or logit regression are standard in the terrorism studies literature for predicting binary outcomes, it is valuable to use QCA when evaluating predictors of a terrorist group's decision to use hoaxes. QCA should be considered "whenever there are good reasons to think that the phenomenon under study is best understood as the result of causally complex structures involving hypotheses about necessity or sufficiency" (Grofman and Schneider 2009, 670). Chapter 5 outlines a wide range of possible logics for hoax activity that are neither mutually exclusive, nor expected to apply across all cases. Furthermore, the indicators incorporated in Chapter 6's hypotheses may point to multiple strategic logics, and there is no existing theory to guide how they might combine to lead a group to hoax in some cases, but not others. QCA can determine inductively which factors are relevant for explaining the hoax outcome *as they occur in combination with other predictors*. It treats cases holistically, rather than isolating independent effects of their characteristics. It is worth noting that this study relies purely on cross-case observational data, and is thus limited to uncovering correlational—not causal—relationships. Any allusion to causality in the discussion of my results should be read as faithfulness to the underlying logic (and linguistic norms) of set theory, not a statement of causal proof.

While the nuanced distinctions between QCA and (logistic) regression are ably summarized elsewhere (e.g. Grendstad 2007, 127; Grofman and Schneider 2009), there are a number of foundational cleavages that explain my choice to use QCA to supplement regression for this group-based analysis. First is QCA's explicitly combinatorial approach

described above: “while standard statistical techniques are good at distilling the net effects of single variables, QCA, by virtue of giving premium to causal complexity, seeks to detect different conjunctions of conditions (configurations) that all lead to the same outcome” (Grofman and Schneider 2009, 663). In some cases, QCA may find a single factor to be relevant in isolation. But whereas regression would conclude that an insignificant coefficient estimate, all else held equal, means the corresponding variable is irrelevant, QCA explicitly checks whether that factor matters when combined with other factors (or their negation).

There are numerous additional features that are useful for distinguishing QCA from regression:

QCA is based on the twin ideas of *necessity* and *sufficiency*. Its motivations include a concern for unraveling causally complex structures in terms of *equifinality*, *multifinality*, and *asymmetric causality* [...] that tend to be omitted or slighted in most discussions of mainstream statistical methods. (Grofman and Schneider 2009, 662–63)

That QCA seeks to uncover necessary and sufficient conditions flows directly from the logic of sets on which it is premised. Necessary causes are required for an outcome to occur (i.e. in all cases of the outcome, the necessary condition is also present). In other words, Y, the dependent variable, is a subset of X, the independent or explanatory variable. Sufficient causes are those that can, on their own, produce an outcome, such that in all cases where the sufficient condition is present, the outcome is also present (i.e. X is a subset of Y).

Figure 7.1 presents a two-by-two illustration of these causal relationships (as in Goertz and Mahoney 2012, 21–23). Whereas set relations identify causal conditions based on empty cells, (logistic) regression “seeks to maximize fit, which in this case means

maximizing the cases that are placed on the main diagonal [from 0,0 to 1,1; the shaded cells in Figure 7.1]. This is the equivalent of weighting deviations from sufficiency and deviations from necessity as equally important” (Grofman and Schneider 2009, 668). Accordingly, regression’s correlational logic obscures any underlying set relations (Ragin 2008a, 22–23).

		X	
		0	1
Y	0	Any value	Empty if X is sufficient
	1	Empty if X is necessary	Any value

Figure 7.1 Two-by-two illustration of necessary and sufficient conditions

In maximizing cases along the (0,0)-(1,1) diagonal, regression not only ignores differences above or below the diagonal; it also assumes symmetric causality, treating cases along the diagonal as logically equivalent, whether they pertain to the absence or presence of both X and Y. For this reason, a probit or logit regression model’s predictive power will remain constant (with reversed signs on the coefficients), whether the dependent variable is specified as 0 or 1. Per the Boolean logic of QCA, however, causal relationships are asymmetrical: to say that X is necessary or sufficient for Y is *not* to say that the absence of X is necessary or sufficient for the absence of Y. Under asymmetric causality, the inverse of what explains an outcome does not necessarily explain the outcome’s negation, and so separate analyses are required to explain the outcome’s occurrence and its non-occurrence (see Grofman and Schneider 2009, 669).

Examination of the group-based hypotheses described in Chapter 6 (section 6.1), highlights how asymmetric causality should be expected in predicting a group’s use of

hoaxes. Using Hypothesis 1, for instance, if groups motivated by single issues have a higher propensity toward hoaxing, correlational logic would suggest that *not* having a single-issue motivation should lead a group *not* to hoax. However, this is neither the case empirically (as shown in the results below and in Appendix Table A.3) nor theoretically. Hypothesis 2 predicts that an ethno-nationalist motivation will also lead groups to hoax. Since a group classified as having an ethno-nationalist motivation would also be classified as being non-single issue, this hypothesis directly contradicts the symmetric correlational logic of Hypothesis 1. For this reason, it is important to interpret regression results in light of set relations between predictive factors and outcomes, as well as among the independent variables themselves.

Finally, whereas regression identifies average conditional correlations for single variables, all else held equal, across the full sample of cases under study, QCA accounts for equifinality and multifinality. Equifinality reflects the notion that many different pathways (whether they be single conditions or combinations thereof) can lead to a given outcome. Alternatively, multifinality recognizes that even similar conditions may generate variation in outcome. As described above and in Chapter 5, there are strong conceptual reasons to expect that both equifinality and multifinality apply with respect to predicting groups' use of hoaxes in their campaign activities.

7.2 Hypotheses revisited

While the present analysis reprises the same theoretical justification for hypotheses given in Chapter 6 (section 6.1), the statements themselves must be reformulated to reflect a

set theoretic approach. The regression's hypotheses were worded in terms of symmetric correlations, to reflect the assumptions of a regression model: if a terrorist group displays X characteristic, then that group is more (or less, if the hypothesized relationship is negative) likely to hoax than groups who do not display X characteristic. From a linear algebraic perspective, the opposite must also be true: if a terrorist group does *not* display X characteristic, then that group is less (or more) likely to hoax than groups who *do* display X characteristic.

Since QCA assumes asymmetric set relations, not correlations, formulating hypotheses in this way is inappropriate. Instead, hypotheses for QCA must be specified to indicate their underlying set relations. In the explicit language of sets, for instance: if a terrorist group belongs to the set of groups displaying X characteristic, then it is likely to also belong to the set of groups who display the outcome (hoaxing). Whereas a correlational hypothesis would assume the opposite to be true (that groups *not* displaying X characteristic are *unlikely* hoaxers), the set theoretic hypothesis makes no claim over groups that do not belong to the set of groups displaying X characteristic. If there is conceptually-sound reason to believe or suspect that a relationship might be symmetrical, then the alternate hypothesis should be explicitly tested.

In the subsections that follow, I present appropriately-specified versions of the hypotheses for set theoretic analysis. Since the purpose of the investigation is exploratory, I do not posit hypothesized combinations of factors. Rather, the hypotheses reflect my directional expectations for each indicator in isolation. As in Chapter 6, each

hypothesis is intended to provide support for one or more of the rationalist mechanisms proposed in Chapter 5.

7.2.1 Hypotheses relating to group motivation

As described in section 7.1 with respect to single-issue and ethno-nationalist motivations, the study's hypotheses relating to group motivation are inherently set theoretical: the relationships are expected to hold when the specified motivation is present, but there is no explicit expectation regarding groups' hoax behaviour when the motivation is absent. Indeed, assuming symmetry in this way would generate conflicting expectations across the motive variables. Accordingly, hypotheses relating group motivation to the incidence of hoaxing are only specified in one direction:

*H₁: If a terrorist group is motivated by **anti-abortionist, animal rights, or environmentalist** sentiments, then that group is a likely hoaxer.*

*H₂: If a terrorist group is motivated by **ethno-nationalism**, then that group is a likely hoaxer.*

*H₃: If a terrorist group is motivated by **religion**, then that group is an unlikely hoaxer.*

7.2.2 Hypotheses relating to group structure

One of the most consistently robust findings from Chapter 6's regression is the negative correlation between groups' hierarchical structure and their propensity to ever use hoaxes, suggesting that groups who are hierarchically structured are less likely hoaxers, and groups who lack hierarchy are more likely hoaxers. Symmetry, in this case, is conceptually founded: non-hierarchical groups will be likely hoaxers, since they are more vulnerable to principal-agent divergences; and hierarchical groups will be unlikely hoaxers, since they are generally more capable of successfully executing sophisticated, serious attacks (and so hoaxes confer a lower relative advantage). Accordingly, two versions of the hypothesis relating group structure to the incidence of hoaxing are specified:

***H_{4A}**: If a terrorist group displays **no hierarchy**, then that group is a likely
hoaxer.*

***H_{4B}**: If a terrorist group displays **hierarchy**, then that group is an unlikely
hoaxer.*

7.2.3 Hypotheses relating to campaign context

The hypothesis relating to groups' radical splinter origins is appropriately specified in only one direction, since the specified mechanism underpinning the hypothesis has no purview over non-splinter groups (see section 6.1.3):

H₅: If a terrorist group emerged as a radical splinter or offshoot of a less violent group or movement, then that group is a likely hoaxer.

Reformulation of the hypothesis relating to the number of serious events perpetrated per group is more complex. Despite the probit estimation's consistently strong positive effect for the logged number of serious events perpetrated per group on its likelihood of ever hoaxing, the correlation between the number of serious events perpetrated and whether the group is a hoaxer is actually very low (only 27.63%). Tests for bivariate consistency between various calibrations and binary cut-offs of the number-of-serious-events variable (and its negation) against the hoax outcome (and its negation) suggest that there is a substantial degree of unimportant variation in the number of serious events perpetrated per group, which may be obscuring a stronger linear correlation. Contrary to the concept underpinning the probit hypothesis—that the likelihood of a group hoaxing increases with the number of serious events that it perpetrates, and vice versa—these tests reveal that the most consistent threshold for predicting the hoax outcome is perpetrating more than three serious incidents.⁷⁰ From a set theoretic perspective, there is no difference in hoax behaviour between a group perpetrating four events or 2,492 (the maximum value in the sample); but groups

⁷⁰ Perpetrating more than three serious events is an almost always necessary condition (consistency = 0.87) for employing hoaxes. Conversely, perpetrating three or fewer serious events is almost always sufficient (consistency = 0.80) for not hoaxing. A similar relationship emerges in the continuous fuzzy set calibrated with 10 serious incidents as the point of full inclusion in the set, 5 as the point of maximum ambiguity, and 3 as the point of full exclusion, but the crisp indicator performs marginally better.

perpetrating four or more serious events are qualitatively different from those perpetrating three or fewer.

Three, clearly, is neither a high threshold nor a reasonable approximation of what would constitute a “prolific” group. Accordingly, it is the low end of the distribution that is more conceptually instructive: the least prolific groups are particularly unlikely to hoax. This low threshold corroborates the theoretical expectation that (among serious terrorism perpetrators) hoaxes only confer strategic value once the group either builds up a degree of reputational capital to have its hoax believed, and/or incorporates hoaxes into a mixed strategy.

*H₆: If a terrorist group perpetrates **three or fewer serious incidents**, then
that group is an unlikely hoaxer.*

7.3 Required modifications to the data

While the basic sample and input data are identical to those used in Chapter 6 (discussed in section 6.2, and presented in Appendix Table A.1), some independent variables had to be transformed in order to satisfy the conditions of crisp and/or fuzzy set modeling.⁷¹

⁷¹ Crisp set elements are completely binary, indicating either full membership in the set (=1) or full exclusion from the set (=0). Fuzzy set elements allow for partial membership in the set. As in crisp sets, a fuzzy set score of 1 reflects full membership in the set and 0 reflects full non-membership. Fuzzy sets also, however, incorporate a range of intermediate values. A score of 0.5 reflects a case that is neither fully in nor fully out of the set. Scores greater than 0.5 but less than 1 reflect cases that are more in than out, and scores less than 0.5 but greater than 0 reflect cases that are more out than in. The most rudimentary fuzzy sets incorporate only these three values: 1, 0.5, and 0. Four-value fuzzy sets include two qualitative states (fully in or out of the set), with one intermediate level between each of these states and the crossover point. Four-value fuzzy sets are helpful when there is a lot of information about the cases, “but the evidence is not systematic or strictly comparable from case to case” (Ragin 2008a, 31). More nuanced is

Table 7.1 presents the guidelines applied for assigning crisp and fuzzy set scores to the terrorist organizational conditions examined.

Apart from the variable measuring the number of serious events perpetrated per group, all the variables used in the probit regression are already binary; they can be incorporated into crisp- or fuzzy-set QCA as-is. A variable combining broadly left- and right-wing motivations (which includes the single issues as well as their economic counterparts) is added to the variable list to probe whether single-issue motivations generate an independent effect, or whether their role is driven by non-hierarchical structure. To clarify the solutions' sensitivity to crisp codings of *purely* religious/ethno-nationalist and *containing* religion/ethno-nationalism, I coded fuzzy set equivalents for religion and ethno-nationalism, which consider the degree to which each ideology is salient in the group's overarching motivation. I similarly tested a fuzzy scoring of organizational structure, to incorporate greater nuance with respect to the structural features present.⁷² However, the basic distinction between hierarchy and non-hierarchy regularly generated greater coverage and consistency than this fuzzy alternative, and so the original crisp indicator was used instead.

the six-value fuzzy set, which includes the two qualitative states of in and out, but with two intermediate levels between them and the crossover point. The most fine-grained type of fuzzy set is "continuous," where values are calibrated (using the fsQCA software) based on existing continuous data using conceptually-defined qualitative anchors to define full membership, full non-membership, and the crossover point of maximum ambiguity.

⁷² The fuzzy set indicator for group structure included six values: 1 = group has a clearly defined centralized command (*leadership* and *command and control* = 1); 0.8 = group has leadership, but no centralized command; 0.6 = group has centralized command, but no leadership; 0.4 = group has functional differentiation (*functional differentiation* = 1) but no leadership or command/control; and 0 = group displays no indication of centralized command or control.

Table 7.1 Guidelines for crisp and fuzzy set scoring of organizational conditions

	FACTOR NAME (SHORT-HAND)	CALIBRATION DEFINITION	SET TYPE
MOTIVE	single-issue (SINGLEISSUE)	1 = anti-abortionist, animal-rights, or environmentalist motivation 0 = motive is not anti-abortionist, animal-rights, or environmentalist	Crisp set
	left or right (LEFTRIGHT)	1 = left- or right-wing motivation, including single issues 0 = motive is not left- or right-wing, nor single-issue	Crisp Set
	purely ethno-nationalist (PUREETHNO)	1 = purely ethno-nationalist motivation 0 = group's ideology contains any non-ethno-nationalist component	Crisp set
	contains ethno-nationalism (CONTAINETHNO)	1 = motivation contains an ethno-nationalist component 0 = motivation does not contain an ethno-nationalist component	Crisp set
	ethno-nationalist (ETHNO)	1 = purely ethno-nationalist motivation 0.67 = strongly ethno-nationalist component 0.33 = slight ethno-nationalist component 0 = no ethno-nationalist component	Four-value fuzzy set
	purely religious (PURERELIG)	1 = purely religious motivation 0 = group's ideology contains any non-religious component	Crisp set
	contains religion (CONTAINRELIG)	1 = motivation contains a religious component 0 = motivation does not contain a religious component	Crisp set
	religious (RELIGIOUS)	1 = purely religious motivation 0.67 = strongly religious component 0.33 = slight religious component 0 = no religious component	Four-value fuzzy set
	religious ethno-nationalism (RELIGETHNO)	1 = motivation is both religious and ethno-nationalist 0 = motivation is purely religious, purely ethno-nationalist, or neither	Crisp set
STRUCTURE	hierarchy (HIERARCHY)	1 = group displays some degree of hierarchy (i.e. <i>leadership or command and control</i> = 1) 0 = group displays no hierarchy; market structure	Crisp set
CAMPAIGN CONTEXT	splinter (SPLINTER)	1 = group originated as a radical splinter 0 = group did not originate as a radical splinter	Crisp set
	committing more than 3 serious events (SERIOUS1)	1 = group committed more than 3 serious events 0 = group committed 3 or fewer serious events	Crisp set
	committing many serious events (SERIOUS2)	1 = group committed 50 or more serious events 0.5 = group committed 10 serious events 0 = group committed 3 or fewer serious events	Continuous fuzzy set

The number of serious events variable is transformed into a crisp score by separating groups responsible for three or fewer serious incidents (=0) from those

responsible for greater than three incidents (=1). A fuzzy set calibration of this variable is also tested, retaining three events as the point of full exclusion from the set (groups responsible for three or fewer events are coded as 0), but testing sensitivity to a higher threshold for full inclusion in the set of conceptually “frequent” perpetrators (groups responsible for 50 or more events are coded as 1), and a higher point of maximum ambiguity (groups responsible for 10 events are neither in nor out of the set).⁷³

7.4 Models and methods

Both csQCA and fsQCA are built upon Boolean algebra and the associated analysis of “truth tables.” In a truth table, columns exist for each of the model’s conditions (k) and the outcome (Y). As many rows exist as there are logically possible combinations of conditions (2^k). Each row thus represents a unique configuration of conditions. Groups that are analytically similar (i.e. they are identical with respect to each of the model’s conditions—although they *may* differ with respect to the outcome) will be assigned to the same row. A row’s outcome is then assigned based on the observed outcome of the cases attributed to that row, given a particular consistency threshold. For instance, with a consistency threshold of 0.8, an outcome of 1 will be assigned to rows where at least 80% of groups displaying that particular configuration of conditions hoax. If fewer than

⁷³ For sensitivity analysis, I tested multiple crisp and fuzzy score calibrations for the number of serious events perpetrated per group. The ones reported in Table 7.1 and included in Models A-F performed most effectively. The greater-than-three threshold for the crisp set division of the number of serious events perpetrated was derived from trail-and-error tests of set relations; it was not theoretically-informed. Since “[c]alibration of degree of membership in sets should be purposeful and thoughtful, never mechanical” (Ragin 2008a, 32), the fuzzy score is more sensitive to what might be meaningfully construed as a large number of events (50) and a very small number of events (3).

80% of the groups in a given row hoax, the outcome will be assigned as 0. Rows that have no associated empirical cases are treated as “logical remainders” (discussed below).

In effect, each truth table row that is associated with the outcome may be construed as a statement of sufficiency. The final step in a truth table analysis is the logical minimization of such expressions. Logical minimization summarizes the primitive (and complex) statements of sufficiency from the truth table into a combined Boolean expression. Per Boolean algebra, if two expressions differ along only one condition and agree on the outcome, the condition on which they differ is irrelevant, and can be eliminated. The result is a more parsimonious expression that is effectively a superset of the two more complex original statements.

In this chapter, I conduct both crisp and fuzzy set QCA, using Ragin’s fsQCA 3.0 software for Mac.⁷⁴ Per the software’s default (Ragin 2008b, 47), I apply a consistency threshold of 0.8 when assigning outcomes. Because the number of cases in the present analysis is relatively large for QCA (Ragin 2008b, 46), I apply a frequency threshold of two (i.e. at least two observed cases must exhibit the same combinations of factors to be considered in the solution). While this more conservative frequency threshold increases the incidence of logical remainders (i.e. truth table rows with no corresponding empirical cases), it is chosen to avoid an excessively complex solution resulting from idiosyncratic outliers.

In terms of model specification, QCA models should balance parsimony and comprehensiveness. On the one hand, given QCA’s exploratory nature, all relevant

⁷⁴ The algorithm applied for all QCA analyses is the software’s default, Quine-McCluskey.

conditions should be present to maximize the likelihood of appropriately specifying sufficient and/or necessary conditions. If a relationship is relevant in a more parsimonious model, it will also be present (although possibly in combination with additional factors) in a more comprehensive model.⁷⁵ On the other hand, the number of conditions evaluated increases exponentially the range of possible combinations of factors incorporated into the truth table. Although many cases are included in the present analysis, they still exhibit limited diversity: not every possible combination of factors corresponds to an empirical case. In truth table analysis, limited diversity generates logical remainders: those truth table rows and their corresponding combinations of factors that are logically possible, but empirically unobserved (or unobservable).

Logical remainders can occur for three main reasons: 1) the number of truth table rows exceeds the number of cases in the sample (arithmetic remainders); 2) the configuration of factors is logically possible, but does not exist in a social reality that is conditioned by various historical, political, social, and other processes (clustered remainders); and 3) the configuration of factors is implausible in the world as we know or imagine it (impossible remainders) (Schneider and Wagemann 2012a, 152–57). The present 81-group sample can easily be tested with a six-condition model without generating arithmetic remainders,⁷⁶ but clustered and impossible remainders abound.⁷⁷

⁷⁵ If an additional condition that is *irrelevant* to the outcome is added to the model, the original statement of sufficiency will exist as a superset over these new, complex, expressions. If the presence or absence of the condition is *relevant* to the outcome, then the new condition (or absence thereof) will simply be incorporated into the original statement using the logical AND.

⁷⁶ There are fewer logically possible combinations ($2^6 = 64$) than there are observations.

⁷⁷ For instance, hierarchical single-issue groups (and their respective subsets) constitute clustered remainders: while it is theoretically possible that they might exist, there are no such examples in the

In applied QCA (and comparative social science research more generally), limited diversity and an associated large number of remainders is common; how they are treated will influence the truth table solution formula (Schneider and Wagemann 2012a, 175–176; 160).

It is problematic to leave a large number of remainders untouched, since doing so will generate QCA solutions that are excessively complex. When there are no empirical cases that differ from observed cases on only one (or more) dimension(s), the algorithm cannot factor out potentially irrelevant conditions from the solution.⁷⁸ This conservative approach yields unwieldy results that are difficult to interpret and generalize.

QCA's Standard Analysis procedure tackles the challenge of logical remainders by supplying "simplifying assumptions" that facilitate logical minimization in the absence of sufficiently diverse observations. These assumptions are based on counterfactual logic: if a group displaying the given conjunction of variables were to exist, would the corresponding truth table row help to minimize the QCA solution? Standard Analysis incorporates two kinds of simplifying assumptions: those based on "easy" counterfactuals, which rely on empirical evidence as well as existing theoretical knowledge regarding each individual condition's expected relationship with the outcome; and "difficult" counterfactuals, which are based on empirical evidence alone (Schneider and Wagemann 2012a, 168). Whether easy or difficult counterfactuals are incorporated,

observed data. Alternatively, impossible remainders are generated by the mutually exclusive motive variables (any combination including purely religious, purely ethno-nationalist, or single-issue motivation).
⁷⁸ For instance, if conditions A, B, and C are shown to generate the outcome, Y, but there are no empirical cases exhibiting A, B, and *not* C, the solution would have to include all three conditions – even if the presence or absence of C is actually irrelevant to the outcome. A two-condition solution (A and B) would be more parsimonious, and more appropriate, in this case.

the simplifying assumptions of the Standard Analysis will *never* generate a solution formula that contradicts the available empirical information. In other words, any solution that incorporates these assumptions “will be a superset of the truth table rows that contain empirical information and that are sufficient for the outcome of interest” (Schneider and Wagemann 2012a, 161). The Standard Analysis procedure for truth table analysis thus generates three different solutions, based on the degree to which simplifying assumptions are accepted. The “complex” solution does not incorporate any simplifying assumptions; the “intermediate” solution only incorporates remainders that would constitute easy counterfactuals, based on the researcher’s theoretical inclinations; and the “parsimonious” solution incorporates “any remainder that will help generate a logically simpler solution” (Ragin 2008b, 81). In other words, the parsimonious solution assumes that all remainders are sufficient for the outcome—as long as they simplify empirically-established conditions.

In my interpretation of the results below, I rely exclusively on the parsimonious solution. In the present study, the parsimonious solution yields a manageable number of solution terms, and is unaffected by the two risks of simplifying assumptions identified by Schneider and Wagemann (2012a): that impossible remainders will be included in the solution, generating “untenable assumptions” (176); and that highly plausible counterfactual claims may be ignored, if they do not contribute to logical minimization based on observed data (177). The risk of untenable assumptions based on impossible remainders is insignificant here, because in counterterrorism applications of these findings, one will always begin with an observed case; a group must exist and carry out

terrorist activity before it becomes important to identify it as a hoaxer or non-hoaxer. It is difficult to imagine a scenario where a practitioner (or scholar, for that matter) would ever need to address an impossible remainder in any meaningful sense. The risk of ignoring “good” counterfactuals is irrelevant here because the purpose of the analysis is to build theory; I do not yet have “strong theoretical and substantive arguments that [any] specific remainder (or set of remainders) would produce the outcome, if only it existed” (Schneider and Wagemann 2012b, 199). In the absence of a (thus-far) empirically-supported theory for hoax behaviour, incorporating easy counterfactuals based only on my hypotheses would introduce unsubstantiated bias. The solutions I present are thus the most logically simple expressions available, considering the limited diversity amongst terrorist groups in my sample.

Although the analysis is exploratory in nature, my model specification is rooted in the hypotheses established above and in Chapter 6. Analogous to probit Model 1, QCA Model A uses basic crisp set scores for each of the six hypotheses: single-issue, ethno-nationalist, and religious motivations; hierarchical structure; radical splinter origins; and the number of serious incidents perpetrated. At a difference to the probit models, QCA models *can* incorporate the single-issue indicator in the analysis; where it is included, it should simply appear independently as a sufficient condition for hoaxing, with 100% consistency. Subsequent models are designed to test the robustness of each Model A indicator, and to bolster my interpretation of Chapter 6’s regression results. Model B, for instance, tests sensitivity to a more theoretically-informed, fuzzy set calibration of the indicator for the number of serious incidents perpetrated. Model C sheds light on whether

it is single-issue motivation, or a broadly rightist or leftist orientation combined with hierarchical structure, that generates hoax behaviour. Models D, E, and F clarify the nature of the effect(s) of religion and ethno-nationalism on hoax behaviour. Because of probable asymmetry in the hypothesized relationships, I run each QCA model to predict both the outcome of hoaxing (i.e. *hoaxer*=1) and its negation (i.e. *hoaxer*=0).

“Goodness of fit,” for a QCA model, is described in terms of solution *consistency*, “the degree to which the cases sharing a given combination of conditions [...] agree in displaying the outcome in question” (Ragin 2008a, 44), and *coverage*, “the degree to which a cause or causal combination ‘accounts for’ instances of an outcome” (Ragin 2008a, 44). Consistency is thus roughly analogous to the previous chapter’s percent of correctly predicted values, whereas coverage indicates the practical significance of a given sufficient condition in the same way that an R-squared measure describes how much variation a given model captures. The solutions presented in Table 7.2 are presented using the variable short-hands given in Table 7.1’s definition of variables. The expressions describing sufficient conditions and their influence on the outcome employ the following notation: “ ~ ” indicates the negation of a factor or the outcome; “ * ” indicates the logical AND; and “ + ” indicates the logical OR.

7.5 Results and discussion

The discussion of results proceeds in two parts. In the subsections that follow, I interpret the QCA results on their own. At risk of redundancy, I discuss each hypothesis (and the associated combinations of factors that emerge in expressions of sufficiency) in turn.

Then, in section 7.6, I evaluate these findings in relation to the probit results from Chapter 6.

Table 7.2 QCA models and their minimized solutions

MODEL A (crisp set)	$HOAXER = f(SINGLEISSUE, PUREETHNO, PURERELIG, HIERARCHY, SPLINTER, SERIOUS1)$	
	$SINGLEISSUE + (PUREETHNO * \sim HIERARCHY) \rightarrow HOAXCHOICE$	
	Solution coverage: 0.304	Solution consistency: 0.875
MODEL A (crisp set)	$\sim HOAXER = f(SINGLEISSUE, PUREETHNO, PURERELIG, HIERARCHY, SPLINTER, SERIOUS1)$	
	$\sim PURERELIG * (HIERARCHY * \sim PUREETHNO + \sim SERIOUS1)$ $+ PURERELIG * SPLINTER \rightarrow \sim HOAXCHOICE$	
	Solution coverage: 0.759	Solution consistency: 0.917
MODEL B (fuzzy set)	$HOAXER = f(SINGLEISSUE, PUREETHNO, PURERELIG, HIERARCHY, SPLINTER, SERIOUS2)$	
	$SINGLEISSUE \rightarrow HOAXCHOICE$	
	Solution coverage: 0.217	Solution consistency: 1
MODEL B (fuzzy set)	$\sim HOAXER = f(SINGLEISSUE, PUREETHNO, PURERELIG, HIERARCHY, SPLINTER, SERIOUS2)$	
	$\sim PUREETHNO * \sim PURERELIG * HIERARCHY$ $+ PURERELIG * SPLINTER \rightarrow \sim HOAXCHOICE$	
	Solution coverage: 0.655	Solution consistency: 0.927
MODEL C (crisp set)	$HOAXER = f(LEFTRIGHT, PUREETHNO, PURERELIG, HIERARCHY, SPLINTER, SERIOUS1)$	
	$\sim HIERARCHY * (PUREETHNO + LEFTRIGHT * SERIOUS1) \rightarrow HOAXCHOICE$	
	Solution coverage: 0.348	Solution consistency: 0.889
MODEL C (crisp set)	$\sim HOAXER = f(LEFTRIGHT, PUREETHNO, PURERELIG, HIERARCHY, SPLINTER, SERIOUS1)$	
	$\sim PUREETHNO * \sim PURERELIG * (\sim LEFTRIGHT + HIERARCHY)$ $+ PURERELIG * SPLINTER \rightarrow \sim HOAXCHOICE$	
	Solution coverage: 0.741	Solution consistency: 0.935
MODEL D (crisp set)	$HOAXER = f(LEFTRIGHT, RELIGETHNO, HIERARCHY, SPLINTER, SERIOUS1)$	
	$LEFTRIGHT * \sim HIERARCHY * SERIOUS1 \rightarrow HOAXCHOICE$	
	Solution coverage: 0.261	Solution consistency: 1
MODEL D (crisp set)	$\sim HOAXER = f(LEFTRIGHT, RELIGETHNO, HIERARCHY, SPLINTER, SERIOUS1)$	
	$HIERARCHY * (SPLINTER + LEFTRIGHT) + RELIGETHNO$ $+ \sim LEFTRIGHT * \sim HIERARCHY * \sim SERIOUS1 \rightarrow \sim HOAXCHOICE$	
	Solution coverage: 0.759	Solution consistency: 0.917

Table 7.2 Cont'd

MODEL E (crisp set)	$HOAXER = f(LEFTRIGHT, CONTAINETHNO, CONTAINRELIG, HIERARCHY, SPLINTER, SERIOUS1)$	
	$\sim CONTAINRELIG * \sim HIERARCHY * SERIOUS1 \rightarrow HOAXCHOICE$	
	Solution coverage: 0.304	Solution consistency: 1
	$\sim HOAXER = f(LEFTRIGHT, CONTAINETHNO, CONTAINRELIG, HIERARCHY, SPLINTER, SERIOUS1)$	
MODEL F (fuzzy set)	$CONTAINRELIG * (CONTAINETHNO + SPLINTER)$	
	$+ LEFTRIGHT * HIERARCHY \rightarrow \sim HOAXCHOICE$	
	Solution coverage: 0.672	Solution consistency: 0.929
	$HOAXER = f(LEFTRIGHT, ETHNO, RELIGIOUS, HIERARCHY, SPLINTER, SERIOUS1)$	
MODEL F (fuzzy set)	$LEFTRIGHT * \sim HIERARCHY * SERIOUS1 \rightarrow HOAXCHOICE$	
	Solution coverage: 0.261	Solution consistency: 1
	$\sim HOAXER = f(LEFTRIGHT, ETHNO, RELIGIOUS, HIERARCHY, SPLINTER, SERIOUS1)$	
	$RELIGIOUS * (ETHNO + SPLINTER)$	
$+ LEFTRIGHT * HIERARCHY \rightarrow \sim HOAXCHOICE$		
Solution coverage: 0.598		Solution consistency: 0.929

7.5.1 Animal rights, environmentalist, and anti-abortionist motivation

Initially, the QCA results appear to provide strong support for Hypothesis 1—that terrorist groups motivated by anti-abortionist, animal rights, or environmentalist (“single-issue”) sentiments are likely hoaxers. In each model where the single-issue indicator is included, it appears as an independent, sufficient condition for hoaxing, accounting for 21.7% of hoaxing groups (with perfect consistency). In other words, the conditional probability of hoaxing if a group is motivated by anti-abortionist, animal rights, or environmentalist objectives equals one; and these motivations collectively explain 21.7% of cases where groups hoax.⁷⁹ These results are consistent with the theory that some groups hoax as a

⁷⁹ As expected, given the hypothesized asymmetric relationship, *not* being a single-issue group is never a sufficient condition for *non*-hoaxing.

normatively-consistent outlet for incomplete radicalization to violence, since hoaxes offer a less-severe alternative to serious activity that does not cause physical harm to human beings.

However, as in Chapter 6, each of the observed single-issue groups also lacks hierarchy, so the solution positing single-issue motivation as a unique predictor of hoaxing is premised on a simplifying assumption: that single-issue groups *displaying hierarchy* would also hoax, if they were to exist. Since hierarchy is hypothesized to be consistent with non-hoaxing, however, this counterfactual is not an easy one. The models incorporating the single-issue indicator are unable to distinguish—based on empirical observations—whether motivation plays an independent role, or whether its seeming effect is a function of limited diversity amongst single-issue groups and the underlying effect of non-hierarchical structure.

In Model C, the single-issue indicator is replaced by a variable combining all broadly left- or right-wing organizations (whether they be motivated by their associated single issues, or economic policies). This broad motivational indicator combines with non-hierarchy and the perpetration of more than three serious events to generate a sufficient condition for hoaxing, with perfect consistency (a combination also sufficient for hoaxing in Models D and F). In other words, all groups displaying these three characteristics use hoaxes at some point in their campaigns. This combination of characteristics also exhibits a higher coverage than the single-issue indicator alone, accounting for 26.1%—as compared with the 21.7%—of hoaxing groups. Given the absence of any theoretical explanation that might suggest why broadly left- or right-wing groups would be inclined

toward hoax use, this increase in coverage suggests that non-hierarchy—not motivation—may be driving the relationship.

However, when the single-issue groups are removed from the sample (as was done for probit Models 1A and 2B), this sufficient combination disappears.⁸⁰ Although there *are* three non-hierarchical, economically-oriented groups who carry out more than three serious attacks (accounting for this combination’s nearly five percent increase in coverage compared with the single-issue indicator alone), they do not hoax with adequate consistency to generate a sufficient condition for hoaxing when their non-economically-oriented, similarly-constituted, counterparts are removed from the sample.⁸¹ Although the effect of single-issue motivation on hoaxing captured using QCA is overdetermined, and at least partially explained by the lack of hierarchy exhibited by the same groups, there is thus evidence to suggest that the single-issue motivations do play a role in increasing groups’ likelihood of hoax use.

In the following sub-section, I address the effect of hierarchy on hoaxing to better explore this dynamic. Subsequent sub-sections address the other hypotheses in their respective order.

⁸⁰ For the reduced sample, Model C generates only one sufficient condition (pure ethno-nationalism combined with non-hierarchy), with very low solution coverage (0.111) and consistency (0.667). Across all of the models, sufficient conditions that do not involve the single-issue or left- and right-wing indicators remain constant (albeit with reduced solution coverage and consistency) when the single-issue groups are removed from the sample.

⁸¹ Although the effect is attributed more narrowly to the single-issue motivations, the combined left-wing/right-wing indicator is used in subsequent models, since it effectively constitutes a superset of the single-issue groups. If a relationship is true of left- or right-wing groups, it must also be true for all single-issue groups; the more general motive category simply casts a wider net and thus leads to generally higher coverage, while also allowing for more parsimonious solutions.

7.5.2 (Non-)hierarchical structure

Hypothesis 4 posits both that non-hierarchical groups are likely hoaxers, and that hierarchical groups are unlikely hoaxers. The QCA solutions predicting both the hoax and non-hoax outcome support this hypothesis, in both directions. Lacking hierarchy (in other words, following a market structure) always appears—or is implied—in the solution for hoaxing, and hierarchy always appears in the solution for non-hoaxing. There is only one inconsistent result—where non-hierarchy appears in a solution for non-hoaxing (Model D)—but the discrepancy does not appear to be theoretically significant.⁸² That said, group structure is never sufficient on its own for predicting whether a group will hoax or not.

As discussed in section 7.5.1, non-hierarchy is implied in solutions for identifying hoaxing groups that include single-issue motivation (Models A and B).⁸³ When single-issue motivation is replaced with a broadly left- or right-wing motivation, the combination of this broad motive with non-hierarchy and a minimum requirement of having perpetrated more than three serious events (Models C, D, and F) is also sufficient for generating the hoax outcome. Similarly, in Model E, non-hierarchy and the minimum threshold of perpetrating more than three serious events combine with *any* non-religious ideology to generate the hoax outcome. As long as non-hierarchical groups commit more than three serious attacks and lack any trace of religious motivation, then, they are likely to use

⁸² While the four cases adhering to this solution are perfectly consistent in displaying the non-hoax outcome, they exhibit an incredibly low raw (6.9%) and unique (5.2%) coverage. Their presence as a unique solution is likely due to their inclusion of two groups lacking structure coding altogether (Black December, and the Justice Guerillas).

⁸³ Non-hierarchy is left out of the parsimonious solution due to the simplifying assumption that single-issue hierarchies would hoax, if they were to exist. The limited diversity of empirical cases cannot rule out that non-hierarchy may be required to generate hoaxing amongst groups that follow an animal rights, anti-abortionist, or environmentalist ideology.

hoaxes at some point. This combination explains a substantial portion (30.4%) of the hoax outcome. Alternatively, if a non-hierarchical group is purely motivated by ethno-nationalism, it is a likely hoaxer (Models A and C); there is no requirement for the number of serious events perpetrated by non-hierarchical ethno-nationalist groups. No matter what the corresponding motive, a lack of hierarchy increases groups' vulnerability to principal-agent problems, and correspondingly increases their propensity toward hoaxing.

To predict the outcome of *non*-hoaxing, hierarchy must be combined with any ideology that is neither purely religious nor purely ethno-nationalist (Models A, B, and C); or, more narrowly, any broadly left- or right-wing ideology (Models E and F), sometimes with a corresponding requirement that the group have radical splinter origins (Model D). Because such hierarchical groups are more capable of sophisticated, serious attacks, hoaxes' relative simplicity and low costs do not offer a substantial relative benefit, and are unlikely to be strategically employed. They are also less vulnerable to principal-agent dynamics that would lead to hoaxing.

That left- and right-wing groups are present in predicting hoaxing (when they are non-hierarchical) and non-hoaxing (when they are hierarchical) indicates that the absence or presence of hierarchy (and sometimes radical splinter origins) can be sufficient to distinguish hoaxers from non-hoaxers. Returning to Hypothesis 1, then: while single-issue motivation may play a role in solidifying groups' propensity toward hoaxing, these groups' non-hierarchical structure—combined with their general left- or right-wing orientation, and campaign context—might otherwise be sufficient to ensure their hoax behaviour.

Purely religious and purely ethno-nationalist ideologies, on the other hand, exert a clear effect that is not overridden by the effects of hierarchy or a lack thereof.

7.5.3 Ethno-nationalist motivation

Hypothesis 2 posits that terrorist groups motivated by ethno-nationalism are likely hoaxers. The QCA solutions illustrate that the presence of an ethno-nationalist ideology only affects hoaxing to a minimal degree: in models A and C, a purely ethno-nationalist ideology combines with a lack of hierarchy to generate the hoax outcome—but in each solution, this expression explains only a small percentage of hoaxing groups (8.7%) and with a low consistency (66.7%). There is thus only a very weak case for claiming that non-hierarchy and ethno-nationalism are sufficient for generating the hoax outcome.

