PLEDGING TO HARM:

A LINGUISTIC ANALYSIS OF VIOLENT INTENT IN THREATENING LANGUAGE

MARLON DAVID HURT

Doctor of Philosophy

ASTON UNIVERSITY

August 2020

M.D.Hurt, PhD Thesis, Aston University, 2020.

Aston University

Pledging to Harm:

A Linguistic Analysis of Violent Intent in Threatening Language

Marlon David Hurt

Doctor of Philosophy

2020

Thesis Summary

Legal systems around the world assume that violent intent is not only real, but that it is also detectable in threatening language. However, empirical studies examining how, or even whether, violent intent is encoded in language are rare, and tend to explore the issue primarily through psychological theory. This linguistic analysis hypothesizes that authorial intent is indeed detectable in the language of threats, if only obliquely, because the functional aim of a threat issued with true violent intent is different than one issued for other communicative purposes, e.g., to cause fear. A novel combination of frameworks is employed to test this hypothesis on a dataset of six realized and eight non-realized threats. First, Audience Design Theory and Speech Act Theory delimit the investigation to the most common kind of threatening language, called 'leakage' in the threat assessment literature and a 'pledge to harm' in Speech Act Theory. Next, the Folk Concept of Intentionality and Biological Naturalism theorize which cognitive elements of intent may be expressed by pledges to harm. Finally, Systemic Functional Linguistics, and the discourse semantic method of Appraisal in particular, identify the various attitudinal and interpersonal meanings in the pledge dataset. Non-realized pledges are discovered to contain significantly more violent ideation, creating a prosody of heightened menace, while the realized pledges are more concerned with ethical evaluations. Hypothetically, these patterns of stancetaking show that the non-realized and realized texts are engaged in divergent 'fields of activity', that of *announcing* and *explaining* respectively. Different communicative purposes point to different psychological intentions spurring the production of each pledge type, potential evidence that violent intent is indeed detectable in the language of pledges to harm.

Keywords: Systemic Functional Linguistics; Appraisal analysis; forensic linguistics; leakage; intention and intentionality

ACKNOWLEDGEMENTS

"Human beings are discourse"

—Rumi

To Erin: your love is everything; I am who I am today because you found me when I was lost.

To Tom and Chris: your steady hands and gentle hearts are a model to me; thank you for the unwavering support which has made so much of this new life possible.

To Cyril and Caliere: I am so lucky to bear witness to the beginnings of your own discourses.

To Mom and Dad: to you both, I owe my love of language and my belief that justice is an ideal worth striving for.

To Tammy: my great fortune is that I've been able to continue learning from you; thank you for taking a chance on me and for keeping me close to the field that I love.

To Melanie and Karla: I couldn't have closed the theoretical circle without your expertise; I look forward to comparing notes again in the future.

To the retired law enforcement officer: this thesis would be so much the poorer without your generosity.

And, of course, to Tim: thank you thank you thank you for the *years* of your time and energy; you've shown me how forensic linguists can make the world a better place, not only through the science of linguistics but by always attending to the human values at stake in what we do; thank you for guiding me through the most difficult thing I've ever done.

TABLE OF CONTENTS

Chapter 1: Intent and the Law 15

1.1: Aims and Approach 20

1.2: Demarcating the Data 23

1.2.1: Speech Act Theory 23

1.2.2: Audience Design Theory 27

1.2.3: Leakage 29

1.3: Outline of the Dissertation 34

Chapter 2: Theorizing Intent 36

2.1: Linguistic Intent 37

2.2: Psychological Intent 40

2.2.1: Components of Intent 42

2.2.2: Reasoning 48

2.2.3: Commitment 52

2.3: Language of Intent 56

2.3.1: Language of Threatening 61

2.3.2: Realized vs Non-realized Threatening 70

2.4: Summary 80

Chapter 3: Data and Method 84

3.1: Ethics and Process of Data Collection 86

3.1.1: Process 90

3.2: Dataset 90

3.2.1: Non-realized Corpus 96

3.2.2: Realized Corpus 100

3.3: Method 101

3.3.1: UAM Corpus Tool 3 103

3.3.2: Appraisal and Systemic Functional Linguistics 104
3.3.2a: Attitude 106
3.3.2b: Engagement 108
3.3.2c: Graduation 109
3.3.3: Risks and Limitations of the Appraisal Method 111
3.4: Summary 114
Chapter 4: Attitude 116
4.1: Attitude Polarity 118
4.2: Attitude Null Results 123
4.3: Summary 128
Chapter 5: Judgement 132
5.1: Tenacity 133
5.2: Capacity 137
5.2.1: Positive Capacity 140
5.2.2: Negative Capacity 141
5.2.3: Violent vs Nonviolent Negative Capacity 143
5.2.4: Parts of Speech 146
5.2.5: Tense 147
5.2.6: Target Specificity 148
5.2.7: Action Specificity 155
5.2.8: Linking in the Rodger NR and Valle NR Texts 160
5.2.9: Summary of Capacity 161
5.3: Propriety 162
5.3.1: Negative vs Positive Propriety 169
5.3.2: Who is Un/Ethical? 171
5.3.3: Parts of Speech 177

5.3.4: Tense 183

5.3.5: Summary of Propriety 185

Chapter 6: Appreciation 187

6.1: Composition 188

6.1.1: Negative Composition 191

6.1.2: Summary of Composition 196

6.2: Valuation 197

6.2.1: Positive Valuation 202

6.2.2: Negative Valuation 206

6.2.3: Valuation and Tense 208

6.2.4: Summary of Valuation 212

Chapter 7: Engagement 215

7.1: Justification 219

7.2: Contraction vs Expansion 221

7.3: Dialogic Expansion: Attribute and Acknowledge 222

7.3.1: Whose Voice is Admitted? 225

7.4: Dialogic Contraction: Proclaim 232

7.5: Dialogic Contraction: Disclaim and Deny 235

7.5.1: Topics of Denial 239

7.6: Dialogic Contraction: Temporal Shifting 245

7.7: A Note on Planning 247

7.8: Summary 248

Chapter 8: Modal Auxiliaries 254

8:1: Overview 258

8.2: Likelihood 261

8.3: Ability 271

8.4: Inclination 277

8.5: Requirement 281

8.6: Summary 287

Chapter 9: Graduation 292

9.1: Summary 298

Chapter 10: A Blind Test 300

10.1: Process and Analysis 301

10.2: Results 307

10.3: Discussion 307

Chapter 11: Discerning Intent 315

11.1: Original Contributions 316

11.2: Difficulties and Limitations 317

11.3: Primary Findings 319

11.4: Discussion 322

11.4.1: Linked Reasoning Processes 324

11.4.2: Register Variation and Communicative Intent 327

11.5: Future Research 333

References 336

Appendices 351

LIST OF FIGURES

Figure 1.1: Pathway Model of Severe Targeted Violence in Schools 32

Figure 2.1: Student Ideologies about Threatening Language 78

Figure 2.2: Distribution of Modals by Threat Realization 79

Figure 3.1: Example of Auto-Normed Corpus Tool Output 103

Figure 4.1: Attitude Polarity by Corpus 119

Figure 4.2: Positive Attitude Dispersion by Corpus 120

Figure 4.3: Negative Attitude Dispersion by Corpus 120

Figure 4.4: Judgement and Appreciation as Institutionalised Affect 124

Figure 5.1: Positive Tenacity Dispersion by Corpus 134

Figure 5.2: Negative Tenacity Dispersion by Corpus 134

Figure 5.3: Tokens of Tenacity by Realization Category 135

Figure 5.4: Tokens of Tenacity by Author 135

Figure 5.5 Positive Capacity Dispersion by Corpus 138

Figure 5.6: Negative Capacity Dispersion by Corpus 138

Figure 5.7: Positive and Negative Capacity by Realization Category 139

Figure 5.8: In/capacity by Author 139

Figure 5.9: Positive and Negative Capacity by Realization Category (Adjusted) 140

Figure 5.10: Violent vs Nonviolent Lexemes of -Capacity by Realization Category .. 143

Figure 5.11: Violent Lexemes of Capacity by Parts of Speech 146

Figure 5.12: Violent Verbs by Tense 148

Figure 5.13: The Nominal Hierarchy 151

Figure 5.14: Positive Propriety Dispersion by Corpus 164

Figure 5.15: Negative Propriety Dispersion by Corpus 164

Figure 5.16: Tokens of Propriety by Realization Category 166

Figure 5.17: Tokens of Propriety by Realization Category (Adjusted) 166

Figure 5.18: Proportion of Propriety by Realization Category 167
Figure 5.19: Directedness of Propriety by Realization Category 171
Figure 5.20: Objects of Propriety by Realization Category 174
Figure 5.21: Objects of Propriety Subcategories 175
Figure 5.22: 'Other' Objects of Propriety by Realization Category 176
Figure 5.23: Parts of Speech in Propriety by Realization Category 179
Figure 5.24: Unethical Processes by Realization Category 181
Figure 5.25: +/-Propriety Tense by Realization Category 183
Figure 5.26: Present Tense +/-Propriety by Realization Category 184
Figure 6.1: Positive Composition Dispersion by Corpus 189
Figure 6.2: Negative Composition Dispersion by Corpus 189
Figure 6.3: Composition Polarity by Realization Category 190
Figure 6.4: -Composition Senses by Realization Category 192
Figure 6.5: Positive Valuation Dispersion by Corpus 197
Figure 6.6: Negative Valuation Dispersion by Corpus 198
Figure 6.7: Valuation and Judgement Variables 200
Figure 6.8: Judgement Polarity by Realization Category 201
Figure 6.9: Valuation Polarity by Realization Category 201
Figure 6.10: Valuation Polarity by Realization Category (Adjusted) 202
Figure 6.11: Judgement Meanings via +Valuation by Realization Category 203
Figure 6.12: +Valuation Meanings by Realization Category (Adjusted) 204
Figure 6.13: -Valuation Meanings by Realization Category 206
Figure 6.14: +/-Propriety via Valuation Sorted by Tense in Realized Texts 209
Figure 6.15: +/-Propriety via Valuation by Tense and Semantics in Realized Texts . 210
Figure 7.1: Dialogic Contraction and Expansion 222
Figure 7.2: Acknowledge and Distance 223

Figure 7.3: Attribute Dispersion by Corpus 223
Figure 7.4: Acknowledge Dispersion by Corpus 224
Figure 7.5: Attributed Voices 225
Figure 7.6: Acknowledged Voices 226
Figure 7.7: Other Voices 231
Figure 7.8: Portion of Proclaim by Realization Category 233
Figure 7.9: Proclaim Dispersion by Corpus 233
Figure 7.10: Portion of Proclaim by Realization Category (Adjusted) 234
Figure 7.11: Disclaim and Deny by Realization Category 235
Figure 7.12: Disclaim Dispersion by Corpus 236
Figure 7.13: Deny Dispersion by Corpus 237
Figure 7.14: Grammatical Markers of Denial by Realization Category 237
Figure 7.15: Denial Topics by Realization Category 241
Figure 7.16: Denial Topic 'Polarity' by Realization Category 243
Figure 8.1: Student Ideologies about Threatening Language 255
Figure 8.2: Modal Auxiliaries 258
Figure 8.3: Modal Meanings 259
Figure 8.4: Modal Meanings by Realization Category 259
Figure 8.5: 1st Person + Future Tense Modal Meanings by Corpus 261
Figure 8.6: Likelihood Dispersion by Corpus 262
Figure 8.7: Likelihood Modals by Realization Category 262
Figure 8.8: Likelihood Modals by Realization Category (Adjusted) 263
Figure 8.9: Ratio of Attack-related Likelihood Modals 265
Figure 8.10: Likelihood Modal Auxiliaries by Time Frame 266
Figure 8.11: Pre-attack and Attack Modals of Likelihood 268
Figure 8.12: Post-attack Modals of Likelihood 270

Figure 8.13: Post-attack Modals of Likelihood (Adjusted) 270
Figure 8.14: Ability Dispersion by Corpus 272
Figure 8.15: Ability Modals by Realization Category 273
Figure 8.16: 1st Person Attack-related (A-R) Modals of Ability by Corpus 275
Figure 8.17: Inclination Dispersion by Corpus 278
Figure 8.18: Modal Auxiliaries of Inclination 278
Figure 8.19: Attack-related Inclination 279
Figure 8.20: Requirement Dispersion by Corpus 282
Figure 8.21: Modal Auxiliaries of Requirement 283
Figure 8.22: Levels of Requirement 283
Figure 8.23: Required Violence by Corpus 284
Figure 8.24: Required Violence by Corpus (Adjusted) 284
Figure 9.1: The System of Graduation 293

List of Tables

Table 1.1: Contrasts among Threatening, Warning, and Promising 25

Table 1.2: Hierarchy of Attributes and Audience Roles..... 27

Table 2.1: Cognition and Volition 43

Table 2.2: Contrasts among Belief, Desire, and Intention..... 46

Table 3.1: Pledge to Harm Dataset 91

Table 4.1: Statistical Significance in Attitude..... 118

Table 4.2: Tokens of Affect 125

Table 5.1: Statistical Significance in Judgement by Realization Category..... 132

Table 5.2: Lemmas of Violent Ideation by Realization Category..... 145

Table 5.3: Targets of Incapacity by Corpus..... 149

Table 5.4: Nominal Hierarchy of Stated Targets 152

Table 5.5: Scope of Proper and Common Noun Phrases of Stated Targets 153

Table 5.6: Troponymy of Future Tense Violent Verbs..... 156

Table 5.7: Violent Verbs and Modifier Types 158

Table 5.8: Primary Target by Text..... 174

Table 5.9: Noun Phrases..... 178

Table 5.10: Process Lexemes of Negative Propriety 180

Table 5.11: Modifiers Marked for Propriety 182

Table 6.1: Statistical Significance in Appreciation by Realization Category..... 187

Table 6.2: Tokens of Plan-related +Composition 190

Table 6.3: Narrative Category of ‘Other’ Tokens of -Composition..... 194

Table 7.1: Statistical Significance in Engagement 218

Table 7.2: Projection Nexuses of Target Voices 228

Table 7.3: Projection Nexuses of Audience Voices..... 230

Table 7.4: Tokens of Temporal Shifting 246

Table 8.1: Modal Auxiliaries and their Meanings 257

Table 8.2: Concordance Lines of Ability Modals 274

Table 9.1: Statistical Significance in Graduation 295

Table 10.1: Blind Test Results 307

Table 10.2: Statistically Significant Appraisal Features Counts 311

Table 10.3: Rodger R and Text2 Feature Count Errors 313

Table 11.1: Communicative Intent by Realization Category 332

List of Appendices

Appendix A: Justin Carter Facebook Threat 351

Appendix B: Ethical Decision-making Tree 352

Appendix C: Jerry Varnell Pledge 353

CHAPTER 1 INTENT AND THE LAW

Legal systems around the world assume that intent is not only a real psychological quality, but also that it can be detected in people's behavior—including their language behavior. Shuy (1981: 115) puts the matter bluntly when he characterizes a criminal court case as "little more than the establishment of intentions and the evidence of having carried them out." In the United States, as in many other countries, criminal intent, or *mens rea*, is often weighed as an equal counterpart to the criminal act, or *actus reus*, being prosecuted. Depending on whether the defendants meant the harm they caused, the law's sensitivity to intent can lead to differing punishments for otherwise identical offenses. Evidence of *mens rea*, for example, may justify the escalation of involuntary manslaughter—an unintentional killing—to the more serious charge of voluntary manslaughter or even murder (18 U.S.C. § 1111-1112).

Similarly, law enforcement agencies operate under the assumption that violent intent can be detected beforehand, with the aim of preventing threatened actions from being committed. In literature devoted to the field of threat assessment, for example, intent is considered a necessary precursor to violence, a required step situated between the idea of violence and the violent act itself (Meloy & Hoffmann, 2014). However, while these assumptions are both commonsensical and commonplace, the question of whether violent intent may indeed be detected beforehand remains fraught. Most efforts to address the issue have focused, with some success, on behavioral cues (Borum, Fein, Vossekuil, & Berglund, 1999; Meloy, 2015; O'Toole & Smith, 2014), e.g., stalking (Meloy, 2015) or the purchase of a weapon (Calhoun & Weston, 2015). Or they have looked to psychological factors such as pathological narcissism (Smith, 2006). "The contemporary research on threats is dominated by case studies" which "identify warning behaviors or risk factors associated with (threats to commit) targeted violence" (Geurts,

Granhag, Ask, & Vrij, 2016: 54). This focus on behavior and biography often relegates the *language* of threatening to a secondary status in these literatures—that is, if language is considered at all. Instead, the *act* of threatening (Meloy & O'Toole, 2011) and the medium through which the threat is conveyed (Scalora, Baumgartner, Zimmerman, Callaway, Maillette, Covell, Palarea, Krebs, & Washington, 2002) are given more predictive weight than any linguistic patterns evident in the words themselves.

And yet, cases where violent intent has been expressed "often start and end with little more than the threatening communication itself" (Smith, 2006: 10). The "lack of empirical guidance" (Borum *et al.*, 1999: 326) concerning which language features may correlate with psychological intent is problematic in forensic contexts, where real-world consequences follow from how a threat is interpreted by authorities. The dangers are not limited to the party being threatened. While threats of violence are themselves illegal, the possible repercussions to the threatener are more severe when he or she is accused of harboring *mens rea*—i.e., of expressing a true intent to harm other people rather than just venting anger (Fraser, 1998). The question of how or whether violent intent is encoded in language, then, is fundamental to the administration of justice.

A few cases highlight the inherent dangers facing authorities, citizens, and the wider public in such situations. The first began in 2013, when an 18-year-old Texas resident named Justin Carter took part in a heated online exchange that included the language shown in Text 1.1, via the social media platform of Facebook (Appendix A).

> Text 1.1: Justin Carter Facebook Threat[1]
> I'm fucked in the head alright, I think I'ma SHOOT UP A KINDERGARTEN AND WATCH THE BLOOD OF THE INNOCENT RAIN DOWN AND EAT THE BEATING HEART OF ONE OF THEM

[1] Following the practice of Gales (2010: 1): "All non-standard language use (e.g. misspellings, incorrect lexical choice, unusual syntax, spacing, and punctuation) has been left intact in all of the example texts herein."

Carter's statements were reported to the Texas state police, who subsequently arrested the young adult and charged him with felony 'terroristic threatening' (Texas Code § 22.07), which carried a possible sentence of 10 years in prison and a $10,000 fine (Sanders, 2018). Carter's lawyer argued that his language was sarcastic, and therefore lacked any true intent to carry out the violence he described (Pinsof, 2013). Along with several external pieces of evidence supporting his lawyer's claim (e.g., Carter did not own, nor had he ever tried to acquire a weapon, police found no other threatening writings, etc.), qualities of the language itself seem to favor such an interpretation. For instance, while shooting up a kindergarten is (sadly) possible, the impossibility of eating a still-beating heart and the near impossibility of causing blood to literally 'rain down' easily qualify as hyperbole, a device often used to signal ironic intent (Kreuz & Roberts, 1995). Texas prosecutors eventually—though only tacitly—agreed to this interpretation as well when, in 2018, they downgraded the charge to a misdemeanor 'false alarm or report' (Sanders, 2018). Where a 'terroristic threat' requires an intent to cause fear and disruption, a 'false alarm' openly states that the content of a communication which threatens a "bombing, fire, offense, or other emergency" is "false or baseless" (Texas Code § 42.06). In Carter's case, the reduced charge was an endorsement by the state of the view that his words did indeed lack a real-world intent to perform the expressed actions. An appeals court agreed, writing that the language of the threat "on its face is so clearly hyperbolic and sarcastic that it would be imprudent to proceed further" (Tex. Ct. App., 2015).

However, it is fair to ask what the court might have said about the darkly hyperbolic writing by Tyrelle Shaw. As part of a longer posting published on his blog in 2015, the 25-year-old New Yorker explained his plan to "hit over a million Asian Women in the face with a stick," which would result in "an independent civil war" and "change history" (Shaw, 2015). The pattern of escalation in Shaw's ideations appears to mirror

Carter's. As with attacking a kindergarten, hitting a woman in the face is also (sadly) possible. However, the numeric impossibility of a single person assaulting a million women, and the chances of this resulting in a full-fledged civil war, seem just as wildly hyperbolic as Carter's ideations, and thus similarly discreditable. Shaw nevertheless followed through on his threat, managing to injure four women—and terrify the larger New York City Asian community—before committing suicide.

Alongside matters of overly severe criminal charges (in the case of Carter) and concern for the safety of potential victims (in the case of Shaw), the issue has potential financial, cultural, and emotional ramifications for the public as well. In the case of threats directed at schools, for instance:

> deciding how to respond usually falls to administrators, who lean heavily on the advice of law enforcement officials, often have little verified information to go on and only a few hours to make the call, and have a sense that they might be second-guessed no matter what. An administrator fears not reacting strongly enough when lives are at stake, but college and school officials say there are costs to overreacting—in policing expenses, lost classroom time, frayed nerves and the danger of encouraging copycat threats. (Pérez-Peña, 2015)

This dynamic was succinctly illustrated in 2015, when an identical email threatening a variety of physical assaults on schools and students was sent to the Los Angeles Board of Education as well as to government officials in New York City. In response, the Los Angeles Unified School District closed all 900 schools, at a potential cost of $29 million. New York chose to keep their schools open. The violence the emails threatened never materialized in either city, and the Federal Bureau of Investigation (FBI) later determined the threat was not credible. "New York officials received the email at roughly the same time, and with three hours less time to assess it, came to a sharply different decision" (Branson-Potts, Ceasar, & Blume, 2015). The two "sharply different" interpretations by authorities of a single threatening text is evidence that confusion still surrounds the detection of violent intent.

These cases, and the many like them which occur annually in the United States alone, illustrate the ongoing need for empirical research focused on the *language* of violent intent. However, threat assessors and other professionals tasked with navigating the types of situations presented by Carter, Shaw, and others are not dealing with a threat as it is prototypically conceived—that is, as a promise of future harm to the addressee (Fraser, 1998). Carter, for instance, was not speaking to the children of the unnamed kindergarten; neither was Shaw writing to the Asian women he imagined assaulting. Even the anonymous LA/NYC emailer, who closed the email by telling the reader *this may be your last day*, was primarily threatening the students and teachers of the two school districts rather than the education and government officials to whom the email was sent.

And yet, communications like these, which target a grammatical 3rd Person, are considered more common occurrences than direct threats to the grammatical 2nd Person of the addressee (Meloy & O'Toole, 2011). In the threat assessment literature, this kind of warning behavior is called 'leakage' (Meloy, Hoffman, Guldimann, & James, 2014). In the linguistic literature, such a communication is called a 'pledge to harm'[2] (Harmon, 2008). (Both will be further theorized in the sections below.) The frequency of pledges to harm means they are a relatively widespread forensic issue. However, pledging's less prototypical nature in comparison to a direct threat means that texts like Carter's and Shaw's are also understudied as a linguistic phenomenon. Combined with the paucity of research that seeks to discern psychological intent in language, this means that little empirical guidance exists for threat assessors and other authorities whose job it is to weigh the risk these kinds of threats may actually pose (Borum *et al.*, 1999). This research works to close this gap.

[2] The reduced form of 'pledge' will be used synonymously with the full term of 'pledge to harm' throughout the thesis.

1.1 AIMS AND APPROACH

The current study attempts to advance the body of knowledge in the field of threat assessment through a novel combination of aim and approach. Other studies aimed at uncovering psychological intent take alternate approaches (e.g., Smith, 2006), or use the experience and intuitions of professional threat assessment practitioners as the foundation of their analyses (e.g., Calhoun & Weston, 2015). Elsewhere, studies with a similar approach have different aims. Gales' (2010, 2011, 2017) groundbreaking linguistic work on threats, for instance, eschews the issue of psychological intent in favor of uncovering how stance operates within threatening as a genre, along with addressing more fundamental questions about "what threatening language actually is" (Gales, 2017: 2).

The current multidisciplinary study focuses on the question of psychological intent but does so through an established framework of linguistic theory—known as Systemic Functional Linguistics (SFL). Underpinning SFL is the idea that language is a meaning-making resource, and that a choice between one option or another within the language system depends on its intended communicative function. That is, "choices among alternatives" at all levels of a grammar may be best understood "from the standpoint of how [a language] creates and expresses meaning" (Halliday & Matthiessen, 2004: 19). In practical terms, this means that the surface form of a text—the words and clauses put in a particular order by a particular writer—is deeply tied to the purpose the author envisaged the text would serve, however consciously. Analyzing these linguistic forms, then, should allow for a better understanding of their communicative function, which in turn could open a window onto the psychological processes spurring their composition. This analysis is operationalized through a discourse analytical method born of SFL called Appraisal (Martin & White, 2005), which

is designed to discover an author's subjective presence in a text. (Appraisal is discussed more fully in Chapter 3.)

By bringing SFL to bear on whether language may potentially express a particular psychological quality, the current investigation seeks to answer the following questions:

- What similarities and differences exist in the language of a realized versus a non-realized pledge? If differences are discoverable between the two text types, what is the extent of this difference?
- What can the application of Systemic Functional Linguistics, and the method of Appraisal in particular, reveal about the mechanics of stancetaking in the two realization categories?
- How well or poorly do any patterns of stancetaking in the two text types comport with the psychological theories of intention formation and the theories of intention detection by outside social perceivers? How do these patterns compare to previous linguistic research on levels of commitment in threatening language? How do they compare to previous research on leakage and language features considered risk-enhancing by threat assessors?
- How well do language markers flagged as risk-enhancing by previous researchers correlate with pledge authors' subsequent behavior? Are patterns or features evident which have not been addressed in previous research on the subject of language and intent?
- What does the triangulation of linguistic and psychological theories used herein contribute to the cross-disciplinary understanding of how—or even whether—psychological intent is encoded in pledges to harm? What does Appraisal, combined with certain corpus linguistic techniques, contribute to this

understanding, i.e., what can be gleaned of language and intent from both qualitative and quantitative methodologies when applied to pledges to harm?

This study hypothesizes that different *systemic* resources are employed in the communication of true violent intent because the *functional* aim of such a communication is different than that of a text created with a different intent (e.g., to vent anger). Furthermore, because these systemic resources are manifested in the tangible realm of language, this study hypothesizes that the difference between violent intent and other intents is visible to scientific tools, i.e., linguistic tools, and is therefore detectable in the texts themselves.

Focusing on violent intent versus different communicative intents invites a comparison between expressions of imagined future violence which were realized in the world-at-large against those where none of the threatened actions were attempted, meaning the threat remained non-realized. And so, to test these hypotheses, a dataset of *authentic* pledges to harm has been curated. By subdividing this collection based on whether the author is known to have acted on the ideations stated in his or her text and comparing the resulting two subcorpora using the tools of Appraisal analysis, this project approaches two related goals, each equal in value to the other. The first is to offer a scientifically sound criticism—one that is theoretically-grounded and empirically-informed—of the assumption underlying so much legal theory, i.e., that violent intent is both detectable and measurable. The second is to potentially offer those tasked with pursuing justice, whether they be threat assessors on the front end or the courts on the back end, with empirically tested linguistic markers of intent to which they may defer in their decision-making, removing some of the 'subjective focus' from their attempts to adjudicate the meaning of a threatening communication (Durant, 2010).

1.2 DEMARCATING THE DATA

Although SFL is the foundation of this study, clearly demarcating the kind of communication authored by Carter, Shaw, and LA/NYC, requires input from two additional theoretical frameworks, as well as the law enforcement literature on threat assessment. The purpose of the subsections which follow are to properly delimit the scope of the inquiry, and therefore determine which data are appropriate for analysis.

To date, the most concise theorizing of threatening as an interpersonal action is found in Speech Act Theory (Austin, 1975; Searle, 1976) and refinements derived from it (e.g., Fraser, 1998; Shuy, 1993; Storey, 1995). Yet, any understanding of the pragmatic forces which shape specific utterances would be incomplete without elements of Audience Design Theory (Bell, 1984). Both SAT and ADT will be treated sequentially, although the text type in question resides in the area where they intersect. Once the language features of possible texts are identified, the final preparatory step will be an overview of what the threat assessment literature refers to as 'leakage,' to better understand just where on a "pathway to violence" (Meloy *et al.*, 2014: 10) this kind of threatening language typically occurs. This is necessary to distinguish which of the qualifying data do, and which do not, arguably encode the intent to act. (By contrast, a confession might explain the speaker's intentions, but only after the fact.) A comparison of the two resulting corpora of realized and non-realized pledges forms the core of this investigation.

1.2.1 SPEECH ACT THEORY

Speech Act Theory grew from Austin's (1975) observation that language is not used in the real world simply to describe certain states of affairs but also to perform certain actions. SAT has been ably summarized in many places (e.g., Gales, 2010; Harmon, 2008), but one of the central ideas is that a single utterance is comprised of "three kinds

of acts—the locutionary, the illocutionary, and the perlocutionary" (Austin, 1975: 103). In his framework, the 'locution' is simply "what is said" (Harmon, 2008: 40). This is the physical aspect of using language which Austin (1975) elsewhere refers to as the "phonetic act" (p. 115), although the visual representations of writing and signing also qualify. Austin's key intuition is that what is *said* as an utterance is separable from what is *meant* by an utterance. This second level he calls the 'illocution,' and its force lies in what the speaker intends the hearer to understand. Finally, the "consequential effects" which this force has on "the feelings, thoughts, or actions of the audience" is termed the 'perlocution' (Austin, 1975: 101). An example of how these three levels interlock is provided by Austin (1975) with the simple sentence 'Shoot her': the locution is the words themselves, with their plain semantic meanings (e.g., *shoot* means to fire a weapon and *her* is a deictic pronoun referring to a single, female entity); the illocutionary force is the urging or command that the hearer 'shoot her'; and the perlocutionary effect is, possibly, the listener being persuaded by this command to fire his or her gun (Austin, 1975: 101-102).

Building on this, Searle (1976) offers a classification system of illocutionary acts which comprises the following five categories:

- representatives "commit the speaker…to the truth of the expressed proposition" (Searle, 1976: 10);
- directives are "attempts…by the speaker to get the hearer to do something" (Searle, 1976: 11);
- commissives are "those illocutionary acts whose point is to commit the speaker…to some future course of action" (Searle, 1976: 11);
- expressives "express [a] psychological state" (Searle, 1976: 12); and

- declarations, whose members "bring[] about the correspondence between the propositional content and reality" (Searle, 1976: 13), e.g., as happens when a couple is declared to be married and is therefore married.

Threatening certainly qualifies as "an illocutionary act, an intentional act of using language to send a message" (Fraser, 1998: 160). Yet, the content of this message must be twofold. First, the speaker must present him- or herself as intending to personally commit some future action, and, second, that this action will result in a state of the world which is unfavorable for the addressee (Fraser, 1998). In Searle's (1976) taxonomy, then, a threat is best understood as carrying the illocutionary point of a *commissive*. In this, it both resembles and differs from the more benign speech acts of promising and warning, with distinctions between the three hinging on 1) who benefits from the outcome, 2) who controls the outcome, and 3) from whose perspective the proposal is being made, as shown in Table 1.1.

Table 1.1: Contrasts among Threatening, Warning, and Promising (Gales, 2010: 10)

	Threatening	Warning	Promising
To the speaker's benefit	X		
To the hearer's benefit		X	X
To the hearer's detriment	X		
From speaker's perspective	X	X	X
Speaker controls outcome	X		X
Hearer controls outcome		X	

In sum, a threat is a commissive from the speaker's perspective, whose proposed outcome is 1) under the speaker's control, 2) to the speaker's benefit, and, perhaps most crucially, 3) to the hearer's detriment (Shuy, 1993). Gales's (2010, 2011, 2017) significant work on this kind of *direct threat*[3]—where the hearer is also the

[3] The broad consensus in the threat assessment and linguistic literatures is that a direct threat is simply one category within a larger paradigm, which also includes conditional and indirect threat types (e.g., Gales, 2010; Napier & Mardigian, 2003). These distinctions do not bear on the current research, and so 'direct threat' is used here and

individual being targeted—characterizes such a text as containing an interpersonal stance of 'disalignment' between a grammatical 1st and 2nd Person ('I' vs. 'you'). In other words, "[a] threat, as opposed to a warning or a promise, is clearly made to the detriment of the hearer[4]" (Gales, 2010: 10). And in this way, "the threatener is naturally poised against his or her intended audience" (Gales, 2010: 214).

However, as noted above, threat assessors and other professionals tasked with navigating the types of situations presented by, e.g., Justin Carter are not analyzing this kind of communication. Carter's 'threat' satisfies Shuy's (1993) first and second condition: the author is personally committing himself to performing a future action, one which presumably will be to his emotional benefit, if nothing else. But it fails the third, in that he is not threatening to harm his reader—the 'hearer' in Shuy's (1993) conception—but rather a third, outside party. To Harmon (2008: 43), failing this final condition means "that what lawyers, jurists, and legislators refer to as a violent 'threat' towards a third party is not really a 'threat' at all." This is because, "[i]n the strict sense, a speaker cannot threaten someone who is not there" (Harmon, 2008: 35). However, while Carter's text may not be a 'threat' definitionally, it is inarguably a *kind* of threatening communication, easily fitting within Durant and Leung's (2016: 159) broader conception of a threat as "a declaration of one's intention to do injury to a person or his or her property." Harmon (2008: 44) concludes that "a proper linguistic classification of the illocutionary act performed in [such a] context would be a 'pledge,' a type of 'commissive,' since the speaker commits himself to harm a third party". Thus, to distinguish texts like Carter's from the more classically formulated, direct threats

---

throughout as an umbrella term for those threatening communications which are addressed directly to the person(s) being threatened.

[4] Gales (2010: 6) notes that "[t]he majority of previous work on threatening language focuses on spoken discourse. However, because threats can be spoken *and* written, for the purposes of this research, the terms: speaker/hearer and writer/recipient, respectively, will be used interchangeably." The same practice is adopted here.

discussed above, Harmon (2008) uses the term pledge to harm, a usage which is adopted here. To understand why the distinction matters—why *who* is being threatened motivates the addition of a separate illocutionary act to the SAT taxonomy—requires understanding how speakers craft their language to suit their intended audience.

1.2.2 AUDIENCE DESIGN THEORY

"In audience design, speakers accommodate primarily to their addressee" (Bell, 1984: 145). In other words, how a speaker crafts a message is highly dependent on who the message is being crafted for. Of course, "an utterance does not carve up the world beyond the speaker into precisely two parts, recipients and non-recipients" (Goffman, 1981: 137). Additional contextual pressures are who else is allowed to participate in the interaction, and, broader still, which non-participants are known to be listening. In ADT, these audience members are ranked in a hierarchy of roles, each with a different saliency for the person speaking, reproduced in Table 1.2 below.

Table 1.2: Hierarchy of Attributes and Audience Roles (Bell, 1984: 160)

	Known	Ratified	Addressed
Addressee	X	X	X
Auditor	X	X	-
Overhearer	X	-	-
Eavesdropper	-	-	-

Although each listener is conceived as simply receiving the utterance of the speaker, none of the hearers' roles is passive[5]. "As in a theatre, the audience is the responsive, critical forum before whom the utterances are performed" (Bell, 1984: 161). And primary among this audience is the 2nd Person of the addressee. A change in the addressee, then, predicts a change in the linguistic style used by the speaker. For

[5] The role of eavesdropper is the noted exception, since this listener's presence is not known, so it cannot be anticipated. It can therefore have no effect on the speaker's linguistic choices. However, in the age of the internet and social media, this role is nevertheless important, and will be discussed again—if only tacitly—in Chapter 3.

instance, it is difficult to imagine that Carter would use the same emphasis and imagery (e.g., *AND EAT THE BEATING HEART OF ONE OF THEM*) if he had been addressing the children of the unnamed kindergarten he allegedly had in mind.

Harmon's (2008) insistence that a pledge to harm is distinct from a direct threat is premised, in large part, on the understanding of context framed by ADT. A direct threat construes an addressee who is very likely to be opposed to the threat's message, i.e., "resistant to the writer's primary argumentative position" (Martin & White, 2005: 125). The audience of a direct threat is highly unlikely to agree that they deserve to be injured. In a pledge to harm, however, the audience and the victim of the threatening language are not conflated this way. Instead, the two are separate, with a stance of disalignment not between a grammatical 1st and 2nd Person ('I' vs. 'you') but between a 1st and 3rd Person ('I' vs 'him/her/them'). There is, therefore, no presumption that the addressee of a pledge is automatically poised against the author's stance. Indeed, an addressee other than the threatened party could be construed by a pledge writer as non-resistant/neutral, or perhaps even as "compliant" (Martin & White, 2005: 62), i.e., as someone who is sympathetic to what the author has to say. Positioning a recipient this way would be a diametric departure from how direct threats are understood to operate.

Meanwhile, the threatened party in a pledge is relegated to the role of auditor at most and overhearer at least, with a correspondingly diminished impact on the writer's choice of language—and thus on the final composition of the pledge itself. For these reasons, the linguistic style of texts which threaten a third, outside party rather than the 2nd Person of the recipient cannot be assumed to follow the patterns discovered in more general studies of threatening communications (e.g., Gales, 2010, 2017; Smith, 2006).

Having isolated the speech act under investigation, certain texts may be judged to qualify as a pledge to harm where other threatening texts might not. However, isolating any linguistic markers which correlate with the presence of violent intent

requires a further subdivision of the resulting dataset, ideally into one corpus of pledges which contain violent intent and one corpus of pledges which were created for other purposes. This division requires concepts from the literature on threat assessment.

1.2.3 LEAKAGE

A central goal of any assessor faced with a communicated threat is to determine "whether the individual is making a threat, poses a threat, or is somewhere in the process of moving from one to the other" (Turner & Gelles, 2003: 93). In Fein and Vossekuil's (1998) understanding, the difference between communicating a threat (i.e., making) and taking real-world steps towards attacking a target (i.e., posing) is a matter of the threatener's intention. A person may make a threat, for instance, with the intent to cause fear in the target, but *without* the intent to proceed to physical violence. By drawing this distinction, Fein and Vossekuil (1998: 14) mean to clarify a misconception, namely that "[p]eople…often associate threateners with attackers, as if the two are the same." Disentangling the two behaviors offers an assessor more clarity on the nature of a situation. Carter, for instance, made a threat to an unspecified kindergarten. Shaw, on the other hand, both made and posed a threat to the Asian women in his community.

In Calhoun and Weston's (2015) more colorful reformulation of this typology, Shaw would be considered a 'hunter' while Carter was merely a 'howler'. "Hunters truly intend to commit an act of violence against whatever target they have selected" (Calhoun & Weston, 2015: 259). By contrast, howlers "communicate inappropriately by making threats or improper suggestions or requests, but they never advance beyond those inappropriate communications" (Calhoun & Weston, 2015: 259). These related conceptions suggest that it is possible, even expected, that an authentic threatening communication nevertheless may *not* include a true intent to harm. Interpreting these communications—to say nothing of adjudicating their meaning—is therefore closely in-

line with Durant's (2010: 138) understanding of 'validity', in that "appealing to a standard of 'validity' in interpretation brings in evaluative scales other than logical demonstration." Threatening language asks assessors, and courts on the back end, to evaluate "claims and expectations that go beyond truth claims" (Durant, 2010: 138) since, logically, claims about the future can be considered neither true nor false. The question, then, is whether statements of future harm are *valid*.

But, of course, hunters might not telegraph their intentions at all before the fact. "Persons who *pose* an actual threat often do not *make* threats, especially direct threats" (Fein & Vossekuil, 1998: 14). In fact, a direct threat "is arguably the least important [warning behavior] in a threat assessment, and least frequent" (Meloy & O'Toole, 2011: 515). Any communications by such individuals are far more likely to be to a third party rather than to the targets themselves. Fein and Vossekuil (1998) continue:

> Two-thirds of assassins and near-lethal approachers were known to have spoken or written in a manner suggesting that they were considering an attack. Would-be assassins told family members, friends, colleagues, and associates about their thoughts and plans, or they wrote down their ideas in journals and diaries. (Fein & Vossekuil, 1998: 15)

In the literature, sharing violent ideation with family, friends, etc., is known as 'leakage' (O'Toole, 2000: 14). Briefly defined, leakage "is the communication to a third party of an intent to do harm to a target" (Meloy & O'Toole, 2011: 514). This may occur through a variety of media, including "songs, drawings, doodles, [and] tattoos" (O'Toole, 2000: 14), as well as through such language-primary genres as "planned or spontaneous utterances…letters, diaries, emails, voice mails, blogs, journals, internet postings, tweets, text messages, video postings, and future means of social communication that are yet to be invented" (Meloy & O'Toole, 2011: 514). Thus, while leakage encompasses more social semiotic resources than just language, when the mode of transmission is linguistic the contours of the term overlay directly with Harmon's (2008) conception of a pledge to harm, i.e., as a speech act whereby "the speaker

commits himself to harm a third party" (p. 44). (Also, viewing leakage as a warning behavior that is categorically distinct from a directly communicated threat, as Meloy and O'Toole (2011) do, implicitly acknowledges the importance of Audience Design Theory concepts—that knowing the intended audience is a crucial component for understanding how an utterance should be construed.)

Hand-in-hand with the idea that 'making' and 'posing' a threat are different behaviors is the theory that threateners travel a pathway to violence (Meloy *et al.*, 2014). Some will leak their violent ideations to third parties and proceed no further down the path, i.e., they will *make* a threat. Others will communicate their intentions and then continue on, sometimes to the point of the attack itself, thus *posing* a threat. The value in pathway models is that they "describe a number of stages that an offender must go through before he attacks, and that this sequential structure can be used as a type of inferred ordinal scale" (Meloy *et al.*, 2014: 12). Calhoun and Weston (2015) offer one such model, wherein a hunter moves from:

- feeling a *grievance* to
- *developing the idea* that only violence can resolve their injury, to
- *researching and planning* the attack, to
- *making preparations* according to the dictates of the plan and the opportunities available, to
- *breaching* the target's security (however primitive or sophisticated that may be), and then to
- *attack*. (Calhoun & Weston, 2015: 259)

If, as noted earlier, it is possible for a 'howler' to make a threat without the real intent to pose a threat, then it is possible to view leakage as a kind of threshold along this pathway to attacking—a line which many threateners come to but do not cross. In other words, those aggrieved parties who howl but stop short of "furthering a plan or building capacity for a violent act" (Borum & Reddy, 2001: 380) may be theorized as producing texts which lack a true intent for future violence. This is a vital criterion for curating the competing subcorpora of pledges produced with violent intent and pledges

produced with other communicative goals in mind. Meloy *et al.* (2014) offer a more in-depth pathway model, one which emphasizes the coaction of internal states with subsequent behaviors. Importantly, this model (Figure 1.1 below) includes leakage (circled in red) as well as concrete signs that a threatener has moved beyond language and into real-world activities which Calhoun and Weston (2015) might begin to recognize as 'hunting.'

Figure 1.1: Pathway Model of Severe Targeted Violence in Schools (Meloy *et al.*, 2014: 11)

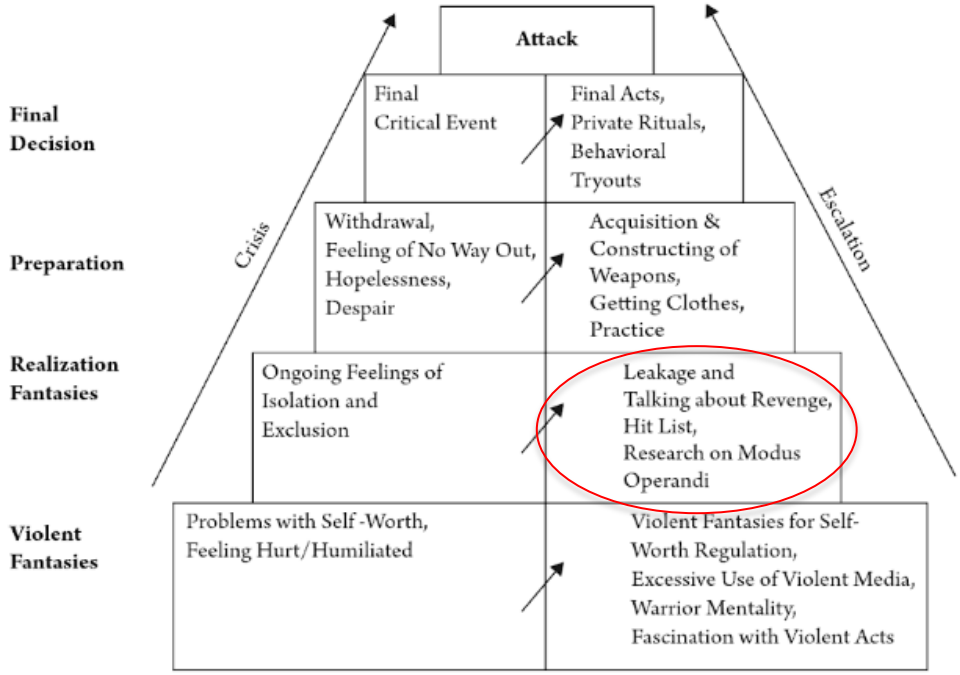

According to this, leakage emerges from the abstract, cognitive world of realization fantasies, but precedes concrete, real-world steps such as acquiring assault-related materials (clothes, weapons, etc.). At this point in the spiral of escalation the threatener is essentially dealing in *information*—both expressing it (e.g., talking to friends, compiling a hit list) and consuming it (e.g., researching the target). For the purposes of this study, then, the author's subsequent behavior is the only observable

indicator of whether his or her pledge was sincere or whether it was produced for some other kind of communicative effect. However, using Meloy *et al.*'s (2014) paradigm means that reported behavior need not include the attack itself. Preparation for the attack, much less a failed attempt at attacking, may be considered enough to classify a pledge as realized, since they signal that an author has moved beyond language. A pledge like Justin Carter's may thus be theorized as producing leakage which lacks true violent intent and be sorted into the 'non-realized' corpus. A text produced by someone like Tyrelle Shaw, on the other hand, would be considered realized, not only because he escalated to the attack itself, but also because he demonstrated preparation behaviors (e.g., acquiring a weapon). Shaw produced language that captured more than a mere fantasy of violence. Shaw's writing would thus be placed with the pledges which were somehow 'realized'.

Using an author's subsequent behavior as an indicator of intent is not a foolproof measure, of course. "Many persons may have been prevented or deterred from taking action because of a prompt response to their threatening communications" (Fein & Vossekuil, 1998: 15). Yet, as Carter's own arrest demonstrates, evidence collected after a threatener has been detained can clarify the author's intentions one way or the other, without the threatened action itself having been carried out. (Issues surrounding the classification of non-realized texts in particular will be explored more fully in Chapter 3.) That said, while the pathways offered by Calhoun and Weston (2015) and Meloy *et al.* (2014) may not contain surefire tests for classifying a text as either realized or non-realized, each still presents benefits. Not least, perhaps, is that legal authorities tasked with determining intent are also likely to be consulting these or similar models.

Unfortunately, it is at this point where the research on leakage stalls. Concrete language features are generally not identified as either risk-enhancing or risk-reducing. Instead, the literature focuses on the simple fact that a third-party communication has

occurred, its mode of transmission (e.g., verbal vs written), and roughly defined semantic 'themes' such as violence, hopelessness, despair, etc. (Meloy & O'Toole, 2011). Lexical items, e.g., "modals of intent" like *must* and *will* (Mardigian, 2008, via Gales, 2010: 26), are considered only rarely, while both systemic and structural elements are ignored almost completely.

To move beyond this point, any effort to discern violent intent in language must first identify the cognitive qualities that are being *expressed* via language—and thus which attitudinal resources an Appraisal analysis should be alert to. Doing so requires better isolating psychological intent, what it is composed of and how it is formed, and how this interrelates with the pragmatic intent inherent to language use.

1.4 OUTLINE OF THE DISSERTATION

The thesis proceeds as follows: Chapter 2 examines the existing literature on the overlapping but ultimately distinct natures of linguistic/pragmatic intent and psychological intent, as well as previous attempts to measure psychological intent's presence in threatening language. Chapter 3 expounds both the method of Appraisal analysis and an explanation of the dataset, its character and collection. Chapters 4 through 7 offer an Appraisal analysis of the data, moving from areas within the system of attitude through to an analysis using the system of engagement. Chapter 8 focuses on modal auxiliaries, a lexical class related to engagement but given its own analysis here due to its longstanding associated with expressions of psychological intent. Chapter 9 briefly examines the null results of the system of graduation. Chapter 10 applies the findings of Chapters 4 through 9 to an outside corpus of six additional pledges in a blind test format, in an attempt to identify which elements of the Appraisal analysis extrapolate more broadly. And finally, Chapter 11 concludes with a discussion of the possible implications

of this analysis on how intent may be encoded linguistically and thus for the detection of *mens rea* in pledges to harm.

CHAPTER 2 THEORIZING INTENT

Threat assessment practitioners Bulling and Scalora (2013: 4) argue that "[o]ne test of whether a threat is credible is the…intent of the entity posing the threat." As grounded in common sense as this may feel, this straightforward formulation nevertheless obscures the theoretical complications which abound in determining whether true violent intent is present in a threat. For example, what constitutes intent? Is intent monolithic or is it composed of more fundamental cognitive materials? Are there different kinds of intent and, if so, is each type constructed by the mind in the same way? Is each of equal strength? Most importantly, how many of these psychological qualities—if any—are detectable solely in the language of a pledge to harm?

These questions and others result largely from the fact that intentions, as internal mental states, are not directly observable (Malle & Knobe, 2001). They are only evident when they affect a person's behavior, including the person's language. In fact, this behavioral view is standard among U.S. legal professionals and reflected in the definition of intent found in Black's Law Dictionary: "being a state of mind, [intent] is rarely susceptible of direct proof, but must ordinarily be inferred from the facts" (Black, 1968: 947). Discerning intent as a means of assessing a threat's credibility, as Bulling and Scalora (2013) suggest, therefore involves two interlocking components: one is behavioral; one is psychological. To do their jobs effectively, threat assessors must first have a clear understanding of the psychological state being measured and, next, empirical evidence that particular behaviors, including language choices, are reliable indicators of this state.

However, this task is complicated by the fact that the term 'intent' is used differently—and not always interchangeably—across the legal, linguistic, and psychological literatures, areas which all intersect within the task of threat assessment.

The goal of this chapter is thus three-fold: first, to disambiguate the usage of 'intent'; second, to delineate which potential 'state of mind' is being tested for when intent is at issue; and third, to explore the current understanding of how this state of mind is—or is not—encoded linguistically.

2.1 LINGUISTIC INTENT

Compared with what will here be called psychological intent, or the "design, purpose, or determination with which a person acts" (Black, 1968: 947), linguistic intent is perhaps the more straightforward. However, the fact that the two are separate phenomena, serving purposes which are not always identical, can make analytical attempts to disentangle them problematic. This is especially true in texts where evidence suggests that the linguistic and psychological intents are somehow disaligned or out of sync, a possibility flagged by Solan and Tiersma (2005: 204) when they note that "to make a threat, the speaker does not *actually* have to be sincere, but need only *appear* sincere." In such cases, the linguistic intent necessary to perform a speech act like pledging may *mask* rather than reveal the psychological intent underlying the pledge's production, a situation which is central to the current research.

Linguistically, to say that a speaker meant something by an utterance is, in the Gricean account, to say that the speaker *intended* to be understood a certain way (Grice, 1957). As Searle (1969) summarizes:

> In speaking, I attempt to communicate certain things to my hearer by getting him to recognize my intention to communicate just those things. I achieve the intended effect on the hearer by getting him to recognize my intention to achieve that effect, and as soon as the hearer recognizes what it is my intention to achieve, it is in general achieved. He understands what I am saying as soon as he recognizes my intention in uttering what I utter as an intention to say that thing. (Searle, 1969: 43)

From the viewpoint of Speech Act Theory, then, linguistic intent is simply what a speaker hopes to accomplish by speaking, i.e., how he or she hopes to be understood.

Within the three-tier abstraction of communication discussed in Chapter 1—that of locutionary act, illocutionary force, and perlocutionary effect—intent therefore operates at the level of the illocution, whereby “we succeed in doing what we are trying to do by getting our audience to recognize what we are trying to do” (Searle, 1969: 47). This could be in the form of a promise a speaker sincerely plans to keep, a request he or she hopes the hearer will fulfill, etc. In this framework, intention is complementary to the semantics of the words and phrases themselves (Searle, 1965). Relying on their “general powers of rationality and inference” (Searle, 1975: 61), among other things, speakers thus use the social context they share with their listeners to layer additional, relevant meanings onto the plain semantics of an utterance. The exhortation to “kill him!” for instance, may be an incitement to murder if the audience is an angry mob, or merely a strongly stated opinion of disapproval if the speaker is a baseball fan dismayed with an umpire’s call (Abramson, 1978). Although the locution is identical, the speaker leading the mob and the speaker watching the game are using it to communicate wildly different propositional content—each *intends* the audience to understand two very different things.

Both to convey and to recover the intended meaning, speakers and audience rely on four conditions by which a speech act is judged to be “happy or felicitous” (Austin, 1962: 42), i.e., successfully performed. One of these is what Searle (1965) calls the sincerity condition, or simply that a speaker means what he or she says. In the case of a promise, for instance, this condition is satisfied if “the speaker intends to do the act promised” (Searle, 1965: 12). Of course, it is possible to speak insincerely. One may promise in “bad faith” (Austin, 1962: 11). But this does not negate the act itself, since “the performance of the [speech] act counts as an *expression* of [whatever] psychological state” is demanded by the sincerity condition (Searle, 1969: 65). To Searle (1969), then, an expression of thanks is understood as signaling a speaker’s

internal feelings of gratitude, whether the speaker is actually feeling thankful or not. "This law holds whether the act is sincere or insincere, that is *whether the speaker actually has the specified psychological state or not* [emphasis added]" (Searle, 1969: 65). In other words, a promise which the speaker has no intention to keep is still considered a felicitous promise. Similarly, uttering the words *thank you* can count as a successful act of appreciation, whether the speaker is truly grateful or not. The *linguistic* intent of the speech act—the speaker's hope that the listener understands a promise has been made or that thanks has been given—thus works independently of the *psychological* intent underlying the promising or the thanking. For social purposes, the speaking is considered evidence of the feeling.

In the case of threatening language, this independence has important theoretical and legal implications. Generally, threats are viewed as "a genre committed to violence and threatener control" (Gales, 2010: ii). For a threat to be judged felicitous, the utterance must communicate "the intention to perform the act" (Fraser, 1998: 162). Since a pledge differs from a direct threat only in its expected audience, there is no reason to believe that a different set of rules applies in this regard. However, Searle's "law" (1969: 65)—that the performance of the speech act counts as the expression of the associated attitude—allows for the production of felicitous threats which nevertheless lack the sincere psychological state of violent intent. More simply, this law allows for the phenomenon of *non-realized threats*, including non-realized pledges to harm. In such texts, the linguistic intent of communicating has a side effect not of revealing but of *obscuring* the psychological intent of the communicator, which is other than to perform the stated violence. Returning to the case of Justin Carter, his linguistic intent was a pledge to harm the children of a kindergarten. In this he was successful, and the felicitous production of this speech act triggered the legal problems which followed. But the psychological intent spurring the speech act was not to open fire on a kindergarten.

Rather, he was perhaps hoping to shock the sensibilities of the person with whom he was arguing. (The appeals court implicitly endorsed this distinction, in fact, when it called his writing "clearly hyperbolic and sarcastic" (Tex. Ct. App., 2015).)

An instance where the speech act is achieved yet somehow still false—i.e., "professed but hollow" (Austin, 1962: 18)—is characterized as an 'abuse' of the procedures. To Austin, a communication like Carter's "is perhaps misleading, probably deceitful and doubtless wrong, but it is not a lie or a misstatement" (Austin, 1962: 11). To better understand how and why linguistic intent may diverge from psychological intent, however, requires an examination of the mental state from which linguistic intent is diverging. As Searle (2004: 11) says, "[o]ur use of language is an expression of our more biologically fundamental mental capacities, and we will not fully understand the functioning of language until we see how it is grounded in our mental abilities."

2.2 PSYCHOLOGICAL INTENT

Intentionality is a "central facet of human cognition" (Malle, Moses, & Baldwin, 2001: 2). As such, it has been treated under a wide array of disciplines, approaches, and traditions. In developmental psychology, for instance, intentionality has been studied within the paradigm of 'theory of mind,' while in social psychology it is examined under 'attribution theory' (Malle *et al.*, 2001: 1). But other psychological approaches to intentionality abound, e.g., the so-called 'theory-theory,' simulation theory, etc. Beyond psychology, the issue has been grist for, among others, philosophers of mind, who have posited a range of occasionally competing frameworks to explain its workings, e.g., functionalism, behavioralism, eliminativism, etc. (Searle, 2004). Some of these theories begin their line of inquiry inside the intender's mind, attempting to understand intentionality as an internal mental state before it becomes visible as action. Others take

an external path, asking instead how people manage to detect this otherwise invisible mental state in the behavior of those around them.

Examining the merits of each and every theory across fields as vast as psychology and philosophy—and from both a cognitive and behavioral standpoint—is well beyond the scope of the current research. Thankfully, such an effort may also be unnecessary. This is because, despite the multiplicity of approaches, "intentionality's constituent components represent basic mental categories, such as belief, desire, and awareness" (Malle *et al.*, 2001: 1). So, rather than needing to wade through each and every '-ism' in turn, this brief foray into the psychology of intent can focus on the cognitive building blocks common to them all. Though, of course, further theorizing will be necessary. The definitions of belief and desire only become useful, for instance, when the interplay between them is conceptualized. For this task, the current research draws primarily on two compatible paradigms, Biological Naturalism and the Folk Concept of Intentionality:

- Biological Naturalism (BN): Searle's (2004) theory is used to clarify intent as an *internal* state preceding action. As one of the fathers of Speech Act Theory, Searle (2004) is ideally suited to identify the borders between a philosophy of language and a philosophy of mind, and how intentionality must be theorized differently in each area. As he himself notes, "we cannot explain the intentionality of the mind by saying it is just like the intentionality of language" (Searle, 2004: 160).
- Folk Concept of Intentionality (FCI): In turn, Malle's (1999) theory is used to understand the *external*, social aspects of intent, i.e., how intent is recognized by social perceivers. FCI has the added benefit of addressing the legal issues surrounding intent's detection in behavior. Malle and Knobe (1997), for instance, assert that jurors—and in many ways the law itself—employ folk concepts of

intentionality to assess responsibility and blame. More pertinently, FCI "promises clarity of *mens rea* concepts and a reconciliation of the legal and layperson's view of human behavior" (Malle & Nelson, 2003: 563).