Ethno-nationalism is more important in generating the non-hoax outcome, but only when co-absent or co-present with religion. In models A, B, and C—where a combined religious ethno-nationalist ideology is logically impossible—their joint effect is implied by a co-absence of *purely* religious and *purely* ethno-nationalist ideologies, plus hierarchical structure, to indicate likely hoaxers.⁸⁴ This joint condition is both highly consistent (91.7%) and with substantial coverage (56.9%). Analogously, where models *allow* for a combined religious ethno-nationalist ideology—whether indicated by the

⁸⁴ In many cases, groups that are neither purely ethno-nationalist nor purely religious are jointly both. A second sufficient condition, including a lack of both purely ethno-nationalist and purely religious ideology, as well as the lack of a broadly left- or right-wing motivation, also emerges in Model C, but with very low unique coverage (only 8.6%). The rest of the cases exhibiting this combination of factors also belong to the solution involving the presence of hierarchy and co-absence of pure religion and pure ethno-nationalism, which has a far greater unique coverage (36.2%) and explanatory power: the presence of hierarchy is far more instructive than a solution indicating only the absence of three alternative ideologies.

jointly religious ethno-nationalist variable (Model D); the intersection of the sets containing religious and ethno-nationalist ideological components (Model E); or the intersection of the fuzzy scores for religion and ethno-nationalism (Model F)—the presence of ethno-nationalism combined with religion is sufficient for predicting which groups will never hoax. Model D's crisp indicator for jointly religious ethno-nationalist ideology performs the best of these combined solutions, accounting for 24.1% of non-hoaxers, with a high (93.4%) consistency.⁸⁵ This result suggests that the presence of any tertiary motive (usually a left- or right-wing economic orientation) dulls the combined religious ethno-nationalist effect.

The religious ethno-nationalist solution for non-hoaxing may be explained by hierarchy,⁸⁶ or by the economic club explanation for religious groups' particular lethality when they resort to terrorist violence (Berman 2009; Berman and Laitin 2008): since religious ethno-nationalist groups are likely social service providers, they are likely to have a committed membership base that is unlikely to defect when carrying out complex and/or lethal attacks. They (along with their religious splinter counterparts) are thus particularly invulnerable to the principal-agent problems that appear to generate hoax activity more generally among non-hierarchical groups.

⁸⁵ The pathways to non-hoaxing in Model F, which uses fuzzy set scores for religion and ethno-nationalism, are identical to those produced in Model E, which uses crisp scores for ideologies containing religious and ethno-nationalist components. However, despite identical solution consistencies, Model F's fuzzy scores generate a 7.4-point reduction in solution coverage; the fuzzy sets' nuance actually diminishes the variables' effect, indicating that it is a crisp, either/or, distinction that matters—not matters of degree.

⁸⁶ All but one of the sample's fifteen religious ethno-nationalist groups are non-hoaxers (93.3%), and all but three are also hierarchies (80.0%).

On their own, then, while ethno-nationalism and the absence of ethno-nationalism do not generate any substantial independent effect on the hoax or non-hoax outcome, ethno-nationalism *does* matter to a substantial degree when it is combined with religion to indicate groups who will never use hoaxes. QCA cannot reject the null hypotheses with respect to ethno-nationalism's individual effect on hoaxing, but it supports a new hypothesis: that religious ethno-nationalists are particularly unlikely hoaxers.

7.5.4 Religious motivation

Hypothesis 3 posits that terrorist groups motivated by religion are unlikely hoaxers. As indicated above, religious ethno-nationalists are particularly unlikely to ever use hoaxes. Religious ideology also exerts an effect on groups' hoax behaviour, independent of ethno-nationalism.

Generally, religion is more effective in explaining why some groups never use hoaxes than explaining why some groups do. The *presence* of a religious ideological component is never included as a predictor of hoaxing, and in only one model does the *absence* of a religion-containing motivation predict hoaxing: Model E's combined solution of any non-religious, non-hierarchical group responsible for more than three serious events. While this solution has perfect consistency, and explains 30.4% of hoaxing groups, by including only the *absence* of a religious ideological component, it neither supports nor contradicts Hypothesis 3. Alternatively, both the presence and absence of religion are included in solutions for non-hoaxing. Both purely religious ideologies (Models A, B, and

C) and ideologies containing religious components (Models E and F) combine with radical splinter histories to explain why some groups never use hoaxes. When a purely religious ideology is *absent*, non-hoaxing may still result under certain conditions,⁸⁷ but these do not logically contradict the unidirectional Hypothesis 3.

Accordingly, Hypothesis 3—that religious groups are unlikely hoaxers—is confirmed only where religion is accompanied by either a radical splinter history, or a jointly ethno-nationalist ideology. Interestingly, religion and religion-containing variables are the only motivations to consistently appear or be implied in any QCA solution without a corresponding requirement for hierarchy or a lack thereof. This finding suggests that the effect of religious splinters and/or religious ethno-nationalists on non-hoaxing is sufficiently strong to override any moderating role that non-hierarchy might otherwise play.

7.5.5 Radical splinter origins

Hypothesis 5 posits that terrorist groups who originated as radical splinters or offshoots of a less violent group or movement are unlikely hoaxers.⁸⁸ The QCA results find no independent effect of splinter history on a group's propensity to hoax, but they do clearly indicate (as discussed in section 7.5.4) that *religious* splinter organizations—whether they

⁸⁷ Non-hoaxing results in the absence of a purely religious ideology: in the mutual absence of ethno-nationalism and the presence of hierarchy (Models A, B, and C); in the mutual absence of ethno-nationalism and left- or right-wing orientation (Model C); and when groups conduct three or fewer serious events (Model A).

⁸⁸ As described in section 6.2.3.3, this dissertation accepts a broad definition for splinter or offshoot organizations. They are not limited to only splinters of existing terrorist groups, but to radical offshoots of non-terrorist movements that enjoy some form of political organization, where there is a degree of continuity in group membership between the “mother” movement and the splinter.

be purely religious (Models A, B, and C) or containing religion (Models E and F)—are particularly *unlikely* hoaxers. In Model D, splinter history combines with hierarchical structure to predict non-hoaxing. Hypothesis 5 thus holds, but only when radical splinter groups are also either hierarchically structured, or at least partially religiously-oriented.

In either case, both components of the sufficient condition are predicted to generate non-hoaxing. In the case of the religion-splinter combination, groups are expected to use particularly extreme forms of violence, and so hoaxes offer a lower relative advantage. In the hierarchy-splinter combination, the processes of group structuring and/or formation would produce a dedicated membership base that is unlikely to defect from the new group when carrying out complex and/or lethal attacks. Whether this commitment flows from a tight command-and-control structure, or the members' desire to form a splinter organization dedicated to more extreme forms of violent activity, hoaxes are unlikely to result from principal-agent breakdowns, as they might amongst non-hierarchical, non-splinter organizations.

Bivariate tests for consistency in set relations between groups' radical splinter origins and their use of hoaxes suggest that there may be an independent effect of splinter origins on non-hoaxing. The presence of a splinter history is inconsistent with the use of hoaxes, but has a 0.76-79% consistency (depending on the variable's specification) as a sufficient condition for non-hoaxing.⁸⁹ In other words, splinter history cannot on its

⁸⁹ Not being a splinter group has no predictive value on the presence or absence of hoaxing. With consistencies ranging from 0.74 to 0.83, depending on the variable's specification, the absence of splinter origins is identified as a necessary condition for both the hoax and non-hoax outcomes, which in effect means it cannot be logically necessary for either.

own explain a group's propensity to hoax, but it comes close to being sufficient, on its own, for explaining non-hoaxing across the entire sample of perpetrators. Despite the over-determination of the combined solutions for non-hoaxing discussed above, it is thus plausible that a group's splinter origins exert an independent (although not uniquely sufficient) effect.

7.5.6 Number of serious events perpetrated

Hypothesis 6 posits that terrorist groups perpetrating three or fewer serious incidents are unlikely hoaxers. Although the number of serious events perpetrated does not figure in the solution for every model, where it *does* appear, it confirms a symmetrical relationship: having committed more than three serious events is sometimes associated with groups' use of hoaxes, and having committed three or fewer serious events is sometimes associated with non-hoaxing. There are no counter-examples.⁹⁰ In neither case, however, is the number of serious events perpetrated a sufficient condition on its own.

Perpetrating more than three serious events becomes a sufficient condition for hoaxing when it is associated with a lack of hierarchy and either lacking a religious ideological component (Model E) or displaying a broadly left- or right-wing motivation (Models C, D, and F). Both combinations align with the expected effect of non-hierarchy on hoaxing. Perpetrating a very small number of events becomes a sufficient condition for identifying groups who *never* hoax when combined with any motive that is not purely

⁹⁰ Having committed three or fewer serious events is never included in a solution for hoaxing, and having committed more than three serious events is never included in a solution for non-hoaxing.

religious (Model A); and any motive that is neither left- or right-wing oriented, when the associated groups also lack hierarchy (Model D). In this latter case, perpetrating three or fewer serious events overrides the otherwise consistent relationship between non-hierarchy and the presence of hoax use. The number of serious events perpetrated thus does not always matter, but when it does, it exerts a powerful effect in the hypothesized direction.

The above analysis is based on a crisp score for the number of serious events perpetrated, distinguishing groups responsible for three or fewer serious events from those responsible for four or more; a threshold determined by trial-and-error tests of basic set relations with the hoax (and non-hoax) outcomes. When a more theoretically-informed fuzzy score for the number of serious events perpetrated is included (Model B), the variable does not figure in any solutions. This (non-)finding confirms that it is only the perpetration of a very small number of incidents that matters when determining whether or not a group will use hoaxes. Anything beyond three serious incidents reflects unimportant variation.⁹¹

7.5.7 Overview of QCA results

Overall, QCA suggests that the incidence of hoaxing is best explained by a lack of hierarchy, when combined with either a purely ethno-nationalist or broadly left- or right-

⁹¹ Indeed, when a lower threshold (10 events) for full inclusion in the set of frequent serious perpetrators is substituted into this model (results not displayed), it never factors into the solution for hoaxing, although being a full non-member of the greater-than-10-events set sometimes does. Compared with the fuzzy score for the number of serious incidents perpetrated tested in Model B, however, this compressed calibration has an even lower bivariate consistency. This comparison confirms the substantial degree of unimportant variation that exists above the three-event threshold.

wing motivation (or the absence of any religious component). Single-issue motivation, on its face, plays a consistently sufficient role in determining whether a group will hoax, but given the sample's lack of hierarchically-structured single issue groups, non-hierarchical organizational structure tells a more compelling story (without relying on counterfactual cases). While these pathways to hoaxing are relatively consistent across models, they account for only a moderate proportion (21.7 – 34.8%) of groups who hoax. Pathways to non-hoaxing are more complex, but have significantly greater coverage (59.8-75.9%). Hierarchical groups with either purely ethno-nationalist or broadly right- or left-wing orientations are likely non-hoaxers, as are religious groups with either a combined ethno-nationalist motivation or a radical splinter history.

That the QCA models are more effective in identifying groups who will never hoax over groups who sometimes do is helpful from a counterterrorism perspective. Since this dissertation aims to distinguish hoaxing groups from non-hoaxing groups, rather than distinguishing hoax events from serious events, models identifying likely hoaxers offer no insight into whether a particular event should be treated seriously or not. Even if a group were a likely hoaxer, every one of the sample's groups also carries out serious attacks; it would be impossible to know based on organizational characteristics alone whether any particular incident were hoaxed or serious. It should be noted that this outcome is partly an artifact of my sample selection, which required as a criterion for inclusion that perpetrators carry out at least one serious incident; pure hoaxers (although they may exist) are excluded from the analysis. Alternatively, in identifying with a rather high degree of confidence which groups *never* use hoaxes, the models can help responders

identify immediately—as long as perpetrators are known—which incidents to always take seriously. The solutions predicting which groups will never hoax have a minimum 91.7% consistency (i.e. they correctly predict the outcome of a serious-event-only group at least 91.7% of the time). This means that nearly 9% of the time, a hoaxing group will be wrongly identified as serious. Still, it is safer to wrongly identify a hoaxing group as serious, than to identify a serious group as a hoaxing one.

Furthermore, the QCA models apply a minimum frequency threshold of two cases, meaning that they exclude logical combinations of factors displayed by only one group. Accordingly, it can be deduced that groups follow more idiosyncratic pathways toward the use of hoaxes than groups whose organizational characteristics prevent them from ever hoaxing.

While it is difficult to derive conclusions regarding groups' intent based on purely observational data, the findings described above suggest that groups' use of hoaxes may result most substantially from principal-agent problems and individual members' cost-benefit analyses, rather than an overarching strategic logic for the group. If groups are hoaxing strategically, their logic is not indicated by broad-based organizational characteristics (as suggested by the preponderance of idiosyncratic combinations of group features in the QCA models predicting the hoax outcome), but by particularities of the group and its circumstances. These same organizational characteristics, however, are sufficient in most cases for precluding groups' incorporation of hoaxes into their campaign activities.

7.6 Putting it all together: Probit, QCA, and the strategic logic of hoax perpetrators

In this section, I consider the full range of evidence—from Chapter 6’s probit regression, and the present chapter’s QCA analyses—for the influence of a range of terrorist organizational characteristics on the binary dependent variable of whether a group ever uses hoaxes, or not. QCA’s combinatorial approach and underlying set relations nuance the net effects of isolated variables uncovered by probit, which assumes symmetrical relationships. I return to each hypothesis in turn, to accept, reject, or modify it in light of the empirical results.

7.6.1 Revisiting hypotheses on group motivation

With respect to the effect of anti-abortionist, animal rights, or environmentalist sentiments on groups’ use of hoaxes, both probit and QCA agree: having a single-issue motivation is, on its surface, sufficient for a group’s decision to hoax. QCA reveals that this relationship is asymmetrical: the presence of single-issue motivation leads to hoaxing, but its absence has no effect on hoaxing or non-hoaxing.⁹² However, this relationship is over-determined. Since all observed single-issue groups also lack hierarchy (Appendix

⁹² Despite embodying a perfect set relation (Appendix Table A.3), the variables for hoaxing and single-issue motivation only have a correlation of 40.73%. This weak correlation confirms Ragin’s (2008a) contention that set-theoretic arguments “should be evaluated on their own terms, that is, as (asymmetric) set relations not as (symmetric) correlational arguments” (3). A basic correlational test would have missed this very important relationship. That there are 18 hoaxing groups who are not single-issue-motivated does not in any logical way challenge the claim that single-issue groups are hoaxers. Those groups undermine the correlation, but not the logical argument for sufficiency.

Table A.4), neither method is able to separate the effect of single-issue motivation from the effect of non-hierarchical structure, which also indicates the use of hoaxes.

When an alternate specification of the basic QCA model—using broadly left- or right-wing ideology instead of single-issue motivation—is tested against a reduced sample of perpetrators, excluding those with single-issue motivation, non-hierarchical economically-oriented groups are neither likely nor unlikely hoaxers. This finding corroborates the lack of significance of the broader left- and right-wing motive controls in probit Model 2B. It thus appears that the single-issue motivations produce an effect on hoaxing that cannot be fully captured by non-hierarchy. If non-hierarchy were the sole motivating factor, then non-hierarchical economically-oriented groups would also consistently hoax.

Accordingly, a unidirectional Hypothesis 1—that single-issue-motivated groups are likely hoaxers—can be tentatively accepted, but with the caveat that the effect may be partially attributed to such groups' lack of hierarchy. The corresponding strategic logics—that single-issue groups hoax because of a normative bias against the use of violence for causing physical harm to sentient beings (section 5.3.1); or because hoaxes can serve as dry runs, offering casualty-averse groups a means to hone discriminant targeting (section 5.2.2)—can also only be tentatively accepted.

With respect to the effect of ethno-nationalism on groups' hoax behaviour, probit and QCA diverge substantially. The probit estimations support the existence of a robust, positive, net effect of purely ethno-nationalist motivation on the use of hoaxes. They also indicate that ethno-nationalism's effect disappears when combined with any other

motivation—whether it be religion, or left- or right-wing economic policies. QCA, however, shows that ethno-nationalism *only* produces an effect in combination with other factors. In some cases, the factors it combines with directly contradict the regression results. In particular, QCA consistently identifies the combination of religious ethno-nationalist ideology as a sufficient condition for never hoaxing—a combination that never even approached significance in the regressions. That probit overlooks a relationship so clearly captured by QCA is a function of its inability to handle asymmetric relationships. Combined religious ethno-nationalism does not matter for predicting hoax use *on average*, because the groups who are *not* religious ethno-nationalists are fairly divided between hoax use and non-use. However, nearly all (93.3%) religious ethno-nationalists are non-hoaxers, accounting for the high consistency of corresponding QCA solutions.

Returning to purely ethno-nationalist motivation, the QCA solutions can shed light on why probit implies a strong relationship that does not appear to hold when cases are examined holistically. While the presence of pure ethno-nationalism appears in some solutions for hoaxing (albeit with very low coverage and consistency), the *absence* of a purely ethno-nationalist ideology appears frequently in solutions for non-hoaxing, in various combinations with substantial coverage and high consistency. Although there is a linear correlation, it is heavily weighted toward co-absence of the ideology and hoax outcome: a nuance which the regression ignores.⁹³ While the regression results are

⁹³ Indeed, a basic cross-tabulation of pure ethno-nationalism and hoaxing reveals that 75% of groups lacking pure ethno-nationalist motivation are *non*-hoaxers, and slightly more than half (55.6%) of purely

robust, since QCA reveals that they are largely driven by the *absence* of purely ethno-nationalist ideology, they are not theoretically instructive. Groups *lacking* purely ethno-nationalist motivation may display any number of alternative ideological orientations, including partial ethno-nationalism. It is not meaningful to interpret the coefficient's large average marginal effect as supporting the hypothesized underlying causal mechanism.

Accordingly, while the hypothesized positive relationship between ethno-nationalism and hoaxing is robust in the regression analysis, its practical interpretation is nonsensical. The unidirectional hypothesis formulated for QCA testing is a more useful alternative, but is not empirically substantiated: QCA reveals it is the absence of pure ethno-nationalism that indicates non-hoaxing, not the mutual presence of either. In terms of hoaxers' strategic logic, then, there is no empirical support for the theory that ethno-nationalist groups hoax in a quest for legitimacy, by signalling restraint and avoiding the imposition of high costs on their constituents (section 5.3.2).

With respect to the effect of religion on groups' hoax behaviour, the probit results indicate a consistently strong positive effect that is the opposite of the hypothesized relationship that religious groups would be *unlikely* hoaxers. While Chapter 6 offered potential explanations for this reversed coefficient sign,⁹⁴ QCA reveals that the effect (as with ethno-nationalism) is likely driven by co-absences of purely religious ideology and

ethno-nationalist groups *are* hoaxers. There is thus consistency between mutual presence and absence of the conditions, but it is heavily weighted to one side.

⁹⁴ For instance, the positive relationship between purely religious ideology and hoaxing may be spurious, due to disproportionate attention paid to religious groups by dataset coders and/or the likelihood of unaffiliated individuals free-riding on the lethal reputations of religious groups, carrying out false flag hoaxes in their name. The result may also be explained by the strategic logic of hoaxing to complement their serious activities (section 5.2).

hoaxing, not their mutual presence; having a religious ideology does not actually increase a group's likelihood of hoax use. In QCA, the presence of religious ideology is never a sufficient condition (alone or in combination) for groups' hoax use. Both its presence and absence, however, play a role in explaining non-hoaxing. While the absence of religion in solutions for non-hoaxing helps explain the regression's positive coefficient, it is not theoretically instructive (for the same reasons discussed above with respect to ethno-nationalism). The *presence* of religious ideology in solutions for non-hoaxing, on the other hand, supports the unidirectional hypothesis that religious groups are unlikely hoaxers.

Per QCA, however, religion is never sufficient on its own to explain non-hoax use. Instead, religion is only relevant when combined with either ethno-nationalist motivation, or radical splinter origins. Given these combinations, not all of the hypothesized mechanisms explaining religious groups' non-use of hoaxes (section 4.2.1.1) appear to hold. The first mechanism posited that religious groups are unlikely to hoax because they find theological justifications for extreme violence and are unconcerned with earthly constituents and their associated norms. For such groups, hoaxes' inability to cause death or bodily harm is not an advantage worth pursuing (and risking undermined estimations of credibility to achieve). While this explanation may explain religious splinters' avoidance of hoaxes, it cannot explain religious ethno-nationalists' non-hoax use, since ethno-nationalists *should* be concerned with earthly constituents and their norms.

Non-hoaxing by both religious splinter groups and religious ethno-nationalists, however, can be explained by a second mechanism: the economic club explanation for radical, religious groups' lethality (Berman 2009; Berman and Laitin 2008). The club-based

mechanism posits that religious groups will not hoax because they maintain a select, highly-committed membership base that is unlikely to defect; their corresponding ability to carry out complex and lethal attacks means that hoaxes do not offer a relative advantage over serious tactics, when audience costs are factored in. According to Berman and Laitin (2008), religious groups are able to guarantee this reliable membership base by requiring that members make substantial sacrifices in order to benefit from the group's social service provision. For religious splinter organizations, members' decision to join the splinter (rather than maintain alignment with the less-severe mother organization) constitutes one such sacrifice. For religious ethno-nationalists, the combined effect may be driven by an increased likelihood of their service provision to encourage group membership in the first place, and corresponding lethality when they eventually resort to violence.⁹⁵ Alternatively, this select membership base—amongst religious ethno-nationalists and religious splinter groups alike—also reduces these groups' vulnerability to principal-agent problems that might otherwise result in hoax activity.

7.6.2 Revisiting hypotheses on group structure

With respect to the effect of hierarchy (and non-hierarchy) on groups' hoax behaviour, both probit and QCA agree: there is a symmetrical, negative relationship between hierarchy and groups' use of hoaxes. All Hypothesis 4 formulations are thus empirically

⁹⁵ The social service providers that Berman identifies (Hamas and Hezbollah) are both religious ethno-nationalists; and the non-social service providers that overlap with the existing sample (al-Aqsa Martyrs; Popular Front for the Liberation of Palestine [PFLP]) are not religious ethno-nationalists (Berman 2009, 141).

supported. While QCA consistently includes (or implies) non-hierarchy in solutions for hoaxing, it also consistently includes the presence of hierarchy in solutions for non-hoaxing, although in neither case are they sufficient on their own. While most of these conditions involve the absence of various ideologies, some of them include requirements for the presence of broadly left- or right-wing ideology.⁹⁶ Importantly, for these left- or right-wing groups, the presence or absence of hierarchy (and sometimes splinter origins) is sufficient for distinguishing hoaxers from non-hoaxers.

That probit and QCA results for the structure indicator so robustly point to a symmetrical relationship is reflected in the strategic logics underpinning the hypothesis. Both the presence and absence of hierarchy are indicative of distinct qualitative states that each follow an independent strategic logic for hoaxing and non-hoaxing. The use of hoaxes by non-hierarchical groups is likely attributed to principal-agent divergences, where individual members defect from their group's *modus operandi* to privately pursue hoaxes' low costs and risks (section 5.1). Non-hierarchical, leaderless resistance groups may further encourage hoax events by inciting individuals who have not fully radicalized to violence to carry out activities in the group's name (section 5.3.1). Alternatively, the propensity of hierarchical groups to never use hoaxes may be explained by such groups' invulnerability to principal-agent breakdown. It may further be explained by their increased capacity for carrying out more sophisticated, lethal attacks (because of their

⁹⁶ For predicting hoaxing, the structure condition is combined with a requirement for motivation (usually the absence of religion or ethno-nationalism, or the presence of broadly left- or right-wing motivation) and sometimes with a requirement of perpetrating more than three serious events. For predicting non-hoaxing, the structure condition is combined with any ideology that is neither purely religious nor purely ethno-nationalist (frequently the presence of a left- or right-wing motivation is required, sometimes in combination with radical splinter origins).

greater financial and human resources). Under such conditions, the low cost requirements of hoaxing engender a correspondingly low relative advantage, thus discouraging their use in place of serious tactics (section 5.1).

7.6.3 Revisiting hypotheses on campaign context

With respect to the effect of radical splinter origins on a group's use of hoaxes, probit and QCA agree: splinter organizations are unlikely hoaxers. However, although the probit results imply that the opposite should also be true (that non-splinter groups are likely hoaxers), QCA—as well as the loss of coefficient significance, under certain regression model specifications—reveals that the relationship is asymmetric: splinter origins matter for predicting the non-hoaxer outcome, but a lack of a splinter history has no bearing on groups' hoax use. This asymmetric relationship supports the strategic logic for splinter groups' non-hoaxing: since there already exists a non- (or lesser-) violent outlet for furthering their cause, individuals who elect to partake in the splinter group's activities are both unlikely to defect from their new organization (and hoax via principal-agent divergences) and unlikely to find hoaxes' lower severity (section 5.3) attractive.

The unidirectional Hypothesis 5 (and corresponding strategic logic), however, cannot be accepted with certainty, because splinter origins are never on their own a sufficient condition for non-hoaxing; it requires either a religious motivation, or hierarchical structure. Since both religion and hierarchy also indicate non-hoaxing, any hypothesized independent effect for splinter origins is over-determined.

Finally, with respect to the number of serious events perpetrated per group, QCA clarifies the regression result. Probit suggests a positive, logarithmic relationship between the number of serious events a group perpetrates and its likelihood of employing hoaxes. Basic set relations, however, show that it is only the perpetration of events above or below a very low threshold that matters: groups committing more than three events are likely hoaxers, whereas groups committing three or fewer events are likely non-hoaxers. This number-of-events indicator is not always included in solutions for hoaxing and non-hoaxing, but it is always consistent with a positive relationship when it is. In some cases, perpetrating very few serious acts is so significant that it even overrides the otherwise-consistent effect of non-hierarchy in predicting hoax use.

Both the probit and QCA results thus support a symmetrical hypothesis pointing to a positive relationship between the number of serious events perpetrated and hoax use, but the substantial degree of unimportant variation above a very low threshold uncovered by set relations indicates that the strategic logic underpinning this hypothesis must be revised. If groups are hoaxing in the context of mixed strategies (section 5.2.1), they only need to “repeat the game” (i.e. attack) three times before they have built up sufficient reputational capital to force counterterrorist response, and hoax without otherwise signalling weakness and undermining their credibility for future attacks.

7.7 Conclusion: On research methods and the characteristics of likely hoax perpetrators

In this chapter, I used QCA to run a parallel empirical analysis to that conducted in Chapter 6 using probit regression. While each method provided substantive insight into the effect of various terrorist organizational characteristics on the use of hoaxes, the comparison of their findings yielded nuanced insight into the limitations of conventional statistics for evaluating such complex phenomena as terrorist groups' use of hoaxes.

Section 7.6's comparison of each method's results shows how regression's assumption of symmetrical net effects can generate misleading results. This is especially clear with respect to the combined religious ethno-nationalist motivation indicator: QCA reveals a consistently strong asymmetric relationship that is missed entirely by the regression. The regression's assumption of symmetry is also problematic with respect to certain binary indicators. Symmetry assumes logical equivalence between the absence and presence of a characteristic, which is not always conceptually reasonable. Regarding this dissertation's motive variables, for instance, the *absence* of an ideology means very little (since any of a wide range of alternative ideologies may be present when one is absent); the mechanism underpinning each motivation-related hypothesis is contingent on the motive's *presence* only. In assuming symmetry, the regression sometimes generates conceptually untenable results, as is the case with religion, when the seemingly counter-intuitive relationship is really driven by the *absence* of an ideology. Furthermore, as evidenced with respect to the number of serious events perpetrated per group, the regression is unable to capture the effects of unimportant variation. Lastly, in evaluating

the net effects of isolated variables, the probit models are unable to reveal under which conditions their variables are more (or less) important—something that QCA’s holistic examination of cases and corresponding causal combinations does well. Accordingly, although the probit results are highly significant and robust across models, QCA offers more reliable results in the present context.

That said, both QCA and probit are limited in similar respects. First, neither method can overcome the limited diversity of observed cases. Without sufficient empirical variation, cross-case comparisons will never be able to disentangle the independent effects of certain conditions. In particular, until a hierarchical single-issue group is observed, it will never be conclusively known whether existing single-issue groups’ hoax use is guided by their motivation and associated norms against the use of violence, or whether it is driven by their non-hierarchical structure and preponderance for principal-agent divergences, or attacks by loosely-connected members who are not fully radicalized to violence. Secondly, neither method can shed light on how groups sequence hoaxes in their broader campaigns, nor can they elucidate strategic differences in sequencing approaches. While a time-series dimension can be added to either methodology, this limitation cannot be remedied in the present context because of the binary dependent variable: it is impossible to know if hoax use increases or decreases given certain organizational conditions (such as the number of serious events perpetrated), because a binary dependent variable can only ever indicate whether hoaxes occur at all.

Irrespective of methodology, the above-described conclusions are subject to the limitations of my sample: serious terrorism groups with a demonstrated interest in CBRN weaponry. Since many studies have sought to identify organizational differences between CBRN-users and those who have not crossed the unconventional weapons-use threshold (section 4.2.3), it is plausible that these groups reflect a biased sample. Furthermore, the dependent variable does not systematically include domestic, non-CBRN hoaxes. This chapter's non-hoaxing groups, then, may actually be hoax perpetrators—they just may do so using conventional weapons and/or against domestic targets. Chapter 9's use of CIDB data corrects for any possible bias to this end.⁹⁷ Finally, the sample's requirement that groups perpetrate at least one serious event eliminates all pure hoaxers from the sample. This group-based analysis is thus unable to shed light on which conditions lead groups to *only* ever use hoaxes.

Although probit and QCA are premised on different understandings of logic and causality, both techniques are designed to identify predictors of the dependent variable. The findings they generate should be interpreted in relation to Chapter 8's statistical cluster analysis, which looks for relationships between independent variables, amongst a sample of groups who all use hoaxes to some degree. Given how, per QCA, most groups follow rather idiosyncratic pathways to hoaxing, this cluster analysis is expected to be of particular value; it will elucidate relationships between factors, even where limited diversity impedes QCA's truth table comparisons.

⁹⁷ Despite a fair degree of sample overlap, the CIDB data does *not* reclassify as hoaxers any of the groups coded here as non-hoaxers. The dependent variable assignments applied here are thus sound.

8. A typology of terrorist hoaxers

While Chapters 6 and 7 used probit regression and QCA, respectively, to identify patterns differentiating hoaxers from non-hoaxers, this chapter uses cluster analysis to look for patterns amongst a sample of groups who all use hoaxes to some degree. These patterns are used as the foundation for a typology of terrorist hoaxers.

Cluster analysis is a statistical technique used to identify natural groupings (clusters) amongst objects in a dataset, based on their coded characteristics. In this chapter, each object is a terrorist group that has been identified in the WMDDDB, ITERATE, or GTD as carrying out both serious terrorism events and hoaxes. They are clustered according to a range of structural, motivational, and contextual characteristics—most of which flow directly from the relationships hypothesized and established in Chapter 6.

Importantly, cluster analysis is an exploratory technique used to generate, rather than to test, hypotheses (StataCorp 2015, 100). Such an inductive approach is appropriate in this context, given that there is no existing hoax literature to inform what a typology of hoaxing groups might look like. Although Chapter 5 presented a range of strategic logics for hoaxing groups and the characteristics that would, theoretically, align most closely with each, the probit and QCA results of Chapters 6 and 7 indicate that the full range of theoretical explanations for groups' strategic deployment of hoaxes is not represented across the available data. A theoretically-derived, deductive, typology of hoaxing groups, then, is unlikely to find empirical substantiation given the limited diversity of observed hoaxing groups.

Building a statistically-derived typology of hoaxing groups via cluster analysis presents two main advantages. First, it ensures that all typological categories correspond with actual, empirical observations—not theoretical “ideal-types” for which there are no real-life examples. Second, it guarantees that natural groupings will be uncovered. Particularly where there exists no *a priori* theoretical reason to expect certain clusters to emerge, identifying these natural divisions is valuable for guiding future evaluations of hoaxers and hoax behaviour. Such a typology is valuable as an empirical guide, since it effectively summarizes holistic features of which groups use hoaxes—not just the characteristics and combinations thereof that generally distinguish hoaxers from non-hoaxers. The extent to which additional cases diverge from the types established here can thus be used as an indication of sample bias; if new cases consistently exhibit alternative combinations of characteristics to the hoaxers included in the present sample, then the existing cross-national datasets that code for hoaxers underestimate such perpetrators’ diversity by virtue of their associated scopes of coverage.

A cluster-derived typology may also shed light on hoax behaviour itself. I leverage this possibility by testing whether the clusters correspond with meaningful patterns in hoax use, and whether a similar clustering emerges from a sample of terrorist non-hoaxers. In the first case, if patterns in hoax use (i.e. event sequencing and/or hoax subtype) align with different hoaxer types, counterterrorism strategies can be effectively targeted based on perpetrator characteristics, before an event has even been concluded. In the second case, if clusters are distinct to their respective hoaxing and non-hoaxing

samples, then the cluster analysis can help frame previous chapters' elaboration of the combinations of factors distinguishing hoaxers from non-hoaxers.

Finally, while previous work has presented typologies of terrorism hoaxers' strategic logics (i.e. Chapter 5; Mahoney 2017), categorizing hoaxers based on their organizational characteristics is entirely new. Even for terrorist groups more broadly, a typology including structural, motivational, as well as contextual group-based characteristics reflects a multidimensional framework rarely seen in the terrorism studies literature. From a methodological standpoint, while cluster analysis has been previously used in terrorism studies to examine patterns amongst terrorist events as well as host country characteristics, using it to identify patterns amongst terrorist *perpetrators* is a substantial innovation. Its success, to this end, contributes meaningful support in favour of further applications of cluster analysis within terrorism and security studies.

8.1 Typologies in terrorism studies

In any field of study, typologies can serve a range of purposes: "(1) to discover new relationships among the things so ordered; (2) to generate hypotheses; (3) to lead on to the development of theories; and (4) to identify areas for investigation" (G.K. Roberts qtd. in Schmid et al. 1988, 56). They thus play a frequently iterative role in theory building.

In terrorism studies, typologies are common, but rarely are applied to great purpose,⁹⁸ incorporated into cumulative theory-building or refinement efforts,⁹⁹ or reflective of empirical data.¹⁰⁰ When designed and applied appropriately, however, typologies have the capacity to contribute substantially to the study of terrorism and the practice of counterterrorism. Importantly, since much empirical terrorism research is siloed within narrow categories (e.g. limited to perpetrators of particular ideologies, or operating within particular geographic areas), theoretically and empirically sound typologies can be useful for identifying areas where extrapolations across experiences and contexts are justified, and where they are not. From a practical perspective, typologies that correspond to meaningful differences in terms of the nature of the threat posed by different “types” can be effective in guiding counterterrorism targeting and response (Marsden and Schmid 2011, 190–91). Thus far, however, few terrorism typologies have been effective in explaining or predicting outcomes (Flemming, Stohl, and Schmid 1988, 155).

This section presents an overview of typological efforts in terrorism studies. It emphasizes prominent reviews of terrorism typologies and efforts to classify terrorism perpetrators (section 8.1.1), as well as uses of cluster analysis in the field (section 8.1.2).

⁹⁸ Schmid et al. (1988) conducted a questionnaire among terrorism studies scholars, who frequently claimed that no existing typologies satisfied Roberts’ conditions, “indicating low regard for or lack of familiarity with typological efforts in the field” (56).

⁹⁹ In their review of terrorism typologies, Flemming, Stohl, and Schmid (1988) find that typologies “frequently lack useful theoretically based delineations of categories [... which] in turn, effectively inhibits interactive sharing of knowledge” (159-60).

¹⁰⁰ Flemming, Stohl, and Schmid (1988) claim that typologies’ deficiencies become particularly apparent “[w]hen matched with empirical data” (160).

8.1.1 Existing typologies of terrorism perpetrators

Although none have been applied to terrorism *hoaxers*, there exist a range of typologies for classifying terrorism perpetrators more broadly.¹⁰¹ There are so many extant typologies that book chapters reviewing them even require their own typology typologies! For instance, Schmid et al. (1988) describe ten “*fundamenta divisionis*, or principles of distinction” that are commonly used as a basis for classification: actors; victims; causes; environments; means; political orientations; motivations; demands; purposes; and targets (40). More parsimoniously, Flemming, Stohl, and Schmid (1988) offer four categories of “classificatory scheme”: those based on groups’ political background or ideology; their underlying motivation or overarching goals; their *modus operandi*; or their theatre of operation and root causes (156). More intuitively, Marsden and Schmid (2011) offer five general classes of terrorism typologies, covering the key questions of who the terrorists are (actors);¹⁰² why they use terrorism (motivation or purpose); where they stage their attacks (location); when they attack (time-frame); and how they carry out their activities (group organization). Given the thorough treatment of each possible classificatory scheme in the three reviews mentioned here, this section emphasizes those classifications that are relevant to this dissertation’s group-based

¹⁰¹ These exist in addition to efforts to classify terrorist activity, broadly writ, in relation to alternate forms of political violence and/or crime (see, for instance, Marsden and Schmid 2011; Schmid et al. 1988); and efforts to classify terrorism events themselves (as in the classificatory schemes used in the event-based ITERATE and WMMDB datasets, discussed in section 3.3.1).

¹⁰² Per Schmid et al. (1988, 43–45), actor-based typologies tend to emphasize violent perpetrators’ relationship with the state. They generally include such types as state and non-state actors; different types of state actors (developing, homogenous nation, old imperial, etc.) and the non-state groups operating within them; or lone-wolf perpetrators. This category of typology is not relevant to the present study, which is limited in scope to defining terrorist perpetrators as non-state, violent, organizational actors.

orientation: typologies of why terrorists act, typologies of how they operate, and typologies based on their historical origins and/or contemporary backgrounds.

8.1.1.1 Typologies of why terrorists act

Marsden and Schmid (2011) collapse into a single category all classifications related to “motivation, cause, purpose, objective and aim” (179). Although they acknowledge that these narrower categories are not synonymous, they are combined for their collective ability to address “the reasons underlying terrorist behavior” (Ibid.). Most terrorist typologies belong to this general category of classification system (Chenoweth and Lowham 2007, 346).

The most common dimension applied in this type of classification scheme is groups’ political orientation or ideology. These typologies include such categories as left- and right-wing, religious, or ethno-nationalist, as well as more nuanced types (e.g. “extreme left”; “Marxist”; “Maoist”; “supremacist”). Flemming, Stohl, and Schmid (1988, 158–61) show that, despite their popularity, ideology-based classificatory schemes are subject to a number of flaws. First, the concepts they encompass are usually specific to particular studies and how they choose to define and label the categories. Groups may be categorized differently in different typologies (e.g. “extreme left” versus “leftist”), due to the often subjective and arbitrary nature of coding. In other cases, classifiers are employed but not defined (e.g. what constitutes “orthodox” communism?). These

problems impede knowledge cumulativeness across studies.¹⁰³ Second, such classifications are generally based on unsubstantiated assumptions. In those rare cases where ideology is linked to particular behavioural outcomes, theoretical explanations (e.g. in differentiating discrimination in targeting logic between left- and right-wing terrorists) do not hold when tested empirically. As such, the authors argue that such classifications on their own are “incomplete” (160) and “of limited utility” (162).

Typologies of why terrorists act may also be based on actors’ intermediate objectives, such as provocation, morale-building, and publicity—in other words, those medium-term goals that connect tactical outcomes with ultimate objectives. While Flemming, Stohl, and Schmid (1988, 162–65) laud such purpose-based typologies as being “more uniform” than those based on ideologies, they find that they suffer from similar shortcomings. Additionally, such approaches fail to give insight into the terrorists’ *modus operandi*, even where long-range goals are incorporated: “the manner in which these goals are classified [...] is far too vague to be of much value in illuminating the relation between terrorist type and terrorist behavior” (Flemming, Stohl, and Schmid 1988, 165).

8.1.1.2 Typologies of how terrorists operate: methods, means, and organizational structure

Typologies of terrorist methods and means tend to emphasize terrorism events themselves, not terrorist actors. They emphasize such categories as: targeting logics (e.g.

¹⁰³ Flemming, Stohl, and Schmid (1988) present a compelling demonstration of these flaws by comparing classifications of a consistent set of groups across four popular typologies. Groups’ allocation across the different typologies diverge substantially.

discriminate versus indiscriminate); weapon types (e.g. bombs or firearms); or tactics (e.g. car bombings or mass shootings). When means and methods are the only dimension along which objects are classified, typologies are unable to contribute to hypothesis development regarding links between terrorism types and known terrorist groups (Flemming, Stohl, and Schmid 1988, 167). Even where objectives are purely descriptive, because terrorist behaviours change over time, these approaches do not offer “an accurate characterization of terrorist activity for all types of terrorists at all times” (Ibid.)

Typologies have also been developed to address a range of emerging threats—such as agro-, cyber-, and CBRN terrorism—as subsets of terrorist activity (Marsden and Schmid 2011, 188). As I argue in section 2.4.1, hoaxes could be typologized in this sense along with suicide terrorism as modes of terrorist activity. Although there have been efforts to classify subsets of terrorist activity, there appears to be only minimal effort to typologize perpetrators within these subsets (as is the goal of the present hoax-centric typology-building effort), despite meaningful expectations that perpetrators of certain types of terrorism will be differently constituted.¹⁰⁴

Unlike the reviews from the 1980s, Marsden and Schmid’s (2011) overview of terrorism typologies includes classifications of groups’ “organizational make-up” in their discussion of typologies that emphasize how terrorist groups operate (186-91). Their inclusion of organizational structure typologies reflects a broader shift in terrorism

¹⁰⁴ The three typological reviews identify only one typological effort designed to unpack terrorist perpetrators’ interest in a single tactical method: Goldaber’s typology of hostage-takers (summarized in Marsden and Schmid 2011, 172–75; and Schmid et al. 1988, 54). Goldaber’s typology incorporates a range of psychological, criminal, and political conditions, as well as the “who, what, where, when, why, and how” of their activities. It has been evaluated as “one of the very best typologies available, offering great usefulness to practicing hostage negotiators” (Marsden and Schmid 2011, 173).

studies toward increased understandings of groups' internal dynamics. The classification of groups according to their organizational structure, the authors argue, is reflective of on-the-ground evolutions in the nature of terrorist groups' organization (186). Arquilla and Ronfeldt (2001) were the first to typologize terrorist groups according to their organizational structure. Their three basic network types (chain; star or hub; and all-channel) are the basis of the four-category typology structure developed by Kilberg (2011; 2012), whose coding methodology I apply in this dissertation.¹⁰⁵

8.1.1.3 Typologies based on groups' origins

A final category of terrorist typologies is that dedicated to classifying groups according to their historical (political) roots, contemporary backgrounds, geographic origins, and/or theatre(s) of operations. While these typologies exhibit "a greater degree of theoretical maturity" than other types (Flemming, Stohl, and Schmid 1988, 167), they also suffer from similar classificatory flaws to those described above. It is frequently unclear where groups should fit, because categories are either too broad, or unrepresentative of reality (Ibid.)

8.1.2 Cluster analysis in terrorism studies

Interestingly, despite their impressive breadth and depth of nuance, none of the three most comprehensive reviews of efforts to typologize terrorists and terrorist activity (Marsden and Schmid 2011; Fleming, Stohl, and Schmid 1988; Schmid et al. 1988)

¹⁰⁵ In Kilberg's (2011; 2012) framework, structure is measured along three dimensions: leadership, functional differentiation, and centralized command and control. These three features combine to reflect four organizational types (see footnote 61).

incorporate examples of cluster analysis, nor do they acknowledge the methodological benefits that cluster analysis can provide for inductive typology-building. The vast majority of terrorism typologies are deductively derived. This emphasis on deductive approaches appears misplaced, given the degree to which inductive typology construction can contribute to theory building, and given that building a typology that encompasses all types—an outcome guaranteed by cluster analysis—may well be “[t]he aspirational outcome of an inductive approach to typology construction” (Marsden and Schmid 2011, 193; emphasis added).

A handful of studies have applied cluster analysis to the study of terrorism. While some examine events data to identify patterns in terrorist means and methods (Adnan and Rafi 2015; Dwivedi et al. 2014), others examine country characteristics to uncover enabling conditions for foreign fighters (Hewitt and Kelley-Moore 2009) and terrorism more broadly (Tikuisis 2009). Another has examined relationships between terrorists’ ideology and their attack methodology (Strang 2015). Unfortunately, these studies are methodologically-oriented, and of limited theoretical utility. Published predominantly in computer science journals, they exhibit inadequate substantive knowledge of terrorism in practice and in theory. They are correspondingly built upon questionable choices in variable selection and measurement; and they rarely interpret findings in any substantive or applied manner. Despite these substantial limitations, however, the mathematical robustness of their approaches indicate that clustering methods have the potential to yield methodologically-sound insights into terrorist activity.

Within the more conventional body of terrorism studies literature, cluster analysis appears to have only been applied once, in Chenoweth and Lowham's (2007) "On Classifying Terrorism." In this article, Chenoweth and Lowham cluster terrorism incidents based on perpetrators' motives, weapons, targets, location, and destructiveness. The primary benefit of this approach, they argue, is that it is able to capture the contextual elements that are frequently omitted in large-sample studies while achieving the replicability and extrapolation that are frequently precluded by small-*n* thick descriptions (Chenoweth and Lowham 2007, 345). In other words, it "provides more comprehensive ways of classifying terrorism incidents across cases while maintaining sensitivity to the context of the cases" (Ibid., 346).

To conduct their cluster analysis, Chenoweth and Lowham use the U.S. State Department's *Chronology of Significant Terrorist Incidents* dataset, which includes 259 incidents of terrorism targeting Americans between 1961 and 2003, and 45 dichotomous variables. Due to reporting bias that emphasizes attacks with an American connection, the authors recognize that this dataset is subpar "from a purely scientific standpoint" (346). Accordingly, they argue that the dataset represents a "hard case for the argument that cluster analysis can provide novel information": if even a biased dataset reveals interesting patterns based on attack characteristics, then "this method should be applied to other, more objective and complete, databases" (Chenoweth and Lowham 2007, 348).

Their dataset passes this hard test. Chenoweth and Lowham's cluster analysis—based on the Jaccard coefficient as their measure of similarity (defined in section 8.3.1), and average between groups linkage as their clustering algorithm (defined in section

8.3.2)—identifies ten core clusters of terrorism event types. They assess core characteristics of each cluster based on the proportion of groups exhibiting each characteristic within each cluster, and find that group type (including a selection of ideologies and possible relationships with the state) were only crucial in two of the ten clusters: a surprising finding, given the prevalence of actor-based typologies in the literature. Means, motives, and methods (in particular, the use of bombings, claims of responsibility, and target choices) were more prominent defining characteristics of each cluster. Although they do not include indicators for group structure or organization, their cluster outputs highlight the importance of typologies that classify terrorism along multiple dimensions.

8.2 The data

This section outlines the classification's unit of analysis, the sample from which objects are drawn (section 8.2.1), as well as the variables that guide the clustering (section 8.2.2).

8.2.1 The unit of analysis and sample: All hoaxing groups

Consistent with earlier chapters, the unit of analysis is again the terrorist group. As in the previous probit and QCA analyses, groups are only included in the sample if they are responsible for at least one serious terrorism event across ITERATE, GTD, or the WMDDB. However, the sample differs from that used in Chapters 6 and 7 along two primary dimensions. First, only serious groups who also carried out hoaxes (as identified via ITERATE, GTD, or the WMDDB) are included in the analysis. The primary clustering sample

thus excludes groups who never hoax. Second, the clustering sample is not restricted to groups with an interest in CBRN terrorism (as defined by their inclusion as perpetrators in the WMDDDB). In addition to the 23 hoax perpetrators counted in the WMDDDB-centric sample of previous chapters, the clustering sample includes an additional 13 groups who committed hoaxes according to ITERATE, and a further 5 groups identified as hoaxers per GTD event descriptions.¹⁰⁶ The total sample of groups ($n=40$)¹⁰⁷ thus encompasses the population of serious terrorism groups identified in any of the three cross-national datasets as also conducting terrorism hoaxes. The names of groups included in the sample are listed in Appendix Table B.1.

The sample is designed in this way to take advantage of the full breadth of cross-national terrorism hoax data. Previous chapters were limited to only those groups demonstrating an interest in CBRN weaponry, because they required a sample of non-hoaxers, too. As described in section 6.2.1, it is infeasible to collect coding details for group structure and campaign context across the full population of terrorist groups present in ITERATE and/or GTD. By focusing only on groups who hoax, I am left with a far more manageable sample size, making it possible to evaluate *all* hoax perpetrators captured by the cross-national datasets, without the limiting requirement that they also

¹⁰⁶ As in Chapter 6, it also includes a handful of groups who committed neither CBRN nor transnational hoaxes, but whose hoax activity is elucidated in event descriptions in the WMDDDB and/or GTD. Five groups are included beyond Chapter 6's original sample for having committed hoaxes identified per the GTD's weapon type category "fake weapons." While not all "fake weapons" incidents are by definition hoaxes, I verified these "fake weapons" incidents' event descriptions for consistency with the definition of hoaxes provided in Chapter 2 to ensure that they align conceptually with the hoax phenomenon.

¹⁰⁷ A 41st group, the Arab Revolutionary Council (ARC) is excluded from the sample because structure details are unknown; the clustering mechanism eliminates objects with missing data.

demonstrate an interest in CBRN weaponry (which was required of the sample used previously).

8.2.2 Clustering variables

While most typological efforts fail to appropriately consider variable selection (Flemming, Stohl, and Schmid 1988, 170), deciding which variables to include can be considered the first step in designing a cluster analysis methodology. This section outlines my decisions regarding variable selection for the cluster analysis. All variables are summarized in Table 8.1, with basic coding decisions presented in Appendix Table B.2.

As a starting point, this typological effort begins with the theoretically-informed group-based characteristics selected for use in previous chapters' regression and QCA models. Some of these variables failed to generate significant results in their respective analyses, indicating that they were unhelpful in distinguishing hoaxers from non-hoaxers. However, they are retained here since they may still be relevant in terms of defining classes of perpetrators within the hoaxing group sample. For the most part, these reprised variables are incorporated into the cluster analysis in their originally-coded form. However, because the proximity coefficient selected (discussed in section 8.3.1) ignores negative matches—where each object scores “0” for the attribute’s presence—in the clustering process, one important variable had to be recoded: hierarchy.¹⁰⁸

¹⁰⁸ “By convention, we shall always code the most important outcome (which is typically the rarest one) by 1, and the other by 0” (Kaufman and Rousseeuw 1990, 26).

For hierarchy, 1s and 0s are reversed, so that the clustering process actually measures the typological relevance of *non-hierarchy*. I chose to reverse the outcomes firstly because non-hierarchy is a less common attribute than hierarchy, and secondly because more nuanced degrees of hierarchy can also be tested here, using the functional differentiation and centralized command and control variables.¹⁰⁹

Coding for the variable indicating the number of serious events perpetrated per group also had to be modified. Although some clustering algorithms are capable of handling mixed types of variables, for the present analysis I preferred to dichotomize this otherwise-continuous variable, so as to work with a full binary set. Converting quantitative data to binary categories is advisable when the original quantitative variable is “noisy” (Legendre and Legendre 1983), and/or when a known threshold value would justify a binary split (Everitt et al. 2011, 68). In the case of the measure of serious events perpetrated, both conditions hold. First, the variable can be considered noisy due to inconsistencies in coding coordinated events that undermine the count’s precision. Second, there is a known threshold to justify dichotomization: Chapter 7’s set analyses show that committing more than three attacks effectively distinguishes hoaxers from non-hoaxers, given other organizational conditions. Hoaxing groups committing three or

¹⁰⁹ The leadership indicator is excluded from the cluster analysis because it is perfectly correlated with the indicator for hierarchy. In a clustering algorithm, when two variables are highly correlated, one should be removed (Kaufman and Rousseeuw 1990, 62). Furthermore, in the present hoaxer-only sample—despite the possibility that a “hierarchical” group might lack leadership, if it is otherwise centrally controlled—there is a perfect correlation between the two indicators: all of the hierarchical groups also have leadership structures, and vice versa. Since only one of the two variables can be retained, non-hierarchy is preferred over leadership because of its stronger performance in the regression and QCA analyses, and important conceptual relevance.

fewer events would thus be outliers within the hoaxing sample, and offer a valuable classificatory dimension.¹¹⁰

I also tested two new variables for describing groups' campaign context. The first indicated whether the group's serious activity was restricted to the domestic sphere (or, its inverse: whether the group ever engaged in serious transnational incidents). The second indicated whether the group ever committed a serious incident involving CBRN agents (indicating capability in particularly extreme modes of terrorist activity). If either of these variables are important differentiators among hoaxing groups, they can shed light on the degree to which relying on hoax data predominantly from the transnationally-oriented ITERATE and/or CBRN-oriented WMDDDB may bias results.

Because cluster analysis is not constricted in terms of degrees of freedom (cf. probit) or exponential combinatorial possibilities and limited diversity (cf. QCA), a clustering model can incorporate far more variables than was possible in either of the two preceding empirical chapters.

¹¹⁰ From a set-theoretic perspective, anything above three events is unimportant variation in terms of determining which groups will hoax. However, there is a wide range in the number of serious events perpetrated by each group in the present sample of hoaxers: 1 - 2,492. Distinctions within this range may be useful for identifying classes of hoaxers, particularly given that there is reason to suspect that groups perpetrating large numbers of attacks use hoaxes differently from those who perpetrate fewer attacks (see section 6.1.3). Accordingly, I tested indicators for the number of events perpetrated at both the low and high ends of the spectrum: very infrequent serious event perpetrators (three or fewer), and very frequent serious event perpetrators. Doing so allowed me to capture groups' important contextual attributes, without the excessive complexity that a mixed variable-type clustering approach would entail. Ultimately, I found that the measures indicating large numbers of serious events were insignificant in each cluster model tested, and so they were not included in the final model presented and discussed here.

Table 8.1 Description of variables for clustering (n=40)

VARIABLE	DEFINITION	DESCRIPTION (% = 1)
IDEOLOGY		
<i>Single-issue</i>	= 1 if the group has an anti-abortionist, animal rights, or environmentalist motivation	12.5%
<i>Purely religious</i>	= 1 if the group is solely motivated by religion; 0 if the group has any other motivation (even if there is a religious component)	20%
<i>Purely ethno-nationalist</i>	= 1 if the group is solely motivated by ethno-nationalism; 0 if the group has any other motivation (even if there is an ethno-nationalist component)	25%
<i>Economically leftist</i> ¹¹¹	= 1 if the group's ideology contains an economically leftist component; 0 otherwise and/or if also <i>single-issue</i>	35%
<i>Economically rightist</i> ¹¹²	= 1 if the group's ideology contains an economically rightist component; 0 otherwise and/or if also <i>single-issue</i>	0%
STRUCTURE		
<i>Functional differentiation</i>	= 1 if there is evidence that the group involves at least one clearly-defined cell or department dedicated to a specialized function (e.g. political wing, bomb center, communications department)	32.5%
<i>Centralized command and control</i>	= 1 if there is evidence that the group's actions are coordinated from a central executive	27.5%
<i>Lack of hierarchy</i>	= 1 if the group lacks any form of hierarchy (leadership or centralized command and control); 0 if there is any evidence of hierarchy	37.5%
CAMPAIGN CONTEXT		
<i>Radical splinter</i>	= 1 if the group originated as a radical splinter/offshoot of a less-violent group or movement	27.5%
<i>Purely domestic</i>	= 1 if the group never committed any transnationally-oriented incidents (i.e. the group is not responsible for any incidents captured by ITERATE)	27.5%
<i>CBRN capable</i>	= 1 if the group ever committed a serious incident involving CBRN weaponry	40%
<i>3 or fewer serious events</i> ¹¹³	= 1 if the group committed 3 or fewer serious incidents	17.5%

¹¹¹ Per CIDB guidelines, economically leftist ideologies emphasize "redistribution of wealth by the government and nationalization of industry" (TSAS 2015, 3).

¹¹² Per CIDB guidelines, economically rightist groups emphasize "extreme laissez-faire policies, small government, or policies that favor status quo economic elites" (TSAS 2015, 4).

¹¹³ The mean value of the number of serious events perpetrated per group is 347.975 events; the median value is 75; and the population standard deviation is 648.4 events. As a further indication of the distribution of these values: the 25th percentile (bottom quartile) is 9 events; the 50th percentile is 75; the 75th percentile (top quartile) is 229.5; the 90th percentile is 1,532; and the 95th percentile is 2,073.5. Three events corresponds to the 17th percentile.

8.3 Methodology

Cluster analysis has been aptly described as “the art of finding groups in data” (Kaufman and Rousseeuw 1990, 1). Its methodological goal is to uncover clusters of objects in a dataset, so that objects within the same cluster are as alike as possible, and as different from the objects in other clusters as possible.¹¹⁴ Clusters are based solely on the data contained in independent variables; unlike previous chapters’ probit and QCA analyses, they are not intended to differentiate variables that are associated (or not) with a dependent variable. The present cluster analysis is conducted using the Stata 14 statistical software package.

In addition to the question of which variables to use (discussed in section 8.2.2), three key choices guide the way in which a cluster analysis is performed: 1) choosing which measure of proximity to employ; 2) choosing which clustering algorithm to apply; and 3) choosing (how to determine) the optimal number of clusters. The following subsections address each choice in turn, with a fourth section dedicated to providing guidelines for cluster interpretation. Broadly speaking, my approach follows the major methodological decisions of Chenoweth and Lowham (2007): I use the Jaccard coefficient

¹¹⁴ Note that, while cluster analysis creates classifications of groups, it is distinct from discriminant analysis, which is also sometimes referred to as “classification.” At a difference to cluster analysis, the classification involved in discriminant analyses is used when group categories are already known. The discriminant analysis then describes the differences between these categories, and assigns objects of unknown group membership to the group whose characteristics most closely resemble its own. In other words, “cluster analysis [...] actually establishes the groups, whereas discriminant analysis assigns objects to groups that were defined in advance” (Kaufman and Rousseeuw 1990, 2).

as my proximity measure; average between-groups linkage as my clustering algorithm; and intra-cluster correlations to identify meaningful patterns between perpetrator types.

As a point of linguistic clarification: the entities that are grouped together in a cluster analysis are called “objects.” In the present chapter, each object to be clustered is a terrorist *group* (elaborated in section 8.2.1). This language is ambiguous, since clusters of objects (as determined by the cluster analysis) may also be referred to as “groups.” To avoid any confusion, I will use “group” or “organization” to refer to the clustered object itself (the terrorist group), and the words “cluster,” “class,” or “type” when referring to the groupings of terrorist organizations revealed by the cluster analysis.

8.3.1 Step 1: Choosing a measure of proximity

“Proximity” is a general term for describing “inter-object *similarities, dissimilarities, or distances*” (Everitt et al. 2011, 5). Once appropriate variables have been selected, the first step in a cluster analysis is deciding which measure of proximity (referred to as a “coefficient”) to employ. Even when applied to the same set of objects, different coefficients often generate different clustering solutions, and so it is important to make an informed choice (Everitt et al. 2011, 69). Since there is no algebraic rule for determining the appropriate coefficient (Everitt et al. 2011, 69; Gower and Legendre 1986, 31), the decision must take into account the nature of the data and types of variables used; theoretical intuition regarding how negative matches (i.e. co-absences) should be handled; and assessments of the appropriate weighting of variables and matches in the calculation.

The types of data within a dataset dictate which types of coefficients are appropriate. Coefficients of similarity—which reflect how closely objects resemble one another—are used when dealing with a dataset comprised only of binary indicators.¹¹⁵ There are many possible coefficients for measuring similarity in binary data (described extensively in Everitt et al. 2011, 47; Gower and Legendre 1986, 23). Here, it is important to understand only the broad dimensions along which they differ: 1) the way in which they handle negative matches (also known as zero-zero [0-0] matches; or co-absences); and 2) the degree to which they weight different types of matches or variables.

Table 8.2 counts the range of possible binary outcomes for two objects, i and j . When both objects display a given characteristic, an a -type (also known as a positive, or one-one [1-1]) match occurs. Analogously, when both objects *lack* a given characteristic, a d -type (or negative) match occurs. Types b and c reflect mis-matches, where one object displays the characteristic and the other does not. The sum calculated in the table's bottom right, p , indicates the total number of variables on which objects i and j may agree (or disagree).

¹¹⁵ When a similarity coefficient for two objects (s_{ij}) equals one, the objects agree on all variables; they are identical. A similarity coefficient of zero indicates that the objects differ maximally on all variables. Similarity measures can be easily converted to (or derived from) dissimilarity measures: the similarity of two objects (usually represented in notation as objects i and j) is equal to 1 minus dissimilarity ($s_{ij} = 1 - d_{ij}$). A small dissimilarity is thus the same as a large similarity, and a large dissimilarity is the same as a small similarity. Stata's algorithms use the language of "dissimilarity," so when a measure of "similarity" is used to measure proximity, the software simply transforms the measure into dissimilarity. Since similarity measures are either represented as 0 to 1, or -1 to 1, corresponding similarities range from 1 to 0 or 2 to 0. Following the cluster analysis, fusion values (the measure on a dendrogram's vertical axis reflecting the proximity at which groups are combined in the hierarchy) are converted back to similarities (StataCorp 2015, 105).

Table 8.2 Counts of binary outcomes for two objects¹¹⁶

		Object <i>i</i>		Total
		1	0	
Object <i>j</i>	1	<i>a</i>	<i>b</i>	<i>a + b</i>
	0	<i>c</i>	<i>d</i>	<i>c + d</i>
Total		<i>a + c</i>	<i>b + d</i>	$p = a + b + c + d$

When variables are symmetric (i.e. one-one and zero-zero matches are logically co-equal), it makes sense to use a similarity coefficient that considers both positive (*a*-type) and negative (*d*-type) matches equally. The similarities of such symmetric variables are said to be “invariant,” meaning that it does not matter which category is coded as zero or one: the result will be the same. While there exists a range of invariant similarity coefficients, the most simple and frequently used is the simple matching coefficient, $s_{ij} = (a + d)/(a + b + c + d)$, which takes the sum of all positive and negative matches, divided by the total number of possible matches and mis-matches (Everitt et al. 2011, 46–47).

When variables are asymmetric, however, the similarities they generate are not invariant: positive and negative matches are of unequal importance in determining how similar objects are. Such asymmetry occurs frequently when dealing with rare attributes: “While it can be said that two people with [the attribute] have something in common, it is not so clear if the same can be said of two people who do *not* have it” (Kaufman and Rousseeuw 1990, 25–26). In other words, while mutual presence of an attribute is instructive, it is up to the investigator to decide whether mutual absence of the same

¹¹⁶ Adapted from Everitt et al. (2011, 46) and Kaufman and Rousseeuw (1990, 22–23).

attribute contributes to the objects' similarity. In many contexts, attributing high similarity to two objects "simply because they both lack a large number of attributes" is not sensible (Everitt et al. 2011, 46–47). For this reason, coefficients exist that weight positive matches more heavily than negative matches; the most common non-invariant coefficients exclude negative matches altogether (Kaufman and Rousseeuw 1990, 26).

Finally, certain coefficients allow variables to be weighted, so as to increase or decrease the role they play in determining objects' similarity. Weighting can be determined by the investigator's subjective judgment, or by features of the data matrix itself (Everitt et al. 2011, 63). Both approaches, however, are unsatisfactory. First, weights based on subjective evaluations may simply reify the investigator's prior beliefs regarding how the data should be classified (Everitt et al. 2011, 66–67). Second, popular strategies for determining variable weights based on features of the data (such as weighting variables based on their standard deviations, or applying equal weighting on a large number of variables in an effort to capture any that are relevant) "appear to be ineffective" (Everitt et al. 2011, 67). There are no currently-identified optimal criteria for using the data matrix to determine weights (Gnanadesikan, Kettenring, and Tsao 1995).

Given the above-described conditions, I adopt the Jaccard coefficient (Jaccard 1901; Jaccard 1908)¹¹⁷ – the most famous, and also most simple, non-invariant coefficient

¹¹⁷ Despite its namesake, the Jaccard coefficient was actually first discovered by Gilbert (1884; StataCorp 2015, 543).

for measuring similarity across binary variables. The Jaccard coefficient is like the matching coefficient, but it excludes type-*d* negative matches:¹¹⁸

$$s_{ij} = \frac{a}{a + b + c}$$

In the context of the present study, a coefficient that excludes co-absences is preferred, due to the theorized and empirically demonstrated preponderance of asymmetric variables described in Chapter 7. Finally, that the Jaccard coefficient does not weight variables is an advantage in the present context, since—although previous chapters’ findings identified factors distinguishing hoaxers from non-hoaxers—I have neither theoretical nor empirical expectations regarding how these factors will differentiate hoaxers from one another; and because there is no agreed-upon method for weighting variables based on their mathematical properties.

Ultimately, the Jaccard coefficient conforms with Everitt et al.’s (2011) advice that investigators select “the simplest coefficient applicable to a data set [...] to ease the possibly difficult task of interpretation of final results” (69). That it also aligns with Chenoweth and Lowham’s (2007) coefficient choice is an added benefit: the authors’ success in applying the Jaccard coefficient on terrorism data bolsters my confidence that it will contribute to meaningful clusters in the present analysis.

¹¹⁸ Note that, in Stata, if both objects are coded as all zeros, the Jaccard binary similarity coefficient goes undefined. In this case, Stata treats the undefined vector as if it were equal to one, meaning that the groups are identical. This only occurs when *both* objects are all-zero; if one of the objects contains any ones, then the objects’ similarity will be equal to zero (StataCorp 2015, 543).

8.3.2 Step 2: Choosing a clustering algorithm

Investigators must also decide which clustering algorithm to apply to the proximity matrix that emerges from their choice in proximity coefficient. Although there are “infinite” ways to perform a cluster analysis (StataCorp 2015, 101), these can be broadly classified into two types: hierarchical or partitional. For partitional algorithms, the investigator provides the desired number of clusters, k , upfront. The cluster analysis then builds a single partition in one step, dividing all objects into k clusters in the best way possible. Conversely, hierarchical algorithms are appropriate when the structure of the data is unknown. A hierarchical clustering considers all possible values of k at each of the analysis’s multiple stages: it considers both $k = 1$ and $k = n$, and every possible value of k in between (Kaufman and Rousseeuw 1990, 44). Since I do not hold *a priori* notions regarding what the clusters will look like or how many there will be, using partitional methods would require me to impose an arbitrary structure that would be antithetical to uncovering true underlying patterns in my data. Hierarchical clustering, on the other hand, will run without input parameters regarding k , and has the added benefit of being better-suited to binary data.

Hierarchical cluster analyses present a further advantage in that they can be easily represented in a two-dimensional diagram called a “dendrogram.” Dendrograms illustrate the process by which groups are split or combined at each stage. Branches of the diagram thus explicitly convey how clusters are nested, and the degree of similarity

or difference between each joined object.¹¹⁹ While a dendrogram may look like a simple tree diagram, where each node represents a single variable's bifurcation, it is not immediately obvious from a dendrogram which variables are responsible for which splits (see section 8.3.4 on interpreting the cluster). This ambiguity is because most clustering algorithms are "polythetic," meaning that they use all variables simultaneously, as opposed to "monothetic" approaches which would use only the single most relevant variable at each step of the clustering.¹²⁰

Within hierarchical cluster analysis are two broad types: agglomerative or divisive. Agglomerative hierarchical clustering is a bottom-up approach, beginning with each case on its own (i.e. $k = n$), successively combining with the next-nearest cluster until all cases belong to the same group ($k = 1$). Divisive hierarchical clustering, on the other hand, begins from the top down: all cases initially belong to the same cluster ($k = 1$), and are successively split until each case exists on its own ($k = n$). While divisive methods offer the benefit of immediately revealing the data's structure, evaluating every possible combination for division at every stage of the analysis is computationally demanding. Accordingly, divisive clustering algorithms are rarely available (Stata does not include any) or applied (Kaufman and Rousseeuw 1990, 49; StataCorp 2015, 103). Accordingly, the present cluster analysis uses an agglomerative clustering approach.

¹¹⁹ The greater the vertical distance between nodes (reflected in what is called "stem length"), the more dissimilar the groups being combined are. Short stem lengths between clusters indicate a low degree of distinction; long stem lengths indicate that the clusters are well separated.

¹²⁰ While monothetic approaches would provide the benefit of easy interpretation (it is clear which variable caused which split) and quick computation on large data sets, they are unable to distinguish symmetric binary variables from asymmetric ones (Kaufman and Rousseeuw 1990, 50), and are thus unhelpful in the present context, since many of the indicators under consideration are asymmetrical.

Even within agglomerative hierarchical cluster methods, however, there are a range of possible algorithms that can be applied. These are known as “linkage methods.” Again, there is no single optimal approach to determining which one should apply. The three most frequently used linkage methods are single linkage, complete linkage, and group average linkage—so named to reflect the way in which inter-cluster similarity is defined.¹²¹ Figure 8.1 presents a visual representation of these three approaches.

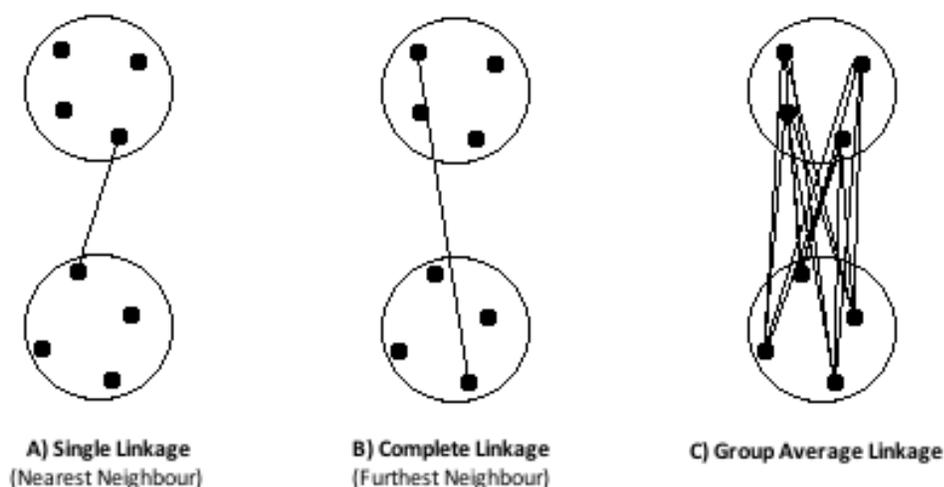


Figure 8.1 Visual comparison of linkage methods¹²²

Single and complete linkage (known as the “nearest” and “furthest” neighbour approaches, respectively) reflect extreme and opposite approaches. In single linkage, inter-group proximity is defined by the greatest similarity between two objects, one from each cluster. While this method can be meaningfully used to identify outliers,¹²³ the single linkage approach’s tendency to result in “chaining”—whereby stringy clusters result from

¹²¹ There are many commonly-used synonyms for each approach, enumerated in StataCorp (2015, 106).

¹²² Figure 8.1 is adapted from Kaufman and Rousseeuw (1990, 47).

¹²³ Outliers “are left as singletons if they are sufficiently far from their nearest neighbor” (Everitt et al. 2011, 81).

noisy data points merging clusters that are otherwise dissimilar—limits its utility in most applications. Complete linkage, on the other hand, defines inter-group proximity by the largest distance between any two objects, one per cluster. It has a tendency to find compact and roughly equally-sized clusters. Neither of these approaches takes the actual cluster structure into account (Everitt et al. 2011, 79).

Conversely, group average linkage represents an intermediate option, defining inter-group proximity as the average (dis)similarity of all objects within each cluster. Accordingly, it avoids the chaining pitfall of single linkage, and is able to handle clusters that are not as ball-shaped as complete linkage would require. Group average linkage is thus fairly robust, and performs well in many situations (Kaufman and Rousseeuw 1990, 47; StataCorp 2015, 103). Given that I hold no preconceived notions regarding the structure of my data nor its likely clusters, group average linkage is an appropriate choice. Furthermore, it has already been applied in terrorism studies to significant effect by Chenoweth and Lowham (2007). The clustering method applied is thus a hierarchical, agglomerative, group average linkage algorithm.

8.3.3 Step 3: Choosing the optimal number of clusters

Since hierarchical clusters are nested, with the possible values for k ranging from 1 to n , determining the optimal number of clusters is important for interpreting the clustering results. Choosing the optimal number of clusters effectively identifies the cutline on the dendrogram, so that it becomes clear which objects should be classified together. Statistical tests can be used to determine the number of clusters. Milligan and Cooper

(1985) describe over thirty such statistically-informed means of selecting the number of clusters, but Stata incorporates only two: the Caliński and Harabasz pseudo-F index (also known as the Caliński-Harabasz test: Stata’s default “cluster stop” command), and the Duda, Hart, and Stork index stopping rule based on pseudo-T-squared (also known as the Duda-Hart $Je(2)/Je(1)$ test).¹²⁴ I use both the Duda-Hart and Caliński-Harabasz tests to ascertain the optimal number of clusters.¹²⁵ Where the tests indicate different optimal numbers of clusters, I make qualitative determinations of the appropriate number of clusters based on visual assessment of the dendrogram.

8.3.4 Step 4: Interpreting the clusters

Following the approach taken in Chenoweth and Lowham (2007), I interpret the clusters based on the characteristics that are most consistently present within each cluster. The “cluster generate” command is used to label each terrorist organization with an identifier for the cluster to which it belongs. I then tabulate the proportion of groups within each cluster displaying each of the input variables, ranging from 0 to 100%. A score of 0% means that no groups within the cluster display the characteristic in question; a score of 100% means that all of them do.

While this chapter’s primary purpose is to classify hoaxers amongst themselves, the final stage in interpreting my results is to check whether these clusters are distinctive

¹²⁴ It is possible to add alternative stopping rules to Stata using the “cluster stop” command.

¹²⁵ The Caliński-Harabasz stopping rule suggests that the optimal number of clusters is indicated by the cluster value with the largest pseudo-F value. The Duda-Hart stopping rule suggests that the optimal number of clusters is indicated where one of the largest $Je(2)/Je(1)$ values corresponds with a low pseudo-*T*-squared value—particularly when this pseudo-*T*-squared value is bookended with larger pseudo-*T*-squareds.

to hoaxers, or whether they are present in a broader sample of terrorist groups. The only feasible way to test the hoaxer typology's distinctiveness is to run the clustering model against the sample of non-hoaxers with an interest in CBRN weaponry (n=54),¹²⁶ whose characteristics were coded for in the previous chapters' analyses. Doing so requires one minor tweak in the clustering variables: the variable measuring groups' serious CBRN capability is switched from indicating the presence of a CBRN capability to a lack thereof, in order to capture the more infrequent outcome (79.63% of groups in the sample are CBRN-capable).¹²⁷ The need to reverse this variable highlights the sample's bias: it is drawn from a sample of groups with a demonstrated interest (serious or otherwise) in CBRN weaponry.¹²⁸ There is thus a higher proportion of CBRN-capable groups within this sample than would be present in a broader sample of all non-hoaxing terrorist groups, since CBRN-interested groups are themselves a minority amongst all terrorism perpetrators. The sample of non-hoaxers is thus neither perfectly analogous to the sample of hoaxers, which includes perpetrators who never demonstrated an interest in CBRN activity, nor to the population of all terrorist perpetrators. Still, following

¹²⁶ While the full, existing, non-hoaxing sample includes 58 non-hoaxing groups, four are culled here because they lack structure coding: Black December; Justice Guerrillas; Commandos for the Defense of Western Civilization [CODECO]; and September 11.

¹²⁷ It should be noted that a very low threshold is used to evaluate CBRN capability. As discussed in section 8.2.2 (and defined in Table 8.1), "CBRN-capable" refers only to whether the group was ever responsible for a serious incident involving CBRN weaponry, per the WMDDDB. There is no requirement that such serious incidents involve a high degree of tactical effectiveness or lethality, only that the group actually possessed CBRN materials and took meaningful steps to somehow integrate them in terrorist activity.

¹²⁸ Since hoax perpetrators are excluded from this sample, the only CBRN-interested perpetrators in the present sample who are *not* CBRN capable are those who threatened or plotted CBRN use, without actual agent possession.

Chenoweth and Lowham (2007), if even a biased dataset reveals interesting patterns, meaningful inferences may still be drawn.

8.4 Results

Results from the cluster analysis are presented in two subsections: 1) describing the clusters that emerged for the sample of terrorist groups; and 2) comparing these clusters to those that emerged amongst the sample of CBRN-interested groups who never hoaxed. In both cases, I evaluate the degree to which the combinations of factors revealed in Chapter 7's QCA analysis—distinguishing hoaxers from non-hoaxers—align with the clustered characteristics revealed in the present analysis, distinguishing amongst classes of hoaxers (and, separately, amongst non-hoaxers). Since QCA only reveals the combinations of characteristics sufficient for distinguishing hoaxers from non-hoaxers, there may be overlap—but there is no guarantee or theoretical expectation. I also evaluate the degree to which clusters appear contingent on their corresponding hoax data sources: the transnationally-oriented ITERATE, and/or the CBRN-oriented WMDDDB.

8.4.1 A statistically-derived typology for hoaxers

Given that the goal of the present chapter's cluster analysis is to inductively construct a typology for terrorist hoaxers, encompassing all possible hoaxing group types, the "best" cluster model was one that yielded relatively balanced group membership with no outliers. This model incorporated the twelve clustering variables summarized in Table 8.1,

and yielded five general classes of terrorist hoax perpetrators, summarized in Table 8.3.¹²⁹ Figure 8.2's dendrogram illustrates the clustering process by which these classes were derived, and shows the breakdown of sub-clusters that emerge when a higher Jaccard similarity threshold is applied.

Table 8.3 Statistically-derived typology of terrorism hoaxers¹³⁰

	Primary Profile (% of groups within cluster displaying the characteristic)	Defining characteristics of sub-clusters¹³¹
	<u>E.g. Direct Action (Squamish Five)</u>	
CLASS 1: LIMITED SERIOUS (n=8)	<ul style="list-style-type: none"> • Non-hierarchical (75%) • Domestically oriented (75%) • Few serious events (87.5%) • Conventional weapons only (0% are CBRN-capable) 	<ul style="list-style-type: none"> A. All-channel structured religious groups (n=2) B. Fully consistent with primary profile (n=4) C. Radical splinter groups who are <i>not</i> domestically-oriented (n=2)
	<u>E.g. Animal Rights Militia (ARM)</u>	
CLASS 2: SINGLE-ISSUE (n=5)	<ul style="list-style-type: none"> • Anti-abortionist, environmentalist, or animal rights ideology (100%) • Non-hierarchical (100%) • Domestically oriented (80%) • CBRN-capable (80%) • Radical splinter origins (60%) 	
	<u>E.g. Liberation Tigers of Tamil Eelam (LTTE)</u>	
CLASS 3: ETHNO-NATIONALIST (n=9)	<ul style="list-style-type: none"> • Transnationally-oriented (100%) • Purely ethno-nationalist ideology (88.89%) • Functional differentiation (66.67%) • Centralized command and control (66.67%) 	<ul style="list-style-type: none"> D. CBRN-capable ethno-nationalist (n=2) E. Ethno-nationalist with leadership but no functional differentiation (n=1) F. Fully consistent with primary profile (n=6)
	<u>E.g. Aum Shinrikyo</u>	
CLASS 4: RELIGIOUS (n=6)	<ul style="list-style-type: none"> • Purely religious ideology (100%) • CBRN-capable (100%) • Not splinter organizations (0% had radical splinter origins) • Mostly transnational (only 16.67% have pure domestic orientation) 	
	<u>E.g. Red Army Faction/Baader Meinhof Group</u>	
CLASS 5: LEFTIST (n=12)	<ul style="list-style-type: none"> • Economically leftist ideology (100%) • Transnationally-oriented (0% are domestic-only) • Mostly lacking CBRN capability (only 8.33% are CBRN-capable) 	<ul style="list-style-type: none"> G. Leftist groups with centralized command and radical splinter origins (n=5) H. Fully consistent with primary profile (n=6) I. CBRN-capable leftist (n=1)

¹³¹ Discussion of sub-clusters is based on the eleven-cluster solution.