2.2.1 COMPONENTS OF INTENT

In their enumeration of the conditions which must be met for a social actor to perceive a particular behavior as intentional, Malle *et al.* (2001) offer a tripartite model:

> First, intentions are directed at the intender's own action whereas desires can be directed at anything. Second, intentions are based on some amount of reasoning whereas desires are typically the input to such reasoning. Third, intentions come with a characteristic commitment to perform the intended action whereas desires do not. (Malle *et al.*, 2001: 4)

As this short schematic implies, an intention to do something is a complex rather than a simplex state. Different, more basic mental components are combined into a larger cognitive structure which is both formal (with the presence of desire and commitment) and processional (through the operation of reasoning). However, the first spark of what may eventually become an intention results from the friction between a person's beliefs and his or her desires. To believe that action A was intentional, for instance, "requires minimally that one grants the agent a *desire* for some outcome O and a *belief* that A will likely lead to O [emphasis added]" (Malle *et al.*, 2001: 4). Understanding psychological intent, then, begins with the delineation of these two states.

Belief and desire both put the thinker in a particular relation with the outside world, though they do so in two distinct ways. For belief, the aim is truth, i.e., to capture some state of affairs as it actually exists[6]. "[I]f I believe that it is raining," for instance, "my belief will be true if and only if it is raining" (Searle, 2004: 168). The matter is therefore one of cognition rather than volition. Desires, on the other hand, are held to be

[6] Whether the belief is objectively true is irrelevant; the stricture here is simply that it is not possible for an agent to knowingly hold a false belief.

neither true nor false. They "do not depict the way the world is; they only represent how the world should be (as judged by the agent's preference)" (Goldman, 2001: 212). To continue with Searle's (2004: 168) example, a desire for rain "will be satisfied or fulfilled if and only if it rains." The wish at the center of a desire is therefore characterized as volitional rather than strictly cognitive, because desires "are not supposed to represent how the world is, but how we would like it to be" (Searle, 2004: 168). Thus, within the framework of biological naturalism (BN), beliefs are the backdrop against which desires operate. The former describe the world while the latter aim to alter the world. This difference is called their *direction-of-fit*[7] with reality (Goldman, 2001; Searle, 2004).

Table 2.1: Cognition and Volition (Searle, 2004: 172)

	COGNITION			VOLITION		
	Perception	Memory	Belief	Intention in Action	Prior Intention	Desire
Causal Self-Reference	YES	YES	NO	YES	YES	NO
Direction of Fit	↓	↓	↓	↑	↑	↑
Direction of Causation	↑	↑	None	↓	↓	None

For a belief, the responsibility is to mirror something true about the world. Belief is thus said to have a *mind-to-world* direction of fit, a relationship Searle (2004) notates using a downward arrow (↓). In contrast, desire wants reality to change to match the content of the desire, a relationship characterized as fitting *world-to-mind*. This is notated using an upward arrow (↑). Searle (2004) describes a third option as well, what

[7] Helpfully, this concept is applied here identically to its use in Speech Act Theory, where direction-of-fit describes the relationship "between the propositional content of an illocution and the world to which the propositional content applies" (Rogers, Wall, & Murphy, 1977: 8). Instead of 'mind,' of course, the unit of measure in SAT is 'words'. An explanation, for example, is said to have a fit of *words-to-world*, while a promise is instead theorized as matching the *world-to-words*. Where the former attempts to map language onto a picture of reality, the latter attempts to influence reality using language.

he calls a null fit (notated as Ø). In basic semantic terms, this is directly akin to presuppositions, which "take for granted that the fit already exists" in one direction or the other (Searle, 2004: 169), such as is evident in the italicized complement clause of "I'm sorry that *I stepped on your foot.*" The formal relations between belief, desire, and other related mental states, such as perception and memory, are set out in Searle's (2004) full chart in Table 2.1 above.

Interestingly, in BN beliefs and desires are both considered inert, i.e., neither has a direction of causation. They may fit the world in one direction or another but by themselves they do not *change* the world. The implications of this for discerning intent in a pledge-to-harm are two-fold. To start, an intention is more than simply a stated belief about the world—for instance, the belief that an assault will occur—though beliefs are the bed from which intent may grow. But neither is intention simply a desire. Threatening language which only takes the form of belief or desire may therefore be hypothesized as lacking other necessary components of true violent intent, such as a world-directed fit of causation. In other words, wanting to change the world is not the same as actually intending to change it. This kind of language is bolded in Text 2.1 below, in an online comment written in 2017.

Text 2.1: Nikolas Cruz Pledge
Im going to watch them sheep fall fuck antifa i **wish** to kill as many as i can

While Cruz did, in fact, proceed to violence, his eventual targets were his classmates—not the 'antifa' protesters he desired harming. However, of the two desires rightly seem closer to the threshold of causation than beliefs. For Searle (2004), the question is one of mode and content. To return to his previous example: believing that it will rain and desiring that it will rain both have the same propositional content, that *it will rain*. They differ simply in their attitude towards this content, i.e., the psychological mode with which they relate to the proposition. Formally, "[w]e can represent this as S(p),

where the S stands for the mode or type of state and the p for the propositional content. Such states are often called 'propositional attitudes'" (Searle, 2004: 166-167). Where a belief is roughly neutral, representing the 'fact' of the content, a desire views the content as somehow favorable or necessary, thereby adopting what Davidson (2001: 3) calls a "pro attitude" towards it. Interestingly, "[t]his distinction" between content and attitude "exactly parallels the same distinction in language" (Searle, 2004: 166), where it is theorized as the difference between the illocutionary *force* and the illocutionary *point* of an utterance. For example, to communicate the illocutionary point of a directive, e.g., 'you leave the room,' a speaker has available a range of modes, or forces, to choose from, such as ordering ("Leave!"), requesting ("Please leave"), etc. The same propositional content (p) may thus be conveyed by different modes (S). Linguistically, this kind of attitude signaling has also been examined under the broader auspices of 'stance' (Biber, 2006; Du Bois, 2007; Jaffe, 2009; Precht, 2003), which is the process of "taking up a position with respect to the form or content of one's utterance" (Jaffe, 2009: 3). In the case of a 'pro attitude', the underlying desire would serve to align the stancetaker *with*, rather than *against*, the propositional content. The intuition that a desire is closer to action than a belief is accurate within both BN and FCI because of the added element of stance alignment involved.

However, bridging the gap between desire and intent also requires examining the content of the proposition. For desires, "many different types of content" may be mentally represented (Malle & Knobe, 2001: 47). There is no proscription against a person desiring something beyond their control. One may desire world peace, for instance, and yet "a desire for world peace does not directly cause the person to perform a particular action" (Malle & Knobe, 2001: 51). The desire itself is invalidated neither by the infeasibility of its own content nor by whoever is responsible for bringing about the desired state of affairs. However, "genuine intentions" are restricted to what the authors

call “action content” (Malle & Knobe, 2001: 48). “A pro attitude has action content when the content of the attitude is an action performed *by the same person who holds that attitude* [emphasis added]” (Malle & Knobe, 2001: 48). This formulation has obvious and direct echoes with the structure of a pledge in SAT. Roughly put, intent is a desire that the desirer imagines fulfilling personally.

A brief summary of the discussion up to this point may be useful. First, both beliefs and desires may have many kinds of content. A desire, however, represents a positive psychological mode towards the content in question, i.e., a positive emotional valence towards the represented material. If the propositional content (p) is violent, then the thinker’s stance or mode (S) toward the imagined event—S(p) in Searle’s (2004) nomenclature—might be represented as D(v), or a desire for violence. Finally, intention, like desire, also features a pro-attitude. But where a desire may hold many kinds of propositional content, an intention is limited to content that is personally actionable by the thinker. These basic contrasts are captured in Table 2.2.

Table 2.2: Contrasts among Belief, Desire, and Intention

	Belief	Desire	Intention
Unrestricted content	X	X	-
Pro-attitude	-	X	X
Restricted (action) content	-	-	X

In the search for where a realized pledge may diverge from a non-realized pledge, then, the question of action content is a strong candidate for further investigation. Yet, how does action content differ from the less restricted representational material of desires and belief? For Searle (2004), the answer involves the intention’s ‘conditions of satisfaction,’ or those things which must come to pass if the intention is to be fulfilled.

> What makes my desire a desire to drink water is that it will be satisfied if and only if I drink water. That is not a psychological remark predicting what will make me feel good, but rather it is the definition of the relevant intentional content. (Searle, 2004: 189)

Or, put more succinctly, “intentionality is representation of conditions of satisfaction” (Searle, 2004: 173). However, if a speaker’s intention is to use the words “I will drink water” only to convince someone that he or she will hydrate in the future, then this condition is satisfied not in the drinking but simply in the speaking. Hypothetically, then, the difference between a realized and a non-realized text is not that a non-realized pledge lacks the action content found in “genuine intentions” (Malle & Knobe, 2001: 48). Rather, *the conditions of satisfaction for a non-realized pledge are either partially or entirely fulfilled by the production of the pledge itself, rather than through actions external to the pledge*.

If this hypothesis is accurate, then a realized text might best be theorized as an expression of the conditions of satisfaction themselves. Exploring Searle’s (2004) scenario further, a person who says, “I will drink water,” while harboring a genuine intention to fetch a glass of water would be verbalizing his or her “relevant intentional content” (Searle, 2004: 189). The condition of satisfaction is directly externalized via language: that personally desiring a drink of water requires personally taking a drink of water. On the other hand, if a speaker uses the same commissive statement to reassure someone that he or she will drink water *but has no psychological intention to actually do so* then the person has issued an ‘insincere promise’ in Austin’s (1962: 18) typology. Here, the response evoked in the listener is the point, rather than any subsequent, real-world action. The conditions of satisfaction themselves are not explicitly encoded; instead, they are fulfilled by the perlocutionary effect of the speech act. If the intent was to reassure a caretaker, and he or she has been reassured, then the intention is realized. Yet actual, physical water plays no part in this second scenario, despite the smokescreen of linguistic intent inherent to the speech act of promising.

The possibility of this kind of bifurcation is highlighted by the linguist Fraser (1998), who points out that threats are produced for a wide variety of reasons: to

intimidate, to challenge authority, to vent anger, to get attention, or even just to be humorous. In this array of motivations, conveying a sincere determination to injure someone in an unlawful manner (Black, 1968) is just one potential goal among many spurring the production of threatening language, including a pledge to harm.

Of course, it is still possible for an agent to hold a pro attitude toward action content—imagining a desired outcome that he or she is personally responsible for—yet fail to act. For instance, in Malle and Knobe's (2001: 46) colorful sketch, "an agent might have a desire to start screaming at her boss even though she has specifically decided not to do so." In all likelihood, the agent's self-censorship is the result of having reasoned through the ramifications of the desired action. Reasoning is thus the penultimate step in the process of "actually deciding to perform the action in question" (Malle & Knobe, 2001: 46) and the next component in Malle *et al.*'s (2001) tripartite model.

2.2.2 REASONING

For Searle (2004), beliefs about the world do not exist in isolation. Rather, each is a node within a larger network of knowledge:

> If I believe I own a car, I must also believe that cars are modes of transportation, that they are used on streets and highways, that they move about, that people can get in and out of cars, that cars are a kind of property that can be bought and sold, and so on. (Searle, 2004: 172-173).

This web of beliefs "is in fact the structure of our conscious life" (Searle, 2004: 174) and forms the cognitive background which gives desires their purpose. A person who wants the world to be different must first *believe* something about the way the world is now. According to FCI, a question naturally arises from the emotional friction created by a desire: can the desire be fulfilled? This may then trigger a series of mental calculations for the agent. "For example, a desire to eat chocolate is not based on any

reasoning, but it may set into motion a reasoning process about how to acquire some chocolate" (Malle & Knobe, 2001: 51). This reasoning process is informed by the agent's network of beliefs, including what the person knows about his or her own "set of abilities, ways of coping with the world, dispositions, and capacities generally" (Searle, 2004: 173), e.g., beliefs that chocolate is edible, that it may be purchased at a store, that the agent is capable of eating it safely, etc.

The echoes between this and the pathway models in the threat assessment literature discussed in Chapter 1 are plain. A person who feels a grievance—the first step in Calhoun and Weston's (2015) journey to targeted violence—*believes* that he or she has been wronged and forms a *desire* to right the perceived injustice. In deciding whether and what to do to restore emotional balance, the threatener "develop[s] the idea that only violence can resolve their injury" (Calhoun & Weston, 2015). Thus, on the path that leads to action "desires stand in the very beginning of the process" (Malle & Knobe, 2001: 46) and function as inputs to reasoning, while "intentions typically function as reasoning output" (Malle & Knobe, 2001: 65). Reasoning is therefore crucial to the formation of intention.

The process is essentially computational, involving the appraisal of sometimes contradictory variables. Like beliefs, desires do not exist in isolation. Unlike beliefs, however, desires may conflict and compete. A desire for chocolate, to use Malle and Knobe's (2001) previous example, may be complicated by, say, a desire to stick to one's diet. "Before making a decision about how to act, the person needs to consider various desires, balancing them against each other and asking which of them can potentially be fulfilled" (Malle & Knobe, 2001: 46). This kind of balancing is evident in Text 2.2, excerpted from a longer pledge to harm which was emailed by Gavin Long to an acquaintance before he shot and killed three policemen in Louisiana.

Text 2.2: Gavin Long Email Pledge (Excerpt)
I know I will be vilified by the media & police, unfortunately, I see my action's as a necessary evil that I do not wish to partake in, nor do I enjoy partaking in, but must partake in, in order to create substantial change within America's police forces and Judicial system.

Here, Long weighs his desire for reform within the American justice system against a second desire to maintain his good name (*vilified*) and a third desire not to personally commit actions which he judges as *evil*. He then briefly explains his reasons for why violence has won out (*necessary*, *must*). This text thus presents evidence, condensed though it may be, that he has worked through the reasoning process by balancing at least three competing desires and arrived at an intention. In this way, Long's writing exemplifies the rudimentary decision-making process described by Malle and Knobe (2001), i.e., that a pro attitude with action content does not rise to the level of intentionality by itself, but only becomes a "full-fledged intention" after the agent answers "(a) whether she is capable of performing the action and (b) whether she has other desires that outweigh her desire to perform the action" (Malle & Knobe, 2001: 55).

This basic calculus is recognized in the threat assessment literature as well. Consideration (a) above is equivalent to what Geurts *et al.* (2016: 55) call 'feasibility,' a variable which "refers to the ease or difficulty of reaching the end state" desired by the agent. Consideration (b), on the other hand, corresponds to 'desirability,' or how positively or negatively a person views the imagined end state, i.e., its valence (Geurts *et al.*, 2016: 55). The difference between people who intend to act and people who are threatening for some other purpose (e.g., to get attention) is that the two "value the *desirability* and *feasibility* of their threat differently" (Geurts *et al.*, 2016: 55). High desirability is evident in Long's email, for instance, through the strong deontic modality of *necessary* and *must*. High feasibility, meanwhile, is indicated by his use of presupposition. He presents his future behavior as taken for granted when he writes *I*

see my action's as…. The violence itself is a given, despite the fact that Long wrote this pledge *before* the shootings. All that Long leaves open to argument is their social value.

For contrast, reasoning which leads an author away from violence is excerpted from a message posted to social media by Ebony Dickens, shown as Text 2.3.

> Text 2.3: Ebony Dickens Facebook Pledge (Excerpt)
> I've thought about shooting every white cop I see in the head until I'm either caught by the police or killed by them. Ha!!!! I think I can pull it off. Might kill at least fifteen tomorrow.

Unlike Long, Dickens's desire to shoot police officers is framed as something she has *thought about* and *might* do rather than as something which she views as necessary or required. This low- or mid-level desirability is then weighed against the two competing desires of preserving both her freedom (*caught*) and her life (*killed*). Through shared world knowledge, a reader would know that the desires for staying both free and alive are very strong indeed, needing little to no extra communicative emphasis or explanation. Along with this, the feasibility of succeeding is presented as less than high through the hedge *I think* and the epistemic modal *can*, an auxiliary which communicates ability but neither inclination (e.g., *will*) nor obligation (e.g., *must*). Finally, the first iteration of the threatening speech act is closed off with the semantically thin but pragmatically loaded exclamation *Ha!!!!*, a tag which does not indicate humor but rather, perhaps, a level of disbelief which might colloquially be glossed as 'could you imagine?' Taken altogether, the low- to mid-level desirability and—at best—middling feasibility are not presented as outweighing the severe consequences of the imagined act (*caught…or killed*). In fact, Dickens's lack of true violent intent was eventually recognized by prosecutors, who dropped a felony charge against her in exchange for a public apology (Torpy, 2016).

If "intentions serve to fulfill desires by identifying a course of action that is feasible to implement for the agent and is compatible with the agent's other desires"

(Malle & Knobe, 2001: 55-56) then reasoning may be characterized as a process of *stancetaking*. For example, an agent recognizes a desire for chocolate—call it D(c)—and subsequently reasons towards a more nuanced stance with regards to obtaining chocolate: is it feasible? does it conflict with other existing desires? etc. As the excerpts from both Long's and Dickens's texts show, this process can indeed be visible in the language of a pledge to harm itself.

Worth noting, the psychological literature reviewed here implicitly argues that the presence of a desire to harm a third party is not equal to true violent intent. Yet, as Dickens's experience illustrates, neither is evidence that an agent has reasoned through the feasibility of committing violence. (Of course, the expression of such desires and such reasoning may be interpreted as threatening and is therefore legally perilous by itself, as Dickens discovered.) For an intention to be judged as "genuine" (Malle & Knobe, 2001: 51), the final test is for the presence of 'commitment.'

2.2.3 COMMITMENT

Dickens's pledge shows that people are capable of walking the pathway to violent intent without actually *arriving* at violent intent. Evidence of reasoning, by itself, is no guarantee that a particular intention has been formed (Malle & Knobe, 2001). Unlike desires, beliefs, or the reasoning process spurred by the conflict of the two, "intentions involve commitments to act" (Moses, 2001: 72). This means that "a third criterion is necessary" (Malle & Knobe, 2001: 56) and this is posited as the addition of 'commitment.' From a threat assessor's standpoint, this final element in in the process is the most important. From a linguistic standpoint, it is unfortunately the most problematic. This is because, of the three stages, language's role in expressing commitment is the most unsure.

The character of commitment as a mental state is one which psychologists often gloss as ‘purposive’ (Gibbs, 2001: 107). It is the ‘settled’ quality which concludes the reasoning process, when a thinker has weighed his or her competing desires and made a decision about how to proceed (Malle & Knobe, 2001: 56). Arguably, detecting the presence of this state is the core of a threat assessor’s responsibility. The very process of behavioral threat assessment, for instance, is “based on the idea that the subject will exhibit identifiable behaviors indicating *the intent to commit* an act of targeted of violence [emphasis added]” (Bulling & Scalora, 2013: 2). Malle and Nelson (2003) posit that this element is a crucial consideration at the other end of the legal process as well, when judges and jurors are asked to weigh questions of *mens rea*. This is because intentions necessarily “involve signs of commitment” (Malle & Nelson, 2003: 571).

This distinction—between a speaker’s intent and his or her level of commitment to the represented action—is also mirrored in the linguistic literature. Fraser (1998: 161), for instance, remarks that it is possible for a speaker to express “an intention though not a commitment to perform an act.” Gales (2010) notes that lexicogrammatical resources can be used by speakers to “demonstrate the stancetaker’s commitment to the mentioned proposition” (p.41) and that these linguistic choices “are oftentimes used by law enforcement practitioners to help determine a threatener’s commitment to carrying out their threatened action” (p.47). As in biological naturalism (Searle, 2004), then, a proposition is seen as separable from the attitude or stance a speaker takes towards it, with the “commitment to carry out the threat” (Gales, 2010: 80) being just one in a range of possible stances. In all areas, commitment is viewed as the final psychological step before real-world action.

However, detecting commitment is problematic. Despite how crucial this mental state is to predicting future behavior, outside observers “cannot directly perceive commitment” (Malle & Knobe, 2001: 56), a sentiment echoed by linguists like, e.g.,

Bojsen-Møller, Auken, Devitt and Christensen (2020) and Muschalik (2018). Instead, social perceivers rely on a set of three behavioral cues to gauge the "degree to which the agent is committed to the represented action" (Malle & Knobe, 2001: 60). If these cues are not apparent, in whole or in part, then the attitude is more likely to be classified not as an intention but merely as a desire (Malle & Knobe, 2001: 60). These three indicators are:

1. making early investments (e.g., buying concert tickets early);
2. accepting opportunity costs (of not choosing alternate options); and
3. a willingness to accept sanctions by announcing intentions publicly.

Malle and Knobe (2001) use the example of a person's commitment to his or her romantic partner. An 'early investment' behavior might be joint purchases made by the couple, while 'accepting opportunity costs' would be forgoing other romantic relationships, and, finally, inviting sanctions might be seen in the introduction of the partner to friends or family (Malle & Knobe, 2001: 57). In terms of extra-linguistic behavior—e.g., buying concert tickets—these criteria make sense. For threat assessors, it is easy enough to argue that a threatener who buys a gun is making an early investment and has therefore demonstrated a level of commitment to the act. Similarly, a threatener who engages in stalking behaviors could be said to be accepting opportunity costs, in terms of the time and energy required to surveil a target.

However, it is difficult to see how these first two criteria may be applied to *linguistic* signals. For example, protestations of love early into a new relationship, to explore Malle and Knobe's (2001) scenario, might indicate both an early investment behavior (indicator 1) and the acceptance of opportunity costs (indicator 2). But in reality, such protestations work as public announcements (indicator 3), in that they imply the speaker's intent to stay with the new romantic partner, and are thus better classified as inviting hearer-imposed sanctions. In general, *every* sample speech act seems to

veer into this third indicator category, and Malle and Knobe (2001) do not offer language examples of their own to act as counterpoints or from which counterpoints may be extrapolated. Thus, when the question is one of language, the first two metrics (early investment; opportunity cost) appear to be unavoidably subsumed by the third (accepting sanctions). For social perceivers looking for clues to commitment in a person's language, the announcement itself would appear to be the only available indicator. This is problematic.

One potential approach to social sanction is through the framework of Politeness Theory (Brown & Levinson, 1987). When agents "put [their] credibility on the line" through "public announcements of intention" (Malle & Knobe, 2001: 57), they engage in a threat to their own positive face, or their "want to be appreciated and approved of by others" (Mao, 1994: 455). In short, they risk rejection from whichever audience they are addressing, however narrowly (e.g., an email to a friend) or broadly (e.g., a public social media post) this audience is construed. This applies to pledge writers as well. By publicly claiming that they will do something violent, pledge authors invite potential ridicule if they fail to act, and thus a diminishment of their status. Simply hitting 'send' or 'publish' is, in this model, a sign of commitment.

The obvious problem with using this as a measure of true malicious intent is that public claims of future violence are made by authors in both realized and non-realized pledge categories. By publicly expressing his desires, Justin Carter accepted just as much risk to his positive face as Tyrelle Shaw, and Ebony Dickens just as much potential loss of face as Gavin Long, yet only Shaw and Long proceeded to real-world violence. Social perceivers attempting to gauge the degree of commitment in each of these four texts through Malle *et al.*'s (2001) model—that the announcement itself is a major sign of commitment—would be forced to conclude that all four authors harbored true violent intent. Yet, in two of the four cases this determination is evidently wrong,

with outsized legal consequences awaiting the misread pledge writers. Gales (2017: 9-10) underlines this danger, writing that "when disputes in the interpretation of meaning arise in forensic contexts, such as those of threatening language, consequences resulting from different intuitions about language use may be more significant."

While the expression of a threat is indeed *socially* significant, Carter's and Dickens's cases demonstrate that its mere existence cannot be the one and only sign of a commitment to act. Beyond the fact of the pledge, then, what is left is the language of the pledge itself. Indeed, "the linguistic information" of a threat is often "the *only* element available for assessment" (Gales, 2010: 2). And anyway, social perceivers do not make such simplistic calculations, mistaking the mere existence of a threat with an actual commitment to act. Readers routinely base judgments about authorial intent on a text's actual content—the words on the page—just as prosecutors eventually did with Dickens and the court eventually did with Carter.

To better understand what in the language of a pledge may signal true malicious intent, the next section moves from the internal world of psychology to the external world of linguistics, reviewing previous studies into how intentions are potentially encoded in language.

2.3 LANGUAGE OF INTENT

Rather than parsing specifically for violent intent, the foundational linguistic literature on threatening is aimed primarily at determining whether a threatening speech act has even occurred. Yamanaka (1995: 37), for instance, concentrates on "what sort of utterances should count as threats as distinguished from warnings, orders, etc." Similarly, Fraser (1998: 169) observes that "a threat need not be in any particular form or phrasing" and that "it may be issued by suggestion or innuendo," which leads him to his main question: "what do we use to ground the conclusion that a serious threat was made?" (Fraser,

1998: 169). Of course, clarifying this bedrock issue—whether locution X qualifies as a threat—has important practical applications, as Shuy (1993) demonstrated in the case of Don Tyner, a businessman who was accused of threatening a former employee's son when he asked, "how's David?" (Shuy, 1993: 109). Shuy's analysis helped in Tyner's eventual acquittal, but the case also lives as a touchstone for the pragmatic issues surrounding threat-making.

However, the current research picks up *after* this crisis point has been passed. From a linguistic perspective, there is no question that a text like Justin Carter's or Ebony Dickens's is a felicitously constructed pledge to harm, leaving the authors open to possible criminal indictment whether they 'meant' to perform the stated violence or not—as both individuals eventually discovered. But the number of studies which begin here—exploring empirical differences not between threats and more innocent forms of writing, but between "threat internal categories (i.e., realized vs. non-realized threats)" (Gales, 2010: 109)—are comparatively few in number. Most of the recent, linguistically-grounded work still seeks to expand on the foundations laid by Shuy (1993), Storey (1995), Yamanaka (1995), Fraser (1998), and others, to refine our understanding of "what threatening language actually is" (Gales, 2017: 2).

This is not to say that investigations into the nature of threatening have no guidance to offer the current study, however. There are several potential language features highlighted by these researchers which might bear on pledging and violent intent, primarily because such studies share certain fundamental premises with the current work. For instance, despite their individual aims, Bojsen-Møller *et al.* (2020), Muschalik (2018), and Nini (2017) all explore the interconnectedness of *formal* realizations with the textual *function* of threatening texts. This approach is consistent with the hypothesis put forward in Chapter 1, i.e., that different systemic resources found in the pledges (their form) are potentially revealing of different communicative aims (their

function). As part of his analysis of register variation in what he calls 'malicious forensic texts' or MFTs, for example, Nini (2017) examines "texts in which the harm was directed to a third party"—i.e., pledges to harm, though he does not use this term. He notes what might at first seem to be an obvious generic distinction, namely that these threats showed a "higher frequency of third person pronouns" (Nini, 2017: 117). However, the larger trend Nini (2017) uncovers is potentially instructive. Pledges in his dataset were distinguished from other MFTs—including direct threats—by their narrative concerns. In other words, the "authors of the texts in which the harmful content is directed towards a third party are likely to recount events to the addressee regarding the third party to whom the threat or abuse is addressed" (Nini, 2017: 117). This finding is especially interesting in light of the attention paid to planning in the writings of threat assessors like Borum and Reddy (2001), Calhoun and Weston (2015), etc., as well as to the supposed centrality of a grievance in spurring violent ideation (e.g., Calhoun & Weston, 2015). Both the plan for an attack and the origin of a grievance could conceivably be communicated in narrative form. The 'story' of a pledge therefore represents a textual dimension worth being attuned to in the analyses which follow[8].

To use the terminology of Bojsen-Møller *et al.* (2020) and Muschalik (2018), the presence of narrative forms might be considered a 'recurrent trait' of pledging. In both studies, recurrent structures are a means of divining whether "speakers always threaten in the same way and for the same reasons" (Muschalik, 2018: 7). Approaching recurrence from the perspective of genre studies, Bojsen-Møller *et al.* (2020) note that the "three crucial aspects of a threat are…futurity, harm, and sender responsibility" (p.16) but that these features may be purposely occluded "because of a greater need for tactical indirectness and silence than for explicitness" (p.38). This invites an

[8] Narrative in the pledges is examined in section 6.1.1 analyzing negative composition.

examination of the stances pledge authors take towards futurity, harm, and responsibility in the current dataset, and whether these stances differ by realization category. Along with this, however, Bojsen-Møller *et al.* (2020), Solan and Tiersma (2005), and Fraser (1998) all raise an interesting functional question, which perhaps points to a deeper asymmetry between pledges and direct threats. Each argues that threatening always involves a speaker's intent to intimidate the addressee. While this may be true of direct threats, the discussion of ADT in section 1.2.2 above throws this functional requirement into sharp theoretical relief. While pledges may indeed be used to intimidate, as the LA Unified text demonstrates—e.g., *It is time to pray to allah, as this may be your last day*—it is an open question whether leakage *always* involves the goal of intimidation. It is decidedly unclear, for instance, that Shaw's blog posts were intended to scare the Asian women he imagined harming. Thus, whether intimidation is also a function of pledges—either primarily or secondarily—and whether it correlates with other potential functions of pledging are issues worth considering.