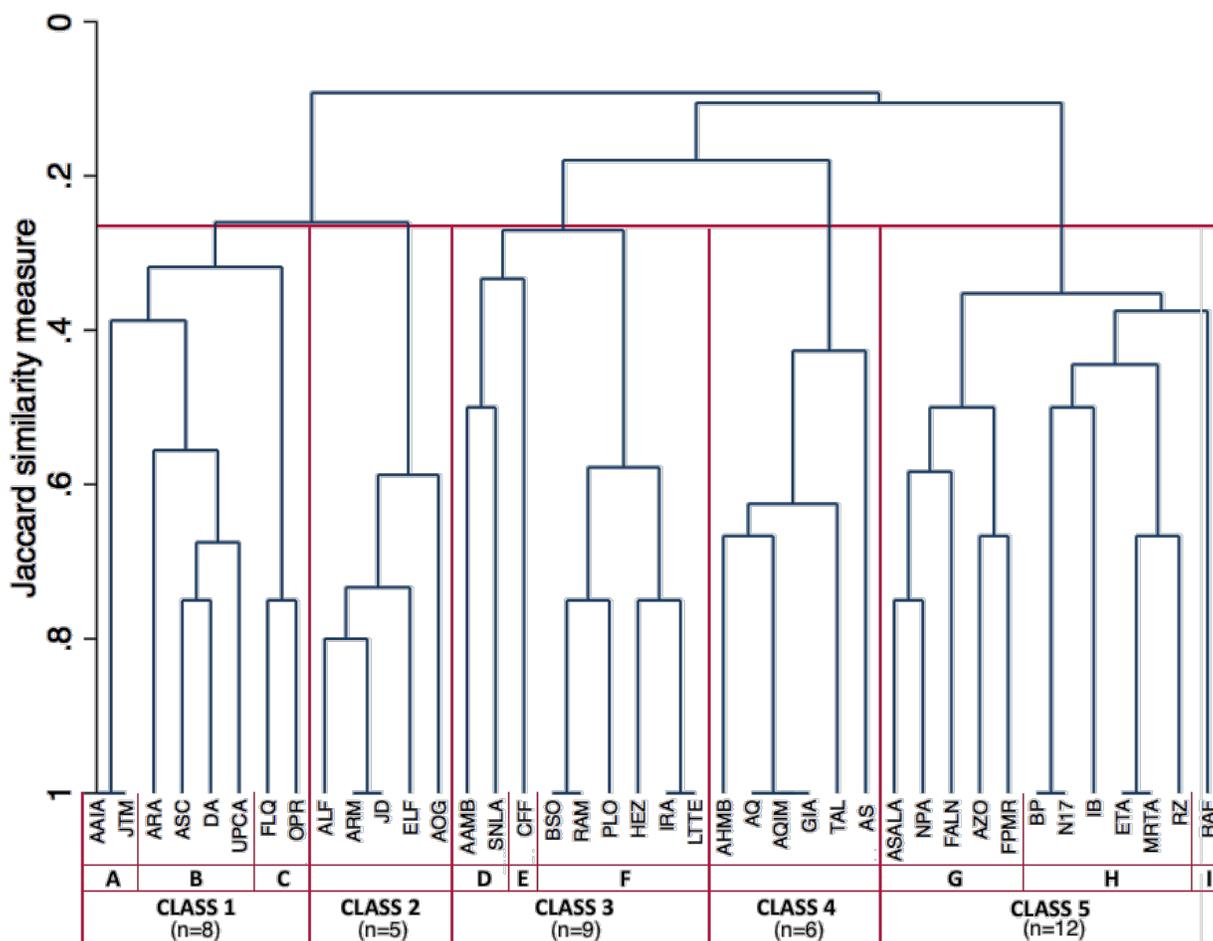


Figure 8.2 Dendrogram for cluster analysis based on Jaccard similarity measure¹³²

Apart from Class 1, which is not defined by ideology, each of the typology's classes aligns with a core ideological orientation: single-issue (animal rights, environmentalist, or anti-abortionist); ethno-nationalist; religious; and economically-leftist.¹³³ For the most part,

¹³⁰ Full details for the clusters on which this typology is based are presented in Appendix Table B.5 and Appendix Table B.6 which indicate the proportion of groups within each cluster displaying given characteristics.

¹³¹ Discussion of sub-clusters is based on the eleven-cluster solution.

¹³² The primary class categories correspond to the 5-cluster solution; the letters indicating sub-classes are based on an 11-cluster solution. These correspond to the classes interpreted in Table 8.3.

¹³³ There are no objects within the sample coded as having an economically rightist orientation, and so right-wing ideology does not affect the clustering mechanism. The variable is retained in the model for theoretical coherence for when it is tested against the non-hoaxing sample, and if ever the sample of groups is expanded.

these classes also align neatly with structural and/or contextual features of the groups' campaigns. The discussion below interprets the five main classes of hoaxers in terms of inter-cluster differences as well as alignment with expected outcomes, given the causal combinations elucidated in Chapter 7's QCA analysis and Chapter 6's discussion of average marginal effects.

8.4.1.1 Limited serious hoaxers

At a difference to the other four classes of perpetrators, Class 1—limited serious—perpetrators are not defined by ideology, but rather by their strong tendency (87.5%) to perpetrate only three or fewer serious events. All of the groups within the sample who perpetrated such small numbers of serious events are included in this class—an unsurprising finding, given that QCA and probit both indicate how this characteristic should be rare amongst hoaxers. These perpetrators also tend to lack hierarchy (75%) and

[Footnote 133 cont'd] I also tested models including additional (or alternate) measures of campaign context and religious orientation. First, I included a variable indicating whether the group was a militant wing of a political party with which it maintained a close alignment, or if the group itself engaged in "legitimate" politics by seeking representation in government (17.5% of groups satisfied this condition). This indicator, however never emerged as a primary cluster descriptor, and its addition to the basic model interpreted here did not affect the clustering solution except for at very low levels of nuance: it was not a relevant trait for distinguishing classes of hoaxers. I had also tested this variable in the earlier regression and QCA analysis, but did not report its findings due to consistent insignificance in distinguishing hoaxers from non-hoaxers. Second, I tested whether perpetrating large numbers of serious events (both in the top quartile and 90th percentile of perpetrators) played a role in distinguishing amongst classes of hoaxers. Neither alternative ever emerged as a distinctive cluster characteristic, and so it was removed from the "best" cluster model, for parsimony. Third, I also tested an indicator for whether a group's religious orientation was Islamic in nature (20% of groups satisfied this condition). While there is no theoretical reason to explain why Islamists might hoax differently from non-Islamists, the frequency with which the "Islamist extremism" label is invoked as a classification for terrorist activity writ large justified testing for whether it was a meaningful differentiator amongst classes of hoaxers. However, the Islamist indicator was highly correlated with both the variables measuring purely religious orientation (0.8447) and containing a religious component (0.8660), and so only one religion variable could be used in the clustering model at one time. From a theoretical perspective, the Islamist indicator was unnecessarily restrictive; and from an empirical perspective it weakened the strength of the clustering.

operate only in domestic spheres (75%). Furthermore, Class 1 perpetrators only ever use conventional weaponry, which is theoretically consistent with their limited number of serious events: they have not built up sufficient expertise to develop more sophisticated means of attack.

In the QCA solutions, a lack of hierarchy generally combines with having perpetrated *more* than three serious events as a sufficient condition for hoaxing. That Class 1 perpetrators are defined by their perpetration of three or fewer events suggests that these groups likely hoax as a function of their market-structure, and corresponding principal-agent problems; they have not perpetrated a sufficient number of serious events to be strategically sequencing their hoaxes as part of a broader campaign. The effect of non-hierarchy on encouraging hoax use, amongst Class 1 perpetrators, is likely stronger than the effect of not having perpetrated a minimum number of events.

8.4.1.2 Single-issue hoaxers

Consistent with the set theoretic relations discussed in Chapter 7, it is unsurprising that Class 2, single-issue, hoaxers are defined also by perfect consistency with non-hierarchical, market structure: all of the groups with animal rights, environmentalist, or anti-abortionist ideologies belong to Class 2, and they all lack hierarchy. Even in this expanded sample of hoaxers, no cases of single-issue hierarchies emerge, confirming that the limited diversity documented in previous chapters is not simply a function of sample selection.

Across the board, Class 2 perpetrators are the most homogeneous perpetrator type; all the groups combine in this class at a Jaccard similarity level only slightly below 0.6. Not only do they agree on motivation and (lack of) structure—characteristics identifying them as “typical” terrorism hoaxers, per the previous analyses—but they also feature shared characteristics that are less common across hoaxing types: they are mostly limited to a domestic sphere of operations (80%); they are mostly capable of employing—however crudely—CBRN weapons (80%); and more than half originated as radical splinters of more legitimate movements or organizations (60%). Class 2 perpetrators are the only ones defined (although only to a moderate degree) by their shared, radical splinter origins—a rare characteristic amongst hoaxers, and one that the predictive analyses suggest should more frequently lead groups to never use hoaxes. This clustering supports the QCA result that splinter histories only prevent groups from hoaxing when they are also religiously-oriented; single-issue, non-hierarchical groups hoax irrespective of their splinter origins.¹³⁴

8.4.1.3 Ethno-nationalist hoaxers

Class 3—ethno-nationalist—perpetrators are all defined by their ethno-nationalist ideology (88.89% are purely ethno-nationalist, although 100% contain an ethno-nationalist ideological component); transnational orientation (100%); and generally

¹³⁴ Additionally, there may be a degree of endogeneity between the non-hierarchy and splinter indicators: when group members splinter from a mother organization, their shared mother group origins and commitment to the use of more radical forms of activity may be sufficient to confer the “entitativity” required to constitute a terrorist group and carry out violent acts, without requiring more formal group structure.

bureaucratic structure (66.67%). That groups only *containing* an ethno-nationalist ideological component were included in the cluster—despite not having included the variable indicating partial ethno-nationalism in the model determining the classes—suggests a strong alignment of any degree of ethno-nationalism with bureaucratic structure and transnational orientation. An exception is where ethno-nationalism is combined with economically leftist orientations (Class 5), in which case it is similarity in economic orientation that overrides structural consistency.¹³⁵ This finding is at a difference to the QCA solution positing a combined sufficient effect of pure ethno-nationalism and non-hierarchy for hoaxing, although it helps to explain that solution's low coverage and weak consistency.

That the clustering identifies so many hoaxing ethno-nationalist groups is reflective of this chapter's expanded sample of hoaxers: in Chapter 7 (section 7.6.1), I showed that probit's strong, positive coefficient for purely ethno-nationalist ideology was actually driven by the frequent non-hoaxing of non-ethno-nationalist groups, not the presence of hoaxing ethno-nationalists. While only slightly more than half (55.56%) of Class 3 perpetrators were responsible for CBRN events (and thus included in the previous chapters' sample), all were responsible for serious transnational activities (as captured by ITERATE). The earlier CBRN-centric sample thus underestimated the incidence of hoaxing ethno-nationalists, explaining the asymmetric result described above. The WMDDB-based sample, however, did *not* underestimate the incidence of hoaxing *religious* ethno-

¹³⁵ While 50% of Class 5 leftist groups also contain an ethno-nationalist ideological element, only one of them (ASALA, included in sub-cluster G) is bureaucratically structured.

nationalists, which are consistently unlikely to hoax, even in this expanded sample. The only hoaxing religious ethno-nationalist in the present sample (Hezbollah; classified with its purely ethno-nationalist counterparts as a Class 3 perpetrator) was already coded as a hoaxer in previous chapters.

8.4.1.4 Religious hoaxers

Class 4—religious—perpetrators are all defined by their purely religious orientation; CBRN capability; and lack of splinter origins. They also tend to operate beyond a domestic sphere (only 16.67% are purely domestic). In the probit analysis, purely religious motivation was a strong positive predictor of hoaxing. In the QCA analysis, religious motivation (pure or otherwise) was never a positive identifier of hoaxers, but it frequently combined with ethno-nationalism and/or radical splinter origins to predict which groups would never use hoaxes. The characteristics of Class 4 hoax perpetrators in the present, expanded, sample are consistent with these results: none of them are co-motivated by ethno-nationalism (by definition; they are all purely religious), and none of them have radical splinter origins.

There are two other purely religious groups included in the sample that are attributed to Class 1 instead of Class 4, by virtue of their domestic orientation and having perpetrated three or fewer serious events. While they differ from Class 4 religious perpetrators along these dimensions, they are consistent in their mutual absence of a splinter history.

8.4.1.5 *Economically leftist hoaxers*

Class 5—leftist—perpetrators are all defined by their economically left-wing ideology, transnational orientation, and general lack of CBRN capability.¹³⁶ Although left-wing orientation did not differentiate hoaxers from non-hoaxers in the probit or QCA results, ideology here is the most distinctive clustering characteristic for defining Class 5 perpetrators. This difference from previous results may again be attributed to the change in sample: all Class 5 leftists are transnationally oriented, and only one (8.3%) is CBRN-capable. The WMDDDB sample thus underestimated the incidence of hoaxing leftists, who happen to be generally uninterested in acquiring a serious CBRN capability, but are consistently transnationally-inclined. This underestimation may help to explain why the broadly left-wing indicator in probit Model 2A only approached significance in predicting the hoax outcome, and why QCA models indicated that left- or right-wing motivation only mattered when combined with non-hierarchical structure (and having perpetrated more than three serious events). While Class 5 leftists all perpetrated more than three serious events, they exhibit a wide range of structural characteristics. Not all of these structural forms may distinguish Class 5 perpetrators from their non-hoaxing counterparts, but the expanded sample does indicate that groups' similarity in exhibiting left-wing orientation is so strong that it overrides wide variation in structural types and splinter histories in the clustering process.

¹³⁶ Only one group is CBRN-capable, and it is isolated as an outlier (RAF; sub-cluster I) in the eleven-cluster solution.

8.4.2 Comparing clusters against the non-hoaxer sample

I tested the same cluster model against the sample of non-hoaxing, CBRN-interest groups (from previous chapters) to evaluate the degree to which the combinations revealed in the clustering of hoaxing groups are distinctive to hoaxers; and the degree to which combinations of factors revealed in Chapter 7's QCA analysis (distinguishing hoaxers from non-hoaxers) are reflected amongst a sample of only non-hoaxers.

When compared with section 8.4.1's clustering of hoaxing groups, the cluster analysis of non-hoaxing groups yields a wholly different set of clusters. Whereas the hoaxers' clusters were neatly defined by alignments of different ideologies combined with various structural and/or campaign contextual features, the seven core clusters that emerged for non-hoaxers were primarily distinguished based on different structural elements combined with features of campaign context.¹³⁷ Ideology was only a core cluster characteristic in two cases: religion and economically leftist orientation. These differing classificatory outcomes suggest that there are fundamental differences between the sample of hoaxers and the sample of non-hoaxers. In other words, the cluster analysis confirms that there are organizational characteristics (and combinations thereof) that differentiate hoaxers from non-hoaxers. Table 8.4 describes the core classes of non-hoaxers, per their most distinguishing features.¹³⁸ Since the purpose of clustering the non-hoaxer sample was to reveal or clarify differences amongst the hoaxing and non-hoaxing

¹³⁷ The number of clusters interpreted here is based on the eleven-cluster solution indicated as being optimal by the Duda-Hart and Caliński-Harabasz stopping rules. Four of these eleven clusters include only one or two group members. I treat these groups as peripheral outliers, and interpret only the remaining seven as core clusters. These seven core clusters account for 90.74% of all groups in the sample.

¹³⁸ Full cluster results are presented in Appendix Table B.7, and represented visually in Appendix Figure B.1's dendrogram.

group samples, I limit my discussion of this classification to only those classes that are relevant for such distinctions.

Table 8.4 Statistically-derived typology of non-hoaxing terrorists

	Profile (% of groups within cluster displaying the characteristic)	Example of typical group
CORE CLASS 1: ALL-CHANNEL (n=5)	<ul style="list-style-type: none"> • All-channel structure (100%) • Not CBRN-capable (100%) • Mostly radical splinter origins (80%) • Mostly transnationally-oriented (80%) 	Real Irish Republic Army (RIRA)
CORE CLASS 2: RELIGIOUS (n=4)	<ul style="list-style-type: none"> • Purely religious ideology (100%) • Domestic-only orientation (100%) • Moderate likelihood of radical splinter origins (75%) 	Ansar al-Islam
CORE CLASS 3: LIMITED SERIOUS (n=11)	<ul style="list-style-type: none"> • Domestic-only orientation (100%) • CBRN-capable (100%) • Mostly few serious events (81.82) 	Babbar Khalsa International (BKI)
CORE CLASS 4: DOMESTIC (n=3)	<ul style="list-style-type: none"> • Functional differentiation (100%) • Domestic-only orientation (100%) • CBRN-capable (100%) • Moderate lack of hierarchy (66.67) 	Khalistan Commando Force
CORE CLASS 5: BUREAUCRATIC (n=11)	<ul style="list-style-type: none"> • Bureaucratic structure (100% - both functional differentiation and centralized command are present in all cases) • Moderately domestically-orientated (63.64%) 	Hizbul Mujahideen
CORE CLASS 6: HUB-SPOKE (n=6)	<ul style="list-style-type: none"> • Hub-spoke structure (100% - all have leadership and functional differentiation) • CBRN-capable (100%) • Transnationally-oriented (100%) 	Hamas
CORE CLASS 7: LEFTIST (n=9)	<ul style="list-style-type: none"> • Economically leftist ideology (100%) • Transnationally-oriented (100%) • Mostly bureaucratic structure (88.89% have functional differentiation; 77.78% have centralized command and control) • Mostly CBRN-capable (88.89%) 	National Liberation Army (ELN)

Classes of both hoaxing and non-hoaxing leftist groups emerge in their respective analyses, with a shared preponderance for transnational orientation. Hoaxing leftists (Class 5) differ from non-hoaxing leftists (Core Class 7), however, in terms of their

structure and CBRN capability. Whereas hoaxing leftists generally lack CBRN capability and exhibit wide variation in terms of organizational structure, non-hoaxing leftists are mostly CBRN-capable and tend toward greater structural sophistication. These differences regarding structure are consistent with the QCA expectation that left- (or right-) wing motivation would combine with non-hierarchy to predict hoaxing, and with hierarchy to predict non-hoaxing.

Interestingly, although there are no cases of hoaxing economically rightist groups (even in the expanded sample of hoaxers), and there *are* four cases of non-hoaxing economically rightist groups, right-wing ideology does not emerge as an important clustering characteristic amongst the non-hoaxer sample. Three of the sample's four economically-rightist groups are clustered together in Core Class 3 (the remaining one is assigned to Core Class 5), but they share sufficient structural and contextual characteristics with the other groups so as to be classified irrespective of their ideology. The lack of a distinct rightist cluster supports the insignificant coefficient for broadly right-wing motivation when tested using probit regression: these groups display other conditions associated with non-hoaxing (namely the perpetration of three or fewer events), such that their motivation has no independent, net effect.

The classes of religious hoaxers (Class 4) and non-hoaxers (Core Class 2) are only similar in that they both contain all purely religious groups. Whereas hoaxing religious groups tend toward transnational orientation, non-hoaxing religious groups are limited to domestic operations. Furthermore, non-hoaxing religious groups are far more likely than their hoaxing counterparts to be radical splinter organizations. This distinction is

consistent with the QCA finding that religious splinter organizations are particularly unlikely to ever use hoaxes.

The QCA analysis predicted that religious ethno-nationalists would also be particularly unlikely hoax users. While the non-hoaxing group sample was not clustered according to any shared motivations, it is interesting that religious ethno-nationalist non-hoaxers clustered together anyway, accounting for 66.67% of groups in each of Core Classes 4 and 6. Both of these core classes align perfectly with respect to their 100% prevalence of functional differentiation and 0% prevalence of command and control structures. They differ, however, with respect to their transnational orientation (Core Class 4 perpetrators are all domestically-oriented, whereas Core Class 6 perpetrators are all transnationally-oriented), and minimum degree of hierarchy (Core Class 4 perpetrators are mostly non-hierarchical, whereas Core Class 6 perpetrators all exhibit a minimum degree of leadership). These core cluster divergences are consistent with the QCA solution treating religious ethno-nationalism as a sufficient condition for non-hoaxing on its own: it does not matter how they are structured, or where their activities are targeted, religious ethno-nationalists are always unlikely to use hoaxes.

8.5 Conclusions from the group-based cluster analysis

This chapter presented a multidimensional typology of terrorist groups engaged in hoax activity, based on their ideological orientations; organizational structures; and various features of their campaign contexts. The typology was inductively constructed using a statistical cluster analysis that generated five classes of hoax perpetrators, predominantly

defined by ideological orientation aligned with particular structural and contextual elements: single-issue groups who are non-hierarchical, CBRN-capable, and domestically-oriented; ethno-nationalist groups who are bureaucratically structured and transnationally-oriented; religious groups who are CBRN-capable and did not originate as radical splinters; economically leftist groups who are transnationally oriented; and limited serious groups—not defined by ideology—perpetrating three or fewer serious events, with a domestic orientation, non-hierarchical group structure, and interest only in conventional weapons. This statistically-derived classification of hoaxers differs from an analogous clustering of non-hoaxing groups, who are more effectively distinguished per their organizational structure, with little regard for ideological orientation.

The two classifications can be contrasted to distinguish hoaxers from non-hoaxers, but only in cases where the group's ideology is religious or leftist, and contextual and structural features are known (as discussed in section 8.4.2).¹³⁹ Beyond this limited extent to which they identify hoaxers and non-hoaxers, the classes do not map on to patterns in hoax behaviour within or across clusters. In an attempt to increase the typology's theoretical leverage, I conducted a fine-grained, qualitative analysis of each group's attack profile, to test for whether certain classes of perpetrator aligned with particular strategic logics or patterns in the sub-category of hoax employed (claim of responsibility,

¹³⁹ For religious groups, those with transnational orientation are more likely hoaxers, whereas domestically-oriented groups—who are also likely to be radical splinter organizations—are more likely non-hoaxers. This pattern bolsters the earlier QCA solution that identified splinter religious groups as likely non-hoaxers. For leftists, non-hoaxing groups tend to be CBRN-capable, and exhibit more sophisticated structures. Hoaxers generally lack a serious CBRN capability, but they do not have any consistent defining structural characteristics. While this characterization of left-leaning groups is hardly conclusive, since the QCA analysis was unable to uncover differences between left-wing and right-wing groups, it offers a further layer of nuance.

warning, or device). Unfortunately, no such patterns emerged. The typology thus falls short of Marsden and Schmid's (2011) marker for a useful terrorism typology: the classification of perpetrator types generated by the cluster analysis does *not* correspond to meaningful differences in terms of the nature of the threat each group poses.

That said, the cluster analysis *did* succeed in identifying clear classes of hoaxing terrorist groups. The result may not be theoretically instructive on its own, but it does indicate avenues for future research. Importantly, the classes can be used as a benchmark for evaluating terrorist groups beyond this sample. If a new hoaxing group does not conform to one of the defined hoax perpetrator classes, it may reveal sample bias in the predominantly transnational- and CBRN-oriented perpetrator sample relied upon here. Conversely, if a group *does* align on the clustered characteristics, but not the expected hoax outcome, it may be meaningfully evaluated in a qualitative paired comparison to identify group characteristics beyond those coded in the present analysis. These new characteristics may in turn be effective in elucidating differences in hoax behaviour.

In their review of terrorism typologies, Schmid, Jongman, Stohl, and Flemming (1988) assert that “[w]hether a typology is functional depends on the needs of the user”

(39). Theory-building is one such need:

“[I]n the absence of a universally accepted definition of terrorism and with the lack of a general theory of terrorism, typology construction can nevertheless be a useful instrument to advance our understanding of terrorism [...]. Theoretical progress in the field of Terrorism Studies will have to be based on typological progress, which, in turn, is based on conceptual progress” (Marsden and Schmid 2011, 193)

In the absence of a universally accepted definition for hoaxing, and the lack of a general theory for hoaxing, this inductively-derived classification contributes a valuable—however limited—early step in building a theory of hoax behaviour by terrorist groups.

PART III: THEORY-TESTING AND CONCLUSIONS

9. Hoaxing in Canada: Testing the theory

On March 1st, 2017, several Canadian news outlets received emailed bomb warnings that “artisanal amateur explosive devices” had been placed in various locations at Concordia University, in Montreal, and that the explosives would be detonated if the university did not halt religious activities (particularly those of Muslim students) on campus. The warning appeared on letterhead of the “Council of Conservative Citizens of Canada (C4),” a supposed splinter group of the U.S.-based, white-supremacist “Council of Conservative Citizens (CCC)” (Frisk 2017). In response, classes were cancelled; thousands of students, faculty, and staff were evacuated; the university increased its private security patrols; and Montreal law enforcement and canine units swept the area for bombs. Ultimately, the threat was deemed a terrorist hoax: no bombs had actually been placed, and there did not appear to be any serious plan to follow through on the warning.

This incident highlights that Canada is a target of hoax activity, and Canadian citizens are victim to its associated costs: lost productivity due to evacuated educational institutions and places of work; wasted emergency response and law enforcement resources; and increased societal fear and suspicion. Canada has a long history of exposure to hoaxes. The CIDB documents a total of 153 terrorism hoax events occurring within Canada between 1960 and 2015, the first one dating back to 1962. Despite this history, however, Canadian hoax events have not received any systematic treatment in the literature. While hoax events are included in Leman-Langlois and Brodeur’s (2005)

evaluation of terrorism linked to Canada between 1973 and 2003 (they account for 73 of the 400 events included in their analysis), they are not treated as a distinct category of terrorist activity; and the first published summary presentation of CIDB data—including hoaxes as well as serious activity—is only descriptive in nature (Tishler, Ouellet, and Kilberg, n.d.).

In this chapter, I use CIDB data to evaluate the history of terrorism hoaxes in Canada, to test the findings gleaned from previous chapters' large-sample probit regression and QCA analyses. CIDB's data presents a valuable test of these earlier findings since it was designed to capture the full scope of terrorists' hoax activity—unlike the CBRN-centric WMMDB, or the transnationally-oriented ITERATE (see Chapter 1, Table 1.1 and section 1.2.1). Although it is limited to events linked to a single country, the CIDB accounts for a more comprehensive documenting of both serious and hoaxed terrorist events within its purview than the previously relied-upon cross-national datasets, and is thus able to shed light on the degree to which previous chapters' findings are extrapolable to all terrorist perpetrators—including those involved only with domestic, non-CBRN activity. Additionally, the CIDB's depth and nuance of event coverage allows me to incorporate a temporal dimension to my analysis, without concern for missed or underreported events. In exploring event sequencing and associated tactical choices, I can thus integrate a more fine-grained approach, and trace processes that were obscured in the earlier cross-sectional analyses emphasizing discrete organizational traits.

The chapter proceeds in six sections. First, I present a descriptive analysis of terrorism hoaxes and hoax perpetrators operating in Canada, including a comparison of

the CIDB-based sample of perpetrators with the theory-building samples used in previous chapters. Second, I review the key findings of previous chapters, and describe to what extent I expect them to hold when tested against CIDB data. Third, I outline my methodology for conducting these tests, incorporating both QCA (for cross-case comparison) and process tracing (for intra-case evaluation of causal mechanisms). Fourth, I present my QCA results and discussion of how the sufficient conditions distinguishing hoaxers and non-hoaxers in the CIDB differ from those generated in Chapter 7. Fifth, I present a narrative overview of hoax use by the FLQ, highlighting how patterns in event sequencing and overall tactical decision-making in their terrorism campaign support (or contradict) the strategic logics of hoaxing presented in Chapter 5. Finally, I conclude with an assessment of how the Canadian data modifies my previously-established understandings of terrorism hoax behaviour, and how these modifications indicate deficiencies in available cross-national terrorism events datasets for studying hoaxes.

9.1 Terrorism hoaxes in the CIDB

In the two subsections that follow, I describe the CIDB's data with respect to hoaxed versus serious terrorism events within Canada; and then compare the sample of perpetrators responsible for these events with the samples of hoaxers (and non-hoaxers) used in previous chapters' analyses.

9.1.1 Hoaxed versus serious terrorism events within Canada

The CIDB covers incidents of violent extremism with a Canadian connection, between 1960 and 2015 (n=1,844). In addition to terrorism events within Canada, the focus of the present analysis, it thus also includes events carried out abroad with Canadian victims, targets, and/or perpetrators (n=394); and violent extremist incidents that do not conform to the definition of terrorism (n=712). When these are culled from the sample, 871 terrorism events remain, of which 153 are hoaxes.¹⁴⁰

Terrorism hoax events in Canada have followed a roughly similar temporal distribution as serious events (see Figure 9.1). The two main exceptions are the 1985-1986 and 2014-2015 periods, where hoax incidents were even more frequent than serious ones. The spike in the 1980s overlaps with a broader spike in transnational hoaxes documented by ITERATE in the same period (see Chapter 3, Figure 3.1). Based on CIDB event descriptions, many of the incidents in this period involved hoaxed bomb warnings against airports and airplanes, and occurred in the year following the June 1985 bombing of Air India Flight 182, on route from Montreal to Delhi. Until September 11th, 2001, the Air India bombing had been the world's deadliest terror incident: all of the plane's 329 passengers—including 280 Canadians—were killed. That many of the Canadian hoax

¹⁴⁰ Additional to these hoaxes are fifteen terrorism hoaxes with a Canadian connection perpetrated abroad. Evaluation of these events is complicated by unidentified perpetrators and motives. They were all, however, classified as threats involving explosives, bombs, or dynamite. Thirteen of them targeted airports and aircraft; the remaining two targeted diplomatic targets (embassies/consulates) and non-aviation transportation (train-related infrastructure), respectively. There were various reasons connecting these hoax incidents to Canada: in five cases, the event involved a Canadian target; in one case, the event involved a Canadian perpetrator; in other cases, Canada was implicated in the event incidentally, due to response activities (e.g. Canadian fighter jets escorted a threatened plane to safety; a Canadian phone operator received/relayed a hoax warning; and—more frequently—a Canadian airport received diverted flights). These fifteen hoax events perpetrated abroad were not concentrated in time: six occurred in the 1970s; four in the 1980s; one in the 1990s; none in the 2000s; and four in the 2010-2015 period.

incidents accounting for this spike mimic the Air India bombers' tactics suggests that hoaxers were exploiting heightened levels of societal fear and sensitivity regarding the terrorist threat to civil aviation, to cause maximum disruption.

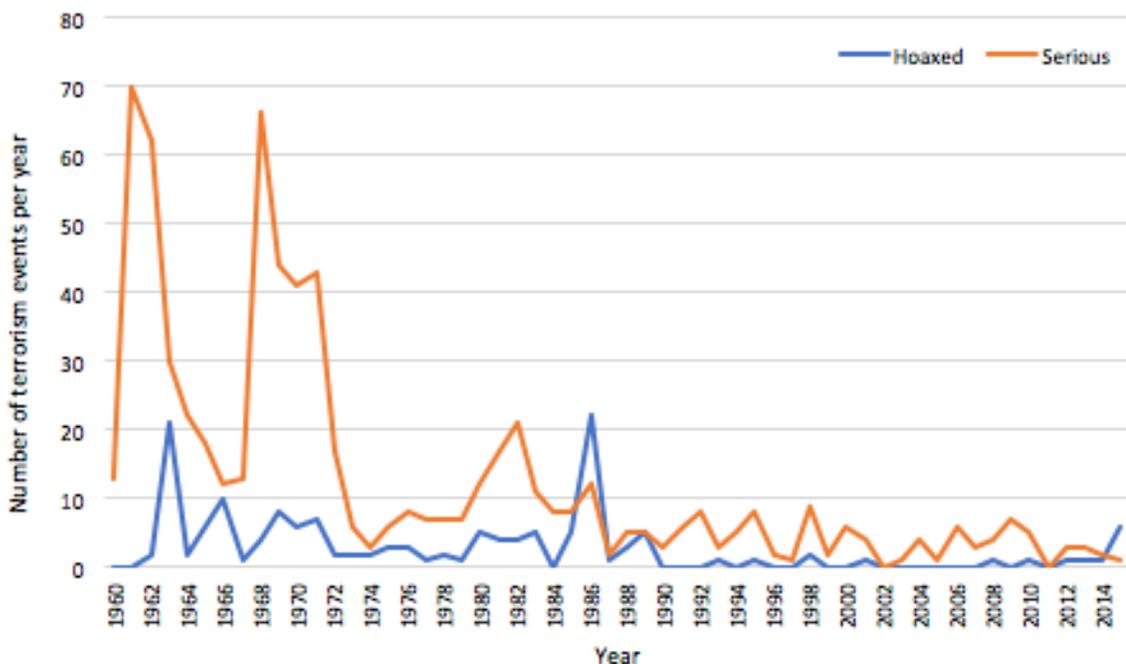


Figure 9.1 Number of hoaxed and serious terrorism events occurring per year in Canada, 1960-2015

The data from the cross-national datasets discussed in earlier chapters do not extend past 2012, and are thus unable to shed light on the observed increase in Canadian-based hoaxes in the 2010s. The CIDB, however, tells us that all of these hoaxes occurred following the fatal October 2014 car-ramming and shooting incidents perpetrated in Saint-Jean-sur-Richelieu and Ottawa.¹⁴¹ They also correspond with a documented shift

¹⁴¹ On October 20, 2014, Martin Couture-Rouleau rammed his vehicle into two Canadian Forces members in Saint-Jean-sur-Richelieu, Quebec, killing one officer and injuring the other. During the ensuing police chase, Couture-Rouleau phoned 9-1-1 to inform police he had been “acting in the name of Allah” (Tishler, Ouellet, and Kilberg, n.d.). The chase ended when Couture-Rouleau was shot and killed by a responding

away from terrorist incidents perpetrated at home to Canadian-affiliated incidents abroad (Tishler, Ouellet, and Kilberg, n.d.). This hoax activity may be explained again by heightened sensitivity to the terrorist threat in Canada following the October 2014 incidents, but also by a corresponding increased difficulty in carrying out serious incidents in the homeland due to augmented counterterrorism efforts (see section 1.1.1); hoaxes are relatively easier to carry out within Canada, and—given the surge in attention paid to the domestic terrorist threat in late 2014 and 2015—have been able to generate a fair degree of “bang” for the buck.

Figure 9.2 illustrates the distribution of hoaxed and serious terrorism events in Canada across various weapons types. For the most part, hoaxers in Canada use the same tactics as their serious counterparts: explosives, bombs, and dynamite are both the most frequently used and hoaxed weapon-type. While serious *and* hoaxed CBRN incidents are incredibly rare in Canada, biological and chemical weapons are the only weapon type (other than fake weapons) to ever be hoaxed more frequently than actually employed. Despite their rarity, this breakdown is consistent with broader CBRN activity as documented by the WMDDDB (Chapter 3, Figure 3.6), where chemicals are the most likely CBRN agents to be used in serious incidents, but biological ones are more likely to be hoaxed. For the CBRN-interested minority of hoax perpetrators in Canada, there is an

police officer. Two days later, Michael Zehaf-Bibeau fatally shot a Ceremonial Guard member at Ottawa’s Cenotaph (the national war memorial), before entering and opening fire inside Parliament Hill’s Centre Block. The incident ended when he was fatally shot during an exchange of gunfire with RCMP officers and House of Commons Security.

aspirational dimension to the tactics employed: they hoax what they cannot carry out in reality.

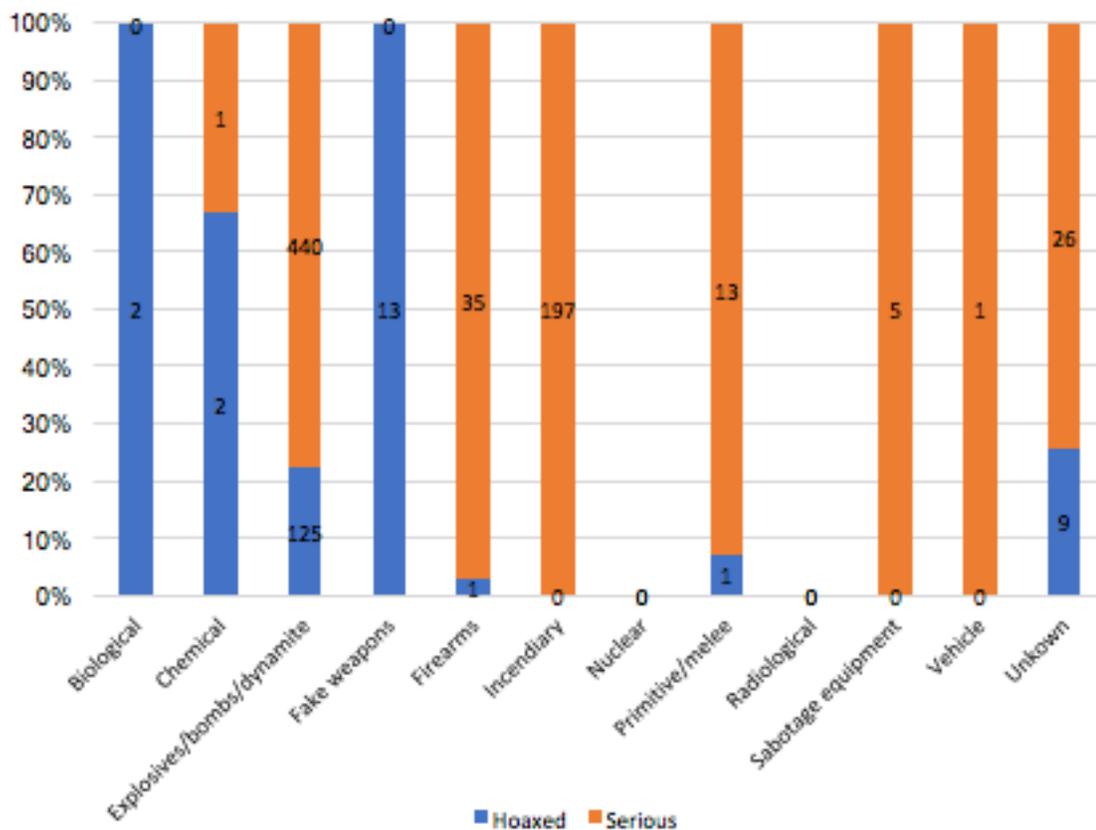


Figure 9.2 Percent likelihood of weapons use being hoaxed or serious for terrorism events in Canada, 1960-2015¹⁴²

9.1.2 The sample: Perpetrator groups responsible for Canadian events

The 447 terrorism events carried out within Canada between 1960 and 2015 were perpetrated by 96 distinctly-identified perpetrators. Many of these, however, are not established terrorist groups. For the present analysis, the perpetrator sample is culled per

¹⁴² Data labels reflect the absolute number of hoaxed or serious events.

the requirements outlined in Chapter 6 (section 6.2.1), with details on sample cleaning presented in Table 9.1. The result is a sample of 42 serious terrorism groups operating in Canada, collectively responsible for a total of 368 serious terrorism events on Canadian soil, and 4,366 serious terrorism events worldwide. Of these groups, fourteen (33.33%) used hoaxes at some point in their campaign, somewhere in the world. These groups are listed in Appendix Table C.1 with coding decisions presented in Appendix Table C.2.

Table 9.1 Cleaning the sample of perpetrators operating within Canada

<i>Total # of perpetrators responsible for terrorism events within Canada =</i>	96
<i># of perpetrators culled for referring to unspecific categories =</i>	11
<i># of "group names" that do not appear to reflect a real organization =</i>	11
<i># of perpetrators culled for referring to "lone wolves" =</i>	6
<i># of perpetrators culled for never having carried out "serious" acts =</i>	10
<i># of perpetrators amalgamated (cells/alternate names of same organization) =</i>	16 ¹⁴³
<i>Total # of groups in sample =</i>	42

9.1.2.1 Comparisons with previous samples of terrorism hoax perpetrators

As demonstrated in Chapter 1, Table 1.1, the cross-national datasets used in previous chapters substantially understate the incidence of terrorism hoax activity. The CIDB captures over ten times more terrorism incidents in Canada than the other datasets, and between 2.5 and 12.9 times more hoax incidents, when time periods are held constant. It also has a stronger track record of identifying terrorism event perpetrators. The cross-national datasets' underestimation of hoax activity is likely attributable to their exclusion of domestic, non-CBRN activity. That over 80% of hoax activity in Canada involved hoaxed

¹⁴³ In none of these cases do various group amalgamations affect the dependent variable outcome of whether a group hoaxes or not, nor its structure coding.

explosives, bombs, or dynamite supports the assessment that much domestic hoax activity is non-CBRN in orientation.

Despite this underreporting of events, there is a fair degree in sample overlap between *perpetrators* who carried out events in Canada, and those perpetrators not limited to any geographic area captured in earlier chapters' samples. The cross-national datasets thus do an acceptable job of identifying relevant actors, even if they do not capture all of the events they perpetrate. Thirteen of the CIDB's 42 perpetrator groups (31.0%) were included in the original sample of perpetrators with an interest in CBRN weaponry used in Chapters 6 and 7's probit and QCA analyses.

Of the sample's fourteen hoaxers, twelve (85.7%) had already been included in Chapter 8's expanded sample of groups who hoaxed in any of the cross-national datasets (WMDDDB, ITERATE, or GTD). Half of these hoaxing groups never carried out *hoaxes* in Canada (they are classified as perpetrators in Canada exclusively as a function of their serious activities), but they are classified as hoaxers by virtue of the full scope of their campaign activities in the world. Importantly, the CIDB does not change the hoaxer status of any groups previously classified as non-hoaxers: the CIDB contributes many additional events, and in some cases substantially modifies the number of serious events perpetrated per group (e.g. it includes 158 serious FLQ events that were not captured in the cross-national datasets), but it never includes additional *hoaxed* events for groups previously classified as non-hoaxers. To the degree that the samples overlap, then, the binary dependent variable assignments in previous chapters are sound.

Table 9.2 compares the incidence of important perpetrator characteristics across this dissertation's various samples. Consistent with previous findings, the CIDB never identifies any single-issue non-hoaxers, nor does it identify any economically rightist hoaxers. It also finds no evidence of religious ethno-nationalist hoaxers, which previous samples had identified as very rare. Conversely, it is substantially more inclined to find serious-only perpetrators who are non-hierarchical and/or who perpetrated a very small number of serious events; as well as hoaxers who originated as radical splinters of more moderate organizations.

Table 9.2 Comparison of perpetrator samples in the present and previous analyses

(% = 1)

		NON-HOAXERS		HOAXERS ONLY		
		PROBIT /QCA	CIDB	CLUSTER ANALYSIS	PROBIT /QCA	CIDB
MOTIVE	<i>Single issue</i>	0%	0%	12.5%	21.7%	21.4%
	<i>Purely religious</i>	24.1%	10.7%	20.0%	34.8%	21.4%
	<i>Contains religion</i>	51.7%	21.4%	25.0%	43.5%	21.4%
	<i>Purely ethno-nationalist</i>	6.9%	21.4%	22.5%	21.7%	21.4%
	<i>Contains ethno-nationalism</i>	43.1%	50%	42.5%	30.4%	50%
	<i>Religious ethno-nationalist</i>	24.1%	10.7%	2.5%	4.4%	0%
	<i>Economically leftist</i>	22.4%	32.1%	35%	13.0%	35.7%
	<i>Economically rightist</i>	8.6%	17.9%	0%	0%	0%
STRUCTURE	<i>No hierarchy</i>	9.3%	28.6%	37.5%	40.9%	42.9%
	<i>Leadership</i>	88.9%	71.4%	62.5%	59.1%	57.1%
	<i>Functional differentiation</i>	53.7%	21.4%	32.5%	27.3%	21.4%
	<i>Centralized command and control</i>	33.3%	21.4%	27.5%	18.2%	21.4%
CAMPAIGN CONTEXT	<i>Splinter origins</i>	27.3%	14.3%	27.5%	18.2%	35.7%
	<i>Three or fewer serious events</i>	27.6%	42.9%	17.5%	13.0%	7.1%
	<i>Number of serious events perpetrated [range]</i>	1 – 1,453	1 – 356	1 – 2,492	1 – 2,492	3 – 2,186

Although the CIDB hoaxers exhibit roughly parallel characteristics, on average, to the cluster analysis sample used to generate Chapter 8’s typology of hoaxers (as illustrated in Table 9.2), the perpetrator samples do not agree with respect to the broad-based *combinations* of factors that previously differentiated amongst classes of hoaxers, nor are hoaxer classes proportionately represented in the CIDB-based sample. Figure 9.3 reflects the distribution of hoax perpetrators per typology class for the sample used in the theory-building Chapter 8, and for the CIDB-based test sample of hoax perpetrators evaluated here.¹⁴⁴ The allocations for the CIDB-based sample includes the two hoaxers that had not been previously classified, and class reassignments required due to new information revealed by the CIDB’s nuanced event coverage.¹⁴⁵

¹⁴⁴ See Table 8.3 for a summary of the five hoax perpetrator types yielded by the cluster analysis.

¹⁴⁵ The two CIDB-based hoaxers that were not included in the cluster analysis are most appropriately assigned to Class 3. The Armenian Revolutionary Army has a purely ethno-nationalist ideology, transnational orientation, and lacks splinter origins—although at a difference to many other Class 3 perpetrators, it also lacks any form of hierarchy. It thus weakens Class 3’s already-mild expectations with respect to group structure. Accordingly, when the left-leaning ethno-nationalist group *Armée de libération du Québec* (ALQ) is considered, it is also grouped with Class 3, since structural features no longer effectively distinguish Class 5 left-leaning ethno-nationalists from Class 3 ethno-nationalists who may or may not have a secondary ideology. For this reason, the ALQ is assigned to Class 3 rather than Class 5, despite its leftist ideological component and unsophisticated structure. Two additional leftist ethno-nationalist groups (ASALA and the Black Panthers) remain assigned to Class 5, per the original clustering, although they may be more appropriately assigned to Class 3 now, given its lowered expectations of structural sophistication.

In the CIDB-based sample, limited serious perpetrators are particularly rare, and ethno-nationalists are substantially more prevalent. Importantly, the CIDB identifies additional events that were not captured in the cross-national datasets, and which change two groups’ “limited serious” classification: the leftist Direct Action (now responsible for four serious events) is assigned to Class 5; and the ethno-nationalist FLQ (now responsible for a substantially-augmented 160 serious events) is assigned to Class 3. CIDB-based perpetrators also espouse more structural diversity than the earlier sample, particularly when they exhibit a combined leftist ethno-nationalist ideology. This diversity blurs the earlier distinctions between Classes 3 and 5 hoax perpetrators regarding structural consistency.



Figure 9.3 Distribution of hoax perpetrators per typology class

9.2 Expected findings

In Chapters 6 and 7, hoaxers were found to be differentiated from non-hoaxers by their:

- Non-hierarchical structure and:
 - a) Single-issue motivation (although this motivation *may* lead to hoaxing even if hierarchy were present);
 - b) Purely ethno-nationalist motivation; or
 - c) Perpetration of more than three serious events (and any non-religious, left- or right-wing ideology).

Conversely, non-hoaxers are likely to be differentiated from hoaxers by their:

- Religious ethno-nationalist ideology;
- Radical splinter origins plus religious ideology *or* hierarchical structure;
- Perpetration of three or fewer serious events; or
- Hierarchical structure (and broadly left- or right-wing ideology).

While the general distribution of characteristics in Table 9.2 does not address combinations of factors, differences in distribution of characteristics across hoaxers and non-hoaxers in the CIDB-based sample (the table's red columns) and the sample used to build these profiles (the table's yellow columns) suggest that these previously-identified (combinations of) factors distinguishing hoaxers from non-hoaxers may not be directly extrapolable to all types of terrorism perpetrators, in all contexts. Of particular suspicion is the divergence in the prevalence of splinter origins amongst CIDB hoaxers, and relative lack thereof amongst non-hoaxers. I thus anticipate that many relationships uncovered in the theory-building chapters will be substantiated by the CIDB, but that the effect of radical splinter histories may not be consistent with previous expectations.

9.3 Methodology

This chapter is designed to test whether the findings derived from my empirical analysis of the CBRN-interested sample of perpetrators are extrapolable to a sample of all terrorism groups operating within Canada; and to test whether the hypothesized mechanisms underpinning these cross-case comparisons hold, when examined at the intra-case level. To do so, I use QCA to differentiate hoaxers from non-hoaxers. I do not repeat the use of probit regression because the present sample ($n=42$) is too small for reliable statistical inference; and because QCA's reliance on asymmetrical relations and causal combinations of factors better approximates terrorist groups' use and non-use of hoaxes than probit's assumed symmetry and net effects of variables in isolation (see section 7.7).

I replicate the two most instructive QCA models from Chapter 7. As described in section 7.4, Model A presents a comprehensive test of all six primary indicators—single-issue, ethno-nationalist, and religious motivations; hierarchical group structure; radical splinter organizations; and the number of serious incidents perpetrated per group—using crisp scores. Model E allows for an alternate examination of the motive variables, using the broadly left- or right-wing indicator instead of single-issue motivation, and the indicator for ideologies *containing* religious or ethno-nationalist components, rather than Model A's requirement that ideologies be *exclusively* religious or ethno-nationalist.¹⁴⁶

As in the original analysis, I test each model against both the outcome of hoaxing (i.e. *hoaxer*=1) and its negation (i.e. *hoaxer*=0), to account for the previously-demonstrated asymmetry in effects. In the present analysis, I apply a consistency threshold of 0.8 (requiring that at least 80% of the groups in a given truth table row agree on the outcome), but I use a frequency threshold of only one group (instead of the previously-required two groups) for determining how many cases must exhibit a given combination of factors before they may be counted as a sufficient condition. I use this lower frequency threshold because of the reduction in sample size ($n=42$, compared with the earlier $n=76$). With the present sample size, there is a guaranteed minimum of 22 logical remainders when six factors are included in the QCA model;¹⁴⁷ applying a lower

¹⁴⁶ I do not present the results of Model D, which uses the combined religious ethno-nationalist indicator, since the mutual presence of each of these ideology is captured by the intersection of the containing religion and containing ethno-nationalism indicators used in Model E. Since there are no religious ethno-nationalist hoaxers (see Table 9.2) including a discrete indicator for religious ethno-nationalism would be meaningless in solutions predicting the hoax outcome.

¹⁴⁷ There are 64 possible logical combinations when six factors are tested ($2^6 = 64$), and so there will be 22 arithmetic remainders with a sample of 42 ($64 - 42 = 22$). These do not include impossible remainders attributable to combinations of mutually exclusive motive variables.

frequency threshold minimizes the effect of additional remainders due to clustering (see section 7.4).

Next, I conduct a qualitative process tracing of hoax use within a broader serious terrorism campaign, to allow for intra-case evaluation of hypothesized causal mechanisms. I focus on the activities of the FLQ, which is both the CIDB's most prolific serious terrorism event perpetrator, and most prolific hoaxer. While the FLQ is an outlier in the Canadian context, due to the large number of events it carries out, it is a "typical hoaxer" per my Chapter 6 and 7 results: it is a non-hierarchical ethno-nationalist group, responsible for many serious incidents. This conformity with previous results is particularly impressive, given that the FLQ was not included in the WMDDDB-based sample of perpetrators used to generate these findings; it thus constitutes a valid test thereof. Since the FLQ is a historical case, I am able to examine the entirety of the group's activities, at all stages of its lifecycle, with the benefit of rich scholarly investigations (and hindsight) to support my analysis. Since it is responsible for a wide range of hoaxes across its lifecycle, it offers ample opportunity to explore several strategic logics—even within a single organization's campaign.

To trace the strategic logics employed in each hoax incident, I pay particular attention to the ways in which the FLQ sequenced its attacks; the tactics it employed; and the victims and targets of its activity. For instance, hoaxes consistently preceding serious attacks using similar weapon-type and delivery patterns indicate the use of dry runs; hoaxes simultaneous to serious activities indicate decoys, used to flood security and emergency response; and hoaxes that diverge substantially from groups' serious activities

in terms of method and geographic locale indicate that they are being used as a low-cost alternative to developing a serious capacity, or as a function of dysfunctional principal-agent dynamics. My analysis is informed by close readings of CIDB event descriptions and additional academic literature, where supplements are available and necessary for understanding the group's behaviour.

9.4 Results and discussion: Comparing Canadian hoaxers and non-hoaxers

In this section, I discuss the results of my QCA analysis of terrorism perpetrators operating in Canada, and the combinations of group characteristics that are sufficient for distinguishing hoaxers from non-hoaxers (full results are presented in Table 9.3). Generally, the models exhibit above-90% solution consistency, and their solution coverages are roughly consistent with what was found for the WMDDDB-based sample: around 30% coverage for identifying hoaxers, and around 70% for identifying non-hoaxers. The discussion below highlights where solutions for the CIDB-based perpetrator sample converge with and/or diverge from the solutions generated by the earlier WMDDDB-based sample.

Table 9.3 QCA models and their minimized solutions (CIDB data)

MODEL A (crisp set)	$HOAXER = f(SINGLEISSUE, PUREETHNO, PURERELIG, HIERARCHY, SPLINTER, SERIOUS1)$	
	$SINGLEISSUE + SPLINTER * (PUREETHNO + \sim HIERARCHY) \rightarrow HOAXER$	
	Solution coverage: 0.357	Solution consistency: 1
	$\sim HOAXER = f(SINGLEISSUE, PUREETHNO, PURERELIG, HIERARCHY, SPLINTER, SERIOUS1)$	
	$\sim PURERELIG * [(HIERARCHY * \sim PUREETHNO * \sim SPLINTER) + \sim SERIOUS1] + (PURERELIG * SPLINTER) \rightarrow \sim HOAXER$	
	Solution coverage: 0.714	Solution consistency: 0.909
MODEL E (crisp set)	$HOAXER = f(LEFTRIGHT, CONTAINETHNO, CONTAINRELIG, HIERARCHY, SPLINTER, SERIOUS1)$	
	$SPLINTER * (\sim HIERARCHY + \sim LEFTRIGHT * CONTAINETHNO) \rightarrow HOAXER$	
	Solution coverage: 0.286	Solution consistency: 1
	$\sim HOAXER = f(LEFTRIGHT, CONTAINETHNO, CONTAINRELIG, HIERARCHY, SPLINTER, SERIOUS1)$	
	$CONTAINRELIG * (CONTAINETHNO + SPLINTER) + \sim CONTAINETHNO * LEFTRIGHT * HIERARCHY + \sim CONTAINRELIG * \sim SERIOUS1 \rightarrow \sim HOAXER$	
	Solution coverage: 0.679	Solution consistency: 1

9.4.1 Findings relating to campaign context

At a difference to the previous sample's findings, radical splinter origins here, under certain conditions, consistently predict the *use* of hoaxes—not just their non-use. Whereas the CIDB-based sample agrees with the theory-building chapters that splinter histories combine with religious motivation (and hierarchy) to explain non-hoaxing,¹⁴⁸ it shows that splinter histories are sometimes also relevant for explaining which groups *will* use hoaxes. In particular, radical splinter groups will use hoaxes when they are motivated by ethno-

¹⁴⁸ There are no cases of religious splinter groups that lack hierarchy; therefore, in every solution for non-hoaxing where splinter combines with religion, the presence of hierarchical structure is implied—it simply is removed from the parsimonious solution due to the simplifying assumption that non-hierarchical religious splinters would never hoax, if they were to exist.

nationalism (as long as they are not also economically left- or right-wing oriented), or when they lack hierarchy.

The CIDB-based solutions thus suggest an amendment to the hypothesis that radical splinter groups are unlikely to hoax. Rather, it appears that the presence of a splinter history exacerbates the effect on hoax use of the corresponding motivational or structural feature. For ethno-nationalist groups and/or non-hierarchical groups, which would be inclined toward hoax use even in the absence of splinter histories, having originated as a splinter group increases the group's likelihood of hoax use even further. For religious groups and/or hierarchical groups, which would otherwise be inclined toward never hoaxing, splinter origins further decrease the group's likelihood of ever hoaxing. Thus, while splinter histories may not always predict sole reliance on more radical violent forms as hypothesized, they are always consistent with more radical adherence to the *violence-related norms* of their associated ideology. In the case of ethno-nationalists, this means that they will be more likely to use hoaxes as a less violent alternative to serious activity.

Regarding the effect of the number of serious events perpetrated on groups' hoax behaviour, the CIDB-based sample is consistent with expected findings. The number of serious events perpetrated is not always relevant for distinguishing hoaxers from non-hoaxers, but where it is relevant in statements of sufficiency, it is the perpetration of a very small number of serious events that is consistent with never hoaxing. While the number of serious events perpetrated indicator never shows up as *sufficient* for explaining hoax use, a basic test for bivariate consistency with the hoax outcome shows

that perpetrating more than three events is almost always *necessary* for hoaxing (92.9% consistency) and perpetrating three or fewer events is almost always *sufficient* for non-hoaxing (92.3% consistency). When groups carry out more than three incidents, then, strategic hoaxing becomes possible in the context of a mixed strategy; they have built up sufficient reputational capital to force counterterrorist response and hoax without otherwise signalling weakness and undermining their credibility for future attacks.

9.4.2 Findings relating to group motivation

With respect to group motivation, single-issue (animal rights, environmentalist, or anti-abortionist) ideologies appear to be sufficient on their own for explaining hoax use, but again limited diversity and the absence of hierarchically-structured single-issue groups means that this result is over-determined. In the Canadian sample, the diversity of single-issue cases is even further limited by a two-thirds overlap between single-issue groups and otherwise non-hierarchical groups with radical splinter origins.¹⁴⁹ The CIDB is thus unable to confirm whether single-issue motivation exerts an independent effect on hoaxing, but neither does it present any contradictory evidence to that effect.

The CIDB-based results are also consistent with the cross-national sample's expectations regarding ethno-nationalism; and even more closely align with the originally-hypothesized relationships than the theory-building results suggested. Consistent with earlier findings, the intersection of religion-containing and ethno-

¹⁴⁹ This set relation explains why the left- and right-wing motive indicator does not combine with non-hierarchy to explain the hoax outcome in Model E; the same groups' hoaxing is already accounted for by their non-hierarchy and splinter origins.

nationalist-containing ideologies (i.e. religious ethno-nationalism) is a sufficient solution for non-hoaxing. In fact, this relationship is even more consistent in the CIDB-based sample, where bivariate tests for consistency reveal religious ethno-nationalism to be sufficient for explaining non-hoax use 100% of the time where it is present. Apart from the combined religious ethno-nationalist motive, the theory-building sample showed that groups lacking ethno-nationalism would be unlikely hoaxers, especially when they are also hierarchical; a combination on which the test sample agrees. Additionally, however, the CIDB-based sample shows that the *presence* of an ethno-nationalist ideology also contributes to the *use* of hoaxes when combined with splinter origins. As discussed above, the splinter histories of such groups may contribute to a more extreme adherence to the group's norms relating to the use of violence. In the case of ethno-nationalist groups, norms tend toward a desire to disrupt—not destroy—society, in order to maintain legitimacy in the eyes of constituents. That the WMDDDB-based sample underestimated this effect adds further support to this hypothesized mechanism: the same norms that would lead ethno-nationalist groups to hoax would also contribute to their disinterest in CBRN weaponry; they are thus far less likely to have been captured in the CBRN-centric WMDDDB than the CIDB, which is not limited by perpetrators' weapons preference (see Table 9.2).

Finally, the CIDB-based sample yields results consistent with prior expectations regarding the effect of religion on hoax (non-)use. As discussed above, religious ideology combines with ethno-nationalism and/or splinter origins as a sufficient condition for non-hoaxing. As for the WMDDDB-based sample, these groups' avoidance of hoaxes is likely a

function of their operation as social service-providing economic clubs, with members who are unlikely to defect during complex and lethal attacks. Hoaxes do not offer these groups much relative advantage.

9.4.3 Findings relating to organizational structure

With respect to organizational structure, the CIDB-based sample again yields results consistent with hypotheses and prior findings: there is a symmetrical, negative relationship between hierarchy and groups' use of hoaxes—although hierarchy is never sufficient on its own to explain hoax use or non-use. For the CIDB-based sample, market structure (that is, non-hierarchy) combines with splinter origins to explain which groups will use hoaxes. Hierarchy, on the other hand, combines with various motivational features and sometimes the perpetration of a small number of serious events to explain hoax non-use. The diversity of conditions involving group structure within and across the CIDB- and WMDDB-based samples of terrorism perpetrators is a function of the clustered diversity of cases included in the sample(s); there is no theoretical reason to explain why those particular combinations should explain hoax use or non-use. The key takeaway regarding group structure is that the absence of hierarchy is consistently relevant in explanations for hoax use, and its presence is consistently relevant in explanations for hoax non-use. With the exception of certain conditions that are sufficient for explaining hoax behaviour irrespective of an associated structural form (i.e. religious ethno-nationalism; religious or ethno-nationalist groups with splinter histories; and possibly

single-issue motivation), the presence or absence of hierarchy generally and effectively distinguishes hoaxers from non-hoaxers.

9.5 Case analysis: The strategic use of hoaxes by the FLQ

Apart from groups' splinter origins, the combinations of factors distinguishing hoaxers from non-hoaxers in the CIDB-based sample are consistent with earlier findings. At a macro-level, then, the QCA analysis supports certain strategic logics for hoaxing, and does not support others (as touched on above, and in greater detail in section 7.6). None of the preceding analyses, however, explored whether these assumptions hold when tested at the micro-level, examining the conformity of individual hoaxed (and serious) events with hypothesized mechanisms. In this section, I present a narrative overview of select FLQ hoax activity, to evaluate the degree to which there is event-level support for the strategic logics for hoaxing proposed in Chapter 5.

The FLQ emerged as a violent faction of (mostly Marxist) Quebec sovereigntists during Quebec's "Quiet Revolution," typically associated with the social and economic reforms of Jean Lesage (elected 1960) and, later, Robert Bourassa (elected 1970)'s Liberal governments.¹⁵⁰ Throughout its violent campaign, the FLQ perpetrated 160 serious terrorism incidents in Canada, concentrated predominantly in the nine-year period between 1963 and 1971. A clear majority of these incidents involved small-scale

¹⁵⁰ See Charters (1997, 134–38) for a thorough treatment of the political culture in which the FLQ emerged. As the movement evolved, the FLQ formally aligned itself as the "covert military wing" of the *Front de libération populaire* (FLP), "much as the Sinn Fein is the political wing and the IRA is the military wing of the Irish national movement" (Crelinsten 1987, 61)—only the FLP never organized as an official political party with electoral candidates.

explosions and incendiary attacks targeting federal government property (postal infrastructure, government offices, and armouries), financial hubs and businesses, and media outlets.

While its first documented serious incident occurred in February 1963, the start of the FLQ's hoax activity coincided with the release of the group's first manifesto, in April 1963. That month, the FLQ made at least three hoax warning calls, each warning of a supposed bombing against a particular target: the Interprovincial Bridge, connecting Ottawa and Hull; and various Montreal locations, including a radio station, hotel, and government offices. With the exception of the Ottawa-Hull bridge—which differed from serious FLQ targets at the time in its geography (located outside of Montreal) and target type (transportation infrastructure)—these hoaxed incidents were sequenced amidst serious FLQ bombings with similar tactics and targets. These serious incidents preceded the hoaxes, were interspersed between them, and continued in the months that followed, indicating that the FLQ was playing a mixed strategy to expend security service response resources without a corresponding investment in resources and risk from the group itself. For the Montreal-based bomb hoaxes, the FLQ had previously demonstrated its ability to carry out serious bombing operations in the area; its hoaxes would have been evaluated as credible, garnering appropriate attention from security and law enforcement.

The following month was also a period of frequent hoaxed and serious activity for the FLQ. In May 1963, the group appeared to hoax in the context of mixed strategies, by mimicking narrower *modus operandi*. On May 20th, the FLQ used its first hoax devices:

dynamite sticks—with no detonators or fuses—wrapped in paper, and placed in a Quebec City post office and various mailboxes. Only three days earlier, the FLQ had carried out a serious coordinated attack using long-delay timers attached to dynamite sticks placed in twelve mailboxes across Westmount, an Anglophone neighborhood of Montreal. The Quebec City hoaxers (whether or not they were affiliated with the cell responsible for the so-called “Westmount Project”) had clearly capitalized on the damage caused by the serious mailbox bombings days earlier, and hoaxed accordingly. While the use of dynamite sticks without a detonator may signal the perpetrators’ incompetence (i.e. that the Quebec City cell simply did not know detonators were required; and that the incident was not a hoax at all, but a failed serious attack), it may also signal that the groups were actively trying to signal their restraint (to demonstrate that while they had the materials to cause destruction, they would rather disrupt postal service with a hoax instead). In either case, the Quebec City mailbox hoaxes played into a mixed strategy, where the counterterrorists were forced into response due to a previously-demonstrated credible capability in the associated tactic. The Montreal-based cell may have picked up on the effect of this mixed strategy: on May 24th, a mailbox bomb hoax device was found in Montreal, echoing the serious attack that had occurred in Westmount only one week prior. This collection of incidents illustrates that multiple strategic logics may be simultaneously at play, where hoax activity is concerned.

The FLQ’s pattern of carrying out hoaxes to mimic previously-demonstrated serious behaviour is also present with respect to bombing attacks targeting Montreal armouries. On May 9th and 20th, 1963, the FLQ conducted serious bombings of

armouries—in spite of heightened surveillance surrounding them. On May 21st, a hoaxed bomb warning was given regarding a supposed bomb at an armoury two blocks from the Montreal Police headquarters. No such bomb was found to be present, but the warning was seen as credible given the success of similarly-designed attacks earlier in the month.

The Ottawa-Hull bridge hoax mentioned above indicates an entirely different strategic logic. When the hoax warning call was made, Canadian Pacific Railway security responded with around-the-clock surveillance, and a serious attack never followed. However, a subsequent event description reveals that an FLQ cell *had* seriously planned to bomb this bridge, but was “deterred by a policeman’s questioning” (CIDB event ID #19630520110310001). The hoax call thus may have been intended as a dry run, but the tactical intelligence gleaned from it regarding security response was sufficient to deter the perpetrators from subsequent action. More plausibly, the call may have been a real (not hoaxed) warning for an intended attack, but the perpetrators had not yet placed the bomb. When they observed the increased security presence, they may have simply decided not to follow through.

If it is interpreted as a hoax, however, the Ottawa-Hull bridge incident fits into a pattern of the FLQ hoaxing as a low-cost means of expanding its geographic spread, to compensate for a lack of operational capacity in an area of symbolic significance: the nation’s capital region. In May 1963, the FLQ had also given a hoax warning regarding a bomb on a plane carrying Prime Minister Pearson, just outside of Ottawa; it also hoaxedly threatened to bomb the North Atlantic Treaty Organization (NATO) ministerial conference in Ottawa later that month. At this point, the FLQ had never successfully

carried out a serious out-of-province attack. It was only years later, in December 1968, that the group ever succeeded in concluding a serious incident in Ottawa.

Following the burst of hoaxes in April and May 1963, the FLQ's hoax activity stopped until October 1965. At that time, the FLQ placed several fake bomb devices—including real sticks of dynamite but no viable means of detonation—at various locations around Quebec City. The “bomb” devices were accompanied by hoax warning calls, alerting police to their presence. In a similar event in July 1967, in Montreal, the FLQ placed an explosive device at City Hall. They simultaneously made hoax warning calls alerting police to the City Hall bomb, as well as bombs supposedly placed in various local mailboxes. An army explosives specialist safely diffused the device found at City Hall, noting that it involved only minimal explosive materials and “probably could not have exploded” (CIDB event ID #19670727110306001). While it is unclear whether the physical device was intended as a hoax or whether it was a poorly-constructed serious device, what *is* clear is that the group exploited its possession of real explosives to give credibility to the hoaxed warnings they called in at the same time. In this case, the hoaxes did not drown out other serious activity (even if it was really intended to explode, the perpetrators had alerted law enforcement to the device's presence), but they *did* constitute a low-cost means of spreading emergency response resources very thin. Furthermore, that they were accompanied by real explosives signalled the group's material (if not operational) capability, and corresponding preference for avoiding serious harm to constituents.

The FLQ's serious violent activity came to a pinnacle in the 1970 October Crisis, precipitated by the abduction of James Cross, a British attaché, and the abduction and eventual murder of Deputy Premier and Quebec Labour Minister Pierre Laporte. This extreme activity coincided with the publication of the FLQ's second manifesto, and a campaign of coordinated bomb hoaxes designed "to support the impression that a widespread revolutionary movement was afoot" (Charters 1997, 155); the hoaxes were a low-cost means of projecting power and force strength, despite a rather weak and disorganized reality. By the following year, the FLQ had been infiltrated by police and security services to the point that any subsequent activity was carried out under police scrutiny, or even covertly facilitated by undercover agents (Crelinsten 1987, 60).

Hoaxes (and serious acts) carried out in the FLQ's name after 1972, when the group effectively ceased to exist, are thus likely indicative of improper attribution.¹⁵¹ When incidents are carried out in the middle of a group's campaign—and *especially* when that group is already characterized by a market structure—it is more difficult to distinguish between group-sanctioned activity (or at least activity perpetrated by true group members), and false flag incidents carried out by unaffiliated individuals.

¹⁵¹ After 1972, the CIDB's first FLQ-attributed incident was a July 1976 hoax bomb device (accompanied by a hoax warning call, to tip off the authorities) placed during a game at the Montreal Olympics stadium. This incident occurred more than four years following the FLQ's last serious bombing; and the group never hoaxed or carried out a serious explosive bombing thereafter. The 1976 hoax was likely perpetrated by an unaffiliated member exploiting the group's name and associated fear. Crelinsten (1987) suggests that events carried out in the FLQ's name following 1972 do not signal a revival of the FLQ, but rather the result of "a few crackpots" adopting the group's name to give the impression of revival: "Only time and research," he concludes, "will tell if any organization really exists behind the façade of these isolated incidents" (Crelinsten 1987, 81). Since the FLQ had ceased formal operations after 1972 (now confirmed with an additional thirty years of hindsight, following Crelinsten's analysis), it is easy to re-assign the FLQ's post-1972 to unaffiliated perpetrators.

The FLQ's (dis-)organization embodies this sort of market structure. The FLQ suffered from frequent generational changes, with as many as fifteen different networks simultaneously operating under the FLQ moniker: "The fundamental point about the FLQ's organizational culture is that *there never was a single organization called the FLQ*" (Charters 1997, 138). Another scholar describes the FLQ's structure as "a centrifugal model minus the hub of the wheel," with no central leadership, and only a "common ideal of a separate Quebec" linking "successive groups of friends who would decide to form a 'cell' and call themselves the FLQ" (Crelinsten 1987, 59). In many cases, there was contact between these cells; if ever they were strictly compartmentalized, this "was more the result of independent initiatives taken by individuals separated by geography or everyday life than a conscious policy emanating from a central command structure" (Ibid.). Importantly, this loose structure and lack of leadership meant that there was no mechanism for controlling which types of activities would be carried out in the FLQ's name. Accordingly, Charters (1997) describes how principal-agent divergences in the FLQ could easily emerge: "Once they were committed to the movement and had created their own networks, it was difficult to prevent [individuals and cells] from launching operations whenever they wished, even if practical or political considerations dictated caution" (Charters 1997, 140). It is thus difficult to attribute with any certainty hoax activity carried out in the FLQ's name to a unitary strategic logic for the group.

However, although different FLQ cells may have diverged with respect to their tactical choice, they would have been subject to the same "ideational oligarchy" (Gray 2013) impressing a common aspiration of establishing an independent Quebec, and a

corresponding desire to be viewed as a legitimate political force. While hoax use is consistent with the hypothesis that such ethno-nationalist groups prefer less severe modes of violent activity to signal restraint and build legitimacy in the eyes of constituents, the FLQ's pattern of serious violent behaviour further buttresses the mechanism underpinning this explanation.

The FLQ network appears to have chosen its targets and tactics so as to avoid fomenting outrage amongst both Francophone and Anglophone populations in Quebec. With the exception of the Westmount mailbox bombings, they were successful: "*the FLQ made no serious attempt to polarize the Anglo-Francophone populations by means of mass murder incidents*" (Charters 1997, 142; emphasis in original). Even their highly symbolic kidnapping incidents—including the capture and subsequent murder of Quebec cabinet minister Pierre Laporte—were not originally designed to threaten Quebecois constituents. The original plan had been to kidnap two *foreign* diplomats—"a British diplomat, as a protest against cultural colonialism, and an American diplomat, as an attack against economic imperialism" (Crelinsten 1987, 69)—and to keep the captives alive; James Cross's captors explicitly told the Quebec Justice Minister that they did not plan to kill their captive (Ibid.). Although Laporte died under his captors' watch, "the circumstances surrounding his death remain unclear" (Ibid., 64); it appears to have been connected with his captors' refusal to bring him to a hospital when, the day prior to his death, he badly injured himself in an attempted escape (Ibid.). Despite the kidnapping's unfortunate end, the FLQ had not intended for the infamous incident to be fatal.

While the hostage scenarios are outliers in the FLQ's broader campaign, casualty-aversion appears to be a consistent feature of their bombing and incendiary activities. Despite having perpetrated 160 serious incidents, only eleven ever resulted in injury, and only two (including the October 1970 hostage scenario) resulted in fatalities. The FLQ's only other fatal incident occurred in May 1966, at the head offices of La Grenade Shoe Company. Perpetrator testimony, however, reveals that the cell responsible had not intended to hurt people. The individual who placed the bomb had supposedly believed it to be only "an inoffensive little bomb" (CIDB event ID #19660505110302001), and the bomb-maker himself had placed a warning phone-call once the bomb was in position; the company dismissed the warning due to pressure to respond to customer calls, which is why individuals were present—despite the perpetrators' efforts—when the bomb exploded. When bomb warning calls were not given, the FLQ sought to diminish harm to humans by timing attacks on property to occur after hours, when people were not expected to be present. FLQ-associated injuries were thus usually incidental, affecting passers-by or night watchmen. Daytime attacks that would result in greater human injury were considered an anomaly within the FLQ's tactical repertoire (CIDB event ID #19690121110314001).

Demonstrating casualty aversion, however, is only half of the mechanism underpinning ethno-nationalists' preference for hoax use. The other half is the expectation that more aggressive violent strategies (particularly against constituents) would erode public support for the group, and reduced estimations of its legitimacy. Public opinion on the FLQ during and after the October Crisis is consistent with this

hypothesis. The high-profile kidnappings of James Cross and Pierre Laporte were designed explicitly to bolster valuations of the group's legitimacy: by exploiting media attention as a platform for making demands, the FLQ guaranteed that negotiations with the government would happen in public, thereby "open[ing] the way to their legitimation as political actors" (Crelinsten 1987, 74–75). During the early phases of the crisis, once the kidnappings were underway and the FLQ's second manifesto had been released, public opinion in favour of the FLQ had accelerated (Crelinsten 1987). It was not until the FLQ crossed a line with Laporte's murder that public opinion turned. The murder "was a propaganda disaster for the FLQ [...]. Thereafter, separatist sentiment was channeled into liberal democratic constitutional politics, through the *Parti Québécois*" (Charters 1997, 162). The FLQ's murderous activity was a clear violation of its legitimacy, and sympathy toward the organization and its revolutionary potential correspondingly evaporated.

The discussion above highlights the substantial degree of equi- and multi-finality predicted at this study's outset with respect to why some terrorist groups turn to hoax activity. Even a single group with relatively constant organizational features appears to hoax for different reasons at different times. Furthermore, it confirms that the same hoax incidents may be plausibly explained by multiple logics, simultaneously. Based on observational data alone, it is impossible to corroborate which (if any) strategy truly caused the hoax outcome.

9.6 Evaluating the theory of terrorism hoaxes in light of Canada's hoax history

In this chapter, I used a new dataset of terrorism incidents in Canada to test my previously-established findings with respect to the factors affecting terrorists' hoax behaviour. The cross-case analyses conducted in Chapters 6 and 7 were limited to a specific subset of terrorist perpetrators: organizations (not individuals) responsible for at least one serious event in ITERATE, the WMDDDB, or the GTD, who also demonstrated an interest in using CBRN weaponry at some point in their campaign. The Canadian-based sample was not restricted by any weapons preferences; it included all serious perpetrators responsible for terrorism incidents in Canada. My analysis thereof sheds light on the degree to which my earlier findings are extrapolable to cases beyond the CBRN-interested sample of perpetrator groups, and whether the mechanisms theorized to underpin my hypotheses hold, when evaluated at the event level.

With respect to the cluster analysis-derived typology built in Chapter 8, the Canadian test sample does not corroborate meaningful distinctions between classes of hoax perpetrators. The typology is undermined by both the varied structural forms of the many left-leaning ethno-nationalists operating in Canada; and the apparent rarity of limited serious hoax perpetrators. It may be that the Canadian sample of hoaxers is simply too small to represent a valid test of the typology's utility. What exists, however, suggests that the typology's classes of hoax perpetrators are not meaningfully extrapolable beyond the sample used to create them.

In terms of differentiating hoaxers from non-hoaxers, however, the Canadian test sample corroborates findings from the earlier theory-building sample. Organizational

hierarchy exhibits a consistently strong, negative effect on whether a group will hoax: hierarchical groups are unlikely hoaxers, and non-hierarchical groups are likely hoaxers. With respect to the effect of ideology on hoax behaviour, the test sample supports the notion that single-issue motivation is a perfect predictor of hoaxing (although, again, the limited diversity of cases means that the single-issue effect cannot be disentangled from that of non-hierarchy). It also corroborates the previously-determined effect of religion: religious ethno-nationalist groups and religious splinter groups are particularly unlikely to ever hoax. With respect to ethno-nationalism, the test sample indicates a positive, symmetrical relationship with respect to hoaxing—a relationship which Chapter 7's QCA analysis only substantiated in one direction. Here, it appears that lacking an ethno-nationalist ideological component is sufficient for explaining groups' non-hoax use, especially when that group is also hierarchical; and *having* an ethno-nationalist ideological component is sufficient for explaining groups' hoax use, especially when that group is also a splinter organization. The test sample also corroborates the finding that hoaxing groups almost always carry out more than three serious events. This relationship is even stronger in the test sample than in the theory-building one: for Canadian-based perpetrators, carrying out more than three serious events is almost always *necessary* for hoaxing, and perpetrating three or fewer is almost always sufficient for non-hoaxing. As in the theory-building sample, predictors of hoax behaviour are better at identifying which groups will never use hoaxes than which groups sometimes will.

The only substantial difference in results gleaned from the test and theory-building samples pertains to the effect of a group's splinter origins on hoax behaviour. In

earlier chapters, splinter history was theorized and empirically-demonstrated to be consistent with hoax non-use: splinter groups were more extreme in their ways, and correspondingly less-inclined to use such a moderate tactic as hoaxing. The test sample corroborates this finding, but also suggests other conditions under which splinter origins lead to both the hoax and non-hoax outcome. Here, splinter origins appear to exacerbate the effect on hoax use of the corresponding motivational or structural feature—in either direction. For ethno-nationalist and/or non-hierarchical groups, splinter histories thus increase the likelihood of hoax use; for religious and/or hierarchical groups, splinter histories increase the likelihood of hoax non-use. That this more multifaceted relationship is evident in the current sample and not the theory-building one is partially a function of the lower relative frequency of ethno-nationalists in the WMDDDB-based sample. Consistent with the hypothesis that ethno-nationalists prefer to use hoaxes because of their concern with audience costs and pursuit of legitimacy, ethno-nationalists are particularly unlikely to be interested in CBRN weaponry. They are thus under-represented in the WMMDB, relative to the broader universe of terrorist perpetrators.

My narrative treatment of the FLQ provides observational evidence that certain strategic logics for hoaxing hold, when examined at the event level. However, this discussion also highlights the inherent equifinality of hoax behaviour. In some cases, the FLQ's event sequencing clearly aligns (particularly in the earlier phases of its campaign) with the mixed strategies explanation for hoaxing: the group made credible hoaxes by mimicking already-demonstrated tactics, thereby forcing a counterterrorism response with minimal perpetrator investment of resources. Its organizational structure—or, more

appropriately, lack thereof—and corresponding implications for attack planning and coordination are simultaneously consistent with a principal-agent explanation for hoaxing. Alternatively, both the FLQ's hoax and serious activities are consistent with the mechanism underpinning the effect of ethno-nationalist ideology on hoax behaviour: that ethno-nationalists are particularly concerned with being perceived as legitimate political actors, and take care to choose less-severe modes of violent activity that do not impose substantial audience costs. Finally, careful examination of the FLQ case reveals evidence of wrongly-attributed hoaxes, which were perpetrated as false flag incidents by unaffiliated individuals.