However, questions of function are parsed most thoroughly by Muschalik (2018), who considers a threat's communicative purpose central to the recurrent structures she finds in her dataset. Importantly for the sake of comparison, her hypothesis that "if the form is somehow determined by the function, it should be possible to predict the function based on the form" (p.183) is very similar to the one pursued here. Although her focus remains on pragmatic intent—i.e., the social work a threat is performing—rather than on the real-world intent of the threatener, her major interpretive contribution to the linguistic literature on threatening is both illuminating and useful. Muschalik (2018) observes that features like conditionality seem to cluster according to whether a threat is uttered "in *prospect* of an action" that the threatener is hoping to incite or prevent, or "in *retrospect* of events that have somehow negatively affected the target" (Muschalik, 2018: 183). The former she terms *manipulative* threats and the latter *retaliative* threats. Whether or

not this classification proves relevant for discerning psychological intent, the precedent set here—that a data-driven taxonomy is supportable from a theoretical standpoint—will prove decidedly useful in Chapter 11, when the clusters of Appraisal features uncovered in Chapters 4 through 9 are holistically examined against the backdrop of the pledges' realization statuses. Finally, Muschalik (2018: 183) concludes that "the forms speakers use to realize their threats are more predictable than it has been argued in the previous studies." The idea that threateners repeatedly choose certain language structures based on the intended social purpose of the threat lends further credence to the hypothesis explored here. For if the communicative aims of realized and non-realized pledges are indeed divergent, then formal differences should also be apparent. And if significant formal differences are discovered, and their communicative function is understood, then both may lead back to the underlying psychological intents spurring each pledge type's production.

That said, studies which *directly* examine psychological intent as it may manifest through language are much rarer. Threat assessors naturally take great care with this question (e.g., Hoffmann & Meloy, 2014), and generally believe that a threatener's language can reveal violent intentions, as when Simons and Tunkel (2014: 207) say that "the offender's *resolution to commit violence* will often become manifest in detectable ways and may be evident in the offender's language." However, almost none of the threat assessment writings are expressly informed by linguistic science. Only one extensive investigation is known to address the difference between realized and non-realized threats with the goal of discerning violent intent, that of Smith (2006). Yet, she does so in reverse procedure from this thesis, using linguistic theory to inform her psychological insights rather than vice versa. Furthermore, like Gales (2010), Smith (2006) considers threats in the broadest sense, making no categorical distinction based

on who is being threatened, i.e., targets in the grammatical 2nd and 3rd Persons are treated equally.

Nevertheless, useful contributions to the current undertaking are found spread across all of these related efforts. What follows will first highlight previous findings relevant to the speech act of pledging (e.g., the language of leakage) before expanding the focus to include discoveries in the language of threats more broadly—from studies like those of Smith (2006) and Gales (2010). The overarching goal is to extract from these literatures any linguistic features and/or constructions which have some level of predictive power, i.e., lexicogrammatical choices which may subsequently be tested for their correlation with one pledge category or the other, whether realized or non-realized.

2.3.1 LANGUAGE OF THREATENING

A 'pledge to harm' as a particular kind of threatening language was first fully defined by Harmon (2008). She singles out pledging as a speech act in order to address the legal pitfalls surrounding what talk therapists call "*Tarasoff*-type situations" (Borum & Reddy, 2001: 375). The case from which this term derives, *Tarasoff v. Regents of the University of California* (1976), is an instructive example of leakage. In 1969, a woman named Tatiana Tarasoff was killed by a student at the University of California, Prosenjit Poddar. Poddar, who harbored a love interest for the woman but had been rebuffed, began to imagine hurting her. He shared his violent ideations with a therapist employed by the school, who was alarmed enough by Poddar's language to contact campus police. During the subsequent interview, Poddar managed to convince the officers that he was rational, in part by promising to stay away from the woman. He was released and no further action was taken by authorities. Two months later, however, Tarasoff returned from a vacation abroad and Poddar followed through on his ideations, stabbing her to death. The woman's parents then sued the Regents of the school, arguing that the

family, or Tarasoff herself, should have been notified by the therapists of the danger Poddar posed. The Supreme Court of California agreed.

The *Tarasoff* case "imposed upon mental health professionals an affirmative duty to warn third parties who are the subject of credible threats uttered by their patients" (Harmon, 2008: 30). These so-called 'duty to warn' or 'duty to protect' laws essentially place therapists in the unlikely role of threat assessors. And while a number of clinical tools have been developed in the forty-plus years since *Tarasoff* to better assist in gauging a patient's risk of violence, such as the HCR-20, COVR, and VRAG, "[n]one of these instruments consider linguistic factors" (Glasgow & Schouten, 2014: 39). As Harmon (2008: 31) notes, a "discussion of exactly what facts are pertinent to the deliberation of whether a threat of the requisite nature has been made, thereby triggering the duty [to warn], is rare."

To be clear, Harmon's (2008) attempt to close this gap does not include the divination of a patient's violent intent. "From the law's standpoint [regarding *Tarasoff*], in a strict sense, whether the speaker *actually intends* the harm is irrelevant" (Harmon, 2008: 45-46). Echoing the therapists whom she hopes to assist, Harmon (2008) rejects the frame that *predicting violence* is a reasonable responsibility to place on mental health professionals. Instead, her focus is "to recognize the locution for what it is: a threat/pledge" (Harmon, 2008: 85). This, of course, is not so different than the goal of Shuy (1993), Fraser (1998), Muschalik (Gales (2017), and others to identify "what threatening language actually is" (Gales, 2017: 2) and whether a threat has even occurred. Nevertheless, Harmon's (2008) program provides an entry point for identifying concrete—and thus testable—linguistic features which may correlate with either a realized or a non-realized pledge.

Her proposal builds on the four felicity conditions put forward by Searle (1969). These are, in Harmon's (2008) words (excluding her examples):

1. Preparatory Conditions: condition(s) that precede the utterance;
2. Sincerity Conditions: conditions that relate to the speaker's state of mind;
3. Essential Conditions: conditions that require the utterance be recognizable as the type of illocutionary act in question;
4. Propositional Content Conditions: conditions that relate to the proper context of the statement. (Harmon, 2008: 42-43)

Harmon (2008: 91) notes that the resulting list of features is "certainly not exhaustive" and that "[o]ne feature might be more important than another" (Harmon, 2008: 91). Both of these equivocations are aspects this research aims to explore. The conditions are treated in the order they are presented.

Preparatory condition: For a pledge to meet this condition requires the completion of two parts: "The speaker would obtain satisfaction from something detrimental to the victim *and* has the capability to accomplish that detriment" (Harmon, 2008: 87). In other words, this condition is focused "on the assessment of the speaker himself; that is, in determining that the speaker can do what is intended" (Harmon, 2008: 88). While Harmon (2008) treats each requirement in a biographical sense—asking who the speaker is and what he or she is feeling—each of the two parts has potential analogs in language.

The first part of this condition, that a speaker would obtain satisfaction from the pledged act, is considered a major, if not primary, motivating factor of targeted violence. By proceeding to action, a person risks "his or her reputation, finances, freedom, and life itself" (Simons & Tunkel, 2014: 206). The reward offsetting this momentous risk is the "emotional profit of redress or feelings of satisfaction not attainable through socially acceptable or legally sanctioned means" (Simons & Tunkel, 2014: 206-207). This would suggest, as Harmon (2008) intimates, that an expression of imagined satisfaction in a pledge to harm should be considered a risk-enhancing factor, and thus more likely to correlate with realized texts.

Harmon (2008: 87) fairly characterizes such feelings of satisfaction as 'subjective,' meaning it is an emotional standard which only the speaker is in a position

to know for certain. She concludes from this that because feelings of satisfaction are subjective their authenticity is "ultimately indeterminable." This may very well be so—Shuy (2014: 4) is certainly not alone among linguists when he warns that "it is impossible for any science to probe accurately into human minds." Yet, this kind of subjectivity, when it is encoded in language, is exactly what Appraisal is built to evaluate. As Martin and White (2005: 1) say in the first sentence on the first page of the Appraisal manual: "This book is concerned with the interpersonal in language, with the *subjective presence of writers/speakers in texts* as they adopt stances towards both the material they present and those with whom they communicate [emphasis added]." More specifically, subjective satisfaction is one of the parameters included within the Appraisal system of attitude, which will be discussed more fully in Chapter 3. If an imagined or anticipated satisfaction is expressed as part of a pledge to harm, then its kind and degree should be measurable through the method of Appraisal.

In contrast Harmon (2008: 87) characterizes the speaker's capability to enact the violence as an 'objective' standard, one which includes the threatener's physical size, strength, and age, as well as access to weapons, history of substance abuse, etc. These historical and behavioral considerations are common—and effective—tools in the world of threat assessment (e.g., Dormond, 2014). Furthermore, this standard echoes the theorizing previously discussed within FCI that a pro attitude with action content does not rise to the level of intentionality until the agent determines "whether she is capable of performing the action" (Malle & Knobe, 2001: 54). Of course, in the context of a *Tarasoff*-type situation where a wealth of personal information about the patient is usually on hand, this expectation is reasonable.

Outside the clinical realm, however, "threateners' identities often remain unknown—which means, unfortunately, that offender-related information is unavailable to investigators" (Smith, 2006: 55). Indeed, Malle and Knobe (2001: 61) counsel that

social perceivers regularly "distinguish…between the agent's belief that the behavior is controllable and the actual controllability of that behavior." Meaning, if an agent believes she can start a car (to use the authors' example) then perceivers are more inclined to judge the agent's behavior as intentional, even if, unbeknownst to the agent, the car's alternator is dead (Malle & Knobe, 2001: 61). Or as Searle (2004: 173) illustrates with a different scenario, "if I form the intention to go skiing I can do so only if I take for granted that I have the ability to ski, but the ability to ski is not itself an additional intention, belief, or desire."

Essentially, an agent's confidence that something can be accomplished counts towards the final tally of intentionality. And where shared world knowledge may insist that a goal is infeasible—such as Shaw's plan to assault one million women—this does not preclude attempts at realizing the imagined violence, as Shaw's case shows[9]. Helpfully, the language of 'capacity' is also addressed within Appraisal, once again through the system of attitude. So, in cases where the biographical details of the threatener may remain unknown (such as with the LA/NYC school pledge mentioned in Chapter 1), the threatener's *opinion* of his or her capacity may still be gleaned.

Sincerity condition: This condition is fulfilled when "[t]he speaker wants to inflict the detriment" (Harmon, 2008: 88). In many ways, this condition is the pivot point between a 'true' pledge to harm and a pledge made for other reasons. As discussed in section 2.1 above, a promise which is uttered with psychological insincerity, i.e., one which the speaker has no intention to keep, may still be *linguistically* sincere, and thus work as a valid speech act "whether the speaker actually has the specified psychological state or not" (Searle, 1969: 65).

[9] Indeed, Appraisal is intentionally blind to such questions, as Martin and White (2005: 33) attest: "from the perspective of interpersonal meaning we are more interested in the rhetorical organisation of a text than its logic."

To determine sincerity, a therapist must look to context. This "includes past communications between the speaker and the hearer" related to, essentially, the speaker's psychological history (Harmon, 2008: 88). For instance, has the speaker responded to treatment? has he or she been hospitalized previously? etc. Despite framing these topics in terms of how they appear in past communications, Harmon (2008) offers no measures for sincerity which are as specific as the satisfaction and capacity of the preparatory condition above. She also judges the "seriousness of the speaker's intentions" to be something which is "ultimately indeterminable" (Harmon, 2008: 88-89)—the exact assessment which, again, is being tested in this research.

Essential condition: For this condition, the measure is not the patient but the therapist—that "[t]he hearer understands the locution to be a threat/pledge" (Harmon, 2008: 89). Additional rules apply to this condition in a *Tarasoff*-type situation, rules which do not bear on the current dataset, e.g., that the threatened party must be clearly or reasonably identifiable. (Despite its vivid violence, for instance, Justin Carter's text would be unlikely to trigger the duties demanded under *Tarasoff* because he offers no details about which kindergarten he imagines attacking.) However, two parts of this condition—the "specificity as to the victim's identity" and "[h]ow involved/detailed the plan to commit the act has become" (Harmon, 2008: 90)—are common foci of the threat assessment literature. For example, Mohandie (2014: 136) says that "evidence of intent" includes the "specificity of [the] plan" which the threatener hopes to enact. Similarly, Borum *et al.* (1999: 330) consider "communications of…plans to attack a target" to be an "attack-related behavior," along with the specificity of the target. If these conclusions are correct, then expressions which feature a higher degree of specificity—about the target and/or the plan of attack—should correlate more with realized pledges than with non-realized.

Propositional content condition: In SAT, where the illocutionary point signals whether an utterance's direction-of-fit is *words-to-world* or *world-to-words*, the propositional content of the utterance is *what* is being fitted (Searle, 1976: 29). Thus, the words of a locution must be generally suitable for the speech act in question. For example, reports and predictions are nearly identical. What separates the two is a difference in temporal content: "a prediction must be about the future, whereas a report can be about the past or present" (Searle, 1976: 30). In terms of a pledge to harm, the content is required to be "a statement of intention to do something detrimental to the victim in the near future" (Harmon, 2008: 90). This is meant to assess the urgency or immediacy of the pledge, i.e., whether the speaker will soon proceed to action. One measure for this is "[w]hether the locution was conditional, and whether the condition is not immediately surmountable" (Harmon, 2008: 91). How easily a condition may be satisfied by a threatener, the threat's audience, or its target is, of course, context-specific and dependent on shared world knowledge. For instance, in the LA Unified pledge, the author tells the school administrators who have received the email, *If you cancel classes, the bombings will take place regardless*. Shared world knowledge and context both support the authority of school administrators to cancel classes as needed. For anyone other than school administrators, though, this condition would not, in fact, be "immediately surmountable" (Harmon, 2008: 91)—the average citizen lacks the necessary authority. Of course, not every condition is as clear cut as this. But that aside, the *presence* of conditionality in threatening language has long been of interest to threat assessors, and is considered one of the three main threat-type categories (shown in bold):

> [Threats] can be **direct**, as in the utterance 'I'm going to kill you tomorrow;' threats can be directly or indirectly **conditional**, as in 'She'll die if you don't pay me $1 million dollars' or 'If you don't leave town, no one knows what might happen,' respectively; and threats can be veiled, or **indirect**, as in 'you'd better watch your back.' [emphasis added] (Gales, 2010: 8-9)

Interestingly, the significance of these categories is debated. In cases of intimate partner violence, for instance, the issuing of conditional threats is singled out as evidence of an 'acute conflict,' meaning a situation in which "a perpetrator is seriously upset" and likely to become violent (Kropp & Cook, 2014: 187). In a different context, studies of threats against public figures found that "those subjects who made conditional threats" were actually "less likely to approach" the politician being threatened (Dietz, Matthews, Martell, Stewart, Hrouda, & Warren, 1991: 1460). The difference here could easily be the personal versus the impersonal nature of the targets, with intimates and public figures on opposite ends of a spectrum. As Muschalik (2018: 57) notes, the use of conditional forms in particular situations might be tied to the situations themselves "and not an indication of the prototypicality of conditional language for threats in general." However, in her examination of threats more broadly, Smith (2006: 78) "found no association between type of threat (direct, conditional, or non-specific/implied) and taking action." In other words, the particular threat category had no apparent bearing on the outcome of the situation in Smith's (2006) dataset. This is despite the fact that "[e]xperienced investigators typically give more credence to direct threats" than to their conditional and indirect counterparts (Smith, 2006: 78).

Findings which run somewhat, though not entirely, counter to Smith's (2006) null results were uncovered by Gales's (2010) more granular, corpus-based work. Approaching conditionality from several grammatical levels, she finds a level of salience not with conditionality as an overarching threat type, but between the functional use of conditionality and realized threats. In her discussion of modality, for instance, she notes that "[i]n realized threats, where *will* makes up approximately 70% of the prediction category, *will* is frequently used in a conditional sense" (Gales, 2010: 184). An example she provides of this particular structure is: *If you do not comply Smith's body <u>will</u> be displayed* (Gales, 2010: 185). This association of conditionality with realized threats

appears buttressed in a separate discussion of verbs controlling *to* clauses. Here, "an added level of conditionality can be seen to accompany the directive nature of the threats" (Gales, 2010: 195), as in *You stand a 99% chance of killing your daughter if you try to out smart us*.

Elsewhere, Fraser (1998) implicitly uses conditionals to distinguish between his conceptions of intention and commitment. For him, commitment is a matter of doing, of asserting that something will come to pass regardless. An intention, meanwhile, is more contingent on circumstances. "We do not find people saying, 'You threatened to fire me if I didn't get the report in on time and you didn't. You lied to me' (Fraser, 1998: 161). Using Fraser's (1998) terminology, the supervisor in this example is guilty of expressing an 'intention' to fire the employee, but not the 'commitment' to do so.

Within the method of Appraisal, conditionals as a kind of syntactic construction are best dealt with via the system of engagement[10], where they are considered dialogically expansive, i.e., as a device which "opens up the dialogic space for alternative positions" (Martin & White, 2005: 103). In other words, by framing future events through a conditional *if-then* clause the authors open the text to the possibility—no matter how small—that the condition will not be met, and that the threatened events will not occur or will occur differently. Whether the current dataset of pledges comports more with: 1) the seeming suggestion by Dietz *et al.* (1991) that this kind of expansion is more likely to be a feature of non-realized threats; 2) with Smith's (2006) null results; or 3) with Gales's (2010) finding that conditionality is likelier to appear in realized threats will be taken up again in the analysis of engagement in Chapter 7.

In sum, Harmon's (2008) theoretical scaffolding helps to identify the atomic components of pledging which are potentially detectable in the pledge's language, e.g.,

[10] Engagement as a system is discussed more fully in Chapter 3.

satisfaction, capacity, conditionality, etc. The next and final section will focus on the two linguistic studies, Smith (2006) and Gales (2010), which work to uncover the differences between realized and non-realized threats in particular.

2.3.2 REALIZED VS NON-REALIZED THREATENING

The primary attributes shared by this research and that of Smith (2006) are the agreement that "delving into threatening communications to find predictive factors for violence" (p.vii) remains a pressing pursuit, and relatedly, that it may be practically possible "to link characteristics of threatening communications with threateners'…intention" (p.75). However, the current research departs from her work generally—though not entirely—in the areas of linguistic theory and methodology. As noted previously, Smith's (2006: 5) study is mainly psychological, aimed at understanding how the "psychopathological, social, demographic, and dispositional characteristics of the threatener, target/victim type and relationship with threatener" affect whether a threat will be carried out. Thus, her interest in language is primarily in its role as a vehicle for transmitting psychological information about the writer, e.g., the writer's level of self-confidence, distrust of others, etc. For this task, she turns to psycholinguistics, "that field of study concerned with psychological aspects of language studies" (Kess, 1992: 1), and specifically to the complementary projects of Gottschalk, Winget, and Gleser (1979) and Hermann (2005), which provide her with a total of fifteen quantifiable language elements.

Briefly, the Gottschalk-Gleser method is a form of content analysis, which seeks to "measur[e] psychological states inferable from the content of communication" (Gottschalk *et al.,* 1979: 2). It attempts this by parsing the semiotic resources a speaker employs in his or her language. "[F]rom the linguistic standpoint, it involves itself with (word) signs to signs (syntactics), the relationship of signs to referents (semantics), and

the relationship of the sender to the sign (pragmatics) and the sign to the sender (pragmatics)" (Gottschalk *et al.,* 1979: 1). These linguistic choices are then used as indicators of "six psychobiological constructs—*anxiety*, *hostility*, *social alienation-personal disorganization*, *cognitive impairment*, *depression*, and *hope*" (Smith, 2006: 50). Essentially, the scales designed to quantify each state rely on a contextual reading of various "references" made by a speaker (Gottschalk *et al.,* 1979: 31). Anxiety about death, for instance, is signaled by statements which refer to "death, dying, threat of death, or anxiety about death" (Gottschalk *et al.,* 1979: 31) such as "I tried to get up the courage to do away with myself" (Gottschalk *et al.,* 1979: 35).

Hermann's (2005: 186) work is also, at bottom, a kind of 'content analysis'. However, where Gottschalk-Gleser attempts to measure psychological *states*, which "may fluctuate considerably from day to day or hour to hour" (Gottschalk *et al.,* 1979: 31), Hermann (2005) uses language to assess the more stable dimensions of personality *traits*. She focuses on seven in particular:

> (1) the belief that one can influence or control what happens, (2) the need for power and influence, (3) conceptual complexity (the ability to differentiate things and people in one's environment), (4) self-confidence, (5) the tendency to focus on problem solving and accomplishing something versus maintenance of the group and dealing with others' ideas and sensitivities, (6) general distrust or suspiciousness of others, and (7) the intensity with which a person holds an in-group bias. (Hermann, 2005: 184)

Hermann's (2005) traits are identified through a range of linguistic forms, like pronoun usage, action verbs, noun phrases which refer to people other than the speaker, etc. High conceptual complexity, for instance, is "coded with words, such as *possibility* and *approximately*, which indicate or suggest the ability to see different dimensions" (Smith, 2006: 47-48). By contrast, allness terms like *without a doubt*, *absolutely*, or *irreversible* are said to signal low conceptual complexity on the part of the speaker (Smith, 2006).

Thus, along with the eight scales of “anxiety, hostility outward, hostility inward, ambivalent hostility, social alienation/ personal disorganization, cognitive impairment, hope, and depression” (Smith, 2006: 53) derived from Gottschalk *et al.* (1979), Smith (2006: 46) also “tests the hypothesis that Hermann’s traits, measured through the threateners’ communications, could be associated with outcome,” i.e., whether or not a threatener proceeded to real-world action.

Smith (2006) divides her findings into factors which are shown to be risk-enhancing and those which are risk-reducing, i.e., variables which correlate more strongly either with realized or non-realized texts, respectively. Several of these are meta-linguistic—such as how a threat was delivered (e.g., phone call vs email)—or paralinguistic in nature—such as the inclusion of inappropriate capitalization. Additionally, two of the threatened actions which Smith (2006) identifies as having predictive value are fear-inducing but not physically harmful. “Threateners were more likely to act if they threatened to stalk and threatened to reveal detrimental information (whether true or false)” (Smith, 2006: 88). However, neither stalking nor the revelation of information detrimental to a target appear as topics in the current pledge dataset, and so these findings will not be discussed further.

Of the six remaining, relevant risk-enhancing factors, two relate to the targets themselves. The first hinges on whether the entity being threatened is human or not. “Although institutions and objects constituted a substantial portion of the targets in this study…nearly all of the action cases,” i.e., realized threats, “involved people” (Smith, 2006: 87). Helpfully, this distinction between human and non-human is built into the Appraisal system of attitude, where evaluations of people other than the author are schematized via the area of judgement, while the area of appreciation covers non-human ‘things,’ including “things we make” such as abstract institutions and concrete objects (Martin & White, 2005: 56). The second risk-enhancing factor is whether the

target is “an acquaintance of the threatener” (Smith, 2006: 85). While there’s no guarantee, of course, that this kind of relationship would be mentioned in a pledge, how a writer evaluates a human target can certainly shed light on this variable, and is therefore also potentially visible to the area of judgement[11].

The final four factors relate to the threatening communication itself. First, “[t]hreateners were significantly more likely to approach/stalk or harm when they used the language strategy of persuasion in their threat communications” (Smith, 2006: 76). Unfortunately, Smith (2006) does not elaborate on just what kinds of persuasive strategies appear in her dataset. Per Sornig (1989), are they aimed at *convincing* the reader? *Seducing* the reader? Are the rhetorical moves achieved primarily through grammatical means (e.g., declarations vs interrogatives) or through lexical choices (e.g., lexemes signaling an informal register)? Elements of persuasion potentially cut across several systems of Appraisal. While an author’s lexical choices (including phrases) typically fall within both the systems of attitude and graduation, Sornig (1989: 100), for instance, notes the effectiveness of quotation, which “carries conviction, or at least feigns credibility by exploiting the prestige of the person by whom a quotation was originated.” Quotative devices would instead be dealt with in engagement under the heading of ‘attribution.’

Smith (2006) also flags politeness as being more highly associated with realized texts but offers similarly little detail. Do authors who proceed to action tend to include negative politeness markers? Positive politeness? Is the politeness hearer- or speaker-directed? As with persuasion, aspects of politeness can appear at different levels of the language and thus be captured by different systems of Appraisal. Nevertheless, Smith

[11] Per Gales (2010: 3): “Due to the fact that the Appraisal framework…possesses a category called ‘Judgement,’ this spelling will be utilized throughout my research as opposed to the American English spelling ‘judgment’ for the sake of consistency.” This research will proceed similarly and for the same reason.

(2006: 98) speculates that both persuasion and politeness "suggest more deliberative and less emotional thinking" on the part of the threatener. This 'coolheaded' quality could be evidence of the 'settled' quality which concludes the reasoning process, when an agent has weighed his or her competing desires and made a decision about how to proceed (Malle & Knobe, 2001: 56). Both variables are thus worth being analytically alert to and both figure in the discussion which concludes this thesis.

The next risk-enhancing factor identified by Smith (2006) is thematic. Authors "were significantly more likely to act when they repeatedly mentioned love, marriage, or romance" (Smith, 2006: 76). Interestingly, Smith (2006) speculates that these topics may be only superficially affective in nature—that the mention of such generally positive states may hide thinking which is more "predatory than pure affect" (Smith, 2006: 99):

> The presence of this particular focus may reveal that the threatener has moved from surface emotion to thinking and planning—cognitive processes consistent with predation, which would increase the likelihood of violence significantly. (Smith, 2006: 99)

However, though Smith (2006) does not specify, these themes may be more closely related to the threats of stalking in her dataset, since a common motivation among stalkers is that their continued contact with a target "will rekindle their romance" (Smith, 2006: 35).

Finally, and perhaps most interestingly, threateners who were more likely to approach their targets also scored more highly in terms of their conceptual complexity. High conceptual complexity is associated with successful adaptive behaviors, in that high-scoring individuals "can entertain differences and are more flexible in their responses to others' ideas or to the objects in the environment" (Smith, 2006: 47). "Conversely, those who are low in this trait tend to categorize things in dimensions of black or white, good or bad, and are less flexible in responding to stimuli" (Smith, 2006: 47). Of Hermann's (2005) seven psychological traits, this is the only one in Smith's

(2006) dataset which correlates with a case's outcome. However, Smith (2006) sees a useful relationship between this and Gottschalk *et al.*'s (1979) conception of ambivalent hostility: in theory, when one is waxing the other is waning. "Ambivalent hostility scales reflect paranoia, which Gottschalk defined as critical, destructive actions or thoughts of others directed toward self" (Smith, 2006: 51). Ambivalent hostility is thus associated with affective (reactive) violence rather than predatory (proactive) violence. This means that threateners scoring high on this scale are more likely to lash out than to plan ahead.

Importantly, for such threateners the act of writing the threat itself "may assist…in defusing their anger" (Smith, 2006: 90). This possibility has compelling points of contact with the idea of the 'safety valve' effect of violent fantasy found in Gellerman and Suddath (2005: 485): "Rather than being predictive of future violence, such fantasies actually serve as a psychological 'safety valve,' permitting the vicarious, but safe and harmless discharge of strong emotions." As ambivalent hostility rises, and with it the chances for sudden, affect-driven violence, "conceptual complexity would logically diminish" (Smith, 2006: 91). "Thus, it appears that the presence of lower ambivalent hostility and higher conceptual complexity together are consistent with predatory thinking" (Smith, 2006: 91). Again, the strength of each of these factors—of nuanced versus black-and-white thinking—is potentially visible to Appraisal, since the method's utility is in how it quantifies the ways in which authors "approve and disapprove, enthuse and abhor, applaud and criticise, and with how they position their readers/listeners to do likewise" (Martin & White, 2005: 1).

Along with such risk-enhancing elements, certain of Smith's (2006) findings also defy the common conception of what makes a communication threatening. One factor of language-use which she found to be risk-*reducing*, for instance, were "words indicating prejudices concerning religion" (Smith, 2006: 77). She also discovered that "expressing

prejudices concerning race, gender, sexual preference, and ethnicity had no relationship to acting" (Smith, 2006: 89).

The relative surprise of these results is that they run counter to many preconceived notions about what is involved in threatening. Such notions have been called "folk linguistic impressions" (Gales, 2010: 267), or "non-linguists' beliefs about language in general" (Preston, 2007: 181), and threat assessors are as susceptible to such impressions as any other language user. Davis (1997), for instance, articulates the commonly held belief that "almost all of those persons who do commit acts of violence use profanity and other offensive language—before, during, and after the act—to describe or discuss both the victim and the violence itself" (Davis, 1997, via Gales, 2010: 33). And indeed, this accords with a student survey performed by Gales (2010) in which 73% of respondents voiced the belief that profanity would be common to threatening language, "including sexist and racist language" (Gales, 2010: 95). But Davis's (1997) and the students' folk linguistic impressions are belied by Smith's (2006) null result. It is also contradicted by Gales' (2010) own analysis. She found that, in fact, only 24% of the threats in her corpus contained such language. Gales (2010) explains the dangers of this disconnect:

> [O]ur folk linguistic (Preston, 2007) ideologies about threatening language *continually mask, or erase, some of the ways in which threateners demonstrate intent*, mitigate claims, and negotiate meaning in threatening language […] [W]e face the risk of misunderstanding the intended stance, and in the case of threats, this misunderstanding may result in dire consequences [emphasis added]. (Gales, 2010: 200)

Comparing folk ideological beliefs about threatening language to empirical evidence is, of course, one of the major strands of Gales's (2010) work in this area, which questions how authorial stancetaking in a genre like threatening can and does differ from the cultural expectations about such communications. Within "a genre committed to violence and threatener control," for instance, our folk ideologies routinely

focus on “functions that *strengthen* the threatener’s stance” (Gales, 2010: ii-iii). Yet, Gales’s (2010) corpus study of 470 threat letters reveals unexpected patterns of stance markers which actually serve to “*weaken* the threatener’s stance” (Gales, 2010: ii-iii). This is different than uncovering psychological intent, of course. (Indeed, Gales (2010: 264) says that her analysis “supports previous studies that have questioned the use of linguistic form as an indicator of behavior.”) Instead, her goal is a clear-eyed view of a threat’s “underlying meaning or pragmatic intent” (Gales, 2010: 6)—in other words, an understanding unencumbered by ideology. Nevertheless, her comparison of realized and non-realized threats brings to the fore an important resource of stancetaking in English, and one whose predictive value warrants examination: the lexical category of modal auxiliaries.

As a class, modal auxiliaries comprise helping verbs like *must*, *can*, *will*, etc., which offer a lexicogrammatical means of signaling a “speaker’s judgments…of the information in the clauses” (Lock, 1996: 192). This includes signaling a writer’s level of commitment towards, or certainty about, a proposition (Gales, 2010: 118). Put differently, this “special class of stance verbs” (Gales, 2010: 118) offers speakers a fully lexicalized resource for evaluating whether something is considered likely or unlikely, obligatory or optional, frequently recurring or rare, etc. For threat assessors tasked with determining whether an author will proceed to violence, these kinds of judgements offer a valuable window into the threatener’s potential commitment to real-world action. Thus, modal auxiliaries are often the unique group of words to find their way into the threat assessment literature. Mardigian (via Gales, 2010: 26), for instance, refers to *must*, *have to*, and *will* as “modals of intent” to which an assessor should be sensitive.

This intuition—that modals bear a special relationship with commitment—is also reflected in folk linguistic ideologies of threatening language. In her student survey, for instance, 27% of respondents identified “forceful modals” such as *will*, *must*, *shall*, and

have got to, as a type of language which they would expect to find in a threat (Gales, 2010: 96), circled in red in Figure 2.1 below.

Figure 2.1: Student Ideologies about Threatening Language (Gales, 2010: 96)

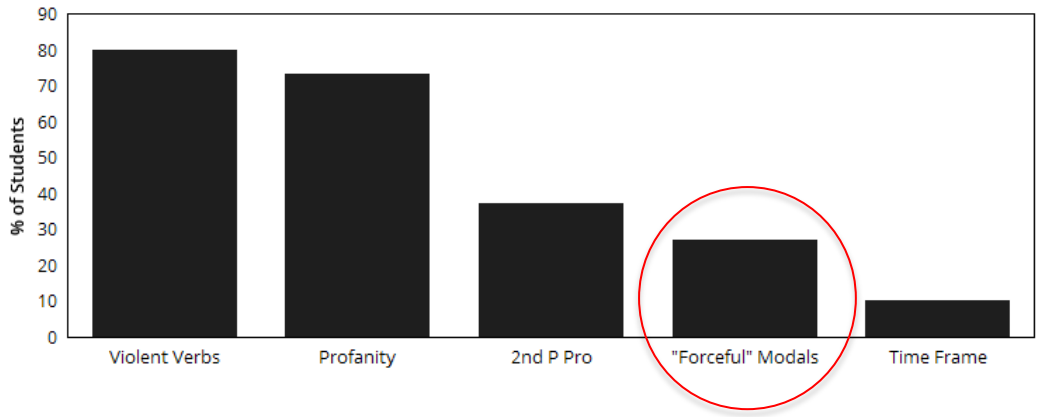

Modals have long been a subject of interest in linguistics—and any area with a long tradition of study tends to be theorized, divided, and subdivided in different, not always compatible, ways. This is certainly the case with modals. O'Donnell (2017), for example, separates this class into three broad categories of meaning, which he calls *likelihood*, *requirement*, and *volition*. Lock (1996), however, employs five headings: *likelihood*, *requirement*, *frequency*, *inclination*, and *potentiality/ability*. Quirk (1985) opts for two broad distinctions: *intrinsic*, which indicates human control over an event (e.g., meanings of *permission*, *obligation*, and *volition*); and *extrinsic*, which indicates human judgement but not human control (e.g., meanings of *possibility*, *necessity*, and *prediction*). For her part, Gales (2010) follows Biber, Johansson, Leech, Conrad, and Finegan (1999) in employing a multilevel approach, in which the three main semantic categories have the capability of expressing both intrinsic (i.e., deontic) and extrinsic (i.e., epistemic) meanings. These three categories are:

> modals of permission, possibility, and ability ("can, could, may, might"); modals of obligation and necessity ("must, should, (had) better, have (got) to, need to, ought to, be supposed to"); and modals of volition, intention, and prediction ("will, would, shall, be going to"). (Gales, 2010: 118)

Of the many linguistic forms which Gales (2010) investigates in her dataset, only the distribution of modals rises to the level of statistical significance. However, instead of being used by realized authors to communicate something like a genuine commitment to act, their significance is instead "specifically in the sub-corpus of *non-realized* threats [emphasis added]" (Gales, 2010: 182). This statistical disparity in distribution is captured in Figure 2.2 below.

Figure 2.2: Distribution of Modals by Threat Realization (Gales, 2010: 183)

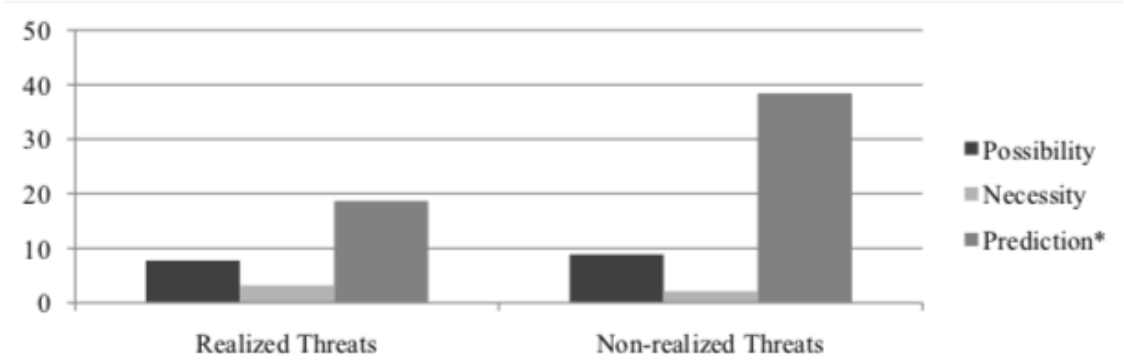

Frequency per 1000 words

Perhaps most surprisingly, modal auxiliaries with meanings related to prediction were significantly more likely to appear in *non-realized* writings. Obviously, a threatener's predictions of future events are an area of great interest to threat assessors, and are believed to tie into other related "warning behaviors" (Meloy *et al.*, 2012: 256). For instance, an "anonymous threatening author who uses language indicating the violent act will occur" is demonstrating one necessary element of "the perceived inevitability of violence" (Simons & Tunkel, 2014: 207). Nevertheless, in Gales's (2010) analysis, modals used to make such predictions are statistically more likely to be a resource for 'howlers' rather than for 'hunters' (Calhoun & Weston, 2015).

2.4 SUMMARY

The goal of this chapter has been to survey both the psychological and linguistic literatures surrounding intent, and to tie the ideas raised by these researchers to the relevant practices and beliefs in the area of threat assessment. The aim has been three-fold: 1) to disambiguate the meaning of 'intent' from its linguistic and psychological uses; 2) to delineate which potential 'state of mind' is being tested for when intent is at issue; and 3) to explore the current understanding of how this state of mind is—or is not—encoded linguistically. Of particular importance to the analysis going forward are the following theoretical concepts and language features:

- The linguistic or pragmatic intent which is necessary for a speech act to be successful—i.e., "happy or felicitous" (Austin, 1962: 42)—is not identical to, and may operate independently from, the psychological intent spurring the speech act.
- "Wherever there is a psychological state specified in the sincerity condition, the performance of the act counts as an *expression* of that psychological state" (Searle, 1969: 65). Accordingly, to "threaten or pledge" to do something "counts as *an expression of intention*" to do that thing (Searle, 1969: 65). This 'law' applies "whether the speaker actually has the specified psychological state or not" (Searle, 1969: 65). This formulation allows in theory for the phenomenon of non-realized pledges, i.e., pledges where the linguistic intent and the psychological intent are not aligned.
- The process of forming an intention begins with the interplay of beliefs and desires (Malle *et al.*, 2001). Where beliefs aim to describe the world, desires instead wish to alter the world (Searle, 2004). Both beliefs and desires may have many kinds of content (Malle & Knobe, 2001). However, a desire represents a positive psychological mode towards the content in question, i.e., a positive

emotional valence towards the represented material. This is known as a pro attitude (Davidson, 2001). Intention, like desire, also features a pro attitude. But where a desire may hold any propositional content, the content of an intention is limited to "an action performed by the same person who holds that attitude" (Malle & Knobe, 2001: 51). Thus, "genuine intentions represent what we call action content" (Malle & Knobe, 2001: 48). In the ensuing comparison of realized and non-realized pledges, devices which may measure action content warrant special attention.

- Further to this, from the perspective of biological naturalism, "*intentionality is representation of conditions of satisfaction*" (Searle, 2004: 173). These conditions are simply what must come to pass if the intention is to be fulfilled. Searle's (2004) formulation raises an interesting possibility: the conditions of satisfaction for a non-realized pledge are either partially or entirely fulfilled by the production of the pledge itself instead of by actions external to the pledge. By contrast, a realized text is, hypothetically, an expression of the conditions of satisfaction themselves.
- Next, arriving at an intention involves a minimal amount of reasoning. This mental calculation 1) begins with desires, which "typically function as reasoning inputs"; 2) weighs the "*desirability* and *feasibility*" (Geurts *et al.*, 2016: 55) between any competing interests; 3) then finally produces an intention, which "typically function[s] as reasoning output" (Malle & Knobe, 2001: 65).
- Finally, the 'settled' quality which concludes the reasoning process characterizes a 'commitment' to act (Malle & Knobe, 2001: 56). Commitment is the crucial consideration for questions of *mens rea* (Malle & Nelson, 2003), although its detection is arguably problematic when the behavior being judged is linguistic.

How these psychological processes may be realized through language varies. Different studies highlight different linguistic features which may be considered risk-enhancing. Risk-enhancing features include:

- An expression of satisfaction at harming the victim (Harmon, 2008);
- An expression of the capability to harm the victim (Harmon, 2008);
- The "specificity of [the] plan" which the threatener hopes to enact (Mohandie, 2014: 136);
- The specificity of the target (Borum *et al.*, 1999);
- Functionally conditional forms (Gales, 2010);
- The targeting of a human rather than an institution or an object (Smith, 2006);
- Attempts by the threatener to persuade the threatened party of something (Smith, 2006);
- The use of politeness markers (Smith, 2006), whose 'coolheaded' quality could be evidence of the 'settled' quality which concludes the reasoning process (Malle & Knobe, 2001); and
- Evidence of higher levels of conceptual complexity allowing the threatener to "entertain differences and [be] more flexible in their responses to others' ideas or to the objects in the environment" (Smith, 2006: 47). Conceptual complexity may be indicative of predatory thinking.

Linguistic features which are considered risk-reducing are fewer in number. Risk-reducing features include:

- Conditional threats as a *category* of threat type (Dietz *et al.*, 1991);
- Higher scores on the scale of ambivalent hostility, which indicates "black and white" thinking and a reduced cognitive flexibility (Smith, 2006: 47). More strongly affective responses may serve not as indicators of commitment but as a

'safety valve,' allowing for "the vicarious, but safe and harmless discharge of strong emotions" (Gellerman & Suddath, 2005: 485);

- "[W]ords indicating prejudices concerning religion" (Smith, 2006: 77);
- The general use of modal auxiliaries as a grammatical class (Gales, 2010); and finally
- The use of prediction modals in particular (Gales, 2010).

As has been noted at various relevant points, nearly all of these features are in some way quantifiable using systems within the method of Appraisal (Martin & White, 2005). The method itself, and the data to which it will be applied, are discussed fully in the next chapter.

CHAPTER 3 DATA AND METHOD

In the roughly two decades since their inception, the use of social media has skyrocketed. In 2005, just 5% of American adults had a presence on one social media platform (e.g., Myspace, Facebook). By 2019, this number had grown to 72% (Pew Research Center, 2019), a more than 1,300% increase. Similarly, internet penetration across the globe has jumped from roughly 15% of the world population in 2005 to more than 45% in 2017 (Clement, 2017). This ubiquity has meant an "increase [in] the number of individuals who can engage in unmediated communication" with the wider world, a situation "which inherently increases the probability of incendiary speech" such as threats (Lidsky, 2012: 149). Indeed, "[o]ur highly digitized societies afford a multitude of easily available channels through which threats can be communicated" (Bojsen-Møller *et al*., 2020: 6). There are no indications that this growth is slowing or reversing; instead, individual usage of social media particularly and the internet more generally is likely to grow.

As ever more people gain access to the internet and the unmediated communication available through social media platforms, the problem of threatening language is also certain to increase rather than diminish. Considering that threatening language is most likely to take the form of a pledge to harm or 'leakage' (Meloy & O'Toole, 2011), understanding whether such language contains evidence of commitment will only become more urgent in the coming years. On the front end, law enforcement personnel and assessors will be tasked with allocating resources to pre-empt genuine threats. After an immediate situation has been resolved, justice systems will be tasked with deciding fair consequences for those ill-advised people who issued a threat[12].

[12] Indeed, the contrast between the Dickens and Carter cases show that fair consequences are still an important question in the United States. Carter, a 19-year-old, was held in county jail for five months with a bond set at a staggering $250,000. During

Yet, despite the growth of the internet and the concomitant rise in public threatening, authentic "[malicious forensic texts] are generally difficult to access" (Nini, 2017: 101). And so, finding appropriate data for research such as this is challenging. Gales (2010: 72) summarizes the hurdles faced by forensic linguists thusly: "performing research within this field is a difficult and oftentimes lengthy process due to two main elements—proprietary methods and/or proprietary data." Regarding data, authentic texts are often unavailable because they are somehow:

> **sensitive in nature** (e.g., threats to national security, which are accessible only to those with top level security clearance), **legally unavailable** (e.g., threats in cases that have not yet gone to trial or are in the process of appeal), or **proprietary** (e.g., threats belonging to corporations or private individuals) [emphasis added] (Gales, 2010: 73)

Even with texts which begin their life in the public sphere—like most of the pledges examined here—availability may *become* an issue as the writings are deleted from the platforms where they were originally hosted. During the collection of this dataset, for example, a common pattern began when news media reported that a pledge to harm had occurred. Usually, the pledge would have been removed in the time it took to search. Generally, the reason for its disappearance was not hard to surmise: it had become evidence in an active investigation and was therefore "legally unavailable" (Gales, 2010: 2).

Of course, active legal disputes can and do affect non-realized writings as well, as authorities work to determine whether any intent was evident (as in Carter's case). And so, these pledges had often been removed as well by the time media reports were

---

this time, he was assaulted by other inmates and then put in solitary confinement—a type of detainment widely viewed as psychological torture (Center for Constitutional Rights, 2017). He was subsequently placed on suicide watch (Pinsof, 2013). The simple public apology demanded of Dickens (Torpy, 2016) appears far more humane in comparison to such an ordeal.

publicizing their existence. The fluidity of the internet—and social media in particular—makes leakage both more ubiquitous and also more difficult to document.

In the choice of method, the path was easier. The framework of Appraisal is not in danger of disappearing, and its public availability has obvious advantages for replicability. As Gales (2010: 72) notes, "many of the *methods* currently used to assess threatening communications are proprietary and require a case-based user fee, a yearly license, or an expensive software purchase, if they are available for public consumption at all." Thus, "these methods cannot contribute to the improvement of the field because they are not available for further scholarly testing, refinement, or general methodological use" (Gales, 2010: 73). Appraisal is not only appropriate as a method for analyzing the pledge dataset—for reasons which are summarized in section 3.3 below—it is also free of such a debilitating drawback. Its application as part of this thesis may thus, it is hoped, contribute to the improvement of the field.

3.1 Ethics and Process of Data Collection

The research design and the data collection method, storage, and analysis were approved by the Ethics Committee of the School of Languages and Social Sciences at Aston University.

The ethical issues surrounding these data are somewhat unusual compared to typical linguistic field work. This stems from the fact that the social action these texts perform is fundamentally to threaten—even though the threat is not directed at the text's intended readers. Pledges to harm are thus instances of an 'illicit genre', i.e., a genre that is "socially and sometimes even legally proscribed" (Bojsen-Møller *et al.*, 2020: 8). The unlawful nature of pledging places their authors in a precarious legal position. For the thirteen authors in this dataset, each was treated by authorities as allegedly guilty of a criminal act. While six have so far escaped subsequent legal action, this is only

because three authors remain anonymous at the time of this writing and the other three died in the course of their attacks. The remaining seven all came under the personal scrutiny of the law and faced consequences of varying severities from the judicial system. For four of the five non-realized authors whose identities are known, police attention was minimally the result of the pledges themselves[13]. This is due to the fact that "even if a threat is not carried out, it is still a verbal act of violence or assault" (Bojsen-Møller *et al.*, 2020: 7). In other words, no matter the author's subsequent behavior, the writings themselves are forensic texts, treated by authorities as evidence of a language crime. Ethical access to these texts, then, hinges in part on the "presumption of access to evidence admitted at trial" (U.S. v. Massino, 2005) practiced in the U.S. and elsewhere. The tradition of publicizing evidence recognizes the interest in public safety inherent to criminal proceedings as well as the role which a public evaluation of such evidence plays in government accountability (Bucqueroux & Seymour, 2009).

That said, this research attempts at all points to abide by "the core principles" of applied linguistics, which De Costa (2015: 245) summarizes as "(1) respect for persons, (2) yielding optimal benefits while minimizing harm and (3) justice." This is somewhat complicated—but not insurmountably so—by the fact that the internet is becoming increasingly domesticated in everyday life (Eynon, Fry, & Schroeder, 2008), which has led to a "debate concerning the ethical implications of online data collection" (Rodham & Gavin, 2006: 92).

However, some ethical concerns were allayed by the reality of the pledges. For instance, there was no need "to work on developing a relationship with [interview] participants […] to ensure that ethical practices are in place" (De Costa, 2015: 249).

[13] The fifth, Elliot Rodger, has one text in each realization category but died in the commission of the realized pledge.

This is because the data is textual, and not the result of a survey or interview. Similarly, because illicit genres like threatening rarely, if ever, emerge from a coherent speech community (Bojsen-Møller *et al.*, 2020), many common ethical concerns are also minimized or eliminated.

For example, Williams (2012) poses a binary set of eight issues which research involving human subjects should consider. Each of the eight has a positive (more ethical) and negative (less ethical) side. Williams' (2012) metrics are addressed in turn below, with negative ethical qualities shown in parentheses. The decision tree itself is reproduced as Appendix B. Generally, though, the collection and treatment of the pledge dataset may be considered ethical because the data and/or research method is:

- public (not private) because news media and the courts made the language evidence publicly available;
- open-access (not restricted) because there is no community or group of pledge authors to join, either online or offline;
- minimally intrusive (not actively intrusive) because the linguistic analysis is passive, with no direct contact between analyst and authors;
- low risk of harm (not high) because there is no community or group to disrupt and any personal identifying information was publicized before this research was undertaken;
- neither broad-based nor intimate because, again, the pledge authors do not constitute either a speech community or a community of practice; and
- not a matter of permission or its refusal since journalists and jurists served the interests of public safety and governmental accountability by putting the language evidence into the public arena.

However, of the eight, the most problematic is Step 4, singled out here for special comment. Williams (2012) asks: “What is the purpose of the group? What is the nature of the discussion and is the group seeking privacy?” This step also considers whether minors are involved in the research, noting that this requires more sensitive research practices than data generated by adults. (Note, Williams (2012) nowhere asserts that research on data created by minors is inherently unethical.) In fact, three of the pledge authors were minor-aged teenagers when they wrote their texts: Alex Hribal (16 years old); Kip Kinkel (15 years old); and the Skyline H.S. pledge author (16 years old). Skyline is considered non-realized, and because no charges were brought, the author remains anonymous (and arguably, anonymous authors are the least vulnerable of all). Both Hribal and Kinkel carried out attacks; both are now adults serving out their sentences. Williams (2012) primary concern in Step 4 is “the vulnerability of the group.” But, again, 1) no cohesive group or community exists between these threateners, and 2) the nature of their cases is criminal, which means that the interest in public safety supersedes privacy concerns. For these reasons, the inclusion of these pledges is considered minimally problematic. Thus, despite the presence of teenaged authors, this research may be considered <u>less sensitive</u> (not more).

Williams’ (2012) final metric asks the researcher to reflect on his or her own ethical stance and whether the risks to the human subjects involved outweigh the benefits. For the pledge authors who are neither anonymous nor dead, this research poses no additional risk to the writers’ reputations, nor does this analysis have any conceivable legal ramifications (e.g., an extended prison sentence for an author already serving time). If anything, every effort is made in this thesis to deal with each pledge fairly and without prejudice, as something constructed for a legitimate communicative purpose, even if the purpose is socially unacceptable (Bojsen-Møller *et al.*, 2020). Finally, the goal of this thesis can be reached only if the motives of the threateners are

also treated as legitimate—not minimized, mocked, or dismissed because they proved disruptive or dangerous. Such a task could only be accomplished by respecting even the most morally reprehensible members of the dataset as, first and foremost, *human beings*.

3.1.1 PROCESS

The general procedure used to compile the dataset was straightforward. Typically, a news report of some kind would indicate that an incident had occurred. These incidents were either actual assaults of some kind or they spotlighted individuals like Carter and Dickens who ran afoul of the law through their language alone. An internet search (e.g., Google) would follow and either the text could be located (occasionally on its original publishing platform, e.g., Dylann Roof's and Tyrelle Shaw's personal blogs, or in subsequent, publicly available criminal complaints, e.g., Gilberto Valle) or the text could not be located. In rare cases, the complete text was featured in the news report itself (e.g., LA Unified). The collection of older texts (e.g., Alex Hribal) was a side effect of searches for more recent pledges.

3.2 DATASET

The dataset itself contains fourteen pledges for a total of 4,190 words. These texts have been divided into two categories: six pledges are classified as realized, meaning the authors attempted to enact the violent ideations outlined in their texts; eight pledges are non-realized, where records indicate that no such attempt occurred. Henceforth, the realized texts will be referred to as 'R' texts and the non-realized as 'NR,' and these designations will be appended to the author names to clarify which corpus they fall within, e.g., Dickens NR.

Methodologically, the two realization types are used essentially as reference corpora for each other, working off the hypothesis stated in Chapter 1 that each type

Table 3.1: Pledge to Harm Dataset

Type	Author	Description	Year	Word Count
Non-realized	Archangel Michael	Anonymous email threatening to bomb official buildings across Wyoming	2016	104
	Brahm, Jake	Chat room post describing the simultaneous bombing of several NFL stadiums	2006	209
	Dickens, Ebony	Facebook post threatening to shoot policemen	2015	120
	LA Unified	Anonymous Email threatening bomb and gun assaults at Los Angeles area high schools	2015	354
	McKelvey, Kayla	Consecutive Twitter posts threatening Kean University students	2015	104
	Rodger, Elliott	Excerpt of self-published 'autobiography' threatening male strangers around his Santa Barbara apartment	2014	230
	Skyline	Anonymous chat room post threatening a gun assault at Skyline High School	2012	248
	Valle, Gilberto	Excerpt of chat messages threatening to abduct and cook a targeted woman	2012	183
			Corpus Total	1,552
Realized	Hribal, Alex	Handwritten essay—subsequent knife attack wounded 21 at Franklin Regional High School	2014	979
	Kinkel, Kip	Handwritten note—subsequent gun attack killed 2 and wounded 25 at Thurston High School	1998	189
	Long, Gavin	Email—subsequent gun attack killed 3 policemen	2016	653
	Rodger, Elliott	Excerpt of self-published 'autobiography' threatening a sorority house near UC Santa Barbara campus—subsequent gun attack killed 6 and wounded 14	2014	183
	Roof, Dylann	Excerpt of a blog post—subsequent gun attack killed 9 church parishioners in Charleston, NC	2015	200
	Shaw, Tyrelle	Blog post—subsequent blunt object attack wounded 4 in New York City, NY	2015	434
			Corpus Total	2,638
			Dataset Total	4,190

was created for a different communicative purpose and will therefore evince distinctive patterns of language usage. The full dataset is shown in Table 3.1 above.

The dataset is heterogeneous in several ways. First, the texts were sent to their intended audiences through various modes of communication. Some are emails, some are social media postings, others were published to online chat forums, etc. Despite such "meta-linguistic" differences (Gales, 2017: 2), however, fully 85% are computer-mediated communications (CMC). The two which are not, Hribal R and Kinkel R, are handwritten, though each is in a style which could easily appear on a modern online platform. Hribal R's essay, for instance, shares many generic hallmarks of a blog post, making it very similar to, e.g., Shaw R's text. Similarly, Kinkel R's handwritten note is not unlike a Facebook post, e.g., Dickens NR's. The methods of distribution are thus of less interest here than the linguistic output of the authors themselves.

Next, four of the fourteen are excerpts of a longer discourse. In some cases, relying on an excerpt was unavoidable. Valle NR's language, for example, was only available through the criminal complaint filed against him. The bulk of his case data remains legally unavailable. In other cases, Appraisal's deep-dive approach made coding and analyzing longer texts unrealistic. Roof R's full blog post, for instance, is over 2,400 words long, which would be a substantial object of Appraisal analysis all by itself. Similarly, Rodger's full autobiography is 105,676 words long—well beyond the scope of this study and the resources of a single Appraisal analyst. Even the shorter fantasy sequence from which his two excerpts are pulled is still an overly bulky 1,031 words. Where excerpts were unavoidable, the general aim was to mimic the length of text which a threat assessor may encounter in a live situation. Based on an average of the 470 authentic threats in Gales's (2017: 8) CTARC corpus, this is roughly 323 words per unique text. The average word count in this dataset is 299 per text. That said, only four excerpts are used, and their boundaries were not decided arbitrarily by 'ideal'

length. Rather, the excerpt captures the beginning and ending of the pledge, i.e., the opening and closing of an imagined violent event as a topic.

Finally, the production of the texts spans a period of eighteen years, with the earliest written in 1998, and the last in 2016. However, ten of the fourteen were authored between 2014 and 2016, and thirteen of the fourteen between 2006 and 2016. Kinkel R's text, written in 1998, remains in the dataset because of its relatively small word count size and because, upon subsequent analysis, its language use did not show any peculiar or 'outlier' tendencies which may have resulted from language change.

Having noted these differences, the dataset is also homogeneous in certain vital ways. First, of course, each shares the generic qualities of a pledge to harm, e.g., somehow threatening future harm against a grammatical 3rd Person.

Second, and most practically, each chosen text is a minimum of 100 words. This floor is admittedly arbitrary. However, "[m]ost of the very frequent linguistic features are already stable in 100 token samples, and previous multidimensional studies have employed the 100-token cut-off with good results" (Nini, 2017). And incidentally, this accords with similar choices made by the creators of two pieces of threat assessment software, Profiler Plus and PCAD (Smith, 2006). Both programs "require a minimum number of words in the text to achieve reliability of the computer-generated scores on each trait (Profiler Plus requires a text with 100 or more words; PCAD, a minimum of 90)" (Smith, 2006: 65).

Third, each is an *authentic* linguistic production, and thus a legitimate forensic text which either: describes violent events which the author later attempted or accomplished (realized); or whose threatening tone caused authorities real concern about the author's intentions (non-realized). Because this study has an element of psychology, authentic texts—meaning, texts produced by authors with some kind of

personally relevant, real-world intention, whether to predict violence, instill fear, etc.—were considered critical for measuring the cognitive quality of intent.

Fourth, each was intended to be received and read by a particular audience, however narrow (one individual) or broad (loved ones, legal authorities, etc.). This means that no text in the dataset was a private production later unearthed against the author's will, such a journal entry. This criterion ensures that the application of Audience Design Theory is appropriate.

Fifth, each was originally conceived in the medium of writing, i.e., none of the pledges are transcribed speech[14]. This uniformity is meant to facilitate comparisons between the findings of this study with other studies like it—Gales (2010: 75), for instance, only examines "written rather than spoken registers"—as well as comparisons with the many CMC pledges that are sure to arise in the future as the use of social media continues to grow.

Sixth, no matter the author's ethnic background (where such information is known), all are written in a relatively informal register of Standard American English. What little variation exists in this regard was further judged to be negligible "because threats are not register-dependent" (Gales, 2010: 76) and cross-register comparisons are therefore both acceptable and warranted (Gales, 2010).