Collectively, while this chapter corroborates much of the theory-building analysis presented earlier, it also highlights deficiencies in applications of cross-national terrorism events datasets for studying the hoax phenomenon. Most importantly, the sheer scope of the cross-national datasets' coverage invariably corresponds with shallower, and less-nuanced coverage. Nowhere is this more evident than in the cross-national datasets' underestimation of serious events perpetrated per group. While in the case of Direct Action the CIDB only contributed one extra group to the analysis, for the FLQ, the CIDB contributed an extra 158 serious events. According to the cross-national datasets, the FLQ had been a limited serious perpetrator, responsible for only two incidents; in Canada, it is the most prolific terrorist group of all time. This problem is partly a function of the cross-national datasets' depth of coverage, but it is also a problem of temporal scope: ITERATE and GTD only capture events beginning in 1968 and 1970, respectively. Unless a group's entire lifecycle is captured within a dataset's timeframe, aggregating event details to the

group level will always be a sub-par approach. The false flag hoax incidents carried out under the FLQ moniker exemplify a further challenge inherent in discretely identifying event perpetrators within large-*n* events databases.

By virtue of their respective scopes, the WMDDB and ITERATE also underestimate the diversity of hoaxers. The CIDB suggests that many hoaxing groups are ethno-nationalists, who are disinterested in CBRN weaponry. A substantial portion of these are also limited to domestic activities. Relying on the WMDDB and ITERATE as sole sources of hoax events, then, systematically underestimates the activity of these groups. Accordingly, when the WMDDB or ITERATE are used to evaluate hoax activity, results should be caveated that they *only* apply with respect to the corresponding CBRN-centric or transnational-oriented sample.

10. Conclusions and steps for future research

The overarching goal of this study has been to present the field's first systematic attempt at evaluating terrorists' hoax behaviour. At a macro-level, the study succeeds in bringing coherence to what has previously been an incomplete and sometimes inconsistent literature regarding hoax use by terrorists.

Given the lack of hoax literature to date, this study was framed as a theory-building exercise. My empirical approach involved harnessing multiple research methods and all publicly-available large- n data on hoaxes, to test for consistency with a rationalist framework outlining how hoax use can be strategic for terrorist groups under certain conditions. Some robust results emerged, based on indicators of group motivation, structure, and campaign context. In Chapter 9, I summarized my empirical findings based on a cross-national sample of terrorism perpetrators (section 9.2), and presented a test thereof using an alternate data source: the Canada-centric CIDB. In section 9.6, I explicitly addressed the degree to which the Canadian data confirm patterns in perpetrators' use of hoaxes, and indicate areas in which the cross-national datasets offer biased representations of hoax perpetrators.

Since these substantive matters have been summarized in Chapter 9, my conclusions here emphasize this study's secondary aim: providing insight into the degree to which existing data can shed light on the hoax phenomenon. In this chapter, I appraise the ways in which my hoax theory, data, and empirics align (and diverge). I emphasize the limited extent to which observational perpetrator data (which forms the bedrock of my

analysis) can corroborate the strategic logic(s) underpinning terrorist groups' hoax activity. With data-related caveats in mind, I go on to discuss how this study still yields substantively and methodologically instructive findings for both policy and research. I conclude with a discussion of possible avenues for further inquiry in a hoax-centric research agenda.

10.1 The limits of hoax data, theory, and empirics

In the subsections that follow, I elaborate how available hoax data—and the limits thereof—qualify the study's findings; and how these findings relate to the rationalist theory for terrorism hoaxes proposed in Chapter 5.

10.1.1 Hoax data and the study of perpetrator groups

In this study's introduction (section 1.2.1) I described in detail the general challenges posed by terrorism events data, and hoax data in particular. Here, I will reprise those elements that are directly relevant to the interpretation of this study's findings, and discuss other aspects revealed by my analysis.

The most significant data limitation in this study is that the scope of terrorism events data coding for hoaxes is imperfect. The available cross-national datasets cover only transnational hoaxes (ITERATE) or hoaxes involving CBRN weapons (the WMDDDB); there is no publicly-available, cross-national data-source that takes stock of domestic hoax activity involving conventional weapons. Comparisons with the Canada-centric CIDB illustrate the degree to which the cross-national datasets underestimate the incidence of

hoax activity, and the diversity of hoax perpetrators. When time periods are held constant, the CIDB captures over ten times more terrorism incidents in Canada than the other datasets, and between 2.5 and 12.9 times more hoax activity. The cross-national datasets thus substantially underestimate some groups' productivity. The FLQ, for instance, was responsible for only two serious terrorism incidents across the WMDDB, ITERATE, and GTD; the CIDB reveals an additional 158 serious FLQ events not accounted for elsewhere. Fortunately—although this may be a matter of chance, not design—the CIDB does not reveal additional *hoaxed* events for groups that were previously classified as non-hoaxers. Dependent variable assignments for the analyses of hoaxers and non-hoaxers in Chapters 6 and 7 are thus reliable.

While the cross-national datasets severely underestimate the incidence of terrorism activity, however, they do an acceptable job of capturing responsible hoax perpetrators: 85.7% of the Canadian hoaxing groups had already been included in the sample of hoaxers examined in Chapter 8's cluster analysis. Many of these groups (and their non-hoaxing counterparts), however, do not ever show an interest in CBRN weaponry, and so they are severely under-accounted-for in the predictive analyses of Chapters 6 and 7. Notably, the WMDDB-based sample of perpetrators underestimates the ideological variety of splinter organizations, structural variations amongst ethno-nationalist groups, and the presence of ethno-nationalist groups altogether. Although the WMDDB-based sample also overestimates the presence of religious groups, such actors are otherwise similarly constituted across the WMDDB-based and CIDB-based samples, such that their over-representation does not skew results.

Relying on the WMDDDB and ITERATE as sole sources of hoax events, then, systematically underestimates the activity of certain groups—namely, domestic-oriented ethno-nationalists who are disinterested in CBRN weaponry. Where the WMDDDB or ITERATE are used to evaluate hoax activity, results should be caveated that they *only* apply with respect to the corresponding CBRN-centric or transnationally-oriented sample. Conversely, although the CIDB is not limited in terms of weapon type or perpetrator orientation, it *is* limited in geographic scope: it is unclear whether the experience of terrorism perpetrators operating within Canada is generalizable to other contexts. In particular, the Canadian sample may over-state the incidence of ethno-nationalism, by virtue of the country's long history of linguistic nationalism, multicultural heritage, and associated prevalence of émigré causes. A more homogenous nation-state may not experience the same variety of terrorism perpetrator types as Canada. Still, the high degree of overlap in hoax perpetrators across the CIDB and the cross-national datasets indicates that—despite their limitations—the cross-national datasets can still offer a reasonably accurate degree of insight into the organizational perpetrators of terrorist hoax activity.

Another source of bias stems from how the perpetrator samples themselves were designed. First, I examine only *organizational* perpetrators. My findings thus say nothing about the motives of individual hoax perpetrators (a choice which I justify on practical and theoretical grounds in section 1.3.1). There are many challenges relating to the operationalization of terrorist organizations (Miller 2016). In particular, there are pitfalls associated with linking attacks to particular organizational perpetrators (LaFree, Dugan,

and Miller 2015, 72–74). Although I endeavoured to implement consistent coding rules,¹⁵² decisions on how to operationalize perpetrator groups and distinguish them from related entities were ultimately subjective. In my sample construction, I erred toward organizational continuity across time. Other datasets preserve the name of a perpetrator at the time of attack, such that when a group stops carrying out attacks under a certain name, they may misleadingly imply that the organization itself has ended (LaFree, Dugan, and Miller 2015, 74; in reference to the GTD's coding practices). Neither approach is perfect, but my preference for continuity is apt in the present context since my study's goal is to uncover strategic logics across the entirety of groups' campaigns.

Second, my sample includes only terrorist organizations *responsible for at least one serious event*. It thus systematically excludes pure hoaxers. However, just as this study identified organizational characteristics associated with terrorist groups who will never use hoaxes, there may equally be characteristics associated with groups who will never carry out serious activity. None of this study's findings can explain the behaviours of these groups, which means I do not address the full breadth of terrorism hoax activity. Since a group engaged in only hoax activity would likely fall short of legal definitions of terrorist actors (unless the hoaxes interfered with the provision of an essential service), however, the exclusion of pure hoaxers from the analysis does not pose a significant theoretical problem with respect to understanding terrorist hoax activity.

¹⁵² For instance, I recognized splinter groups as distinct entities, wherever the split reflected a permanent schism; and I retained group coherence when name changes were not accompanied by substantial shifts in membership or organizational structure.

Third, the bulk of the study's empirical analysis was conducted on a sample of perpetrators *with a demonstrated interest in CBRN weaponry*. I selected this sample because it was feasible to collect data on the required independent variables for this manageable sample size, but also because CBRN-interested groups presented a hard case for hoaxing: they are otherwise interested in the most serious of terrorism tactics, so why would they also engage in fake terrorism? The good news is that, as a hard test, if patterns in hoax use are visible here—and they are—they should also be visible elsewhere. The bad news is that we know there are meaningful differences between groups who are interested in CBRN weaponry and those who are not (see section 4.2.3), and so it is unclear to what extent their experience is extrapolable to perpetrators who are only interested in conventional weaponry. As discussed above, the test of findings against the CIDB-based perpetrator sample partially compensates for the WMDDDB-based sample's deficiencies. By virtue of focusing only on a single country context, however, the CIDB-based sample also produces results that are not necessarily generalizable. Fortunately, findings across the two samples are relatively consistent, suggesting that the patterns uncovered are attributable to the variables themselves, and not features of their respective samples.

Finally, there are limitations with respect to the subjectivity inherent in assigning discrete codes to characteristics of terrorist organizations, which are fundamentally complex social phenomena; my coding to that end is not infallible.¹⁵³ Part of the problem

¹⁵³ The IRA and RIRA may be used as examples to illustrate these coding challenges, particularly with respect to motive and splinter origins. Fortunately, although alternate classifications for variables (especially motive) can have substantial implications for analysis (see Flemming, Stohl, and Schmid (1988);

here lies in the requirement that static indicators be assigned to groups who naturally evolve over time. Because my purpose in this study was to examine hoax use within the context of broader terrorism campaigns, my coding decisions reflect trends that are both salient and persistent; they are not reflective of organizations' temporal whims, or features that do not pertain to the entire period under study.¹⁵⁴ While incorporating such

footnote 103), my study's key findings do not appear to be sensitive to alternate codings for motive and splinter origins for the IRA and RIRA (see footnotes 168 and 172).

The IRA and RIRA are each coded as purely ethno-nationalist, although it could be argued that both also, at times, espoused an economically leftist dimension. That the IRA at its founding was politically Marxist is without doubt. However, whether this Marxism endured across the full time period studied here is up for debate (McKinley 1992). Throughout the conflict, the IRA and its associated political body (Sinn Fein) became increasingly pragmatic in their pursuit of a resolution to the conflict (Moloney 2002). This pragmatism required a re-evaluation of certain foundational tenets of the group's ideology. John Morrison's (2016) interview with Danny Morrison illustrates aspects of this balancing act. Importantly, the salience of (democratic) socialism in the group's ideology appears to have declined over time, to the point of practical insignificance by the time they reached the crux of a resolution. Most public data on the IRA (and RIRA)—particularly those discrete indicators coded within existing datasets—focuses exclusively on the groups' ethno-nationalist dimensions; any mention of a leftist economic orientation is clearly secondary or, more frequently, omitted altogether (as in TOPs and Jane's, as well as BAAD-1 and Jones and Libicki's (2008) data—the primary sources consulted in this dissertation's coding process). Within the present study, only the primary ethno-nationalist aspect of the IRA and RIRA has been retained, but it is necessary to note that both could be coded as also being economically leftist in future research. While doing so would undoubtedly enrich the quality of the data, it would require a robust justification for breaking with past coding trends, and collaboration amongst regional and group-specific experts across the full sample of perpetrators. Future large-*n* data coding efforts would do well to incorporate the fine-grained nuance that qualitative approaches and detailed case-based knowledge can provide. In the present context, however, a primary aim was to test the reliability of existing data, which involved retaining many coding decisions from existing data.

The question of groups' splinter histories is also open to debate. Whereas RIRA is unambiguously a splinter of the Provisional IRA (PIRA), assigning the splinter indicator to the IRA itself was problematic, due to the attribution of events to groups in the events-based datasets that form the basis of my sampling. Since events are not consistently differentiated between the founding IRA group and PIRA, I erred on the side of conservatism in assigning the splinter variable (for instance, the GTD only assigns events to the IRA, not PIRA, even though its earliest events occur in 1970—well after the founding IRA's supposed dissolution). Because of this inability to distinguish the events of the founding IRA from those of PIRA, and the implied organizational continuity between them, I treated the IRA and PIRA as a single organization (reflected in my dropping of "Provisional" in the group's name). In my framing of the splinter history hypothesis, a group cannot be a "splinter" if it is effectively a continuation of the same organization. It is worth noting that the IRA is coded in the present dataset as a militant wing (of Sinn Fein), although this variable never figured significantly in the study's analysis or results.

¹⁵⁴ For Chapters 6, 7, and 8, this means that the coded characteristics describe all (or most) of the groups' activity between 1968 and 2012. For Chapter 9, the characteristics must apply to all (or most) activity between 1960 and 2015. The exceptions, of course, are those groups who underwent such permanent structural change that it justified splitting the group into two distinct observations, identified in the

nuanced differences—particularly with respect to ideology—would have improved my results’ theoretical leverage, doing so in the context of a cross-sectional analysis would have artificially inflated the number of cases in my sample. Furthermore, the required level of deep, case-based knowledge for all groups within the sample would have been beyond the realm of feasibility for a single researcher. Maintaining coding consistency with extant data sources and their coarse indicators of organizational characteristics—however flawed—was both feasible, and purposefully congruous with my aim of testing the limits of existing data.

The data limitations discussed above are not unique to the present study. Difficulties in collecting data on terrorism perpetrators are well-documented, and largely a function of the research subjects’ clandestine nature (Ackerman and Pinson 2014, 230; LaFree and Dugan 2007, 189; LaFree, Dugan, and Miller 2015, 23; Silke 2001, 2, 12; Young 2016, 13). A reliance on secondary data (namely media-based, open-source events databases)—however flawed—is inevitable in terrorism research. Other, more traditional, sources of crime data are deficient for dealing with terrorism (LaFree and Dugan 2007, 182–83; LaFree 2010, 24; LaFree, Dugan, and Miller 2015, 14–15).¹⁵⁵ Events

appendices by their periods of coverage (e.g. al-Qaeda, pre- and post-9/11; Khmer Rouge, pre- and post-1996).

¹⁵⁵ The arguments for which traditional crime data sources are not used to study terrorism are summarized by LaFree, Dugan, and Miller (2015): “Although terrorist event databases such as the GTD are imperfect, they persist because they are useful. They rely on media reports because the more traditional sources of data on illegal events fall short of providing the information needed for a comprehensive picture of worldwide terrorist attacks. We are unable to depend on official sources, because terrorism data are not uniformly collected by police and legal authorities within individual countries. We cannot conduct victimization surveys because victims of terrorism are often chosen randomly and report little useful information. Self-reported data are problematic because terrorist offenders are often unwilling to talk to researchers and unlikely to be objective when they do talk. In the absence of these more traditional data sources terrorist event databases provide both researchers and analysts with an alternative method for studying characteristics of attacks” (226).

databases are frequently applied in large-sample terrorism research, with findings caveated by recognition of the data's relevant limitations (see Ackerman and Pinson 2014, 236; Asal, Ackerman, and Rethemeyer 2012, 230, 249; Findley and Young 2012, 709; LaFree and Dugan 2007, 198; LaFree, Dugan, and Miller 2015, 25, 227).

Even for the harshest of terrorism research critics, "researchers can still take steps to address [serious concerns with the manner in which data is collected] when they progress to analysing the gathered data" (Silke 2001, 9).¹⁵⁶ This study's use of multiple samples drawn from the best data available to me, and a multi-method research design, help to hedge against the data's limitations.¹⁵⁷ My findings thus have sufficient internal validity (and a degree of external validity, accorded by the check against CIDB data) to draw preliminary conclusions on the factors affecting terrorist groups' hoax use.

Furthermore, there is value in knowing the extent to which existing data can shed light on the hoax phenomenon. Given the challenges of terrorism data collection, if large-sample research is to be conducted on hoaxes, using existing data as a starting point is a reasonable choice. First, this data exists and is in wide circulation, yet nobody has tested it. Second, the research community's understanding of hoaxes will progress more quickly

¹⁵⁶ Silke (2001) identifies the reliance on secondary data sources as a significant impediment to the advancement of knowledge in terrorism research. He pessimistically concludes that "[o]ur knowledge of terrorism most certainly is deficient but the field shows no clear ability to improve this situation. After 30 years of study, we simply should know more about terrorism than we currently do. That we continue to languish at this level of ignorance on such a serious subject is a cause of grave concern" (Silke 2001, 13).

¹⁵⁷ The validity of WMDDDB events data for CBRN terrorism has been questioned. As a result, its creator went on to develop the POICN database, which explicitly incorporates measures for source reliability and data validity. Of the 449 CBRN events initially coded for in POICN in the 1990-2010 period, 17% displayed some event uncertainty, and 9% displayed considerable uncertainty (Ackerman and Pinson 2016, 628). Approximately one quarter of incidents thus display a considerable amount of noise. Unfortunately, the POICN database is inapplicable to the present study, since its coders did not receive funding support to include information on hoax incidents.

if researchers' time and effort are not expended on parallel data collection efforts. Third, we cannot refine or improve in data collection or the empirical analysis thereof if we do not first know the limits of what we already have. My findings are thus a valuable first step, substantively and methodologically.

10.1.2 The (mis-)alignment of theory and empirics

Although the data discussed above generate some consistently robust relationships relating the hypothesized indicators with the hoax (or non-hoax) outcome, the relationship between these empirics and the strategic framework outlined in Chapter 5 is frequently unsubstantiated.

This misalignment between the study's overarching theory and its empirics stems from the nature of the theory and of available data. To directly test a rationalist theory would require knowledge of the groups' decision-making processes, and likely first-hand accounts of why the groups behaved the way they did:

“In the absence of extensive information on a terrorist group's internal deliberations and attack planning, statements about the intent of particular operations [derived from events datasets] will invariably be educated guesses rather than certain statements of fact. These realities mean that there is a great deal of noise in the data, some of which originates from limits in the coding efforts, but more from the nature of the data on which it is based.”
(Jackson and Frelinger 2008, 588)

Observational events data are thus a subpar means of evaluating perpetrator logic and intent, but—given terrorist groups' clandestine nature—they are frequently the best (and only) data available to researchers.

The specific hypotheses I test in Chapters 6, 7, and 9 are at best proxies for elements of group strategy—determined as a function of which indicators are readily

observable. They can indirectly corroborate the presence of certain logics, but they do not purport to represent direct tests thereof. As a result, this study consistently finds that the same empirical results simultaneously corroborate multiple strategic logics. It is impossible to disentangle—based on observational events and perpetrator characteristics alone—which logics are truly driving group behaviour. Although I vary my sample and variable selection in an effort to achieve greater nuance to this end, even the variables that I test directly are sometimes over-determined due to the limited diversity of cases (in particular, the effect of groups' animal rights, environmentalist, or anti-abortionist motivations is confounded with the effect of non-hierarchical structure). All of this, however, is consistent with my theoretical prediction that the causes of hoax activity are complex and multiple.

Still, while this study cannot determine which strategic logics *are* motivating groups' hoax behaviour, it can to a degree rule out logics for which there is no empirical evidence. In particular, I did not uncover any "smoking gun" evidence for hoaxes being used as dry runs of future serious attacks. This non-finding perhaps indicates why we do not observe hoaxes more frequently: even though a dry run hoax could strategically be used to improve the effectiveness of a future serious attack, if the hoax were discovered and its perpetrators arrested, whatever serious plot was being pursued would be foiled. Hoaxes still involve operational risks for perpetrators. Additionally, my findings with respect to the number of serious events perpetrated per terrorist group have clear implications for understanding terrorists' mixed strategies. The CIDB data confirms that terrorist groups only use hoaxes when they have carried out more than three serious

incidents. In this context, it appears that this minimum threshold of three events accords the perpetrators sufficient reputational capital to force counterterrorist response and hoax without otherwise signalling weakness and undermining their credibility for future attacks.

The disjoint between the study's theory and empirics is compounded by the cross-sectional nature of the analysis and binary dependent variable used predominantly throughout. By aggregating all events to their organizational perpetrators, I lose important temporal details that would otherwise help clarify the perpetrators' underlying logics—a problem that is “true of much quantitative research on terrorism” (Asal, Ackerman, and Rethemeyer 2012, 230–31). This limitation is why I conduct a more fine-grained analysis of the FLQ's campaign activities in Chapter 9. Although this narrative overview is intended for illustrative purposes only, it does show how examining event sequencing—and detailed characteristics of attacks and their perpetrators' sociopolitical context—can reveal important clues regarding hoax strategies.¹⁵⁸

Despite these limitations—which are common to the entire field of empirical terrorism studies—this study is a valuable first step for understanding terrorism hoax behaviour. It does not definitively provide an answer to why some terrorist groups use

¹⁵⁸ Findley and Young's (2012) empirical tests of the strategic logic of outbidding for explaining general and suicide terrorism presents an alternative means of using events data to evaluate perpetrators' strategic logic. There, the authors identify two testable implications of the logic (linked to observable indicators) and conduct several clever tests thereof. The authors do not, however, find cross-national support for the link between outbidding and terrorism, which, they argue, “suggests a need to examine the logic of outbidding more closely” (Findley and Young 2012, 719); the effect of outbidding may be context-dependent, and cross-national observational data cannot reveal whether this is the case, or whether the logic is simply unsubstantiated beyond the case of Israel (on which the theory was designed). The authors also suggest that the “tendency towards identifying monocausal arguments about the onset of suicide terrorism [...] might be misguided” (Findley and Young 2012, 719). Statistical inference, as illustrated in this study's Chapter 7, does not handle multiple causality well.

hoaxes, but it does offer a preliminary analysis of claims and hypotheses regarding terrorism hoax use that have never previously been subjected to any rigorous tests. Its substantive findings can be used to iteratively guide future conceptual and empirical research on hoax activity.

10.2 Implications for policy and research

10.2.1 Implications for counterterrorism and national security policy

This study's substantive findings may be applied to counterterrorism applications, with the caveat that they are based on imperfect data and a limited research design. As discussed in Chapter 7, this study is only capable of distinguishing groups who will never hoax from those who sometimes do; it cannot identify which *events* are serious or hoaxed, which would be of greater practical relevance. Although my empirical chapters identify certain markers of groups who are likely to hoax—having a single-issue motivation; being non-hierarchically structured; being an ethno-nationalist splinter organization; and having carried out more than three events—these indicators say nothing of whether a particular incident will be hoaxed. For this reason, the indicators of groups who never hoax are of more value in counterterrorism applications. The activities of religious ethno-nationalists, religious splinters, and groups who are hierarchically structured, or responsible for between one and three serious events, will almost always be serious; their activities should always be treated with extreme caution.

Fortunately, the QCA solutions describing which groups never hoax have greater coverage than the solutions describing which groups sometimes do; whereas hoaxing

groups tend to follow more idiosyncratic pathways to hoaxing, serious-only groups consistently align with the identified sufficient conditions. Although sometimes these indicators will wrongly identify a hoaxing group as a serious one, for counterterrorism applications it is a safer bet to rely on a framework that wrongly assumes hoaxers are serious, than wrongly assuming a serious group is hoaxing.

This study's findings have greater value as an argument in favour of investing in terrorism events data that includes hoaxes. As the significant disparity in events coverage between the CIDB and the cross-national datasets demonstrates, having government funding support and a team of coders intimately familiar with the language and historical context of the country they are evaluating increases the data's validity. Unfortunately, large-scale, cross-national data collection is a multi-million dollar endeavour: "For governments going through budget crises and cost-cutting pressures, data collection can no doubt look like an attractive target" (LaFree, Dugan, and Miller 2015, 24). The impact of valid and reliable data for generating practical insights for national security, however, cannot be overstated (see Ackerman and Pinson 2016, 618). If the law enforcement and security communities are interested in reducing their resource drain attributed to hoax activity, in intervening "left of boom" when hoaxes are used as facilitators of serious activity, or in mitigating unwarranted public fear surrounding hoax events, practitioners would benefit substantially from stronger data on which one could conduct events-based predictive analyses distinguishing hoaxed events from serious ones.

10.2.2 Implications for the study of hoaxes (and terrorism-related phenomena)

10.2.2.1 Large-sample research

Moving the large-sample study of hoaxes forward, then, requires stronger data. While others (see Ackerman and Pinson 2016; LaFree 2010, 43–44; LaFree, Dugan, and Miller 2015, 232–34) present suggestions for improving terrorism events data in general, here I offer specific recommendations for the future study of hoaxes.

First and foremost: data collection efforts must include hoaxes in the analysis! Including hoaxes is required if large-sample hoax-related knowledge is to advance, but it may simultaneously improve the study of serious terrorism and related phenomena. Scholars have already advocated expanding events data beyond completed serious attacks. For instance, LaFree, Dugan, and Miller (2015) suggest that future databases will integrate information on terrorist attacks with “data on perpetrator organizations, failed and foiled plots, contextual data on other types of events, and countermeasures” (234).¹⁵⁹ Hoax data should be included alongside these other valuable dimensions.

Second, the external validity of hoax-related inferences would be augmented in datasets that are not limited to a particular subset of terrorist activity—as with the CBRN-centric WMDDb, transnationally-oriented ITERATE, or Canada-focused CIDB. If hoaxes cannot be directly incorporated into cross-cutting events datasets, then hoax-specific

¹⁵⁹ Recent data collection efforts have succeeded to this end. Crenshaw, Dahl, and Wilson’s database of foiled and failed plots (2017); and Brodeur’s *Équipe de recherche sur le terrorisme et l’antiterrorisme* (ERTA), which includes—in addition to terrorist attacks—details on terrorism support activities, foiled and abandoned plots, and characteristics of surveillance and response (data collection details and scope of coverage summarized in Leman-Langlois and Ouellet 2009) are two such examples.

datasets should be designed to be easily merged with serious events data; otherwise, the data will lack the variation along the dependent variable required for inferential statistics.

Finally, no matter the context, when hoaxes are included in the analysis, they should be coded as a mode of terrorism, not a discrete, alternative tactic (see sections 2.4.1 and 3.3.1). Details of hoaxes' implied weapon type and delivery mechanism should be included, with a secondary variable indicating their hoaxed quality. Doing so would provide extra detail on the hoaxed events themselves, while retaining the capacity to assess the implications of each tactic (whether hoaxed or serious) for responders and the public. Future coding efforts might also consider assigning hoaxes to one of their subtypes: claim of responsibility, warning, or device. Doing so would alleviate concerns regarding over-aggregation—a problem which Crenshaw (2007) elaborates with respect to suicide terrorism:

“All types of suicide attacks are merged together, despite their serving different instrumental purposes [...]. Is their common feature, the death of the perpetrator (more or less willing or intended, depending on the definition), sufficient to produce conceptual coherence? What explains variations within the phenomenon? [...] It is by no means clear that all varieties of suicide attack would have the same origins and outcomes.” (Crenshaw 2007, 162)

In the case of hoaxes, is their hoaxed quality sufficient for considering all hoax incidents together? Future research might explore variations within the phenomenon to probe whether all types of hoaxes have common origins and outcomes. In the present study, hoax subtypes did not vary consistently with respect to perpetrator types (Chapter 8), but patterns in their use may vary along alternate lines not explored here.

10.2.2.2 Small- and medium-sample research

There are numerous feasible approaches to hoax research beyond the large sample. One approach is that favoured by Mahoney (2017, 22–23) in his research on hoax warnings. Mahoney recommends constructing small datasets from the bottom-up, including events data only for specified individual campaigns. He advocates creating campaign-specific databases of threats, using LexisNexis and other searches of open source news reporting during the narrow time periods associated with the campaign in question. He then recommends checking these data against events listed in the GTD and ITERATE to capture whether the group followed through on its threat, or whether it was indeed a “bluff.” He demonstrates the feasibility of this approach in his analysis of Boko Haram’s campaign against Nigerian educational institutions, beginning 2009. Although this approach can generate reliable data for hypothesis formation and testing, Mahoney acknowledges that it does not offer any external validity. Still, in the absence of substantial resources to develop a dataset encompassing the universe of hoax cases, focusing at the campaign-level is a good place to build more nuanced knowledge of groups’ event sequencing—as evidenced by my illustrative treatment of the FLQ in Chapter 9.

While data collected via Mahoney’s method may be suitable to quantitative analysis (if a campaign is sufficiently prolific), it is likely more appropriate as a foundation for qualitative case analysis and process tracing. Case studies will be particularly valuable in furthering hoax research, since they overcome the limitations posed by “the aggregative nature of quantitative analysis” (Asal, Ackerman, and Rethemeyer 2012, 249). In-depth studies of particular campaigns can shed more light on perpetrators’ motives

and operations than the group-based analyses conducted in this study. It is only via such in-depth qualitative analyses that research on terrorism hoaxes will side-step the “interpretive impasse” that has been documented with respect to WMD terrorism research (Ackerman 2005, 140).

Such case-based analyses could be meaningfully supplemented with interview data gleaned from past hoax perpetrators. Terrorists, however, frequently do not want to speak with researchers; and even when they do, their clandestine nature usually presents logistical challenges in conducting interviews (LaFree, Dugan, and Miller 2015, 14). Furthermore, although interview data can in theory speak more directly to perpetrators’ strategic logic, they are still victim to threats to inference. Internal validity of interview data is threatened by terrorists unreliable and biased accounts (LaFree, Dugan, and Miller 2015, 14; 226). External validity is threatened by interview sampling approaches. In terrorism studies, most interviews are carried out via opportunity sampling: “the interviews are carried out with conveniently available groups or individuals and little effort is made to sample systematically” (Silke 2001, 8). Even then, the interviews usually only follow “a semi-structured approach at best” (Ibid., 7), which further limits generalizability.

No matter what the source of data, empirical and theoretical progress in the study of hoaxes requires multiple methods and alternate models of inference.¹⁶⁰ Boolean

¹⁶⁰ In his article on measurement issues related to terrorism, Young (2016) recommends that progress in the field be based on “different statistical models,” as well as “different models of statistical inference” (12). I take this argument a step further, since I believe hoax research will be best continued using a mixed-methods approach, not limited to quantitative techniques.

approaches, and other techniques that account for complex causality, appear to have particular value. A comparison of my statistical and QCA results in Chapters 6 and 7, respectively, reveals the limitations of conventional statistics for evaluating such complex phenomena as terrorist groups' use of hoaxes: net effects of single variables—even statistically robust ones—may be misleading. In particular, my probit results indicate strong, symmetric correlational relationships with respect to religious and ethno-nationalist motivation that do not hold up to scrutiny of their underlying set relations. Furthermore, QCA captures combinations of variables—namely, religious ethno-nationalism as a robust combinatorial predictor of non-hoaxing—that are entirely overlooked by the regression. Methods that assume a single causal pathway leading to hoaxing will always be deficient, since hoaxers' strategic logics appear to be both over-determined and equifinal.

10.3 Conceptual extensions

An important “next step” for research on terrorism hoaxes is to shift the object of analysis from the terrorist organization to individual perpetrators of hoax events. Such a focus would better substantiate any links between hoax behaviour and theories of radicalization to violence. It would also contribute to understanding how and why hoax incidents attributed to groups result from principal-agent dynamics—a key element of this study's hypothesis linking non-hierarchical group structure to hoax activity.

Court documents and commentary regarding prosecuted hoax perpetrators offer a feasible means of conducting preliminary exploratory and descriptive work in this

direction. More advanced explanatory models might feasibly contextualize individual hoax perpetrators relative to unaffiliated perpetrators of serious terrorism attacks, in the same way that the present study compared groups who hoaxed with groups who only ever carried out serious activity. Given that open source materials on individual perpetrators are frequently more credible than those regarding organizational perpetrators (Ackerman and Pinson 2014, 240), this type of study may be less burdened by the limitations of observational secondary data than the present study.

Alternatively, it may be valuable to explore terrorism hoaxers relative to hoaxers in non-terrorist domains. Laura Pellegrini's (2008) doctoral research, for instance, examines criminal hoaxes that are explicitly *non*-terrorist in nature. She identifies four categories of criminal hoax—monetary, attention-seeking, racial, and hate crime—which indicate the range of possible non-terrorism hoax activity.¹⁶¹ Since not all hoax legislation is limited to *terrorism* hoaxes, a study of prosecuted hoax perpetrators might reveal valuable insights in the differences between perpetrators who hoax for financial gain or otherwise personal purposes, versus those who hoax for terroristic aims.¹⁶²

¹⁶¹ Pellegrini (2008) refers to these four categories as a "typology" of criminal hoaxes. As a typology, however, these categories are deficient: she describes racial hoaxes as a subset of hate crime hoaxes, and uses both monetary gain and attention-seeking as motives to explain hoaxes that are racially- or otherwise hate-driven. Still, they are valuable in illustrating the range of possible non-terrorism hoax activity.

¹⁶² For instance, hoax legislation in the United States and Australia is not limited to terrorist motivations, whereas in Canada and the United Kingdom it is. In Canada, hoaxes are addressed in a 2004 amendment to the *Criminal Code*. Article 83.231, "Hoax Regarding Terrorist Activity" (c. 14, s. 32), describes a hoax offender as one who a) conveys—or causes to be conveyed—information about current or future terrorist activity that he or she knows to be untrue; or b) commits an act that is likely to be viewed as terroristic, without believing that this activity is occurring or will occur. In both cases, the terrorist hoaxer's intent must be "to cause any person to fear death, bodily harm, substantial damage to property or serious interference with the lawful use or operation of property."

It may similarly be valuable to explore terrorism hoaxes perpetrated by actors other than the supposed perpetrators themselves: both purported victims and state actors may create fake terrorism attacks.

Pellegrini's research introduces the term "viper" to describe "a person who claims to be the victim of a crime, when in reality he/she is the actual perpetrator of the crime/event" (Pellegrini 2008, 5). She describes how vipers—Victim Perpetrators—are frequently members of a protected racial, religious, or ethnic class, who stage (or report) non-existent hate crimes to gain attention and manufacture victim sympathy. As a recent example, the wave of more than one hundred bomb threats against Jewish community centres worldwide in February and March of this year was ultimately attributed to an American-Israeli youth (Bob 2017). Whether or not it was the perpetrator's intent, his bomb hoax warnings led to a perception of rising anti-Semitism in the United States that was so strong, it forced a formal statement by President Donald Trump (Levenson and Stapleton 2017).¹⁶³ There is a fine line between viper-perpetrated hate crime hoaxes and possible false flag terrorism hoaxes. Either may be perpetrated by actors who wish to stoke public fear regarding the terrorism threat, and demonize out-groups.

In addition to "victim" perpetrators, recent events show that governments have also been progenitors of fake terrorism. In an essay provocatively entitled "From Andijon to Bowling Green: Fabricated Terrorism in Uzbekistan and the United States," Sarah

¹⁶³ A list of fake hate crimes and the news reports that document them is compiled at www.fakehatecrimes.org (Wilcox 2017). Although the website exhibits little academic merit or rigour (its manager is a self-styled academic who self-published a book on fake hate crimes in 1994), its links to news articles from established fora are useful in illustrating the prevalence of such hoax incidents.

Kendzior (2017) argues that state governments (usually autocratic ones) have invented terrorist events and perpetrators to justify their policy decisions and to consolidate state power. In the American context, she is referring to Trump advisor Kellyanne Conway's now-infamous February 2017 claim that two Iraqi citizens had been responsible for a "massacre" in Bowling Green, Kentucky, years prior. Of course, this massacre—and then-President Obama's alleged response thereto of instating a temporary ban on Iraqi refugees—never really happened. Conway fabricated it to justify Trump's own ban on immigration from seven Muslim-majority countries (Kendzior 2017, 10).

Whether created by victims, governments, or the terrorists themselves, "fake terrorism" is very much *du jour*. Hoaxes carried out by terrorist actors, however, are a longstanding phenomenon, with especially pernicious effects—particularly when sequenced in broader campaigns along with serious terrorism attacks. The present study has laid the theoretical and empirical groundwork for future research into terrorism hoaxes, in all its forms. It is my hope that future studies will extend and revise the findings described herein as new and better data become available.

When I first began writing this dissertation, I created a Google e-mail alert for the word "hoax." Apart from the odd celebrity death hoax à la "Paul is dead," this basic set-up generated a stream of fairly relevant results for my purposes. Over the past year and a half, however, my alerts have been increasingly filled with news reports on topics ranging

from climate change to “fake news,” and everything in between.¹⁶⁴ This shift aligns with the spike in popular usage of Oxford Dictionaries’ (2016) word of the year: “post-truth,” defined as an adjective “relating to or denoting circumstances in which object facts are less influential in shaping public opinion than appeals to emotion and personal belief.”¹⁶⁵ Hoaxes—as “fake terrorism” or “post-truth terrorism”—appear to be caught up in the zeitgeist of it all, although whether their seeming increase in prevalence (which post-dates any systematic data collection effort) is a function of a true surge in hoax frequency or the term’s lexicographic popularity is unclear.¹⁶⁶

What *is* clear, however, is that in an era of so-called “post-truth” politics, terrorist perpetrators and those seeking to emulate them may learn the lesson that hoaxed terrorism pays. Although it would be a societal “good” for serious terrorists to substitute into hoaxed activity, the more likely outcome of this lesson is that a great deal of individuals who would never otherwise have engaged in serious terrorism will now conduct hoax acts. The result will be a drain on security, law enforcement, and emergency response services, and a substantial increase in the “noise” that these agencies must filter through to uncover serious security threats. Terrorism hoaxes may be “fakes,” but their indirect costs and consequences are very real.

¹⁶⁴ According to Merriam-Webster (2017), “*Fake news* is frequently used to describe a political story which is seen as damaging to an agency, entity, or person.” The dictionary agrees that “*Fake news* is having a bit of a moment.”

¹⁶⁵ In its justification for selecting “post-truth” as the word of the year, Oxford Dictionaries (2016) describes how the term spiked in 2016 in the context of the United Kingdom Brexit referendum and the United States presidential election.

¹⁶⁶ CIDB data appears to substantiate—at least in the Canadian context—a recent uptick in the use of hoaxes; in the 2014-2015 period, hoaxes in Canada were more prevalent than all forms of serious terrorism combined (Figure 9.1; Tishler, Ouellet, and Kilberg, n.d.).