A final, important note on the characteristics of these pledges: as is by now apparent, the texts were not sampled according to sociolectal considerations like age, gender, class, or ethnicity. Despite all sharing the same informal register of SAE, as stated immediately above, these demographic elements could very well behave as intervening variables in the study, potentially influencing the distribution of whichever

[14] Many threateners also create video material for platforms like YouTube (as Elliot Rodger did). For some, this is their medium of choice. While these materials are certainly informative of motive and motivation, such speech-based pledges are excluded here.

language resources are identified by Appraisal. For instance, both of the texts known to be authored by female writers—Dickens and McKelvey—are in the non-realized corpus. Genderlectal variation, i.e., "actual language performance differences" between men and women (Thimm, Koch & Schey, 2003: 531), could conceivably mediate between Dickens's and McKelvey's violent ideations and their *expression* of these ideations, in comparison to their male counterparts in the dataset. It is not difficult to imagine, for example, how the more frequent presence of softeners or hedges—features often associated with female speech (Thimm *et al.*, 2003)—might affect an analysis of modal auxiliaries like the one undertaken in Chapter 8. (Furthermore, both women are African American, raising similar questions about the potential effects of ethnolectal variation.) Also, despite the fact that it "may not be empirically provable nor falsifiable" (Grant & MacLeod, 2018: 83), idiolectal variation is a potential concern, i.e., how the "language patterns associated with a particular individual" (Grant & MacLeod, 2018: 83) may influence which resources are called on to express psychological intent. In comparative analyses as fine-grained as this, individual usage has the potential to loom large.

However, for such a small study, and under the conditions of data collection outlined above, controlling for these variables would have proved debilitative. Instead, certain guardrails are employed. For example, the overarching hypothesis put forward here posits that each of the two author types will draw from different *systemic* resources. Examining the texts at this broader level should, in theory, blunt the distorting effect of idiolectal variation (Grant & MacLeod, 2018). The subsequent analyses also make use of, e.g., dispersion plots for each Appraisal variable under examination, to better identify outliers of usage no matter their potential sociolectal source. Nevertheless, variation remains a concern and will likely warrant more attention in future iterations of this research.

These caveats in place, the contextual determinations used to categorize each pledge by realization type are discussed on a text-by-text basis in the next two sections.

3.2.1 NON-REALIZED CORPUS

While populating both corpora presented unique challenges, the psychology-based criterion that a non-realized pledge should, ideally, *not* contain true violent intent is, and will likely remain, theoretically fraught—both for this study and for any investigation which attempts a comparison between "threat internal categories" (Gales, 2010: 109). Because intent is a private psychological process, as discussed in Chapter 2, accessing it in other people's minds in its *pre-behavioral* forms is impossible (Muschalik, 2018; Singer, 2014). The analytical danger is that a text classified as non-realized does indeed encode a true commitment to violence, yet the author never attempted the actions due to some kind of practical constraint. For example, suicide prevention literature suggests that even for people with a strong intent to realize suicidal ideations, restricting access to lethal methods reduces successful attempts (Mann, Apter, Bertolote, Beautrais, Currier, Haas, ... & Mehlum, 2005). Whether or not this extrapolates to pledges to harm, this invisibility of intent means that the absence of action on the part of an author can never, with full certainty, be conflated with a true absence of violent commitment in the moment when a pledge was written.

In the end, all that is available for analysis is the language used to express the violent ideations and any correlation these ideations have with the subsequent behaviors of the author. And indeed, it was typical to find that the non-realized pledge writers themselves publicly "declared that they never had the intention, the means, or the commitment to carry out the threat" (Gales, 2010: 181). In these cases, reading through additional available media sources suggested that no real-world steps were taken to commit violence, such as buying bomb materials or procuring a weapon. For

anonymous authors, the two criteria were that none of the threatened events ever came to pass and no evidence was reported that the attacks had been prepared. Contra the LA Unified NR writer, for example, no bombs were ever discovered in LA county or NYC schools and no *jihadist cell* took their *guns to the streets and offices of Los Angeles, San Bernardino, Bakersfield, and San Diego* on the date promised in the pledge— December 15, 2015—or on any other date, for that matter.

A brief summary of the context surrounding each non-realized pledge is provided below.

Archangel Michael: On September 22, 2016, a sender calling him- or herself 'Archangel Michael' emailed Wyoming state officials threatening "state facilities and schools across the state and the Cheyenne Regional Airport" with napalm, pipe bombs, and even a strategically placed weapon of mass destruction (Morton, 2016). Closures across the state were patchy. However, no materials were discovered in any schools or state buildings and no violence occurred. As of this writing, the sender has not been identified.

Jake Brahm: The 20-year-old Wisconsin man posted over 40 identical messages to several internet forums in which he claimed that seven National Football League (NFL) stadiums would be destroyed by truck bombs, for which the terrorist group Al Qaeda would take credit. After his arrest, Brahm claimed that the messages were written as part of a contest, or a "writing duel" with an unidentified Texas man "to see who could come up with the scariest threat" (Smothers, 2006). In literary terms, then, he aspired to realism without harboring any genuine intent for violence. Brahm also expressed consternation that anyone would misunderstand his message as being anything other than a hoax or invention, i.e., a fantasy (The Associated Press, 2008). No additional warning behaviors, e.g., attempts to procure bomb materials, were ever alleged by authorities.

Ebony Dickens: As discussed in Chapter 2, Ebony Dickens, writing on Facebook under the name Tiffany Milan, posted a call in April of 2015 for African Americans to *rise up and shoot at every white cop in the nation starting NOW.* Dickens said she wanted "to write something preposterous in hopes of starting a conversation during protests and riots in Baltimore after the death of 25-year-old Freddie Gray" (Salinger, 2016). Authorities dropped charges of terroristic threats and inciting a riot "in exchange for her offering [her] remorse while surrounded by law enforcement officials at a news conference" (Salinger, 2016).

LA Unified: As discussed in Chapter 1, on December 14, 2015, an anonymous sender claiming to be an Islamic radical emailed the Los Angeles Board of Education as well as New York City officials threatening shootings and bombings in their respective school districts. "By midday, elected officials briefed by law enforcement said the threat did not appear to be credible" (Branson-Potts *et al.*, 2015). The LA Unified School District opted to close all 900 of their schools; NYC law enforcement chose to keep theirs open. No bomb materials were ever discovered and, as of this writing, the threatener has not been identified.

Kayla McKelvey: On November 7, 2015, Kayla McKelvey—then a 25-year-old student at Kean University in New Jersey—was attending a student rally addressing racial issues at the school. Disappointed in the turnout, she is reported to have gone to the campus library, created an anonymous Twitter account, and sent successive messages to the Kean University official account threatening both a bombing and that African American students on campus would be shot. She then returned to the rally to tell the other students about the threats. She pleaded guilty to creating a false alarm, and was sentenced to 90 days in jail and five years' probation (Rahman, 2016).

Elliot Rodger: On May 23, 2014, the 22-year-old man sent a 137-page autobiography-cum-manifesto, titled *My Twisted World* (2014), to select family, friends,

and acquaintances. Towards the end of this document, Rodger discusses at length "every single fantasy I had about how I would punish my enemies" (Rodger, 2014: 132), essentially detailing his ideal agenda for the day of the murders. The non-realized excerpt details his desire to trick male strangers into entering his apartment, where he hoped to torture and kill them. Police reconstructions of the day in question indicate that Rodger never attempted this (Naranjo, 2015), although he successfully carried out other portions of his longer 'fantasy'.

Skyline: On September 19, 2012, an anonymous user on 4chan, a site which allows users to share images as well as to interact through writing and comments, posted a picture of a gun with accompanying text vowing to *open fire on the people in the commons in the morning* at Skyline High School in Washington state. Classes were canceled at the school although no attacks materialized. While the author's identity was eventually uncovered by authorities, he remains unnamed publicly because the then 16-year-old was still legally a minor. After searching the teen's home, officials seized computers and cell phones "but no guns were found" (KOMO Staff, 2012). Investigators also stated they did "not believe the suspect had 'active plans or means' to carry out the attack" (KOMO Staff, 2012).

Gilberto Valle: Over the course of roughly ten months in 2012, New York police officer Gilberto Valle engaged in chat room discussions on a site called Dark Fetish Network about his desires to kidnap, murder, and eat particular women, including his own wife (Moyer, 2015). After discovering some of this content on a shared laptop, his wife alerted authorities. Valle was subsequently convicted of a conspiracy to commit kidnapping and jailed for 20 months until the convictions were overturned by two higher courts. The judge in the final appeal wrote that the case was "about the line between fantasy and criminal intent" and that the "chats are replete with references to fantastical

elements such as a human-sized oven, a spit, and a remote cabin in the woods, none of which Valle owned or made any effort to acquire" (Moyer, 2015).

3.2.2 REALIZED CORPUS

The classification of these texts is more surefooted than their non-realized counterparts because (unfortunately) the authors left measurable marks of their violent intentions on the real world. However, realized pledges in general were more likely to be hidden behind legal challenges, making this corpus the most difficult to populate.

A brief summary of the context surrounding each realized pledge is provided below.

Alex Hribal: In a handwritten essay dated April 6, 2014, the high school student in Pennsylvania explained his reasons for wanting to hurt his classmates using religio-philosophical themes, e.g., *I could say life is evil, because it blocks everyone in the world of heaven where all is good*. On April 9, he slashed or stabbed 21 people at Franklin Regional High School (Routh, 2018).

Kip Kinkel: On May 20, 1998, the day he was expelled from Thurston High School in Oregon, Kinkel shot and killed his parents at their home. He then composed a short, handwritten text explaining that *I have to kill people*. The next morning, May 21, Kinkel drove to the high school and opened fire on students in the cafeteria, killing two and wounding 25 (PBS, 2014).

Gavin Long: On July 17, 2016, a 29-year-old former Marine Corps sergeant, Gavin Long, emailed an internet acquaintance—an Ohio resident with whom he had communicated online. The email was addressed to people he called *FAMILY*, and in it Long argued that *I must bring the same destruction that bad cops continue to inflict upon my people*, i.e., African Americans. An hour after sending the message, Long

ambushed six policemen in Baton Rouge, killing three. Long himself was killed in the altercation (Berlinger, 2016).

Elliot Rodger: The realized portion of Rodger's writing contains ideas for his assault on one of the U.C. Santa Barbara sororities. A timeline of the day of the attacks shows that he tried to gain access to the sorority house and, failing that, shot three young women who were nearby, killing two (Naranjo, 2015). Although the violence he planned for the sorority itself was thwarted, this excerpt nevertheless represents an attempt on his part to carry through with his ideations.

Dylann Roof: 22-year-old, South Carolina resident, Dylann Roof, maintained a personal website devoted to white supremacist propaganda titled The Last Rhodesian. On June 17, 2015, Roof used the site to publish a 'manifesto,' at the end of which was a 196-word section which he called *An Explanation*. In this, he argued that *someone has to have the bravery to take it to the real world, and I guess that has to be me*. Hours later, Roof walked into an historically Black church in Charleston and shot to death nine of the parishioners (Hersher, 2017).

Tyrelle Shaw: In a bog post internally dated June 7, 2015, but published on June 18, the 25-year-old Bronx native announced his intentions to *hit over a million Asian Women in the face with a stick*. Between June 10 and 15, he used a hammer to assault four separate Asian women in various Manhattan neighborhoods. On June 22, his body was discovered in the bottom of an elevator shaft and ruled a suicide (Barker, Wright, & Goodman, 2015).

3.3 METHOD

"Judgments of intentionality set the course of social interactions" (Malle & Knobe, 1997: 101). A comment may be taken either as an insult or merely as a social misstep, for instance, depending on if the speaker is perceived as intending the insult (Malle &

Knobe, 1997). To decode the actions of others, including their speech acts, social perceivers try to glean the intent behind them—what the doer or speaker is trying to accomplish by acting or speaking. And so, once intentions have emerged from the invisible world of cognition and become visible in the realm of behavior, they take on a "social role" (Malle & Knobe, 1997: 102). Thus, in social contexts, intentions function *interpersonally*.

To understand how intentions may be encoded in language, then, requires a method that can "reveal much about the author's underlying position, emotion, and intent" (Gales, 2010: 90). Appraisal analysis is ideally suited to this task, first because it is expressly built to measure "the interpersonal in language" (Martin & White, 2005: 1). However, it does so by attending to a panoply of functional resources authors may call on to achieve their communicative purpose, e.g.: their stances towards the content of their message as well as toward their audience; how they administer praise and blame and position their readers to do the same; the personae they construct for themselves; and the "linguistic mechanisms" they call on "for the sharing of emotions, tastes and normative assessments" (Martin & White, 2005: 1). While "[n]obody can know for sure, of course, what another person's intentions are" (Shuy, 2005: 15), Appraisal's elaborate, interlocking systems and subsystems are an excellent tool for measuring what an author is trying *to do* with language, and thus on the intentions underlying the language itself.

The method itself is expounded more fully below. In terms of process, analysis was performed as follows: after the fourteen texts were collected, each was coded by hand according to the Appraisal scheme laid out by Martin and White (2005). Next, this manual coding was replicated in UAM Corpus Tool (O'Donnell, 2013)—a specialized concordancer discussed in section 3.2.1 below. Corpus Tool provided both raw counts as well as automatically normed frequencies for the various Appraisal variables in each text and across each corpus. The question at this stage was one of occurrences, i.e.,

how often does a particular class of tokens appear in one corpus versus the other. For "frequency data" of this kind, "chi square analysis is the best one to use" (Grant, Clark, Reershemius, Pollard, Hayes & Plappert, 2017: 96). Thus, normed tallies were entered into Microsoft Excel, and the chi-square statistical test was applied to the counts.

Once areas of statistical significance had been identified, these became the focus of various grammatical and semantic investigations. Broadly speaking, then, this approach employed *quantitative* measures to guide subsequent *qualitative* analyses.

3.3.1 UAM CORPUS TOOL 3

As shown above, the fourteen pledges combine for a total of 4,190 words. While a word count this small would usually be odd fodder for procedures drawn from corpus

Figure 3.1: Example of Auto-Normed Corpus Tool Output

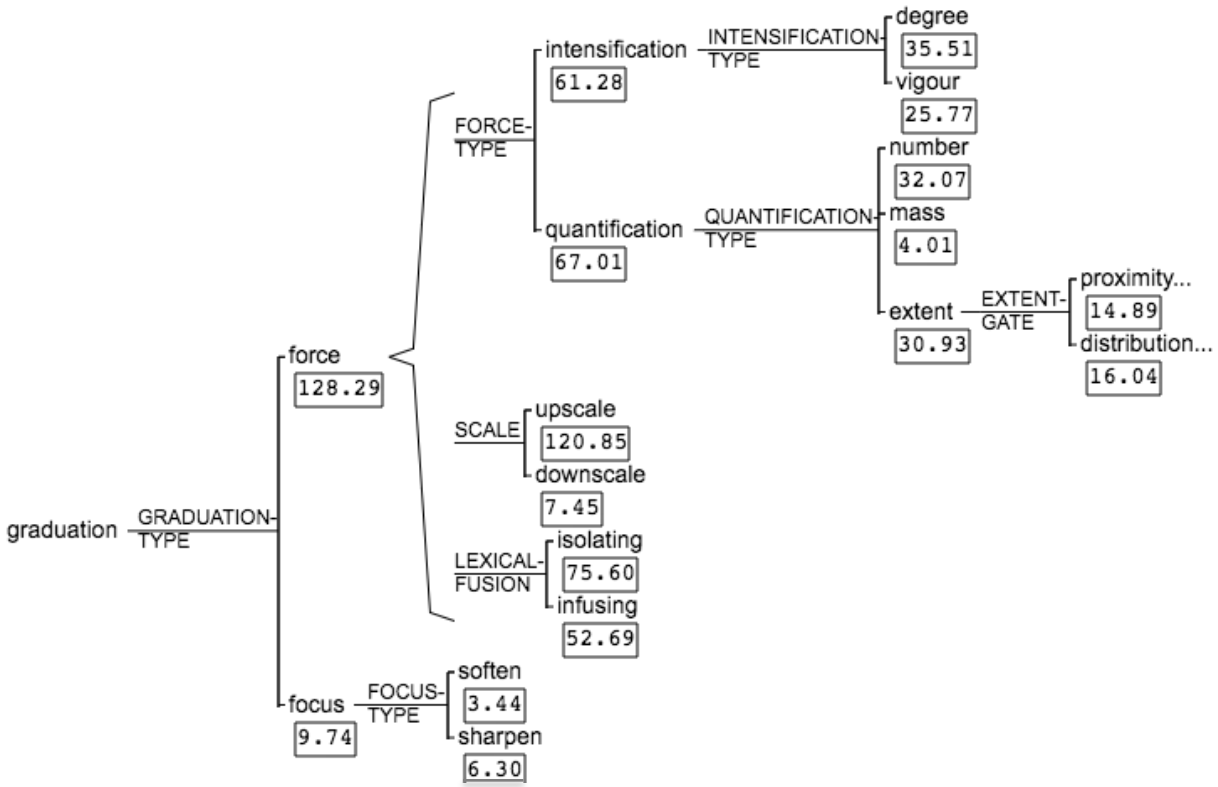

Frequency per 1000 words

linguistics—a field which often analyzes word counts in the millions—Appraisal's close-reading methodology provides ample analytical grist from a minimum of language, so much so that a specialized corcordancer, called UAM Corpus Tool (O'Donnell, 2013), was enlisted to help quantify the data. At the time of this writing, Corpus Tool is a free software application[15] whose interface allows users to annotate for Appraisal, as well as offering ordinary concordancing functions, e.g., n-grams, collocations, etc.

The program also provides a range of output options, including Appraisal variables automatically normed to 1,000 words, for easier cross-corpora comparisons. As an example, the normed counts for the system of graduation in the non-realized pledge corpus are shown in Figure 3.1.

3.3.2 APPRAISAL AND SYSTEMIC FUNCTIONAL LINGUISTICS

In Systemic Functional Linguistics, language is theorized as construing human experience across three simultaneously operating 'metafunctions' (Halliday & Matthiessen, 2004). Lexicogrammatical resources which are used to transform experience into meaning, for instance, are said to serve the 'ideational' metafunction. This captures the fact that speakers use language representationally, to communicate a process of "doing or happening, saying or sensing, being or having—with its various participants and circumstances" (Halliday & Matthiessen, 2004: 29). Or, put more simply, this metafunction conveys "what the clause is about" (Halliday & Matthiessen, 2004: 309).

But language is not just representational; it is also social. Speakers present listeners with propositions and proposals which "inform or question, give an order or make an offer" (Halliday & Matthiessen, 2004: 29). They also "express [their] appraisal

[15] http://www.corpustool.com/index.html

of and attitude towards whoever [they] are addressing and what [they] are talking about" (Halliday & Matthiessen, 2004: 29). Where the ideational metafunction *represents* experience, these propositions and proposals instead actively "negotiate[e] social relations" (Martin & White, 2005: 7). This is the 'clause as exchange' (Halliday & Matthiessen, 2004: 59), and because it is both "interactive and personal" (Halliday & Matthiessen, 2004: 30) this second metafunction is referred to as 'interpersonal.'

Finally, both the ideational and interpersonal motifs are facilitated by a third, 'textual' metafunction, which is "concerned with information flow" (Martin & White, 2005: 7) and is a means for speakers to "creat[e] cohesion and unity" across spans of discourse (Halliday & Matthiessen, 2004: 30). In any SFL-based analysis, focusing on one metafunction does not exclude the others; it merely backgrounds them. This is because systemic theory's approach to language is *comprehensive*, i.e., "it is concerned with language in its entirety, so that whatever is said about one aspect is to be understood always with reference to the total picture" (Halliday & Matthiessen, 2004: 19). As a methodological extension of systemic theory, Appraisal is therefore sensitive to each of the three metafunctions to some degree. However, its main purpose is developing and extending an account of the *interpersonal* in language (Martin & White, 2005: 1).

Essentially, Appraisal is a discourse semantic method for analyzing how language users evaluate themselves and their surroundings, thus providing a framework for classifying the "subjective presence" of a writer in his or her text (Martin & White, 2005: 1). Appraisal's emphasis on interpersonal meaning-making encompasses every area this study hopes to illuminate. Linguistically, pledging to harm is an inherently interpersonal action, both in regard to the target(s) of the writer's antagonism and to how this antagonism is presented to the pledge's audience. The reasons for a pledge—whatever intent may underlie it and infuse it—is also fundamentally a relational issue.

As Malle *et al.* (2001: 1) state, intentionality is "a tool with manifold functions, ranging from the conceptual to the interpersonal and even to the societal." Appraisal is thus an ideal method for arriving at a *linguistic* understanding of the *psychologically* generated intentions which spur the creation of a pledge text, as well as the writer's level of affective and epistemic commitment to those intentions.

Appraisal itself is divided into three separate but interlocking systems, each aimed at uncovering a different set of prosodic meanings: attitude, engagement, and graduation. Each system is discussed in turn in the next few subsections.

3.3.2a **ATTITUDE**

The first system of Appraisal, *attitude*, "highlights how feelings are mapped within texts" (Gales, 2011: 30). The "kinds of feelings" a writer may seek to communicate are further divided into "three semantic regions covering what is traditionally referred to as emotion, ethics and aesthetics" (Martin & White, 2005: 42), concepts which map to the areas of *affect*, *judgement*, and *appreciation*, respectively.

Affect, refers to the "emotive dimension of meaning," or, more specifically, how a writer communicates a personal feeling, e.g., "do we feel happy or sad?" (Martin & White, 2005: 42). This is further delineated into four subtypes, or variables, of emotions: un/happiness, in/security, dis/satisfaction, and dis/inclination. Each set, as indicated by the forward slash, contains both the positive and negative instantiations of the feeling. According to the authors:

> the un/happiness variable covers emotions concerned with 'affairs of the heart'—sadness, hate, happiness and love; the in/security variable covers emotions concerned with ecosocial well-being—anxiety, fear, confidence and trust; the dis/satisfaction variable covers emotions concerned with telos (the pursuit of goals)—ennui, displeasure, curiosity, respect. (Martin & White, 2005: 48)
>
> Finally, dis/inclination covers the range of mental processes and their behavioral counterparts on the cline of fearing to desiring. (Martin & White, 2005: 48)

Where affect concerns the inner experience of an author, judgement addresses the writer's "attitudes to people and the way they behave" (Martin & White, 2005: 52). These feelings are directed outward by the writer, appraising the social and moral behavior of external human actors. The area is comprised of two subcategories—social esteem and social sanction—which are differentiated according to 1) the personal qualities being assessed and 2) by the severity of the consequences which may follow a perceived breach. Social esteem concerns more quotidian evaluations of character, such as how special a person may be ('normality'), how capable they are ('capacity'), and how resolute ('tenacity'). These three values are "critical to the formation of social networks (family, friends, colleagues, etc.)" (Martin & White, 2005: 52). Social sanction, on the other hand, is potentially the more serious, encompassing a person's penchant for truthfulness ('veracity') and ethical behavior ('propriety'), values which "underpin[] civic duty and religious observances" (Martin & White, 2005: 52). Where too much negative social esteem may lead to ostracization from a social network, more extreme transgressions of social sanction are instead handled through the legal system and corrected through state-decreed punishments. "[T]oo much negative esteem, and we may need to visit a therapist; too much negative sanction, and a lawyer may need to be called in" (Martin & White, 2005: 53). Or, to use a different analogy, the difference is "comparable to that [of] venial [and] mortal sins" (Martin & White, 2005: 53).

The final category of the system of attitude, appreciation, speaks to "meanings construing our evaluations of 'things,' especially things we make and performances we give, but also including natural phenomena—what such things are worth (how we value them)" (Martin & White, 2005: 56). These objects may be either concrete (e.g., a gun), or abstract (e.g., an idea). Appreciation is subdivided into three further metrics regarding the inanimate entity's: 1) 'value' (its worthiness); 2) 'composition' (its balance and complexity); and 3) what 'reaction' it elicits (its impact and its quality). Interestingly, this

area shows evidence of the comprehensive character of systemic theory, in that appreciation has echoes of affect within its 'reaction' variable (Martin & White, 2005: 57). Similarly, "positive and negative valuations of something imply positive and negative judgements of the capacity of someone to create or perform" (Martin & White, 2005: 58). Along with these correlations with capacity, valuation is also especially sensitive to the goals of the text's ideational metafunction (Martin & White, 2005: 57). The practical manifestations of these overlaps and echoes are explored in the analysis of appreciation in Chapter 6.

3.3.2b ENGAGEMENT

The resources of stancetaking in the system of attitude work hand-in-hand with the system of engagement, which situates a given text within a larger dialogue of other texts and voices. The heart of this system is the Bakhtinian idea that every utterance stands in some relation to the many utterances which have come before it. Those which stand in isolation from other texts, such as bare assertions which do not "referenc[e] other voices or recogniz[e] alternative positions" (Martin & White, 2005: 99), are conceived as being 'monoglossic.' However, truly monoglossic utterances are rare to the point of vanishing. In fact, Gales (2010: 229) argues that *no* utterance is free from the subjective presence of the speaker, and thus "no utterances can be monoglossic under this framework." Nevertheless, the very few instances of monoglossia which are identified in the dataset are treated in Chapter 7.

The main thrust of engagement is rather towards what the authors call 'heteroglossia,' i.e., the "backdrop of prior utterances, alternative viewpoints and anticipated responses" (Martin & White, 2005: 97) which shape the current verbal performance. Either the utterance closes the text to these outside voices, *contracting* the dialog in a way which prohibits or excludes viewpoints other than the author's, or it

opens the text to these alternate positions, thereby *expanding* the dialog to include room for opinions which the author may not explicitly share (Martin & White, 2005: 102). In essence:

> all verbal communication, whether written or spoken, is 'dialogic' in that to speak or write is always to reveal the influence of, refer to, or to take up in some way, what has been said/written before, and simultaneously to anticipate the responses of actual, potential or imagined readers/listeners. (Martin & White, 2005: 92).

Writers may contract their texts, for example, through devices like pronouncements (e.g., *the facts of the matter are…*) and denials (e.g., *no, didn't, never*). Conversely, they can expand their texts by entertaining alternatives (e.g., *this may be*) or by explicitly acknowledging prior utterances (e.g., *Halliday argues that…*) (Martin & White, 2005: 134). In all, the system of engagement recognizes nine complementary rhetorical actions available to an author.

In terms of threatening, both dialogic motions are of interest, as Gales's (2010) work has shown. For example, dialogic contraction serves to "strengthen the threatener's commitment towards, role in, or responsibility for the threatened action," while expansion "weakens each of those interpersonal functions" (Gales, 2010: 157). Thus, where the system of attitude provides a means of understanding how threateners feel about themselves and their targets, engagement offers a way to gauge both how threateners view the content of their own texts and how they anticipate their audience will react to this content.

3.3.2c GRADUATION

Finally, the meanings within both attitude and engagement are scalable, i.e., speakers may choose between various degrees of positivity and negativity when making an assessment. Tokens of attitude, for instance, may be intensified or de-intensified, as in the choice between someone being judged as *good* versus *brilliant*, or an author

describing their current affective state as *happy* versus *ecstatic* (Martin & White, 2005: 136). Options like this are similarly available at the syntactic level within engagement, as in the choice between the two disclamations *I didn't hurt him* and *I never hurt him* (Martin & White, 2005: 136). "It might be said that attitude and engagement are domains of graduation which differ according to the nature of the meanings being scaled" (Martin & White, 2005: 136). The semantics of graduation are therefore considered key to the Appraisal method as a whole.

This system is divided along two axes, one called 'force' and the other 'focus.' These attend to two broad resources which speakers have available for up-scaling and down-scaling meanings. Force applies to "categories which involve inherently scalar assessments" (Martin & White, 2005: 137), such as the intensity of a quality (e.g., *slightly foolish* versus *extremely foolish*) or a process (e.g., *slightly hindered* versus *greatly hindered*). Force also covers assessments of amount which "apply to entities, rather than to qualities and processes" (Martin & White, 2005: 141), such as number (*a few*), mass/presence (*small*), distribution (*long-lasting*/*widespread*), and proximity (*nearby*/*recent*) (Martin & White, 2005: 154). Furthermore, this axis is sensitive to whether the de/intensification is accomplished through 'infusing' a lexeme, wherein "the scaling is conveyed as but one aspect of the meaning of a single term" (Martin & White, 2005: 143)—e.g., *happy* versus *joyous*—or "through an isolated, individual item which solely, or at least primarily, performs the function of setting the level of intensity"—e.g., the freestanding adverbials *somewhat miserable* and *very miserable* (Martin & White, 2005: 141).

Focus, on the other hand, applies to meanings which are not inherently scalar but rather operate prototypically. "These are the clearly bounded, either–or categories which operate in experiential taxonomies where category membership is more or less precisely determined by some combination of sufficient and necessary conditions"

(Martin & White, 2005: 137). In such cases, graduation situates the assessment either closer to the semantic prototype or further into the fuzzy boundaries of the category. For instance, "[f]rom an experiential perspective, jazz music is a distinct category," and yet a speaker may evaluate a specimen of music as being more prototypical of the genre (*real jazz*) or less (*sort of jazz*) (Martin & White, 2005: 137).

In all, graduation offers a means of coding positive and negative instantiations of scaling across roughly eight variables. Thus, where attitude addresses an author's stance towards the people and things discussed in the text, and engagement addresses his or her stance towards the content and audience of the text, graduation parses the *intensity* of these stances, opening a window on the relative 'loudness' and 'temperature' of a particular communication.

3.3.3 Risks and Limitations of the Appraisal Method

Every method of linguistic analysis offers advantages and disadvantages to a researcher. By illuminating some patterns in a dataset, others are necessarily obscured or overlooked. In this respect, Appraisal is no different. Despite being "the most fully developed model of evaluation" (Thompson, 2014: 48), the method still carries with it a few risks and limitations.

The first risk arises from Appraisal's nature as a discourse semantic method. Martin and White (2005: 70) note that—as is common with discourse analysis—one may "start[] with prosodies and work[] down to their realisations or start[] with realisations and work[] back to the 'mood' of a text." However, a temptation particular to Appraisal is to view the realizations through the lens of the prosody, concentrating on the implicit meanings of a text without accounting for the explicit language it contains. In other words, there can be a strong temptation to conflate the *coding* process—which Gales (2020: 679) characterizes as "steps to interpretation"—with the process of *analysis*—i.e.,

the “revealed patterns of prosodic meaning” (Gales, 2020: 687). But doing so minimizes an author’s black-and-white lexicogrammatical choices by ‘reading ahead’ to the interpersonal effects the author was perhaps trying to achieve.

This danger is both explicit and implicit in several of the issues that, e.g., Thompson (2008, 2014) raises about Appraisal in practice. For one, he flags the potential conundrum posed by nominalizations of behavior, as in “His catching was brilliant” (Thompson, 2008: 178). Even where such nominalizations are clearly presented by an author as *products*, and thus as tokens of appreciation, “the evaluative terms chosen can sometimes be associated with JUDGEMENT” (Thompson, 2008: 179). The issue is also apparent in what he calls, alternately, the ‘Chinese box syndrome’ (2008) and the ‘Russian doll dilemma’ (2014). This is when “one kind of appraisal is nested inside another kind” (Thompson, 2008: 183), e.g., how the negative appreciation of *I found his notes antipathetic* is used to provoke in the reader a negative judgement of the notes’ author (Thompson, 2008: 182). Both examples invite the analyst to blur the semantic boundaries within attitude. But the temptation would result in basic coding choices that are “dangerously impressionistic” (Thompson, 2008: 169).

In both cases, Thompson (2008: 180) rightly concludes that “the wording should be taken as the basis for the initial assignment of categories” when coding. Or, more simply, analysts should “trust the text” (Thompson, 2008: 180). This is because, in practice, nominalizations of behavior may indeed feel like judgements of the person behaving. Yet, jumping past appreciation would ignore the potential significance of the author’s focus on behavior, especially if an explicit judgement of the person might have been expected. Similarly, if an analyst tracks the layers of Russian doll nesting outward, “it can at least be made explicit that each step represents a further move into interpretation” (Thompson, 2008: 184), i.e., a move from coding into analysis. Both issues require recognizing, as Gales (2020) does, that coding and interpretation are

interdependent but nonetheless *different* stages in the process, and that an analyst should be wary of conflating the prosody with how it is realized. Great care has been taken here to do just that.

A second risk lies in what Martin and White (2005) call 'invoked evaluations.' These stances are communicated implicitly, without overt attitudinal lexis, thus asking the reader to recover the stancetaking through contextual signals (Martin & White, 2005). For instance, in the sentence *That child tears the wings off butterflies*, a listener would need a measure of shared world knowledge to know that the underlined verb phrase should be read as a token of negative judgement of the child in question, an interpretive process which Thompson (2008) is suspicious of. And arguably, his suspicion is warranted. Martin and White (2005) themselves recognize the potential pitfall, acknowledging that coding for invoked evaluations may introduce "an undesirable element of subjectivity into the analysis" (p.62). Nevertheless, they see a greater danger in ignoring them, and so they remain part of Appraisal's model of attitude.

However, Martin and White (2005) offer no similar defense of lexical infusion in the system of graduation, which arguably poses a similar risk of impressionistic coding. In Appraisal, infusion occurs when there is "no separate lexical form conveying the sense of up-scaling or down-scaling" (Martin & White, 2005: 143); instead, the de/intensifying is an aspect of a single term. Thus, *joyous* is considered an up-scaled infusion of the term *happy* (Martin & White, 2005: 144). While this may make intuitive sense, no baseline is offered. If *joyous* is an infused version of *happy*, should *happy* be coded as an infused version of *contented*? Is there a set of bedrock lexis that may be viewed as un-infused? If not, where does the process of lexical infusion begin? And if there is no identifiable beginning then might not all stance markers be treated as infused? In which case, would the concept not lose all analytical utility? For these reasons, both invoked evaluation and infusion have been treated here with great care

and are often flagged for the reader when they enter into the analyses (e.g., in the discussion of explicitness in section 4.2).

Finally, although “interpersonal meanings are distributed in waves of semiosis, including interconnections among waves” (Martin & White, 2005: 7), the main limitation of Appraisal is that it identifies the crests of these waves, but not the interconnections. Actual prosodic patterns are left for the analyst to piece together in the interpretive phase of the analysis. (This is likely for the best, since every text is different.) Still, this occasionally makes it necessary to step *outside* the method if the significance of the patterns which Appraisal has highlighted are to be fully understood. For instance, according to Nini (2017), threats directed at third parties tend to be more narrative in form. But despite the rhetorical focus in the system of engagement, Appraisal is surprisingly ill-suited to analyze narrative structures. Thus, to understand whether the variable of negative composition explored in section 6.1.1 is somehow related the ‘story’ of a pledge, the analysis reaches beyond Appraisal to the tools of Narrative Analysis (e.g., Labov, 2003, 2010, 2013). Elsewhere, Appraisal may approach but fail to directly touch categories important to the threat assessment literatures, e.g., the specificity of imagined targets. Any departures from Appraisal’s framework are explicitly noted and justified when they are deemed necessary. Generally, though, they are few in number.

All that said, even with these risks and limitations Appraisal Analysis remains “a rich methodological resource” (Gales, 2020: 691) as “the most fully developed model of evaluation” (Thompson, 2014: 48).

3.4 SUMMARY

The dataset under examination consists of fourteen texts for a total of 4,190 words. These texts have been divided into two categories: six pledges are classified as realized (R texts), meaning the authors attempted to enact the violent ideations outlined in their

writings; eight pledges are non-realized (NR texts), where records indicate that no such attempt occurred. As human-produced data, the pledges were collected and analyzed in as ethical a manner as possible, with an understanding of their illicit nature and the role which such a nature plays in concerns about authorial privacy.

Once the dataset was compiled, each pledge was hand-coded according the scheme laid out in the Appraisal method, as summarized briefly in subsections 3.2.2a through 3.2.2c above. This coding was then replicated in UAM Corpus Tool, a specialized concordancer built to accommodate Appraisal variables. Corpus Tool provided both raw and automatically normed feature frequencies, to which the chi-square statistical test was applied. After areas of significant distributional difference were identified (e.g., meanings of capacity and propriety within the subsystem of judgement), various grammatical and semantic investigations were undertaken. The results of these investigations are discussed in the chapters of analysis that follow.

CHAPTER 4 ATTITUDE

As noted in Chapter 3, meanings which construe a speaker's feelings are grouped in Appraisal within the system called attitude. This is delineated into three semantic areas based on their objects of focus: affect concerns the speaker's own, personal emotions; judgement deals with a speaker's attitudes towards other people, particularly their behaviors; and appreciation addresses the speaker's view of 'things,' i.e., "semiotic and natural phenomena" (Martin & White, 2005: 43). As shown in Table 4.1 below, each of these areas is comprised of a small set of code-able variables, with each variable capturing both positive and negative instantiations of the feeling. So, for instance, when the LA Unified NR author complains about *the loneliness* of high school this is logged as a token of negative security (-sec), the variable within affect covering "ecosocial well-being—anxiety, fear, confidence and trust" (Martin & White, 2005: 49). And when the same author later asserts that s/he is a *devout Muslim*, this is considered a positive instantiation of the same variable, and thus coded as +sec.

The importance of the evaluated object means that attitude is also sensitive to who or what is being appraised, not just to the appraiser. In a subgenre of threats targeting a grammatical 3rd Person, analyzing an author's feelings towards various stance objects can be illuminating. For example, is the target of the pledge the sole focus of the text or are other parties present as well? How are the various participants evaluated? Along with this, distinguishing between *appraisers* can be valuable. Pledge authors may choose to quote external voices, for instance, thereby "co-opting some authoritative second party to the current rhetorical cause" (Martin & White, 2005: 103). Understanding the stance work which an outside voice is forced to perform may shed light on the attitudinal meaning under examination.

Finally, attitude also asks whether the feeling is encoded explicitly, via attitudinally-laden lexemes and phrases, or whether the feeling is left implicit, requiring the reader to recover the meaning through contextual signals. The prior is referred to in Appraisal as 'inscribing,' while the latter is called 'invoking' (Martin & White, 2005: 62). For example, Hribal R employs both tactics to communicate judgements of negative tenacity (-ten), the variable which asks how resolute or dependable the speaker considers other people. Hribal R 'inscribes' his attitude towards the teachers of his high school—i.e., he tells the reader "directly how to feel" (Martin & White, 2005: 62)—when he characterizes them as *Laziness* incarnate. However, he communicates a similar disdain towards his fellow students more subtly when he refers to them as a *herd of sheep*. At the time of Hribal R's writing in 2014, this was a fairly common metaphor in American English for denigrating a group of people as mindless followers or automatons. Yet, there is nothing inherently attitudinal or negative denoted by either of the nouns (*herd*, *sheep*) in particular. Instead of an inscribed judgement, then, Hribal R uses the linguistic context (helped along by shared cultural knowledge) to 'invoke' the ideational meaning of 'mindless followers.'

In these various ways, Appraisal facilitates a deep-bore view of attitudinal data across a range of intersecting metrics: what kind of emotion is being expressed? to whom or at what is it directed? is it positive or negative? is it explicitly stated (inscribed) or is some shared cultural experience necessary to decode the evaluation successfully (invoked)?

The cross-corporal statistical probabilities of each variable are shown in Table 4.1. Areas which registered as significant through chi-square are shown in white, while areas which fell below significance are shaded.

Table 4.1: Statistical Significance in Attitude

Attitude Category		Tokens		Probability (p)
		Non-realized	Realized	
Affect	Un/happiness	9.16	9.73	> .05
	Dis/satisfaction	5.15	8.79	> .05
	In/security	6.30	7.22	> .05
	Dis/inclination	13.75	9.73	> .05
Judgement	Normality	9.74	17.26	> .05
	Capacity	36.08	13.49	< .01
	Tenacity	1.15	6.90	< .05
	Propriety	15.46	29.49	< .05
	Veracity	1.15	2.20	> .05
Appreciation	Reaction	5.15	4.39	> .05
	Composition	26.35	5.02	< .001
	Social-valuation	9.16	26.67	< .01
Polarity	Positive	49.26	67.15	< .05
	Negative	88.20	73.74	< .05
Explicitness	Inscribed	118.56	111.39	> .05
	Invoked	18.90	29.49	> .05
Appraiser	Writer-appraiser	134.59	168.06	> .05
	Other-appraiser	2.86	2.82	> .05
Appraised	Self	29.78	31.69	> .05
	Other	107.67	109.19	> .05

Frequency per 1000 words

Plainly, both corpora are dominated by expressions of judgement and, to a lesser extent, appreciation. Indeed, these two areas are discursively rich enough to warrant their own freestanding chapters of analysis (Chapters 5 and 6, respectively). This chapter instead examines the distributional significance of polarity in the pledges and briefly analyzes the potentially interesting null results found within the attitude system as a whole.

4.1 ATTITUDE POLARITY

"Opinions about entities," of the kind that attitude is built to catalog, "tend to be realized lexically" (Martin & White, 2005: 38). Realizations may either be isolated attitudinal markers (e.g., *make them <u>suffer</u>* [Rodger R]) or as an aggregate of meanings in the form of phrases (e.g., *no one <u>doing anything but talking on the internet</u>* [Roof R]). While a single marker may be of interest by itself, such as the dehumanizing process inherent to

calling his targets a *herd of sheep* (Hribal R), more interesting still are the patterns of stancetaking which these individual tokens enact across a text. As Martin and White (2005: 10) note, "the realisation of an attitude tends to splash across a phase of discourse, irrespective of grammatical boundaries" (Martin & White, 2005: 10). Such splashing forms a 'prosody' of attitude (Halliday & Matthiessen, 2004; Martin & White, 2005), and is central to the way interpersonal meanings are realized.

As noted above, each Appraisal variable recognizes the polarity of positive and negative evaluations, i.e., "traits we admire alongside those we criticize" (Martin & White, 2005: 52). Approaching the polarity of these tokens collectively—negative grouped with negative and positive with positive—it is possible to survey the prosody of this elemental stancetaking choice in the realized and non-realized corpora, both alone and in relation to each other. The resulting landscape, shown in Figure 4.1, is unsurprising in one way and fascinating in another.

Figure 4.1: Attitude Polarity by Corpus

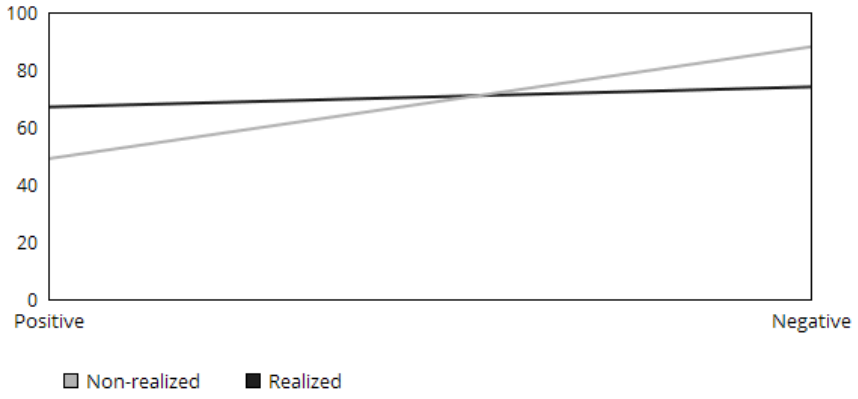

Frequency per 1000 words

In texts devoted to grievance and violent ideation, it is no surprise that pledge authors, as a whole, more consistently express negative feelings. Across the dataset, negative attitudinal markers appear at a rate of 162 tokens per 1000 words, while positive tokens appear at the lesser rate of 116 per 1000. This is in keeping both with

the general character of threatening language as a genre “committed to violence and threatener control” (Gales, 2010: ii) and to leakage in particular, which often features negatively-charged themes of “violence, hopelessness, despair, hatred, isolation, loneliness, nihilism, or an ‘end of the world’ philosophy” (Meloy & O’Toole, 2011: 516).

Figure 4.2: Positive Attitude Dispersion by Corpus

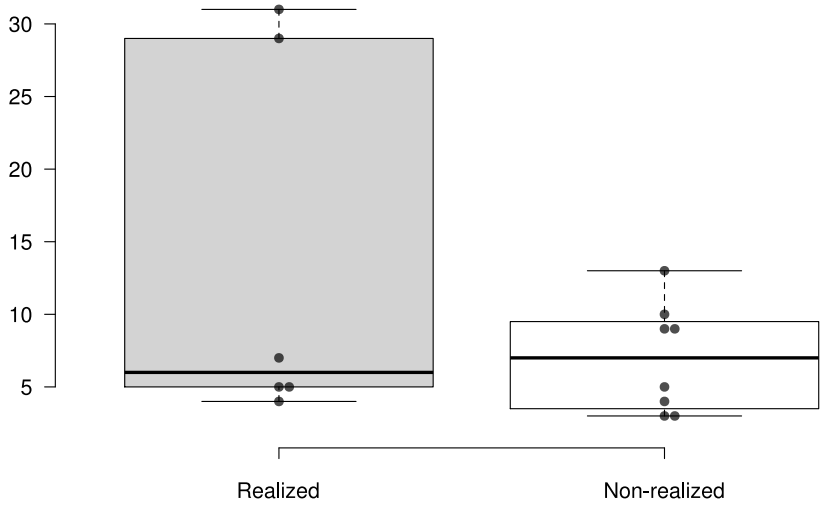

Frequency per 1000 words. No outliers (>1.5IQR) are found.

Figure 4.3: Negative Attitude Dispersion by Corpus

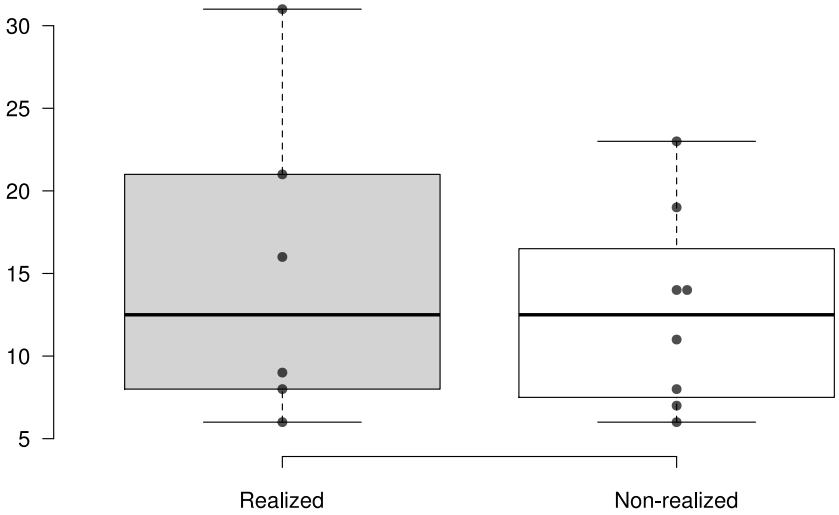

Frequency per 1000 words. No outliers (>1.5IQR) are found.

However, the distribution across the corpora themselves is intriguing, as the above Figures 4.2 and 4.3 show in conjunction with Figure 4.1. From a folk linguistic

perspective, one might expect that texts written by people who would later realize their ideations—i.e., those whose convictions were strong enough to attempt the violence they described—that these writings would tilt decisively towards a negative attitudinal prosody. But this is not the case. The R texts operate within fields of positive emotion (67 per 1000 words) almost as much as negative (74 per 1000 words). Meaning, realized pledge authors are nearly as likely to profess feelings of *satisfaction* as they are of *dissatisfaction*, and to evaluate certain people/things as *good* as well as *bad*. In a genre purportedly born of grievance (Calhoun & Weston, 2015), such emotional balance is noteworthy.

This stands in stark contrast to the NR texts, which are more lopsidedly negative (49 positive versus 88 negative per 1000 words). Non-realized authors are far likelier to build a prosody of criticism across their pledges, one which is *not* equally offset by positive feelings of praise. This heightened negativity is interesting primarily because such relentless disparagement did not, by any account, lead to further real-world actions on the part of these authors.

This potentially cuts against a folk ideology of threatening language, which, as Gales (2010) has shown, typically focuses on those functions which *strengthen* authorial stancetaking. Folk linguistically, realized pledges might be expected to contain stronger stance functions like, e.g., criticism of others. In this way, the im/balance of polarity across the two realization types comports with Gales's (2010) own findings: "threats that have been carried out *and* those that have not been carried out are composed of a combination of functions that both strengthen and weaken the threatener's stance" (Gales, 2010: 263).

Both trends also potentially fit within a broad categorical distinction theorized in the threat assessment literature, that of 'targeted' versus 'affective' violence (Meloy, 2016). Individuals who engage in targeted violence, i.e., whose main aim is to harm a

victim (Meloy, 2016)—rather than unpremeditatedly reacting to a heightened emotional state—tended to score higher on measures of 'conceptual complexity'. This cognitive characteristic "involves the ability to see that other people or places might have different positions, values, ideas or policies" (Smith, 2006: 47). Concomitantly, "lower scores on ambivalent hostility," a trait which was *also* associated with threateners taking action, "indicate lack of paranoia" (Smith, 2006: 91). The two qualities seem to work hand-in-hand, since "it appears that the presence of lower ambivalent hostility and higher conceptual complexity together are consistent with predatory thinking" (Smith, 2006: 91). The emotional balance found in the realized texts could show this relative coolness of affect, or level-headedness, and thus be interpreted as an indicator of "deliberative thinking" (Smith, 2006: 91).

Concomitantly, the imbalance found in the non-realized pledges could be evidence of the opposite. That is, the heightened negativity could be evidence of *affective* ideations which are "driven by emotions of fear and often anger" (Meloy, 2016: 232), rather than as well-considered lead-ins to targeted violence. This would go some way to explaining why the emotional asymmetry of this corpus is more skewed than its realized counterpart. Affective responses are typically "defensive, reactive, emotional, or impulsive" (Meloy, 2016: 232). Smith (2006: 90) theorizes that higher scores on ambivalent hostility would be "more consistent with the thinking of paranoid threateners who respond to perceived threats to self." She then speculates that "[t]he act of writing the threatening communications may assist these threateners in defusing their anger" (Smith, 2006: 90). In other words, the prosody of negativity infusing the non-realized pledges may be working as an affective "'safety valve,' permitting the vicarious, but safe and harmless discharge of strong emotions" (Gellerman & Suddath, 2005: 485).

Such a wide-angle view explains little about the semantics or structures at work within the two threat types, of course. Still, it is intriguing to see potential correlations

between the data and known threat categories—realized with targeted violence; non-realized with affective violence—at such a general level. More fine-grained investigations of prosody and attitudinal polarity are undertaken in several of the analyses which follow—particularly in the discussions of capacity and propriety in Chapter 5.

4.2 ATTITUDE NULL RESULTS

The more meticulous analyses in this study concentrate on Appraisal areas which are shown to be significantly different between the two corpora. However, null results—areas where a distributional difference might have been found but was not—can be of interest as well. Four such regions are discussed here.

Affect: In Appraisal, affect is concerned with the expression of personal feelings, broken into four major sets of meanings: 'affairs of the heart' like sadness and love (un/happiness); ecosocial well-being (in/security); telos, or the pursuit of goals (dis/satisfaction); and finally, desires and intentions (dis/inclination) (Martin and White, 2005). Thus, for instance, when Kinkel R writes *I love my mom and dad so much*, the verb *love* is registered as encoding positive happiness (+hap), i.e., a positive instantiation of a feeling that is considered an 'affair of the heart,' e.g., "sadness, hate, happiness and love" (Martin & White, 2005: 49).

Affect, as an encompassment of "our everyday common sense world" (Martin & White, 2005: 45), is thus a key expressive resource, one which is theorized as underpinning the larger system of attitude. Both judgement and appreciation are viewed, essentially, as personal feelings modified to suit institutional settings, i.e., "the uncommon sense worlds of shared community values" (Martin & White, 2005: 45). Judgement, for instance, "reworks feelings in the realm of proposals about behavior," while appreciation "reworks feelings as propositions about the value of things" (Martin &

White, 2005: 45). A sketch of how affect sits "at the heart of institutionalised feelings" (Martin & White, 2005: 45) is presented in Figure 4.4.

Figure 4.4: Judgement and Appreciation as Institutionalised Affect (Martin & White, 2005: 45)

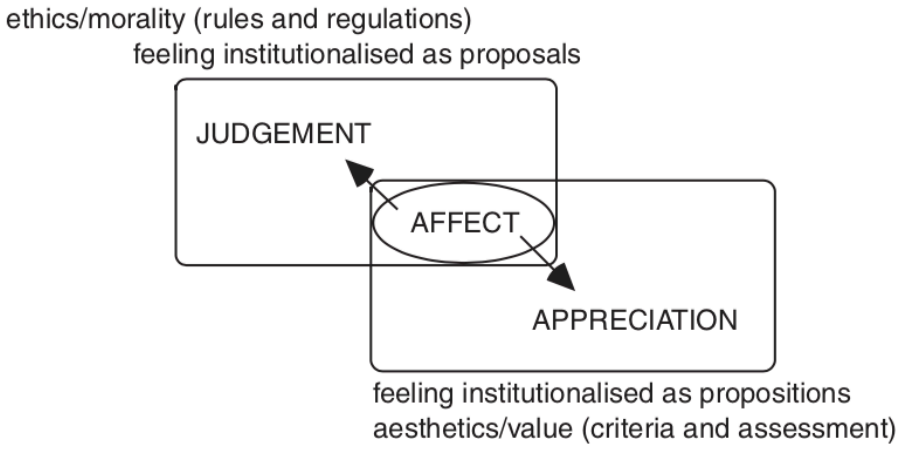

Given this relationship between the three, it is interesting to discover that the pledge authors show no significant difference in their use of the *personal* resources of affect—differing instead quite widely in how they employ the *institutionalized* meanings of judgement and appreciation. This is notable because, if a grievance is the first step on the pathway to targeted violence (Calhoun & Weston, 2015), then this negatively-charged emotion includes "a highly personal meaning for the subject" (Bulling & Scalora, 2013: 6) which can lead to similarly negative stances like a desire for revenge. As Gales (2010: 2) says, threats "are the manifestation of personal feelings, emotions, and intentions." Yet, the personal meanings expressed by both author categories deviate only minimally in their use of affective tokens-by-type, as shown in table 4.2 below.

The parity of the two corpora in this area may thus be due to nothing more complicated than the generic demands of threatening language—something common to the nature of airing a grievance, say, rather than as a resource for indicating a true commitment to violence.

Table 4.2: Tokens of Affect

Affect Variable		Tokens	
		Non-realized	Realized
Un/happiness	misery/cheer	7.45	6.59
	antipathy/affection	1.72	3.14
Dis/satisfaction	ennui/interest	1.15	4.39
	dis/pleasure	4.01	4.39
In/security	dis/quiet	4.58	4.71
	dis/trust	1.72	2.51
Dis/inclination		13.75	9.73

Frequency per 1000 words

However, two null results stand out as worthy of comment. The first is dis/inclination. This variable catalogs such overt desiderative expressions as *I <u>want</u> to use it for the good of society* (Roof R) and *As with any crime, family will be at the center of it all, something I <u>wish</u> was avoidable* (Hribal R). Also captured in this set of meanings is the modal of inclination *will*, which, in the grammatical context of a 1st Person Subject and an action verb, communicates a speaker's 'determination' to do something (Lock, 1996), e.g., *i <u>will</u> shoot any black person i see at kean university* (McKelvey NR). As far as null results go, the case of *will* is especially interesting, since this auxiliary has been called one of just three "modals of intent" by the noted threat specialist, Mardigian (2008, via Gales, 2010: 26)—along with *must* and *have to*—and thereby singled out as a lexical item warranting an assessor's special attention. Indeed, given that realized authors proceeded to action, tokens of inclination in whatever form might be expected to appear more frequently in these texts. But in fact, inclination is used at a higher rate in the *non-realized* pledges, though only slightly—14 words per 1000 here versus 10 words per 1000 in the realized corpus.

Of course, modal *will* cuts across several grammatical categories in English and plays a distinct role in signaling certain rhetorical stances in the system of engagement, and so the behavior of this lexeme will be explored more in the analyses which follow[16].

[16] E.g., Chapter 7 (engagement) and Chapter 8 (modal auxiliaries).

It is nevertheless interesting to find that such overt markers of desire like *will*, *wish*, *want*, etc., are not employed predominantly by the authors who theoretically intended to act.

The second notable null result is in the area of dis/satisfaction. As discussed in section 2.3.1, by committing violence, a threatener risks "his or her reputation, finances, freedom, and life itself" (Simons & Tunkel, 2014: 206). The reward offsetting such a risk is the "emotional profit of redress or feelings of satisfaction not attainable through socially acceptable or legally sanctioned means" (Simons & Tunkel, 2014: 206-207). Thus, as Harmon (2008) intimates, an expression of imagined satisfaction in a pledge to harm should be considered a risk-enhancing factor, and more likely to correlate with realized texts. With the slight exception of dis/interest, though, that is clearly not the case here—personal feelings of satisfaction are relatively infrequent as a whole and do not differentiate the two pledge types.

Explicitness: Explicitness is the cover term for how an author communicates attitude. Is the stance 'inscribed,' i.e., declared openly through "attitudinal lexis that tells us directly how to feel" (Martin & White, 2005: 62)? Or is it rather 'invoked,' which does not declare a stance so much as suggest it? The difference, essentially, hinges on the "degree of freedom allowed readers in aligning with the values naturalised by the text" (Martin & White, 2005: 67). For instance, lexis which directly signals authorial attitude—such as when Rodger R calls a group of women *heartless*—is considered inscribed because the adjective allows for very little interpretive freedom on the part of the reader.

Altogether, the non-realized pledges are slightly more likely to insist on an interpretation through inscription (at a rate of 119 words per 1000 versus 111 per 1000 for the realized pledges) while the realized texts are slightly more likely to depend on contextual readings to convey their stances (29 words per 1000 versus 19 per 1000 in the non-realized pledges). That the two corpora do not diverge significantly in this regard is perhaps interesting from the standpoint of Audience Design Theory. If

“speakers accommodate primarily to their addressee” (Bell, 1984: 145), then realized writers feel less need to ‘sign-post’ their emotions for their intended audiences (Martin & White, 2005: 63). This tendency could be evidence that realized authors view their addressees as closer to in-group than to out-group members and may thus rely on shared cultural values for the recovery of their intended meanings. Or, returning to Appraisal terminology, the realized authors are marginally more confident—and the non-realized authors marginally less—that their intended audiences will read their texts ‘compliantly,’ meaning as sympathetic rather than resistant to the ideations the authors present (Martin & White, 2005: 62).

Appraiser/Appraised: Lastly, Appraisal is sensitive both to who is feeling the emotion and the stance object at which the emotions are directed. In the case of pledges, the two findings—that the appraiser, by a wide margin, tends to be the author, and that the appraisals tend to be outwardly directed, also by a wide margin—are less than surprising in light of the generic features of pledging.

Pledges, like threats, involve an author who personally commits him- or herself to performing a future action (Shuy, 1993). And so, little interpretive effort is needed to understand why the pledge authors would be overwhelmingly the source of attitudinal evaluations in both corpora. If anything, it is more surprising that other voices and their appraisals are admitted into these texts at all (a point which is taken up in Chapter 7 in the analysis of engagement).