Appendices

Appendix A Supplementary materials for Chapter 6

Appendix Table A.1 Organizational characteristics of terrorist groups in the sample¹⁶⁷

Group Name	hoaxer	single-issue	purely religious	contains religion	purely ethno-nationalist	contains ethno-nationalism	religious ethno-nationalist	economically leftist	broadly left	economically rightist	broadly right	leadership	functional differentiation	command and control	hierarchy	splinter	number of serious events
2 nd of June Movement	0	0	0	0	0	0	0	0	1	0	0	1	0	0	1	0	14
Abu Hafs al-Masri Briagde	1	0	1	1	0	0	0	0	0	0	0	0	0	0	0	0	13
Abu Sayyaf Group (ASG) (2007 and earlier)	0	0	0	1	0	1	1	0	0	0	0	1	0	0	1	1	240
Abu Sayyaf Group (ASG) (2008 and onward)	0	0	0	1	0	1	1	0	0	0	0	0	0	0	0	1	240
Aden-Abyan Islamic Army (AAIA)	1	0	1	1	0	0	0	0	0	0	0	1	0	0	1	0	3
Afrikaner Weerstandsbeweging (AWB)	0	0	0	0	0	1	0	0	0	0	1	1	1	1	1	0	7
al-Qsa Martyrs Brigades	1	0	0	0	1	1	0	0	0	0	0	1	0	0	1	1	178
al-Qaeda (after 9/11)	1	0	1	1	0	0	0	0	0	0	0	1	0	0	1	0	105
al-Qaeda (before and including 9/11)	0	0	1	1	0	0	0	0	0	0	0	1	1	0	1	0	27
al-Qaeda in Iraq (AQI)	0	0	1	1	0	0	0	0	0	0	0	1	1	0	1	0	559
al-Qaeda in the Arabian Peninsula (AQAP)	0	0	1	1	0	0	0	0	0	0	0	1	1	1	1	0	338
al-Qaeda in the Islamic Maghreb (AQIM)	1	0	1	1	0	0	0	0	0	0	0	1	0	0	1	0	419
al-Tawhid wa al-Jihad	0	0	0	1	0	0	1	0	0	0	0	1	0	0	1	0	54
Amanat	0	0	0	1	0	1	1	0	0	0	0	1	0	0	1	0	2
Animal Liberation Front (ALF)	1	1	0	0	0	0	0	0	1	0	0	0	0	0	0	0	121
Animal Rights Militia (ARM)	1	1	0	0	0	0	0	0	1	0	0	0	0	0	0	1	4
Ansar al-Isam	0	0	1	1	0	0	0	0	0	0	0	1	1	0	1	1	32
Arab Revolutionary Army	1	0	0	0	1	1	0	0	0	0	0	10
Army of God	1	1	0	1	0	0	0	0	0	0	1	0	0	0	0	0	28
Aryan Nations	0	0	0	0	0	0	0	0	0	0	1	1	1	0	1	0	8
Aryan Republican Army (ARA)	1	0	0	0	0	0	0	0	0	0	1	0	0	0	0	0	16

¹⁶⁷ Variables are defined and described in Table 7.1.

Group Name	hoaxer	single-issue	purely religious	contains religion	purely ethno-nationalist	contains ethno-nationalism	religious ethno-nationalist	economically leftist	broadly left	economically rightist	broadly right	leadership	functional differentiation	command and control	hierarchy	splinter	number of serious events
Aum Shinrikyo (1995 and earlier)	1	0	1	1	0	0	0	0	0	0	0	1	1	1	1	0	33
Aum Shinrikyo (1996 and later)	0	0	1	1	0	0	0	0	0	0	0	1	0	0	1	0	3
Babbar Khalsa International (BKI)	0	0	0	1	0	1	1	0	0	0	0	1	0	0	1	0	12
Basque Fatherland and Freedom (ETA) (2002 and earlier)	1	0	0	0	0	1	0	1	1	0	0	1	1	0	1	0	1962
Basque Fatherland and Freedom (ETA) (2003 and onward)	0	0	0	0	0	1	0	1	1	0	0	1	0	0	1	0	130
Black December	0	0	0	0	1	1	0	0	0	0	0	2
Commandos for the Defense of Western Civilization (CODECO)	0	0	0	0	0	0	0	0	0	1	1	0	1
Devrimci Halk Kurtulus Partisi/Cephesi (DHKP/C)	0	0	0	0	0	0	0	1	1	0	0	0	1	1	1	0	286
Direct Action (Squamish Five)	1	0	0	0	0	0	0	1	1	0	0	0	0	0	0	0	1
Earth Liberation Front (ELF)	1	1	0	0	0	0	0	0	1	0	0	0	0	0	0	1	84
East Turkistan Liberation Organization (ETLO)	0	0	0	1	0	1	1	0	0	0	0	1	0	0	1	0	2
German Taliban Mujahideen	0	0	1	1	0	0	0	0	0	0	0	1	0	0	1	0	1
Groupe Islamique Armé (GIA)	1	0	1	1	0	0	0	0	0	0	0	1	0	0	1	0	252
Hamas	0	0	0	1	0	1	1	0	0	0	0	1	1	0	1	0	356
Harakat-ul-Mujahideen (HuM)	0	0	0	1	0	1	1	0	0	0	0	1	1	0	1	0	23
Hezbollah	1	0	0	1	0	1	1	0	0	0	0	1	1	1	1	0	412
Hizbul Mujahideen (HM)	0	0	0	1	0	1	1	0	0	0	0	1	1	1	1	0	107
Irish Republican Army (IRA)¹⁶⁸	1	0	0	0	1	1	0	0	0	0	0	1	1	1	1	0	2185
Jaish al-Ta'ifa al-Mansura	1	0	1	1	0	0	0	0	0	0	0	1	0	0	1	0	2
Jaish-e-Muhammad (JeM)	0	0	0	1	0	1	1	0	0	0	0	1	1	0	1	0	50
Jamaat-ul-Mujahideen Bangladesh (JMB)	0	0	1	1	0	0	0	0	0	0	0	1	1	1	1	0	21
Japanese Red Army (JRA)	0	0	0	0	0	0	0	1	1	0	0	1	1	1	1	1	23

¹⁶⁸ As discussed in footnote 153, it may be justifiable to consider the IRA and RIRA as having an economically (and broadly) leftist orientation, and the IRA as having originated as a splinter organization. When these alternate codings, however, are taken into account, the regression results do not substantively change. All variables that were significant at the $p < 0.05$ level in Table 6.3 continue to be significant at the same threshold. In terms of marginal effects, all but the splinter variable increase in magnitude when the alternate codings are tested. Furthermore, while there are minor reductions in the models' pseudo R-squareds, the percent of correctly predicted values remain constant. The only substantive difference is that the broadly left indicator in Model 2A loses all traces of significance; however, this reduction in significance is consistent with the hypothesized relationships, and bolsters the interpretation of findings presented in section 6.4.1.

Group Name	hoaxer	single-issue	purely religious	contains religion	purely ethno-nationalist	contains ethno-nationalism	religious ethno-nationalist	economically leftist	broadly left	economically rightist	broadly right	leadership	functional differentiation	command and control	hierarchy	splinter	number of serious events
Jemaah Islamiyah (JI)	0	0	1	1	0	0	0	0	0	0	0	1	1	1	1	1	81
Justice Guerrillas	0	0	0	0	0	0	0	0	0	0	0	1
Kach	0	0	0	1	0	1	1	0	0	0	0	1	0	0	1	0	14
Khalistan Commando Force (KCF)	0	0	0	1	0	1	1	0	0	0	0	0	1	0	0	0	21
Khmer Rouge (1996 and onward)	0	0	0	0	0	0	0	1	1	0	0	1	0	0	1	0	44
Khmer Rouge (pre-1996)	0	0	0	0	0	0	0	1	1	0	0	1	1	1	1	0	183
Kurdistan Workers Party (PKK)	0	0	0	0	0	1	0	1	1	0	0	1	1	1	1	0	1453
Lashkar-e-Jhangvi (LeJ)	0	0	1	1	0	0	0	0	0	0	0	1	0	0	1	0	72
Lashkar-e-Taiba (LeT)	0	0	0	1	0	1	1	0	0	0	0	1	1	0	1	0	132
Liberation Tigers of Tamil Eelam (LTTE)	1	0	0	0	1	1	0	0	0	0	0	1	1	1	1	0	1635
Minutemen	0	0	0	0	0	0	0	0	0	1	1	1	0	0	1	0	3
Movement for the Restoration of the Ten Commandments of God (MRTCg)	0	0	1	1	0	0	0	0	0	0	0	1	1	1	1	1	1
Mujahidin e-Khalq (MeK)	0	0	0	0	0	1	0	1	1	0	0	1	1	1	1	1	123
National Liberation Army (ELN)	0	0	0	0	0	0	0	1	1	0	0	1	1	0	1	0	1388
O'odua People's Congress (OPC)	0	0	0	0	1	1	0	0	0	0	0	1	1	1	1	1	3
Oromo Liberation Front (OLF)	0	0	0	0	1	1	0	0	0	0	0	1	1	1	1	0	13
Patriots Council	0	0	0	0	0	0	0	0	0	1	1	0	0	0	0	0	1
Popular Front for the Liberation of Palestine (PFLP)	0	0	0	0	0	1	0	1	1	0	0	1	1	0	1	0	223
Popular Liberation Army (EPL)	0	0	0	0	0	0	0	1	1	0	0	1	0	1	1	0	208
Rajneeshees	0	0	1	1	0	0	0	0	0	0	0	1	1	1	1	0	6
Real Irish Republican Army (RIRA) ¹⁶⁹	0	0	0	0	1	1	0	0	0	0	0	1	0	0	1	1	265
Realm of Chaos	0	0	0	0	0	0	0	0	1	0	0	1	0	0	1	0	1
Red Army Faction / Baader-Meinhof Group (RAF)	1	0	0	0	0	0	0	1	1	0	0	1	0	0	1	0	145
Republic of Texas (RoT)	0	0	0	0	0	1	0	0	0	1	1	1	0	0	1	1	2
Revolutionary Armed Forces of Colombia (FARC)	0	0	0	0	0	0	0	1	1	0	0	1	1	1	1	0	696
RISE (Reconstruction [unknown] Society Extermination)	0	0	0	0	0	0	0	0	0	0	0	1	0	0	1	0	1
Salafia Jihadia	0	0	1	1	0	0	0	0	0	0	0	1	0	0	1	0	5
Salah al-Din al-Ayoubi Brigades (Islamic Iraqi Resistance Front)	0	0	0	1	0	1	1	0	0	0	0	0	1	0	0	0	1

¹⁶⁹ See footnote 168.

Group Name	<i>hoaxer</i>	<i>single-issue</i>	<i>purely religious</i>	<i>contains religion</i>	<i>purely ethno-nationalist</i>	<i>contains ethno-nationalism</i>	<i>religious ethno-nationalist</i>	<i>economically leftist</i>	<i>broadly left</i>	<i>economically rightist</i>	<i>broadly right</i>	<i>leadership</i>	<i>functional differentiation</i>	<i>command and control</i>	<i>hierarchy</i>	<i>splinter</i>	<i>number of serious events</i>
Scottish National Liberation Army (SNLA)	1	0	0	0	1	1	0	0	0	0	0	0	0	0	0	0	11
September 11	0	0	0	1	0	0	0	0	0	0	0	5
Students Islamic Movement of India (SIMI) (2008 and earlier)	0	0	1	1	0	0	0	0	0	0	0	1	0	0	1	1	14
Students Islamic Movement of India (SIMI) (after 2008; Indian Mujahideen)	0	0	1	1	0	0	0	0	0	0	0	0	0	0	0	1	3
Symbionese Liberation Army (SLA)	0	0	0	0	0	1	0	1	1	0	0	1	0	0	1	1	7
Taliban	1	0	1	1	0	0	0	0	0	0	0	1	1	0	1	0	2492
The Covenant, the Sword, and the Arm of the Lord (CSA)	0	0	0	1	0	0	0	0	0	0	1	1	1	1	1	0	5
The Justice Department (JD)	1	1	0	0	0	0	0	0	1	0	0	0	0	0	0	1	16
Union for the Total Independence of Angola (UNITA)	0	0	0	0	0	1	0	0	0	1	1	1	1	1	1	1	450
Weather Underground	0	0	0	0	0	0	0	1	1	0	0	1	0	0	1	1	46

Appendix Table A.2 Average marginal effects for probit models testing alternate indicators of organizational structure

	Model 1 [<i>hierarchy</i>]	Model A1.1 [<i>leadership</i>]	Model A1.2 [<i>functional differentiation</i>]	Model A1.3 [<i>centralized command and control</i>]
<i>purely ethno-nationalist</i>	0.426*** (1.334)	0.428*** (0.131)	0.408*** (0.132)	0.474*** (0.135)
<i>purely religious</i>	0.263*** (0.074)	0.263*** (0.076)	0.198*** (0.076)	0.217*** (0.079)
<i>hierarchy</i>	-0.344*** (0.112)	--	--	--
<i>leadership</i>	--	-0.277*** (0.106)	--	--
<i>functional differentiation</i>	--	--	-0.261*** (0.088)	--
<i>command and control</i>	--	--	--	-0.177* (0.100)
<i>splinter</i>	-0.350*** (0.134)	-0.341** (0.136)	-0.435** (0.179)	-0.405** (0.175)
<i>number of serious events, logged</i>	0.050*** (0.018)	0.044** (0.017)	0.055*** (0.019)	0.036** (0.018)
Observations (<i>n</i>)	71	71	71	71
Pseudo R-squared	0.3530	0.3262	0.3527	0.2940
Percent correctly predicted values	83.10	81.69	85.92	80.28

Note: Coefficients and standard errors (reported in parentheses) have been transformed to reflect average marginal effects. The constant is not shown. The dependent variable is the dummy *hoaxer*, indicating whether (=1) or not (=0) a group ever uses hoaxes.

***, **, * indicate statistical significance at the 0.01, 0.05, and 0.10 levels, respectively, for a one-tailed test.

Appendix Table A.3 Cross-tabulation of single-issue motivation as a sufficient condition for hoaxing

		<i>single-issue</i>		<i>Total</i>
		<i>0</i>	<i>1</i>	
<i>hoaxer</i>	<i>0</i>	58	0	58
	<i>1</i>	18	5	23
<i>Total</i>		76	5	81

Appendix Table A.4 Cross-tabulation of single-issue motivation as a subset of non-hierarchy

		<i>single-issue</i>		<i>Total</i>
		<i>0</i>	<i>1</i>	
<i>non-hierarchy</i>	<i>0</i>	62	0	62
	<i>1</i>	9	5	14
<i>Total</i>		71	5	76

Appendix Table B.1 List of groups in hoaxer sample and associated ID labels ¹⁷⁰

AAIA*	Aden-Abyan Islamic Army	FLQ	Front de libération du Québec
AAMB*	al-Aqsa Martyrs Brigades	FPMR	Manuel Rodriguez Patriotic Front
AHMB*	Abu Hafis al-Masri Brigade	GIA*	Groupe Islamique Armé
ALF*	Animal Liberation Front	HEZ*	Hezbollah
AOG*	Army of God	IB	Islamic Brotherhood
AQ*	al-Qaeda (after 9/11)	IRA*	Irish Republican Army
AQIM*	al-Qaeda in the Islamic Maghreb	JD*	The Justice Department
ARM*	Animal Rights Militia	JTM*	Jaish al-Ta'ifa al-Mansura
ARA*	Aryan Republican Army	LTTE*	Liberation Tigers of Tamil Eelam
AS*	Aum Shinrikyo (1995 and earlier)	MRTA	Tupac Amaru Revolutionary Movement
ASALA	Armenian Secret Army for the Liberation of Armenia	N17	Revolutionary Organization 17 November
ASC	Arsonists for Social Cohesion	NPA	New Peoples Army
AZO	Al-Zulfikar	OPR	Organization of the Popular Revolution (OPR-33)
BP	Black Panthers	PLO	Palestine Liberation Organization
BSO	Black September	RAF*	Red Army Faction / Baader-Meinhof Group
CFF	Croatian Freedom Fighters	RAM	Reform of the Armed Forces Movement
DA*	Direct Action (Squamish Five)	RZ	Revolutionary Cells
ELF*	Earth Liberation Front	SNLA*	Scottish National Liberation Army
ETA*	Basque Fatherland and Freedom (2002 and earlier)	TAL*	Taliban
FALN	Armed Front for National Liberation	UPCA	Union of Peaceful Citizens of Algeria

¹⁷⁰ While ID labels frequently correspond to the acronyms commonly applied to each group, in cases where acronyms are not in common usage they were assigned for ease of reference.

* indicates groups that were included in the sample for the probit and QCA analyses of Chapters 6 and 7.

Appendix Table B.2 Basic characteristics of serious terrorist groups who hoax¹⁷¹

ID Label	Hierarchy	Leadership	Functional differentiation	Command and control	Splinter	Militant wing	Contains religion	Purely religious	Islam	Contains ethno-nationalism	Purely ethno-nationalist	Single-issue	Economically leftist	Broadly left	Economically rightist	Broadly right	Number of serious events perpetrated	Transnationally-oriented	CBRN-interested	CBRN-capable	CBRN hoax	ITERATE hoax	GTD hoax
AAIA*	1	1	0	0	0	0	1	1	1	0	0	0	0	0	0	0	3	0	1	0	1	0	0
AAMB*	1	1	0	0	1	0	0	0	0	1	1	0	0	0	0	0	178	1	1	1	1	0	0
AHMB*	0	0	0	0	0	0	1	1	1	0	0	0	0	0	0	0	13	1	1	1	0	0	0
ALF*	0	0	0	0	0	0	0	0	0	0	0	1	0	1	0	0	121	0	1	1	1	0	1
AOG*	0	0	0	0	0	0	1	0	0	0	0	1	0	0	0	1	28	1	1	1	1	0	0
AQ*	1	1	0	0	0	0	1	1	1	0	0	0	0	0	0	0	105	1	1	1	1	0	0
AQIM*	1	1	0	0	0	0	1	1	1	0	0	0	0	0	0	0	419	1	1	1	0	0	1
ARM*	0	0	0	0	1	0	0	0	0	0	0	1	0	1	0	0	4	0	1	1	1	0	0
ARA*	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	16	0	1	0	0	0	1
AS*	1	1	1	1	0	0	1	1	0	0	0	0	0	0	0	0	33	0	1	1	1	0	1
ASALA	1	1	1	1	0	0	0	0	0	1	0	0	1	1	0	0	197	1	0	0	0	1	0
ASC	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	2	0	0	0	0	0	1
AZO	1	1	0	0	1	1	0	0	0	1	0	0	1	1	0	0	10	1	0	0	0	0	1
BP	1	1	0	0	0	0	0	0	0	1	0	0	1	1	0	0	24	1	0	0	0	0	1
BSO	1	1	1	1	1	1	0	0	0	1	1	0	0	0	0	0	159	1	0	0	0	1	0
CFF	1	1	0	0	0	0	0	0	0	1	1	0	0	0	0	0	8	1	0	0	0	0	1
DA*	0	0	0	0	0	0	0	0	0	0	0	0	1	1	0	0	1	0	1	0	1	0	0
ELF*	0	0	0	0	1	0	0	0	0	0	0	1	0	1	0	0	84	0	1	0	1	0	0
ETA*	1	1	1	0	0	0	0	0	0	1	0	0	1	1	0	0	1962	1	1	0	0	1	0
FALN	1	1	0	1	0	0	0	0	0	1	0	0	1	1	0	0	136	1	0	0	0	1	0
FLQ	0	0	0	0	1	0	0	0	0	1	0	0	1	1	0	0	2	1	0	0	0	1	0
FPMR	1	1	0	1	1	1	0	0	0	0	0	0	1	1	0	0	849	1	0	0	0	1	0
GIA*	1	1	0	0	0	0	1	1	1	0	0	0	0	0	0	0	252	1	1	1	1	1	0
HEZ*	1	1	1	1	0	1	1	0	1	1	0	0	0	0	0	0	412	1	1	1	1	0	0
IB	0	0	0	0	0	0	0	0	0	0	0	0	1	1	0	0	5	1	0	0	0	1	0
IRA* ¹⁷²	1	1	1	1	0	1	0	0	0	1	1	0	0	0	0	0	2185	1	1	1	0	1	0
JD*	0	0	0	0	1	0	0	0	0	0	0	1	0	1	0	0	16	0	1	1	1	0	0
JTM*	1	1	0	0	0	0	1	1	1	0	0	0	0	0	0	0	2	0	1	0	1	0	0
LTTE*	1	1	1	1	0	0	0	0	0	1	1	0	0	0	0	0	1635	1	1	1	1	1	0
MRTA	1	1	1	0	0	0	0	0	0	0	0	0	1	1	0	0	573	1	0	0	0	1	0

¹⁷¹ These are the basic characteristics used to compute the full range of clustered variables.

* indicates groups that were included in the sample for Chapter 6 and 7's probit and QCA analyses.

¹⁷² As discussed in footnote 153, it may be justifiable to code the IRA as having an economically and broadly leftist orientation, as well as radical splinter origins. However, when these alternate coding decisions are tested, the cluster results remain identical; the IRA continues to be clustered as a Class 3 – Subtype F, ethno-nationalist hoaxer. While the clustering does not change, the alternate coding does provide one example (contrary to the general findings discussed in section 8.4.1.3) of an economically leftist ethno-nationalist group, where structural consistency overrides similarity in economic orientation in terms of generating the classification.

ID Label	Hierarchy	Leadership	Functional differentiation	Command and control	Splinter	Militant wing	Contains religion	Purely religious	Islam	Contains ethno-nationalism	Purely ethno-nationalist	Single-issue	Economically leftist	Broadly left	Economically rightist	Broadly right	Number of serious events perpetrated	Transnationally-oriented	CBRN-interested	CBRN-capable	CBRN hoax	ITERATE hoax	GTD hoax
N17	1	1	0	0	0	0	0	0	0	1	0	0	1	1	0	0	119	1	0	0	0	1	0
NPA	1	1	1	1	1	1	0	0	0	0	0	0	1	1	0	0	1429	1	0	0	0	1	0
OPR	0	0	0	0	1	0	0	0	0	0	0	0	0	0	0	0	2	1	0	0	0	1	0
PLO	1	1	1	1	0	1	0	0	0	1	1	0	0	0	0	0	207	1	0	0	0	1	0
RAF*	1	1	0	0	0	0	0	0	0	0	0	0	1	1	0	0	145	1	1	1	0	1	0
RAM	1	1	1	1	1	0	0	0	0	1	1	0	0	0	0	0	13	1	0	0	0	1	0
RZ	0	0	1	0	0	0	0	0	0	0	0	0	1	1	0	0	66	1	0	0	0	1	0
SNLA*	0	0	0	0	0	0	0	0	0	1	1	0	0	0	0	0	11	1	1	1	1	0	0
TAL*	1	1	1	0	0	0	1	1	1	0	0	0	0	0	0	0	2492	1	1	1	1	0	0
UPCA	0	0	0	0	0	0	0	0	0	1	1	0	0	0	0	0	1	0	0	0	0	0	1

Appendix Table B.3 List of groups in non-hoaxer sample and associated ID labels¹⁷³

2JM	2 nd of June Movement	JMB	Jamaat-ul-Mujahideen Bangladesh
AI	Ansar al-Islam	JRA	Japanese Red Army
AM	Amanat	KACH	Kach
AN	Aryan Nations	KCF	Khalistan Commando Force
AQ1	al-Qaeda (before and including 9/11)	KR1	Khmer Rouge (pre-1996)
AQAP	al-Qaeda in the Arabian Peninsula	KR2	Khmer Rouge (1996 and onward)
AQI	al-Qaeda in Iraq	LEJ	Lashkar-e-Jhangvi
AS2	Aum Shinrikyo (1996 and later)	LET	Lashkar-e-Taiba
ASG1	Abu Sayyaf Group (2007 and earlier)	MEK	Mujahideen e-Khalq
ASG2	Abu Sayyaf Group (2008 and onward)	MIN	Minutemen
AWB	Afrikaner Weerstandsbeweging	MRTCG	Movement for the Restoration of the Ten Commandments of God
BKI	Babbar Khalsa International	OLF	Oromo Liberation Front
CSA	The Covenant, the Sword, and the Arm of the Lord	OPC	O'odua People's Congress
DHKP	Devrimci Halk Kurtulus Partisi/Cephesi	PC	Patriots Council
ELN	National Liberation Army	PFLP	Popular Front for the Liberation of Palestine
EPL	Popular Liberation Army	PKK	Kurdistan Workers Party
ETA2	Basque Fatherland and Freedom (2003 and onward)	RAJ	Rajneeshees
ETLO	East Turkistan Liberation Organization	RC	Realm of Chaos
FARC	Revolutionary Armed Forces of Colombia	RIRA	Real Irish Republican Army
GTM	German Taliban Mujahideen	RISE	Reconstruction [unknown] Society Extermination
HAM	Hamas	ROT	Republic of Texas
HM	Hizbul Mujahideen	SIMI	Students Islamic Movement of India
HUM	Harakat-ul-Mujahideen	SJ	Salafia Jihada
IIRF	Salah al-Din al-Ayoubi Brigades	SLA	Symbionese Liberation Army
IM	Indian Mujahideen	TJ	al-Tawhid wa al-Jihad
JEM	Jaish-e-Muhammad	UNITA	Union for the Total Independence of Angola
JI	Jemaah Islamiyah	WU	Weather Underground

¹⁷³ While ID labels frequently correspond to the acronyms commonly applied to each group, in cases where acronyms are not in common usage they were assigned for ease of reference.

Appendix Table B.4 Basic characteristics of serious terrorist groups who never hoax¹⁷⁴

ID Label	Hierarchy	Leadership	Functional differentiation	Command and control	Splinter	Militant wing	Contains religion	Purely religious	Islam	Contains ethno-nationalism	Purely ethno-nationalist	Single issue	Economically leftist	Economically rightist	Number of serious events perpetrated	Transnationally-oriented	CBRN-capable
2JM	1	1	0	0	0	0	0	0	0	0	0	0	0	0	14	1	0
AI	1	1	1	0	1	0	1	1	1	0	0	0	0	0	32	0	1
AM	1	1	0	0	0	0	1	0	1	1	0	0	0	0	2	0	1
AN	1	1	1	0	0	0	0	0	0	0	0	0	0	0	8	0	1
AQ1	1	1	1	0	0	0	1	1	1	0	0	0	0	0	27	1	1
AQAP	1	1	1	1	0	0	1	1	1	0	0	0	0	0	338	1	0
AQI	1	1	1	0	0	0	1	1	1	0	0	0	0	0	559	1	1
AS2	1	1	0	0	0	0	1	1	0	0	0	0	0	0	3	0	1
ASG1	1	1	0	0	1	0	1	0	1	1	0	0	0	0	240	1	0
ASG2	0	0	0	0	1	0	1	0	1	1	0	0	0	0	240	1	1
AWB	1	1	1	1	0	1	0	0	0	1	0	0	0	0	7	0	1
BKI	1	1	0	0	0	0	1	0	0	1	0	0	0	0	12	0	1
CSA	1	1	1	1	0	0	1	0	0	0	0	0	0	0	5	0	1
DHKP	1	0	1	1	0	0	0	0	0	0	0	0	1	0	286	1	1
ELN	1	1	1	0	0	0	0	0	0	0	0	0	1	0	1388	1	1
EPL	1	1	0	1	0	1	0	0	0	0	0	0	1	0	208	1	1
ETA2	1	1	0	0	0	0	0	0	0	1	0	0	1	0	130	1	1
ETLO	1	1	0	0	0	0	1	0	1	1	0	0	0	0	2	0	1
FARC	1	1	1	1	0	0	0	0	0	0	0	0	1	0	696	1	1
GTM	1	1	0	0	0	0	1	1	1	0	0	0	0	0	1	0	1
HAM	1	1	1	0	0	1	1	0	1	1	0	0	0	0	356	1	1
HM	1	1	1	1	0	1	1	0	1	1	0	0	0	0	107	0	0
HUM	1	1	1	0	0	0	1	0	1	1	0	0	0	0	23	1	1
IIRF	0	0	1	0	0	0	1	0	1	1	0	0	0	0	1	0	1
IM	0	0	0	0	1	0	1	1	1	0	0	0	0	0	3	0	0
JEM	1	1	1	0	0	0	1	0	1	1	0	0	0	0	50	1	1
JI	1	1	1	1	1	0	1	1	1	0	0	0	0	0	81	1	0
JMB	1	1	1	1	0	0	1	1	1	0	0	0	0	0	21	0	1

¹⁷⁴ These are the basic characteristics used to compute the full range of clustered variables. The indicators for broadly left and right ideology are not included, as there were no single-issue groups to be re-distributed across the broadly left- or right-wing categories. The CBRN-interested variable is not included, because it by definition applies to all groups in the sample. The three hoax indicators (CBRN, ITERATE, or GTD) are similarly excluded, since they by definition do not apply to any groups in the sample.

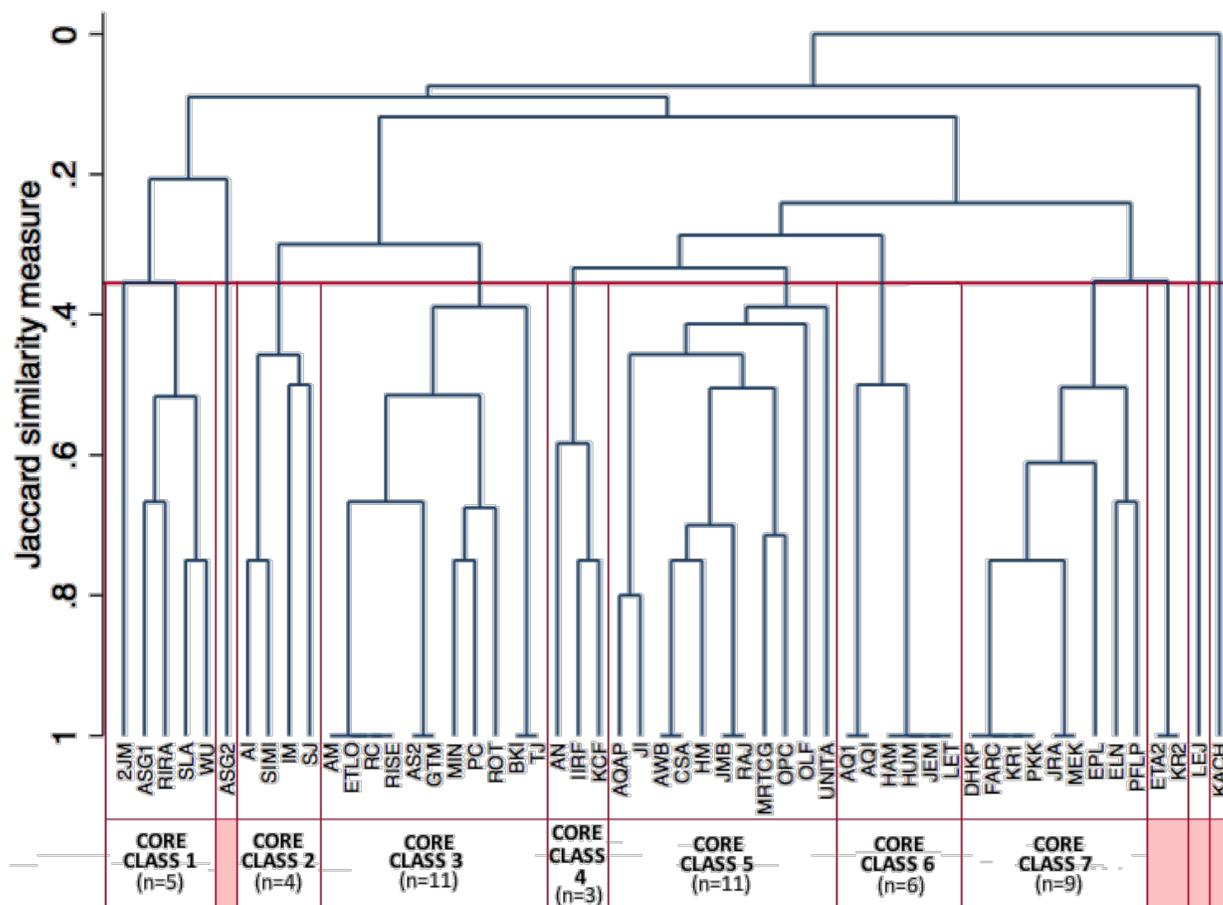
ID Label	Hierarchy	Leadership	Functional differentiation	Command and control	Splinter	Militant wing	Contains religion	Purely religious	Islam	Contains ethno-nationalism	Purely ethno-nationalist	Single issue	Economically leftist	Economically rightist	Number of serious events perpetrated	Transnationally-oriented	CBRN-capable
JRA	1	1	1	1	1	0	0	0	0	0	0	0	1	0	23	1	1
KACH	1	1	0	0	0	0	1	0	0	1	0	0	0	0	14	1	1
KCF	0	0	1	0	0	0	1	0	0	1	0	0	0	0	21	0	1
KR1	1	1	1	1	0	1	0	0	0	0	0	0	1	0	183	1	1
KR2	1	1	0	0	0	1	0	0	0	0	0	0	1	0	44	1	1
LEJ	1	1	0	0	0	0	1	1	1	0	0	0	0	0	72	1	1
LET	1	1	1	0	0	1	1	0	1	1	0	0	0	0	132	1	1
MEK	1	1	1	1	1	0	0	0	0	1	0	0	1	0	123	1	1
MIN	1	1	0	0	0	0	0	0	0	0	0	0	0	1	3	0	1
MRTCG	1	1	1	1	1	0	1	1	0	0	0	0	0	0	1	0	1
OLF	1	1	1	1	0	0	0	0	0	1	1	0	0	0	13	1	1
OPC	1	1	1	1	1	1	0	0	0	1	1	0	0	0	3	0	1
PC	0	0	0	0	0	0	0	0	0	0	0	0	0	1	1	0	1
PFLP	1	1	1	0	0	1	0	0	0	1	0	0	1	0	223	1	0
PKK	1	1	1	1	0	0	0	0	0	1	0	0	1	0	1453	1	1
RAJ	1	1	1	1	0	0	1	1	0	0	0	0	0	0	6	0	1
RC	1	1	0	0	0	0	0	0	0	0	0	0	0	0	1	0	1
RIRA ¹⁷⁵	1	1	0	0	1	0	0	0	0	1	1	0	0	0	265	1	0
RISE	1	1	0	0	0	0	0	0	0	0	0	0	0	0	1	0	1
ROT	1	1	0	0	1	0	0	0	0	1	0	0	0	1	2	0	1
SIMI	1	1	0	0	1	0	1	1	1	0	0	0	0	0	14	0	1
SJ	1	1	0	0	0	0	1	1	1	0	0	0	0	0	5	0	0
SLA	1	1	0	0	1	0	0	0	0	1	0	0	1	0	7	0	0
TJ	1	1	0	0	0	0	1	0	1	0	0	0	0	0	54	0	1
UNITA	1	1	1	1	1	0	0	0	0	1	0	0	0	1	450	1	1
WU	1	1	0	0	1	0	0	0	0	0	0	0	1	0	46	1	0

¹⁷⁵ As discussed in footnote 153, it may be justifiable to code RIRA as having an economically leftist orientation. RIRA consistently clusters with Core Class 1, however, irrespective of the coding for its economic orientation. Consistent with the description of results presented in section 8.4.2, the clustering of non-hoaxing groups is thus more a function of groups' structural features than their motivational orientation.

Appendix Table B.5 Core cluster results for hoaxing groups, by proportion of terrorist groups within each cluster displaying given characteristics ($\geq 60\%$ highlighted)¹⁷⁶

	CLASS 1: NON-SERIOUS DOMESTIC (n=8)	CLASS 2: SINGLE-ISSUE (n=5)	CLASS 3: ETHNO- NATIONALIST (n=9)	CLASS 4: RELIGIOUS (n=6)	CLASS 5: LEFTIST (n=12)
<i>Single issue</i>	0	100	0	0	0
<i>Purely religious</i>	25	0	0	100	0
<i>Purely ethno-nationalist</i>	12.5	0	88.89	0	0
<i>Economically left</i>	25	0	0	0	100
<i>Economically right</i>	0	0	0	0	0
<i>Functional differentiation</i>	0	0	66.67	33.33	41.67
<i>Command</i>	0	0	66.67	16.67	33.33
<i>No hierarchy</i>	75	100	11.11	16.67	16.67
<i>Splinter</i>	25	60	33.33	0	25
<i>Domestic</i>	75	80	0	16.67	0
<i>CBRN-capable</i>	0	80	55.56	100	8.3
<i>3 or fewer serious</i>	87.5	0	0	0	0

¹⁷⁶ The inverse of certain variables may contribute meaningful information for cluster interpretation. These include, in particular, the measures for domestic orientation (transnational orientation); and CBRN capability (conventional-weapons-only capacity).



Appendix Figure B.1 Dendrogram for cluster analysis based on Jaccard similarity

measure, sample of non-hoaxers¹⁷⁷

¹⁷⁷ The horizontal red line indicates the solution resulting in seven core clusters with more than two group members (accounting for 90.74% of all groups). This number of clusters (eleven total) is strongly preferred by the Duda-Hart stopping rule, and a close second under the Caliński-Harabasz rule.

Appendix Table B.7 Core cluster results for non-hoaxing groups, by proportion of terrorist groups within each cluster displaying given characteristics ($\geq 60\%$ highlighted)

	CORE CLASS 1 (n=5)	CORE CLASS 2 (n=4)	CORE CLASS 3 (n=11)	CORE CLASS 4 (n=3)	CORE CLASS 5 (n=11)	CORE CLASS 6 (n=6)	CORE CLASS 7 (n=9)
<i>Single issue</i>	0	0	0	0	0	0	0
<i>Purely religious</i>	0	100	18.18	0	45.45	33.33	0
<i>Purely ethno-nationalist</i>	20	0	0	0	18.18	0	0
<i>Economically left</i>	40	0	0	0	0	0	100
<i>Economically right</i>	0	0	27.27	0	9.09	0	0
<i>Functional differentiation</i>	0	25	0	100	100	100	88.89
<i>Command</i>	0	0	0	0	100	0	77.78
<i>No hierarchy</i>	0	25	9.09	66.67	0	0	0
<i>Splinter</i>	80	75	9.09	0	36.36	0	22.22
<i>Domestic</i>	20	100	100	100	63.64	0	0
<i>Not CBRN-capable</i>	100	50	0	0	27.27	0	11.11
<i>3 or fewer serious</i>	0	25	81.82	33.33	18.18	0	0

Appendix Table C.1 List of groups in sample of Canadian perpetrators and associated

ID labels ¹⁷⁸

<u>AAIA</u> ^{*°}	Aden-Abyan Islamic Army	<u>GIA</u> ^{*°}	Groupe islamique armé
<u>ALF</u> ^{*°}	Animal Liberation Front	HAM *	Hamas
<u>ALQ</u>	Armée de libération du Québec	<u>IRA</u> ^{*°}	Irish Republican Army
<u>AQ</u> ^{*°}	al-Qaeda (after 9/11)	IRI	Initiative de résistance internationaliste
<u>ARA2</u>	Armenian Revolutionary Army	ISYF	International Sikh Youth Federation
<u>ASALA</u> [°]	Armenian Secret Army for the Liberation of Armenia	JCAG	Justice Commandos for the Armenian Genocide
BAF	La brigade d'autodéfense du français	<u>JD</u> ^{*°}	The Justice Department
BKI *	Babbar Khalsa International	JDL	Jewish Defence League
BP [°]	Black Panthers	JRA *	Japanese Red Army
<u>BSO</u> [°]	Black September	KKK	Ku Klux Klan
CA	Cuban Action	MDATF	Militant Direct Action Task Force
CDI	Les chevaliers de l'indépendance	MEK *	Mujahideen e-Khalq
CHFF	Canadian Hungarian Freedom Fighters	MWS	Mohawk Warrior Society
CNM	Cuban Nationalist Movement (anti-Castro)	NLF	Canadians for the National Liberation Front (of South Vietnam)
CSG	Cuban Secret Government	O7	Omega 7
<u>DA</u> ^{*°}	Direct Action (Squamish Five)	OST	Order of the Solar Temple
<u>ELF</u> ^{*°}	Earth Liberation Front	SOF	Sons of Freedom
FEC	Force étudiante critique	SSNP	Syrian Social Nationalist Party
FFFC	Fighting for Freedom Coalition-Ottawa	WFB	Wimmins' Fire Brigade
FLNC	Cuban National Liberation Front	WG	Western Guard
<u>FLQ</u> [°]	Front de libération du Québec	WU *	Weather Underground

¹⁷⁸ * indicates groups that were included in the sample for Chapter 6's probit and QCA analyses.

[°] indicates groups that were included in the sample for Chapter 8's cluster analysis.

 [underlining] indicates hoaxing groups.

While ID labels frequently correspond to the acronyms commonly applied to each group, in cases where acronyms are not in common usage they were assigned for ease of reference.

Appendix Table C.2 Organizational characteristics of terrorist groups in the sample¹⁷⁹

ID Label	Hoaxer	Hierarchy	Leadership	Functional differentiation	Command	Splinter	Militant wing	Contains religion	Purely religious	Islam	Contains ethno-nationalism	Purely ethno-nationalist	Single issue	Economically leftist	Broadly left	Economically rightist	Broadly right	Number of serious events perpetrated
AAIA	1	1	1	0	0	0	0	1	1	1	0	0	0	0	0	0	0	3
ALF	1	0	0	0	0	0	0	0	0	0	0	0	1	0	1	0	0	123
ALQ	1	1	1	0	0	1	1	0	0	0	1	0	0	1	1	0	0	4
AQ	1	1	1	0	0	0	0	1	1	1	0	0	0	0	0	0	0	105
ARA2	1	0	0	0	0	0	0	0	0	0	1	1	0	0	0	0	0	4
ASALA	1	1	1	1	1	0	0	0	0	0	1	0	0	1	1	0	0	200
BAF	0	1	1	0	1	0	0	0	0	0	1	1	0	0	0	0	0	2
BKI	0	1	1	0	0	0	0	1	0	0	1	0	0	0	0	0	0	15
BP	1	1	1	0	0	0	0	0	0	0	1	0	0	1	1	0	0	25
BSO	1	1	1	1	1	1	1	0	0	0	1	1	0	0	0	0	0	160
CA	0	1	1	0	0	0	0	0	0	0	0	0	0	0	0	1	1	18
CDI	0	1	1	0	0	0	0	0	0	0	1	1	0	0	0	0	0	1
CHFF	0	0	0	0	0	0	1	0	0	0	1	1	0	0	0	0	0	2
CNM	0	1	1	0	0	0	0	0	0	0	1	0	0	0	0	1	1	5
CSG	0	1	1	0	0	0	0	0	0	0	1	0	0	0	0	1	1	16
DA	1	0	0	0	0	0	0	0	0	0	0	0	0	1	1	0	0	4
ELF	1	0	0	0	0	1	0	0	0	0	0	0	1	0	1	0	0	86
FEC	0	0	0	0	0	0	0	0	0	0	0	0	0	1	1	0	0	1
FFFC	0	0	0	0	0	0	0	0	0	0	0	0	0	1	1	0	0	2
FLNC	0	1	1	0	0	0	0	0	0	0	0	0	0	0	0	1	1	28
FLQ	1	0	0	0	0	1	0	0	0	0	1	0	0	1	1	0	0	161
GIA	1	1	1	0	0	0	0	1	1	1	0	0	0	0	0	0	0	252
HAM	0	1	1	1	0	0	1	1	0	1	1	0	0	0	0	0	0	356
IRA ¹⁸⁰	1	1	1	1	1	0	1	0	0	0	1	1	0	0	0	0	0	2186
IRI	0	0	0	0	0	0	0	0	0	0	0	0	0	1	1	0	0	5
ISYF	0	1	1	1	1	0	0	1	0	0	1	0	0	0	0	0	0	2
JCAG	0	0	0	1	0	0	1	0	0	0	1	1	0	0	0	0	0	50
JD	1	0	0	0	0	1	0	0	0	0	0	0	1	0	1	0	0	16
JDL	0	1	1	1	0	0	0	1	1	0	0	0	0	0	0	0	0	106
JRA	0	1	1	1	1	1	0	0	0	0	0	0	0	1	1	0	0	23
KKK	0	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	1	24
MDATF	0	0	0	0	0	0	0	0	0	0	0	0	0	1	1	0	0	3
MEK	0	1	1	1	1	1	0	0	0	0	1	0	0	1	1	0	0	123
MWS	0	0	0	0	0	0	0	0	0	0	1	1	0	0	0	0	0	1
NLF	0	1	1	0	0	0	1	0	0	0	1	0	0	1	1	0	0	1

¹⁷⁹ These are the basic characteristics used to compute the full range of variables.¹⁸⁰ See footnote 153 for a discussion of alternate coding possibilities for the IRA as an economically (and broadly) leftist splinter organization.

ID Label	<i>Hoaxer</i>	<i>Hierarchy</i>	<i>Leadership</i>	<i>Functional differentiation</i>	<i>Command</i>	<i>Splinter</i>	<i>Militant wing</i>	<i>Contains religion</i>	<i>Purely religious</i>	<i>Islam</i>	<i>Contains ethno-nationalism</i>	<i>Purely ethno-nationalist</i>	<i>Single issue</i>	<i>Economically leftist</i>	<i>Broadly left</i>	<i>Economically rightist</i>	<i>Broadly right</i>	<i>Number of serious events perpetrated</i>
O7	0	1	1	0	0	0	0	0	0	0	1	1	0	0	0	0	0	57
OST	0	1	1	0	1	0	0	1	1	0	0	0	0	0	0	0	0	1
SOF	0	1	1	0	0	1	1	1	1	0	0	0	0	0	0	0	0	133
SSNP	0	1	1	0	1	0	1	0	0	0	1	0	0	1	1	0	0	9
WFB	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	0	0	3
WG	0	1	1	0	0	0	0	0	0	0	0	0	0	0	0	1	1	3
WU	0	1	1	0	0	1	0	0	0	0	0	0	0	1	1	0	0	47

Appendix D List of acronyms

9/11	September 11, 2001
ALF	Animal Liberation Front
ALQ	<i>Armée de libération du Québec</i>
ASALA	Armenian Secret Army for the Liberation of Armenia
ASG	Abu Sayyaf Group
BAAD	Big Allied and Dangerous Database
BKI	Babbar Khalsa International
CBRN	Chemical, biological, radiological, and/or nuclear
CIA	Central Intelligence Agency
CIDB	Canadian Incident Database
csQCA	Crisp set Qualitative Comparative Analysis
EIAD	Energy Infrastructure Attack Database
ELF	Earth Liberation Front
ELN	National Liberation Army
ERTA	<i>Équipe de recherche sur le terrorisme et l'antiterrorisme</i>
ETA	Basque Fatherland and Freedom (<i>Euskadi ta Askatasuna</i>)
FBIS	Foreign Broadcast Information Service
FLP	<i>Front de libération populaire</i>
FLQ	<i>Front de libération du Québec</i>
FPMR	Manuel Rodriguez Patriotic Front
fsQCA	Fuzzy set Qualitative Comparative Analysis
GSPC	Salafist Group for Preaching and Fighting (<i>Groupe salafiste pour la prédication et le combat</i>)
GTD	Global Terrorism Database
HAZMAT	Hazardous materials
IRA	Irish Republican Army
ITERATE	International Terrorism: Attributes of Terrorist Events
LTTE	Liberation Tigers of Tamil Eelam
MIPT	Memorial Institute for the Prevention of Terrorism
NATO	North Atlantic Treaty Organization
NCTC	National Counterterrorism Centre
PFLP	Popular Front for the Liberation of Palestine
PGIS	Pinkerton Global Intelligence Services
POICN	Profiles of Incidents Involving CBRN Use by Non-State Actors
QCA	Qualitative Comparative Analysis
RIRA	Real Irish Republican Army
SIMI	Students Islamic Movement of India
SNLA	Scottish National Liberation Army
START	National Consortium for the Study of Terrorism and Responses to Terrorism
TKB	Terrorism Knowledge Base

TOPs	Terrorist Organization Profiles (TOPs)
TSAS	Canadian Network for Research on Terrorism, Security and Society
VDR	Violent Dissident Republican
WITS	Worldwide Incidents Tracking System
WMD	Weapons of Mass Destruction
WMDDb	Monterey WMD Terrorism Database

References

- "A World of Terror." 2017. Periscopic. Accessed July 14. <http://terror.periscopic.com/>.
- Abrahms, M. 2006. "Why Terrorism Does Not Work." *International Security* 31 (2): 42–78.
- . 2008. "What Terrorists Really Want: Terrorist Motives and Counterterrorism Strategy." *International Security* 32 (4): 78–105.
- Abrahms, M., and J. Conrad. 2017. "The Strategic Logic of Credit Claiming: A New Theory for Anonymous Terrorist Attacks." *Security Studies* 26 (2): 279–304.
- Abrahms, M., and P. B. K. Potter. 2015. "Explaining Terrorism: Leadership Deficits and Military Group Tactics." *International Organization* 69 (Spring): 311–42.
- Ackerman, G. 2005. "WMD Terrorism Research: Where to from Here?" *International Studies Review* 7 (1): 140–43.
- Ackerman, G., and L. Pinson. 2014. "An Army of One: Assessing CBRN Pursuit and Use by Lone Wolves and Autonomous Cells." *Terrorism and Political Violence* 26 (1): 226–45.
- . 2016. "Speaking Truth to Sources: Introducing a Method for the Quantitative Evaluation of Open Sources in Event Data." *Studies in Conflict & Terrorism* 39 (7–8). Taylor & Francis: 617–40.
- Adnan, M., and M. Rafi. 2015. "Extracting Patterns from Global Terrorist Dataset (GTD) Using Co-Clustering Approach." *Journal of Independent Studies and Research - Computing* 13 (1): 7–13.

- Aguirre, A. P. 2009. "The Philippine Response to Terrorism." *Journal of Policing, Intelligence and Counter Terrorism* 4 (1): 47–63.
- Anderton, C. H., and J. R. Carter. 2009a. *Principles of Conflict Economics: A Primer for Social Scientists*. New York: Cambridge University Press.
- . 2009b. "Terrorism." In *Principles of Conflict Economics: A Primer for Social Scientists*, 126–54. New York: Cambridge University Press.
- Animal Liberation Front. 2015. "The ALF Credo and Guidelines." Accessed April 9. http://www.animalliberationfront.com/ALFront/alf_credos.htm.
- Arquilla, J., and D. Ronfeldt, eds. 2001. *Networks and Netwars: The Future of Terror, Crime, and Militancy*. Santa Monica: RAND Corporation.
- Asal, V., G. Ackerman, and R. K. Rethemeyer. 2012. "Connections Can Be Toxic: Terrorist Organizational Factors and the Pursuit of CBRN Weapons." *Studies in Conflict & Terrorism* 35 (September 2012): 229–54.
- Asal, V., M. Brown, and A. Dalton. 2012. "Why Split? Organizational Splits among Ethnopolitical Organizations in the Middle East." *Journal of Conflict Resolution* 56 (1): 94–117.
- Asal, V., M. Brown, and M. Schulzke. 2015. "'Kill Them all—Old and Young, Girls and Women and Little Children': An Examination of the Organizational Choice of Targeting Civilians." *Political Science Research and Methods* 3 (3): 589–607.
- Asal, V., and R. K. Rethemeyer. 2008. "The Nature of the Beast: Organizational Structures and the Lethality of Terrorist Attacks." *The Journal of Politics* 70 (2): 437–49.

- . 2009. "Islamist Use and Pursuit of CBRN Terrorism." In *Jihadists and Weapons of Mass Destruction*, 335–58. Boca Raton: CRC Press.
- Asal, V., R. K. Rethemeyer, and I. Anderson. 2011. "Big Allied and Dangerous (BAAD) Database 1 - Lethality Data, 1998-2005." Harvard Dataverse. doi:hdl/1902.1/16062.
- Banlaoi, R. C. 2009. "Media and Terrorism in the Philippines: The Rajah Solaiman Islamic Movement." *Journal of Policing, Intelligence and Counter Terrorism* 4 (1): 64–75.
- Beam, L. 1992. "Leaderless Resistance." *The Seditonist* 12 (February).
- Becker, G. S., and Y. Rubinstein. 2011. "Fear and the Response to Terrorism: An Economic Analysis." 1079. CEP Discussion Paper. London.
<http://cep.lse.ac.uk/pubs/download/dp1079.pdf>.
- Beckman, R. L. 1986. "International Terrorism: The Nuclear Dimension." *Terrorism* 8 (4): 351–78.
- Beres, L. R. 1998. "Israel, the 'Peace Process,' and Nuclear Terrorism: Recognizing the Linkages." *Studies in Conflict & Terrorism* 21 (1): 59–86.
- Berman, E. 2009. *Radical, Religious, and Violent: The Economics of Terrorism*. Cambridge: MIT Press.
- Berman, E., and D. D. Laitin. 2008. "Religion, Terrorism and Public Goods: Testing the Club Model." *Journal of Public Economics* 92: 1942–67.
- Bob, Y. J. 2017. "Will the JCC Fake Bomb Threat Suspect Be Extradited?" *The Jerusalem Post*, March 23. <http://www.jpost.com/Israel-News/Analysis-Will-the-JCC-fake-bomb-threat-suspect-be-extradited-485022>.
- Bremer, J. S. 1983. "Offshore Energy Terrorism: Perspectives on a Problem." *Terrorism*

6 (3): 455–68.

Brown, J. M. 2015. "The Bomber Who Calls Ahead: Terrorism, Insurgency, and the Politics of Pre-Attack Warnings." Columbia University.

Bueno de Mesquita, E. 2008. "Terrorist Factions." *Quarterly Journal of Political Science* 3 (4): 399–418.

Cameron, G. 1999. "Multi-Track Microproliferation: Lessons from Aum Shinrikyo and Al Qaida." *Studies in Conflict & Terrorism* 22 (4): 277–309.

Cameron, G., J. Pate, D. McCauley, and L. DeFazio. 2000. "1999 WMD Terrorism Chronology: Incidents Involving Sub-National Actors and Chemical, Biological, Radiological, and Nuclear Materials." *The Nonproliferation Review* Summer: 157–74.

"Canadian Incident Database [CIDB]." 2016. Canadian Network for Research on Terrorism, Security and Society. www.extremism.ca.

Charters, D. A. 1997. "The Amateur Revolutionaries: A Reassessment of the FLQ." *Terrorism and Political Violence* 9 (1): 133–69.

Chenoweth, E., and E. Lowham. 2007. "On Classifying Terrorism: A Potential Contribution of Cluster Analysis for Academics and Policy-Makers." *Defense & Security Analysis* 23 (4): 345–57.

Cragin, K. R. 2006. "Learning to Survive: The Case of the Islamic Resistance Movement (Hamas)." In *Teaching Terror: Strategic and Tactical Learning in the Terrorist World.*, edited by J. J. F. Forest, 189–204. Lanham: Rowman and Littlefield.

Crelinsten, R. 1987. "The Internal Dynamics of the FLQ during the October Crisis of

- 1970." *Journal of Strategic Studies* 10 (4): 59–89.
- Crenshaw, M. 2007. "Explaining Suicide Terrorism: A Review Essay." *Security Studies* 16 (1): 133–62.
- Crenshaw, M., E. Dahl, and M. Wilson. 2017. "Beneath the Tip of the Iceberg: A New Data Set of Failed and Foiled Terrorist Plots." Paper presented at the 58th Annual Convention of the International Studies Association (ISA) in Baltimore, MD: February 22nd-25th, 2017.
- Crenshaw, M., and G. LaFree. 2017. "The Tip of the Iceberg Accounting for Failed and Foiled Terrorist Plots." In *Countering Terrorism: No Simple Solutions*, 69–98. Washington, DC: Brookings Institution Press.
- Crenshaw Hutchinson, M. 1972. "The Concept of Revolutionary Terrorism." *The Journal of Conflict Resolution (Pre-1986)* 16 (3): 383–96.
- Cronin, A. K. 2003. "Behind the Curve: Globalization and International Terrorism." *International Security* 27 (3): 30–58.
- Dahl, E. 2011. "The Plots That Failed: Intelligence Lessons Learned from Unsuccessful Terrorist Attacks against the United States." *Studies in Conflict & Terrorism* 34 (8): 621–48.
- Dishman, C. 2001. "Understanding Perspectives on WMD and Why They Are Important." *Studies in Conflict & Terrorism* 24 (4): 303–13.
- Dixit, A. K., and B. J. Nalebuff. 1991. *Thinking Strategically: The Competitive Edge in Business, Politics, and Everyday Life*. New York: W. W. Norton & Company.
- Donohue, L. K., and J. N. Kayyem. 2002. "Federalism and the Battle Over

- Counterterrorist Law: State Sovereignty, Criminal Law Enforcement, and National Security." *Studies in Conflict & Terrorism* 25 (1): 1–18.
- Dougherty, T. J., T. F. Green, and T. Harrington. 2001. "Anthrax: A Biologic or Pyschologic Threat?" *Journal of Toxicology - Clinical Toxicology* 39 (5): 526–27.
- Drake, C. J. M. 1998. "The Role of Ideology in Terrorists' Target Selection." *Terrorism and Political Violence* 10 (2): 53–85.
- Dwivedi, S., P. Pandey, M. S. Tiwari, and M. A. Kalam. 2014. "Comparative Study of Clustering Algorithms Used in Counter Terrorism." *IOSR Journal of Computer Engineering* 16 (6): 13–17.
- Editor, Y. A., and J. M. Post. 1990. "The George Washington University Conferences on Terrorism." *Terrorism* 13 (2): 165–68.
- Enders, W., G. F. Parise, and T. Sandler. 1992. "A Time-Series Analysis of Transnational Terrorism: Trends and Cycles." *Defence Economics* 3 (4): 305–20.
- Enders, W., and T. Sandler. 1991. "Causality between Transnational Terrorism and Tourism: The Case of Spain." *Terrorism* 14 (1): 49–58.
- . 1993. "The Effectiveness of Antiterrorism Policies: A Vector-Autoregression-Intervention Analysis." *The American Political Science Review* 87 (4): 829–44.
- . 2012. *The Political Economy of Terrorism*. 2nd ed. Cambridge: Cambridge University Press.
- Enders, W., T. Sandler, and J. Cauley. 1990. "UN Conventions, Technology, and Retaliation in the Fight against Terrorism: An Econometric Evaluation." *Terrorism and Political Violence* 2 (1): 83–105.

- Erickson, C. W., and B. A. Barratt. 2004. "Prudence or Panic? Preparedness Exercises, Counterterror Mobilization, and Media Coverage - Dark Winter, TOPOFF 1 and 2." *Journal of Homeland Security and Emergency Management* 1 (4).
- Everitt, B., S. Landau, M. Leese, and D. Stahl. 2011. *Cluster Analysis*. 5th ed. West Sussex: Wiley.
- Findley, M. G., and J. K. Young. 2012. "More Combatant Groups, More Terror?: Empirical Tests of an Outbidding Logic." *Terrorism and Political Violence* 24 (5): 706–21.
- Flemming, P. A., M. Stohl, and A. P. Schmid. 1988. "The Theoretical Utility of Typologies of Terrorism: Lessons and Opportunities." In *The Politics of Terrorism*, edited by M. Stohl, 3rd, revis ed., 153–96. New York and Basel: Marcel Dekker.
- Fodeman, A. D. 2015. "Safety and Danger Valves: Functional Displacement in American Anti-Abortion Terrorism." *Behavioral Sciences of Terrorism and Political Aggression* 7 (3). Taylor & Francis: 169–83.
- Frisk, A. 2017. "Concordia Bomb Threat: Who Is the Group Purportedly behind the Threat?" *Global News*, March 1. <http://globalnews.ca/news/3281430/concordia-bomb-threat-who-is-council-conservative-citizens>.
- Gilbert, G. K. 1884. "Finley's Tornado Predictions." *American Meteorological Journal* 1: 166–72.
- Gill, P., J. Horgan, and P. Deckert. 2014. "Bombing Alone: Tracing the Motivations and Antecedent Behaviors of Lone-Actor Terrorists." *Journal of Forensic Sciences* 59 (2): 425–35.
- Giroux, J., P. Burgherr, and L. Melkunaite. 2013. "Research Note on the Energy

- Infrastructure Attack Database (EIAD).” *Perspectives on Terrorism* 7 (6): 113–25.
- Giuffrida, L. O. 1987. “Selected Papers from: International Terrorism: Threats and Responses.” In *Selected Papers from: International Terrorism: Threats and Responses.*, edited by Y. Alexander and J. Suchlicki, 71–76. *Terrorism* (Volume 10, pp. 51-81).
- Gnanadesikan, R., J. R. Kettenring, and S. L. Tsao. 1995. “Weighting and Selection of Variables.” *Journal of Classification* 12: 113–36.
- Goertz, G., and J. Mahoney. 2012. *A Tale of Two Cultures: Qualitative and Quantitative Research in the Social Sciences*. Princeton and Oxford: Princeton University Press.
- Gower, J. C., and P. Legendre. 1986. “Metric and Euclidean Properties of Dissimilarity Coefficients.” *Journal of Classification* 3 (1): 5–48.
- Gray, P. W. 2013. “Leaderless Resistance, Networked Organization, and Ideological Hegemony.” *Terrorism and Political Violence* 25: 655–71.
- Grendstad, G. 2007. “Causal Complexity and Party Preference.” *European Journal of Political Research* 46 (1): 121–49.
- Grofman, B., and C. Q. Schneider. 2009. “An Introduction to Crisp Set QCA, with a Comparison to Binary Logistic Regression.” *Political Research Quarterly* 62 (4): 662–72.
- Heger, L., D. Jung, and W. H. Wong. 2012. “Organizing for Resistance: How Group Structure Impacts the Character of Violence.” *Terrorism and Political Violence* 24 (5): 743–68.
- Hewitt, C., and J. Kelley-Moore. 2009. “Foreign Fighters in Iraq: A Cross-National

- Analysis of Jihadism." *Terrorism and Political Violence* 21 (2): 211–20.
- Hoffman, B. 2006. *Inside Terrorism*. Revised an. New York: Columbia University Press.
- Horgan, J, and J. F. Morrison. 2011. "Here to Stay? The Rising Threat of Violent Dissident Republicanism in Northern Ireland." *Terrorism and Political Violence* 23: 642–69.
- "If a Tree Falls: A Story of the Earth Liberation Front." 2011. New York: Oscilloscope Laboratories.
- Iglarsh, H. J. 1989. "Inter-institutional Conference on Terrorism Research: A State of the Art." *Terrorism* 12 (1): 64–65.
- ITV Report. 2014. "Teenage Girl Arrested after Twitter Terror 'Hoax' Goes Wrong." *ITV News*, April 14. <http://www.itv.com/news/2014-04-14/teenage-girl-arrested-after-twitter-terror-hoax-goes-wrong/>.
- Ivanova, K., and T. Sandler. 2006. "CBRN Incidents: Political Regimes, Perpetrators, and Targets." *Terrorism and Political Violence* 18 (3): 423–48.
- . 2007. "CBRN Attack Perpetrators: An Empirical Study." *Foreign Policy Analysis* 3: 273–94.
- Jaccard, P. 1901. "Distribution de la Flore Alpine dans le Bassin des Dranses et dans Quelques Régions Voisines." *Bulletin de la Société Vaudoise des Sciences Naturelles* 37: 241–72.
- . 1908. "Nouvelles Recherches sur la Distribution Florale." *Bulletin de la Société Vaudoise des Sciences Naturelles* 44: 223–270.
- Jackson, B. A. 2001. "Technology Acquisition by Terrorist Groups: Threat Assessment Informed by Lessons from Private Sector Technology Adoption." *Studies in Conflict*

- & Terrorism 24 (3): 183–213.
- Jackson, B. A., and D. R. Frelinger. 2008. "Rifling through the Terrorists' Arsenal: Exploring Groups' Weapon Choices and Technology Strategies." *Studies in Conflict & Terrorism* 31 (7): 583–604.
- Jenkin, C. M. 2006. "Risk Perception and Terrorism: Applying the Psychometric Paradigm." *Homeland Security Affairs* 2 (2): 1–14.
- Jenkins, B. M. 1983. "International Cooperation in Locating and Recovering Stolen Nuclear Materials." *Terrorism* 6 (4): 561–75.
- . 1987. "The Future Course of Terrorism." *The Futurist*, no. July/August.
- Jenkins, P. 1988. "Under Two Flags: Provocation and Deception in European Terrorism." *Terrorism* 11 (4): 275–87.
- Jones, S. G., and M. C. Libicki. 2008. *How Terrorist Groups End: Lessons for Countering Al Qa'ida*. Santa Monica: RAND Corporation.
- Joyner, C. C. 1990. "The Rabta Chemical Factory Fire: Rethinking the Lawfulness of Anticipatory Self-Defense." *Terrorism* 13 (2): 79–87.
- Juergensmeyer, M. 1988. "The Logic of Religious Violence." In *Inside Terrorist Organizations*, edited by D. C. Rapoport, 172–93. London: Frank Cass.
- Kahneman, D., and A. Tversky. 1979. "Prospect Theory: An Analysis of Decision under Risk.pdf." *Econometrica* 47 (2): 263–92.
- Kashubsky, M. 2011. "A Chronology of Attacks on and Unlawful Interferences With, Offshore Oil and Gas Installations, 1975 – 2010." *Perspectives on Terrorism* 5 (5–6): 139–67.

- Kaufman, L., and P. J. Rousseeuw. 1990. *Finding Groups in Data: An Introduction to Cluster Analysis*. Hoboken: Wiley.
- Kearns, E. M., B. Conlon, and J. K. Young. 2014. "Lying about Terrorism." *Studies in Conflict & Terrorism* 37 (5): 422–39.
- Kendzior, S. 2017. "From Andijon to Bowling Green: Fabricated Terrorism in Uzbekistan and the United States." *World Policy Journal* 34 (1): 9–12.
- Kilberg, J. 2011. "Organizing for Destruction: How Organizational Structure Affects Terrorist Group Behaviour." Carleton University.
- . 2012. "A Basic Model Explaining Terrorist Group Organizational Structure." *Studies in Conflict & Terrorism* 35: 810–30.
- Koblentz, G. D. 2011. "Predicting Peril or the Peril of Prediction? Assessing the Risk of CBRN Terrorism." *Terrorism and Political Violence* 23 (4): 501–20.
- Kydd, A. H., and B. F. Walter. 2006. "The Strategies of Terrorism." *International Security* 31 (1): 49–79.
- LaFree, G. 2010. "The Global Terrorism Database (GTD): Accomplishments and Challenges." *Perspectives on Terrorism* 4 (1): 24–46.
- LaFree, G., and L. Dugan. 2007. "Introducing the Global Terrorism Database." *Terrorism and Political Violence* 19 (2): 181–204.
- LaFree, G., L. Dugan, and E. Miller. 2015. *Putting Terrorism in Context*. London and New York: Routledge.
- Lee, J. E. C., L. Lemyre, and D. Krewski. 2010. "A Multi-Method, Multi-Hazard Approach to Explore the Uniqueness of Terrorism Risk Perceptions and Worry." *Journal of*

- Applied Social Psychology 40 (1): 241–72.
- Legendre, P., and L. Legendre. 1983. "Partitioning Ordered Variables into Discrete States for Discriminant Analysis of Ecological Classifications." *Canadian Journal of Zoology* 61: 1002–10.
- Leman-Langlois, S., and J. P. Brodeur. 2005. "Terrorism Old and New: Counterterrorism in Canada." *Police Practice & Research* 6 (2): 121–40.
- Leman-Langlois, S., and G. Ouellet. 2009. "L'évolution du Terrorisme au Canada, 1973-2006." In *Terrorisme et Antiterrorisme au Canada*, 58–72. Montreal: Les Presses de l'Université de Montréal.
- Lentini, P., and M. Bakashmar. 2007. "Jihadist Beheading: A Convergence of Technology, Theology, and Teleology?" *Studies in Conflict & Terrorism* 30 (4): 303–25.
- Levenson, E., and A. Stapleton. 2017. "Jewish Center Bomb Threats: 100 since January." CNN, March 13. <http://www.cnn.com/2017/02/28/us/bomb-threats-jewish-centers-jcc/index.html>.
- Leventhal, P., and B. Chellaney. 1988. "Nuclear Terrorism: Threat, Perception, and Response in South Asia." *Terrorism* 11 (6): 447–70.
- Mahoney, C. W. 2017. "Empty Threats: How Organizations Use Bluffing in Terrorist Campaigns." Paper presented at the 58th Annual Convention of the International Studies Association (ISA) in Baltimore, MD: February 22nd-25th, 2017.
- Marsden, S. V. 2012. "Successful Terrorism: Framework and Review." *Behavioral Sciences of Terrorism and Political Aggression* 4 (2): 134–50.
- Marsden, S. V., and A. P. Schmid. 2011. "Typologies of Terrorism and Political Violence."

- In *The Routledge Handbook of Terrorism Research*, edited by A. P. Schmid, 158–200. London and New York: Routledge.
- Matesan, I. E., and R. Berger. 2017. “Blunders and Blame: How Armed Non-State Actors React to Their Mistakes.” *Studies in Conflict & Terrorism* 40 (5). Taylor & Francis: 376–98.
- McAllister, B., and A. P. Schmid. 2011. “Theories of Terrorism.” In *The Routledge Handbook of Terrorism Research*, edited by A. P. Schmid, 201–71. London and New York: Routledge.
- McCauley, C., and S. Moskalenko. 2014. “Toward a Profile of Lone Wolf Terrorists: What Moves an Individual from Radical Opinion to Radical Action.” *Terrorism and Political Violence* 26: 69–85.
- McKinley, M. 1992. “‘Dangerous Liaisons?’: The Provisional Irish Republican Army, Marxism, and the Communist Governments of Europe.” *History of European Ideas* 15 (1–3): 443–49.
- Merriam-Webster. 2017. “How Is ‘Fake News’ Defined, and When Will It Be Added to the Dictionary?” Merriam-Webster. <https://www.merriam-webster.com/words-at-play/the-real-story-of-fake-news>.
- Mickolus, E. F. 2002. “How Do We Know We’re Winning the War against Terrorists? Issues in Measurement.” *Studies in Conflict & Terrorism* 25 (3): 151–60.
- Mickolus, E. F., T. Sandler, J. M. Murdock, and P. A. Flemming. 2011. *International Terrorism: Attributes of Terrorist Events, Data Codebook*. Dunn Loring: Vinyard Software.

- . 2012. "International Terrorism: Attributes of Terrorist Events (ITERATE: 1968-2012)." Dunn Loring: Vinyard Software.
- Miller, E. 2016. "Patterns of Collective Desistance from Terrorism: Fundamental Measurement Challenges." *Perspectives on Terrorism* 10 (5): 5–21.
- Milligan, G. W., and M. C. Cooper. 1985. "An Examination of Procedures for Determining the Number of Clusters in a Data Set." *Psychometrika* 50: 159–79.
- Moloney, E. 2002. *A Secret History of the IRA*. New York and London: W. W. Norton & Company.
- Monaghan, R. 2000. "Single-Issue Terrorism: A Neglected Phenomenon?" *Studies in Conflict & Terrorism* 23 (4): 255–65.
- . 2013. "Not Quite Terrorism: Animal Rights Extremism in the United Kingdom." *Studies in Conflict & Terrorism* 36: 933–51.
- Monterey Terrorism Research and Education Program. 2012. "Monterey Weapons of Mass Destruction Terrorism Database." Monterey: Monterey Institute of International Studies.
- Morrison, J. F. 2016. "An Interview with Danny Morrison." *Terrorism and Political Violence* 28 (3). Taylor & Francis: 620–35.
- Mueller, J. 2005. "Simplicity and Spook: Terrorism and the Dynamics of Threat Exaggeration." *International Studies Perspectives* 6: 208–34.
- . 2006. *Overblown: How Politicians and the Terrorism Industry Inflate National Security Threats, and Why We Believe Them*. New York: Free Press.
- Mueller, J., and M. G. Stewart. 2012. "The Terrorism Delusion: America's Overwrought

- Response to the Terrorism Delusion." *International Security* 37 (1): 81–110.
- Munroe, H. D. 2009. "The October Crisis Revisited: Counterterrorism as Strategic Choice, Political Result, and Organizational Practice." *Terrorism and Political Violence* 21: 288–305.
- Murray, R. 1984. "Killings of Local Security Forces in Northern Ireland 1961-1981." *Terrorism* 7 (1): 11–52.
- National Consortium for the Study of Terrorism and Responses to Terrorism [START].
- 2013a. *GTD Codebook: Inclusion Criteria and Variables*. College Park: START.
- . 2013b. "GTD Frequently Asked Questions." *Global Terrorism Database*.
<http://www.start.umd.edu/gtd/faq/>.
- . 2017. *GTD Codebook: Inclusion Criteria and Variables*. College Park: START.
- National Counterterrorism Centre [NCTC]. 2012. *2011 Report on Terrorism*. Washington, DC: Office of the Director of National Intelligence.
- Nehorayoff, A. A., B. Ash, and D. S. Smith. 2016. "Aum Shinrikyo's Nuclear and Chemical Weapons Development Efforts." *Journal of Strategic Security* 9 (1): 35–48.
- Nesser, P. 2012. "Research Note: Single Actor Terrorism: Scope, Characteristics and Explanations." *Perspectives on Terrorism* 6 (6): 61–73.
- Oxford Dictionaries. 2016. "Word of the Year 2016 Is... Post-Truth." *Oxford Dictionaries*. Accessed December 12. <https://en.oxforddictionaries.com/word-of-the-year/word-of-the-year-2016>.
- Pape, R. A. 2003. "The Strategic Logic of Suicide Terrorism" 97 (3): 343–61.
- Pellegrini, L. A. 2008. "An Argument for the Criminal Hoax." *University of Southern*

California.

Pittel, K., and D. Rübhelke. 2011. "Characteristics of Terrorism." In *Handbook on the Economics of Conflict*, edited by D. L. Braddon and K. Hartley, 143–71. Cheltenham and Northampton: Edward Elgar.

Prunckun Jr., H. W., and P. B. Mohr. 1997. "Military Deterrence of International Terrorism: An Evolution of Operation El Dorado Canyon." *Studies in Conflict & Terrorism* 20 (3): 267–80.

Ragin, C. C. 2008a. *Redesigning Social Inquiry: Fuzzy Sets and Beyond*. Chicago and London: University of Chicago Press.

———. 2008b. "User's Guide to Fuzzy-Set/qualitative Comparative Analysis." Tuscon. <http://www.u.arizona.edu/~cragin/fsQCA/download/fsQCAManual.pdf>.

Rapoport, D. C. 2002. "The Four Waves of Rebel Terror and September 11." *Anthropoetics* 8 (1).

Reporters Without Borders. 2012. "World Press Freedom Index 2011-2012." Paris: Reporters Without Borders.

Richards, A. 2014. "Conceptualizing Terrorism." *Studies in Conflict & Terrorism* 37 (3): 213–36.

Rome, H. 2013. "Revisiting the 'problem From Hell': Suicide Terror in Afghanistan." *Studies in Conflict & Terrorism* 36 (10): 819–38.

Ross, J. I. 1988. "Attributes of Domestic Political Terrorism in Canada, 1960-1985." *Terrorism* 11 (3): 213–33.

Rowlands, D., J. Littlewood, and J. Kilberg. 2012. "CBRN Terrorism: Assessing the Threat

- of CBRN Terrorism by Groups and Individuals in Canada and Worldwide.” Ottawa: Norman Paterson School of International Affairs.
- Ryder, C. 2005. *A Special Kind of Courage: 321 EOD Squadron-Battling the Bombers*. London: Metheun.
- Schmid, A. P. 2011. “The Definition of Terrorism.” In *The Routledge Handbook of Terrorism Research*, edited by A. P. Schmid, 2013, 39–157. London and New York: Routledge.
- Schmid, A. P., A. J. Jongman, M. Stohl, and P. A. Flemming. 1988. “Terrorism and Related Concepts: Typologies.” In *Political Terrorism: A New Guide to Actors, Authors, Concepts, Data Bases, Theories, & Literature*, edited by A. P. Schmid and A. J. Jongman, 39–60. New Brunswick and London: Transaction Publishers.
- Schneider, C. Q., and C. Wagemann. 2012a. “Limited Diversity and Logical Remainders.” In *Set-Theoretic Methods for the Social Sciences: A Guide to Qualitative Comparative Analysis*, 151–77. Cambridge: Cambridge University Press.
- . 2012b. “Potential Pitfalls in the Standard Analysis Procedure and Suggestions for Improvement.” In *Set-Theoretic Methods for the Social Sciences: A Guide to Qualitative Comparative Analysis*, 197–219. Cambridge: Cambridge University Press.
- Sheehan, I. S. 2012. “Assessing and Comparing Data Sources for Terrorism Research.” In *Evidence-Based Counterterrorism Policy*, edited by C. Lum and L. W. Kennedy, 13–40. Springer.
- Siegel, D. A., and J. K. Young. 2009. “Simulating Terrorism: Credible Commitment, Costly

- Signaling, and Strategic Behavior." *PS: Political Science & Politics* 42 (4): 765–71.
- Silke, A. 2001. "The Devil You Know: Continuing Problems with Research on Terrorism." *Terrorism and Political Violence* 13 (4): 1–14.
- Slovic, P. 2002. "Terrorism as Hazard: A New Species of Trouble." *Risk Analysis* 22 (3): 425–26.
- StataCorp. 2015. *Stata Multivariate Statistics Reference Manual: Release 14*. College Station: Stata Press.
- Stern, J. E. 2000. "Larry Wayne Harris (1996)." In *Toxic Terror: Assessing Terrorist Use of Chemical and Biological Weapons*, edited by J. B. Tucker, 227–48. Cambridge: MIT Press.
- Strang, K. D. 2015. "Exploring the Relationship between Global Terrorist Ideology and Attack Methodology." *Risk Management* 17 (2): 65–90.
- Sunstein, C. R. 2003. "Terrorism and Probability Neglect." *The Journal of Risk and Uncertainty* 26 (2/3): 121–36.
- Tikuisis, P. 2009. "On the Relationship between Weak States and Terrorism." *Behavioral Sciences of Terrorism and Political Aggression* 1 (1): 66–79.
- Tishler, N. 2013a. "C , B , R , or N : The Influence of Related Industry on Terrorists ' Choice in Unconventional Weapons." *Canadian Graduate Journal of Sociology and Criminology* 2 (2): 52–72.
- . 2013b. "C , B , R , or N : The Influence of Related Industry on Terrorists ' Choice in Unconventional Weapons." 13–01. TSAS Working Paper Series.
http://www.tsas.ca/wp-content/uploads/2013/10/13-01_Tishler.pdf.

- . 2016. "Taking Hoaxes Seriously: Characteristics of Terrorism Hoaxes and Their Perpetrators." 16-03. TSAS Working Paper Series. http://tsas.ca/wp-content/uploads/2016/04/TSASWP16-03_Tishler.pdf.
- . 2017. "Trends in Terrorists' Weapons Adoption and the Study Thereof." *International Studies Review* 0 (0): 1-27. doi:<https://doi.org/10.1093/isr/vix038>.
- Tishler, N., M. Ouellet, and J. Kilberg. n.d. "A Survey of Terrorism in Canada: 1960-2015." In *Terrorism and Counterterrorism in Canada – Canada Among Nations*, edited by J. Littlewood, L. Dawson, and S. Thompson. Toronto: University of Toronto Press.
- TSAS. 2015. "CIDB Codebook." (March 14). http://extremism.ca/Content/CIDB_Data_codebook.pdf.
- Tucker, J. B. 2000a. "Introduction." In *Toxic Terror: Assessing Terrorist Use of Chemical and Biological Weapons*, edited by J. B. Tucker, 1-14. Cambridge: MIT Press.
- . 2000b. "Lessons from the Case Studies." In *Toxic Terror: Assessing Terrorist Use of Chemical and Biological Weapons*, edited by J. B. Tucker, 249-70. Cambridge: MIT Press.
- . ed. 2000c. *Toxic Terror: Assessing Terrorist Use of Chemical and Biological Weapons*. Cambridge: MIT Press.
- Veness, D. 2001. "Terrorism and Counterterrorism: An International Perspective." *Studies in Conflict & Terrorism* 24 (5): 407-16.
- Wallace Jr., P. S. 2002. "Anti-Hoax Legislation in the 107th Congress: Addressing Problems since September 2001." RL31314. CRS Report for Congress.
- Weinberg, L., A. Pedahzur, and A. Perliger. 2009. *Political Parties and Terrorist Groups*.

2nd ed. London and New York: Routledge.

Weinstein, J. M. 2007. *Inside Rebellion: The Politics of Insurgent Violence*. New York: Cambridge University Press.

Wilcox, L. 2017. "Fake Hate Crimes." [Fakehatecrimes.org](http://www.fakehatecrimes.org).
<http://www.fakehatecrimes.org>.

Wilkinson, P. 2011. *Terrorism versus Democracy: The Liberal State Response*. 3rd ed. New York: Routledge.

Young, J. K. 2016. "Measuring Terrorism." *Terrorism and Political Violence*, 1–23.

Zulaika, J., and W. A. Douglass. 2008. "The Terrorist Subject: Terrorism Studies and the Absent Subjectivity." *Critical Studies on Terrorism* 1 (1): 27–36.