And of course, pledges, like threats, require a role complementary to the threatener—that of the threatened (Gales, 2010). This is because, while a grievance is deeply personal, it also involves a cause (Bulling & Scalora, 2013). The other-oriented prosody of both corpora thus comes as little surprise, and is arguably of a piece with the dynamics noted under affect above, i.e., that resources expressing institutionalized,

outward-directed feelings of judgement and appreciation are, by and large, preferred over personal, inward-directed emotions.

4.3 SUMMARY

In Appraisal, meanings which construe a speaker's emotions are grouped within the system called attitude, which comprises affect (personal feelings), judgement (attitudes towards other people), and appreciation (attitudes towards 'things'). Along with attitude's sensitivity to where the emotions are directed, it also asks whether the feeling is encoded explicitly, via attitudinally-laden lexemes and phrases (i.e., inscribed), or whether the feeling is instead stated implicitly, requiring the reader to recover the meaning through contextual signals (i.e., invoked).

Of the twenty attitude variables tested in the pledge dataset, a significant difference in usage between the realized and non-realized authors was discovered in seven, spread across three areas: judgement; appreciation; and overall attitudinal polarity. Judgement and appreciation are analyzed separately, in Chapters 5 and 6 respectively.

However, the results within polarity are interesting when viewed in terms of semantic prosody, a key means by which interpersonal meanings are realized in language. The dataset as a whole skews towards negative attitudinal meanings, at a rate of 162 negative tokens per 1000 words versus 116 positive. This means that pledge authors reliably opt for markers communicating criticism rather than admiration. In texts devoted to grievance and violent ideation, such a result might be expected.

Yet, the distribution of praise and blame across the two text types is less expected. The realized texts—writings by people who felt aggrieved enough to move from language to action—are almost as likely to operate within fields of positive emotion (67 per 1000 words) as negative (74 per 1000 words). By contrast, the non-realized

texts are much more negative (49 positive versus 88 negative per 1000 words). Both trends echo a broad categorical distinction discussed in Chapter 2, that of 'targeted' versus 'affective' violence (Meloy, 2016). Individuals who pursued targeted violence tended to score higher on measures of conceptual complexity and lower on ambivalent hostility, leading Smith (2006: 91) to conclude that "the presence of lower ambivalent hostility and higher conceptual complexity together are consistent with predatory thinking" (Smith, 2006: 91). The emotional balance found in the realized texts could show this coolness of affect, or level-headedness, and thus be interpreted as evidence of "deliberative thinking" (Smith, 2006: 91).

Conversely, the heavily critical prosody of the non-realized pledges could be evidence of heightened hostility, an affective reaction "driven by emotions of fear and often anger" (Meloy, 2016: 232). If so, "[t]he act of writing the threatening communications may assist these threateners in defusing their anger" (Smith, 2006: 90)—thus working as an emotional 'safety valve' (Gellerman & Suddath, 2005: 485)—and help explain why these authors did not proceed to real-world violence.

Along with the significance of attitudinal polarity, the statistical null results in this system were briefly explored. For instance, all four variables in the area of affect, which is concerned with the expression of personal feelings, are used at nearly the same rates by both author types. This is notable because a grievance is considered the first step on the pathway to targeted violence (Calhoun & Weston, 2015), and this negatively-charged emotion is understood to include "a *highly personal* meaning for the subject [emphasis added]" (Bulling & Scalora, 2013: 6). One possible explanation for this failure in differentiation may be that it is generic, i.e., that personal meanings are used with the frequency discovered here as a common way to air grievance, and that affect itself is not a resource for communicating a commitment to act.

However, the null result of one variable in particular stands out as worthy of comment—that of dis/inclination. Given that realized authors proceeded to action, tokens of inclination might be expected to appear more frequently in these texts. But in fact, inclination is used at a slightly higher rate in the *non-realized* pledges—at 14 words per 1000 versus 10 words per 1000 in the realized corpus. This is notable especially in the case of the modal auxiliary *will*, a lexeme which has been called one of just three "modals of intent" (along with *must* and *have to*) by Mardigian (2008, via Gales, 2010: 26). As a means of communicating determination (Lock, 1996), modal *will* would theoretically be a prime resource, but this is not so in terms of its usage within affect. Nevertheless, the behavior of this lexeme in the dataset will be explored more in later analyses (especially Chapter 8's direct investigation of modal auxiliaries).

Next, the explicitness with which evaluations were made also failed to approach significance. Realized and non-realized author employed in roughly equal measure both inscribed attitude, i.e., "attitudinal lexis that tells us directly how to feel" (Martin & White, 2005: 62), and invoked attitude, which suggests rather than declares a stance. In other words, if "speakers accommodate primarily to their addressee" (Bell, 1984: 145), then neither author type was more likely to craft their pledge using signals of in-group values. Instead, both realized and non-realized corpora relied heavily on lexical 'sign-posts' (Martin & White, 2005: 63) to convey their intended attitudinal meanings, inscribing their emotions directly into the text for their audience.

Finally, both pledge types overwhelmingly placed the writer in the position of the appraiser, and an outside entity in the position of appraised. This is not surprising, since pledges, like threats, involve an author who personally commits him- or herself to performing a future action (Shuy, 1993). Also, taking on the role of threatener assumes a complementary role of the threatened (Gales, 2010). Thus, the other-oriented prosody

of both corpora is, from the standpoint of pledges as a kind of threatening, also to be expected.

In sum:

- pledge writers express personal feelings at roughly the same rates, but differ significantly in how they use institutionalized feelings (judgement, appreciation);
- the non-realized texts are far more negatively oriented than the realized texts, which are more attitudinally balanced;
- there is no functional difference in how attitude is presented to the pledges' intended audiences, i.e., whether the evaluations are signaled overtly through inscribing or covertly through invoking;
- the variables of inclination and satisfaction, despite being highlighted in the threat assessment literature, are not sources of distinction between the two pledge types; and finally,
- who is appraising and who or what is appraised both show generic qualities, and are not a point of distinction between the two pledge types.

With the analysis of judgement, the next chapter will begin the exploration of those areas of Appraisal that were identified as statistically significant by chi-square.

CHAPTER 5 JUDGEMENT

In the Appraisal framework (Martin & White, 2005), the subsystem of judgement addresses an author's feelings towards the behavior of others. Five variables are available for coding within this subsystem: *normality* (how special?); *capacity* (how able?); *tenacity* (how resolute?); *propriety* (how ethical?); and *veracity* (how honest?). Of these, a statistically significant difference in usage between the R and NR corpora was detected across three: capacity, tenacity, and propriety. The two areas which fall below significance (normality, veracity) are shaded in Table 5.1.

Table 5.1: Statistical Significance in Judgement by Realization Category

Judgement Variable	Tokens		Probability (p)
	Non-realized	Realized	
Normality	9.74	17.26	< .1
Capacity	36.08	13.49	< .01
Tenacity	1.15	6.90	< .05
Propriety	15.46	29.49	< .05
Veracity	1.15	2.20	< .5

Frequency per 1000 words

A brief word on the two null results is warranted. Normality is found in statements like *I am not like my <u>family</u> or <u>friends</u> at all* (Hribal R), while veracity is seen in Long R's assessment that *good cops* are punished when they *<u>speak out</u>, & <u>point out</u> the wrong's & criminal acts of a bad cops*. Judgements of whether third parties are special (or not) do not have a direct correspondence with any language features flagged in previous research on threatening. (Keeping in mind that 'how special' is a different question from 'how specific'. Specificity of both the target and the violent act is addressed in the analysis of capacity below.) Normality does not, therefore, call for further comment. Of the two, then, the low frequency of veracity is perhaps the more surprising. The null result with this variable means, quite simply, that the matter of dis/honesty is not of special concern to this set of pledge authors. In particular, the "mortal sin" of lying (Martin & White, 2005: 53) is not one of the "behaviors [that] are

punishable in the eyes of the threatener" (Gales, 2010: 225) since the realized and non-realized pledges are found to focus on truthfulness and lying to a roughly equal degree. While questions of veracity do not differentiate the corpora, however, other ethical concerns do. These are explored under propriety below.

Turning to the three areas of statistical significance—capacity, tenacity, and propriety—the sections that follow investigate each from a *quantitative* standpoint first (e.g., a comparison of normed token counts). Numeric differences and similarities are then used as guideposts for asking *qualitative* questions (e.g., are discourse semantic patterns apparent). The least semantically rich area, tenacity, is addressed first, followed by analyses of the far richer meanings within capacity and propriety.

5.1 TENACITY

Tenacity corresponds to the grammatical modulation of inclination, specifically an author's opinion of another's "resolve" to do something (Martin & White, 2005: 54). Examples of tenacity found in the two corpora include:

- *They worked hard to achieve freedom in heaven.* (Hribal R)
- *I understand that Asian Women are afraid of African American Men* (Shaw R)
- *All Black ppl should rise up and shoot at every white cop in the nation starting NOW.* (Dickens NR)

Obviously, this encompasses a different set of meanings than those found in the variable of dis/inclination within the subsystem of affect. Tenacity does not cover *personal* feelings of resolve (which are discussed in section 4.2 above). Rather, meanings in this area relate to the resolve authors see in the people around them. Of the two pledge types, tokens of tenacity are far more likely to appear in the pledge of an author who would later proceed to act, as the dispersion values shown in Figures 5.1

Figure 5.1: Positive Tenacity Dispersion by Corpus

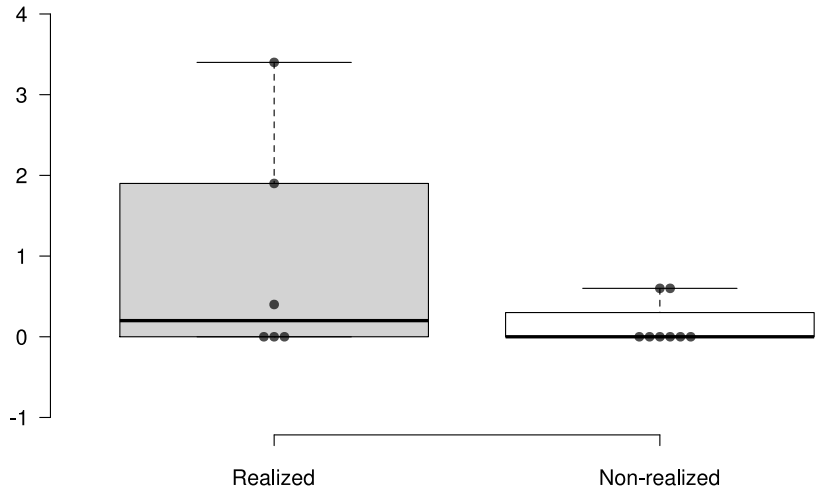

Frequency per 1000 words. No outliers (>1.5IQR) are found.

Figure 5.2: Negative Tenacity Dispersion by Corpus

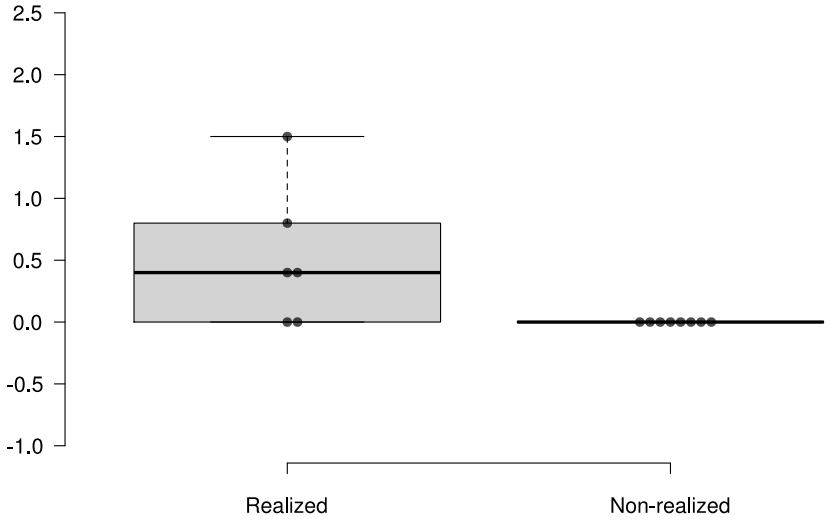

Frequency per 1000 words. No outliers (>1.5IQR) are found.

and 5.2 demonstrate. Not only that, realized authors are also more likely to admire rather than disparage the tenacity of the people they cite, signaling their heightened interest in the will to action of others.

Along with highlighting the significant difference in usage between the pledge types, Figure 5.3 also shows that tokens of tenacity are employed relatively infrequently in either corpus. Only six of the 14 authors in the dataset employ tenacity at all—two

Figure 5.3: Tokens of Tenacity by Realization Category

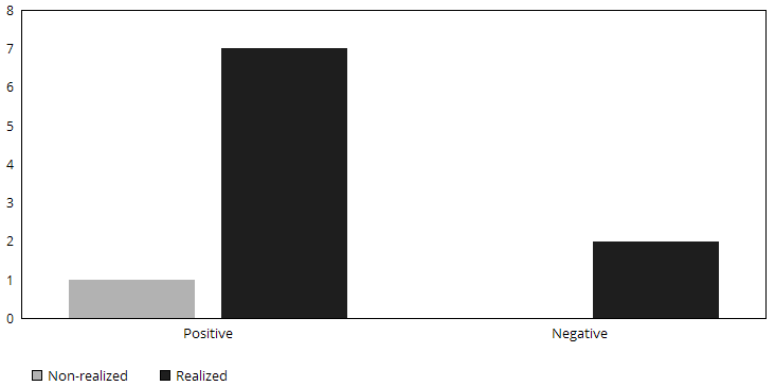

Frequency per 1000 words

non-realized authors and four realized authors. The distribution of token counts by author is shown in Figure 5.4 below. (Note that the figure has been normed to 100 rather than 1000 to more accurately capture just how rare a resource tenacity proves to be across the pledges.)

Figure 5.4: Tokens of Tenacity by Author

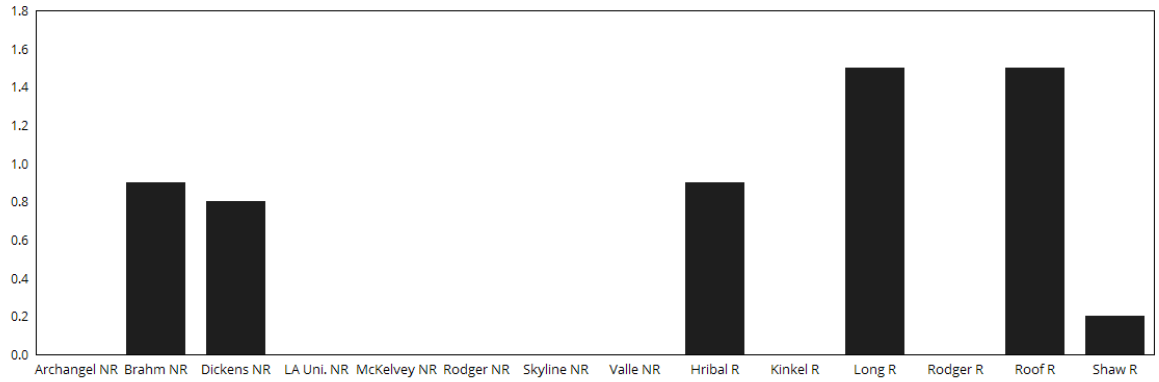

Frequency per 100 words

Other than the disparity in frequency—that a token of tenacity is nine times more likely to appear in a realized pledge—there is unfortunately too little data here to extrapolate more meaningful differences between the pledge types. Nevertheless, one functional pattern may be faintly apparent. In the three pledges where more than a single token appears, tenacity is used to negotiate the 'solidarity' and 'status' (e.g., Brown & Gilman, 1960) of the author in relation to groups the author either admires or

condemns. Hribal R, for example, praises Eric Harris and Dylan Klebold, the notorious Columbine High School shooters, saying they *worked hard to achieve freedom in heaven*, while simultaneously dismissing their victims as a *herd of sheep*. Harris and Klebold's will to action is discursively framed as a model for Hribal R's own.

A push and pull between action and inaction is also apparent in Roof R's pledge. He employs -tenacity to note that possible allies in his community are not *doing anything* except *talking on the internet* and contrasts this communal inaction with *skinheads* and the *KKK*—groups Roof R believes would take real-world action if they were present in his part of South Carolina. He concludes by shifting to +tenacity, positing that *someone*—meaning Roof R himself—must have the *bravery* to act in their absence. Similarly, Long R argues that the *current system* of law enforcement allows *bad cops to flourish, excell, & go unpunished*, and that until this changes *my people, & the people in general will continue to strike back*. A more or less vigorous condemnation of those who could act but have not is central to the use of tenacity in each of these three realized texts:

- *I am not like my family or friends at all. I don't belong with them. I belong with people that developed quickly, that changed the world when everyone else were just sheep in a herd.* (Hribal R)
- *If not, my people, & the people in general will continue to strike back against all cops until we see that bad cops are no longer protected & allowed to flourish.* (Long R)
- *We have no skinheads, no real KKK, no one doing anything but talking on the internet. Well someone has to have the bravery to take it to the real world, and I guess that has to be me.* (Roof R)

In the few pledges where more than one token of tenacity is available for analysis, solidarity with groups who share the author's underlying grievance but have yet

to rectify it is an expressly stated justification by authors who would proceed to violence. No such nuanced assessments of action versus inaction appear in the non-realized texts using this resource of appraisal.

5.2 CAPACITY

As conceived in Appraisal, capacity speaks to qualities like 'competency' and 'productivity' (Martin & White, 2005: 53). This variable also addresses judgements of a person's physical wellbeing ('unsound, sick, crippled') and mental health ('insane') (Martin & White, 2005: 53). However, the "range of meanings listed" as examples for a variable like capacity "is not exhaustive" (Martin & White, 2005: 52). And indeed, despite its more prosaic conception, capacity turns out to be an excellent tool for capturing the violence at the heart of a pledge to harm. This is because judgements of another person's physical injury and death represent evaluations of their physical capacity and its diminishment, in part or in whole. Similarly, assessing others as fearful or confused is, fundamentally, an evaluation of their *mental* capacity.

Examples of **physical** capacity from both corpora include:

- *i will shoot any black person i see at kean university* (McKelvey NR)
- *Your security will not be able to stop us.* (LA Unified NR)
- *I will attack the very girls who represent everything I hate in the female gender: The hottest sorority of UCSB.* (Rodger R)

Examples of **mental** capacity include:

- *Humans don't understand me.* (Shaw R)
- *The embarrassment would be too much for them.* (Kinkel R)

The dispersion of this resource across the dataset is shown below in Figures 5.5 and 5.6, first according to its positive polarity and then by its negative.

Figure 5.5: Positive Capacity Dispersion by Corpus

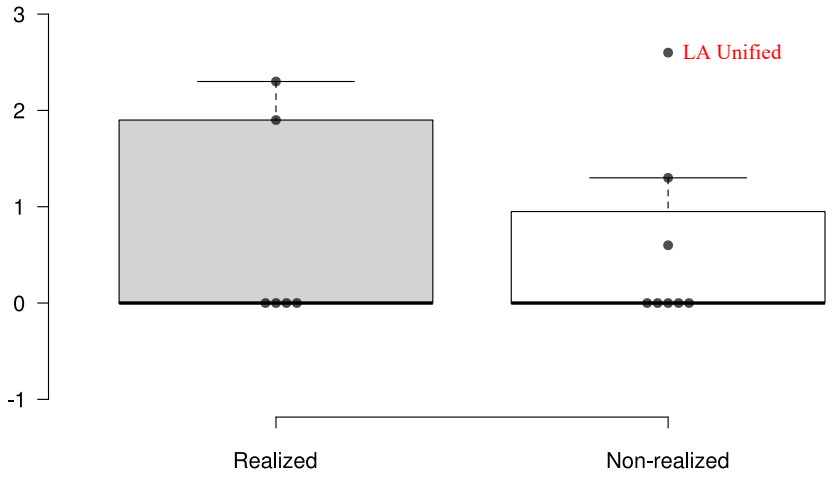

Frequency per 1000 words. One outlier (>1.5IQR) is found.

Figure 5.6: Negative Capacity Dispersion by Corpus

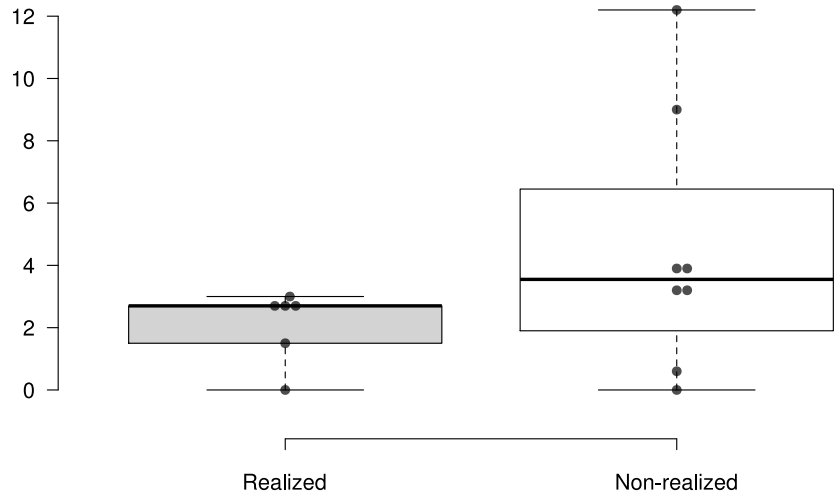

Frequency per 1000 words. No outliers (>1.5IQR) are found.

From a folk linguistic perspective, “a genre committed to violence and threatener control” (Gales, 2010: ii) might be expected to make liberal use of such a resource. And this indeed proves to be the case in this dataset. At a rate of 50 tokens per 1000 words, capacity is the most common attitudinal resource in these pledges (followed by the variable of propriety, at 45 tokens per 1000 words). However, as Figures 5.5 through 5.7 show, the two text types do not make use of this resource equally.

Figure 5.7: Positive and Negative Capacity by Realization Category

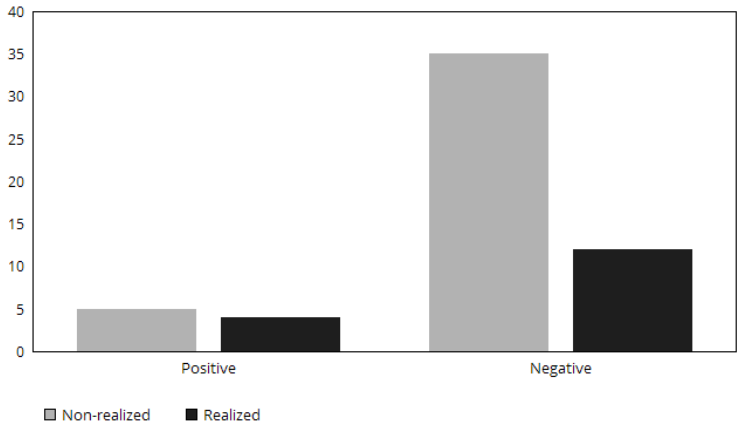

Frequency per 1000 words

While judgements of capacity are likely to be negative in both corpora, tokens of *in*capacity are far more likely to appear in a non-realized text, at a ratio of 3:1. Before the significance of this finding is explored further, however, two things are worthy of note. First, LA Unified qualifies as an outlier in its use of +capacity. This will be addressed below. Second, although Rodger NR's pledge is not technically an outlier, as Figure 5.6 shows, this text makes outsized use of -capacity (as does Valle NR's, though to a lesser degree), demonstrated in Figure 5.8. This text will therefore be treated somewhat carefully, and removed from the discussion except at key points when its re-inclusion will be made explicit.

Figure 5.8: In/capacity by Author

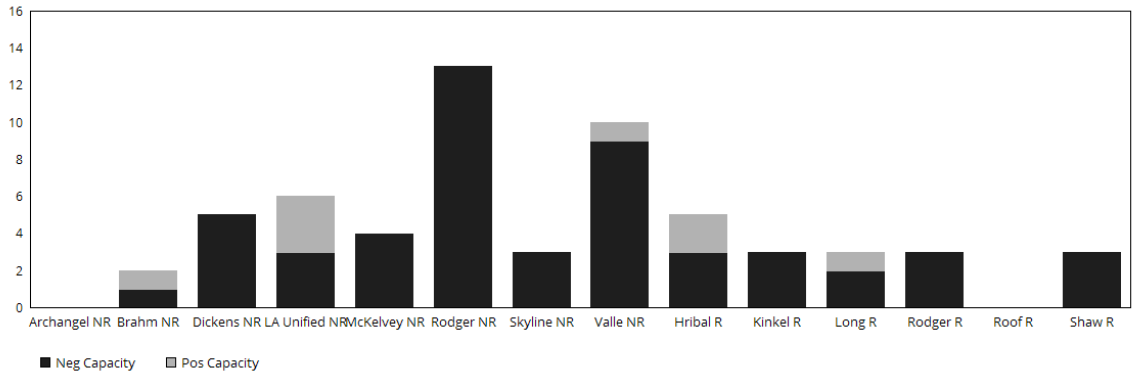

Frequency per 1000 words

Without Rodger NR's pledge, the 3:1 ratio tightens between the two realization categories. However, the overall pattern remains true: the adjusted NR corpus features the meanings of incapacity at over twice the rate compared to the realized pledges, as the Figure 5.9 demonstrates below.

Figure 5.9: Positive and Negative Capacity by Realization Category (Adjusted)

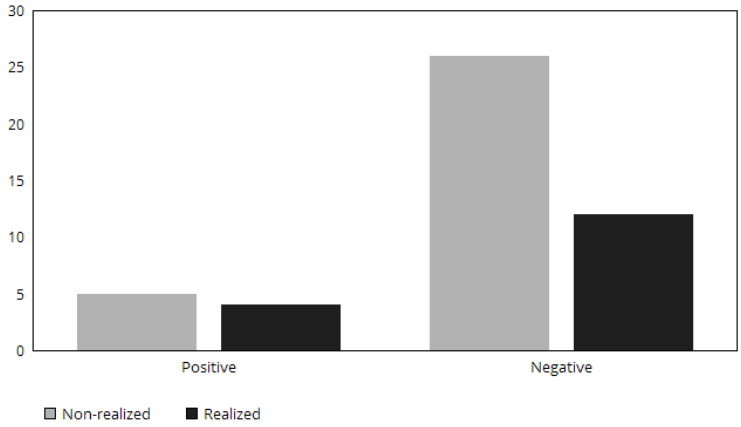

Frequency per 1000 words

This disparity is, of course, of a piece with the significant difference in attitudinal polarity, as discussed in section 4.1 above, of which +/-capacity is a part. It is interesting that even such a wide-angle view shows the non-realized texts—those featuring actions the authors would *not* later attempt—employ a semantic resource useful for expressing violent ideation at an appreciably greater rate. However, the full implications only become apparent when the monolith of 'negative capacity' is broken into its component pieces. First, though, the fewer tokens of +capacity are addressed.

5.2.1 POSITIVE CAPACITY

Positive judgements of capacity are much less frequent in the dataset, and appear at a near parity in the two pledge types (per Figure 5.5, above). Indeed, LA Unified NR qualifies as an outlier because this text features a mere four tokens of positive capacity. The analysis of this area is therefore brief. What is most notable about this small set of

attitudinal markers is that they are not employed simply as inversions of their negative analogues. No party in these pledges is deemed to be *capable* in a textual vacuum, say, simply as the opposite of being *incapable*. Instead, positive capacity is employed strategically by both author types, generally to reinforce more negative meanings in the token's immediate contextual vicinity.

For example, approving statements of a person's or organization's capabilities are used to underscore the futility of resisting the author's will, as when LA Unified NR says that *beefing up security…won't matter*. Elsewhere, positive capacity heightens the cruelty of an action designed to cause harm, like when Valle NR muses that he would like to *cook her over a low heat, keep her alive as long as possible*, or Hribal R's hope that his victims will *realize their previous lives are going to be taken by the only one among them that isn't a plebeian*. Even for the most civic-minded author in the dataset, Long R, concern for the *welfare* and *well-being* of the people in his community draws much of its emotional resonance from the *destruction* being wrought by *bad cops*. In these ways, positive capacity is far from being merely a nod to the estimable qualities of others. Instead, like the interplay of positive and negative tenacity discussed above, +capacity draws its full, discursive meaning from the more prominent themes of *in*capacity against which it is set. The positivity thus serves to undergird and reinforce—rather than lighten or counter—the overwhelming negativity of capacity meanings in the dataset, a rhetorical strategy that is employed in a similar manner by both author types.

5.2.2 NEGATIVE CAPACITY

The use of negative tokens is far richer in the dataset, covering a range of semantic and discursive meanings, e.g., to denote violent acts (as shown in the examples at the start of this chapter), and to derogate or infantilize classes of people (e.g., *faggots*, *cunts*, *girls* [for adult women], *kids* [for teenaged peers]). Incapacity also appears in more anodyne

constructions, such as *your security will not be able to stop us* (LA Unified NR). In this section and the ones which follow, the many and varied meanings in this area are analyzed along a host of semantic axes: whether the tokens assess physical or mental capacity; are violent or non-violent; are placed in the grammatical past or future; etc.

Beginning with physical versus mental capacity, a sizable portion of the R texts' non-violent negative judgements are concerned with the mental abilities of their stance objects, such as when Shaw R writes *Humans don't understand me* and Kinkel R laments that *They couldn't live with themselves!* This shows a certain comfort with, if not simply a "capacity to infer mental states" in other people (Malle & Nelson, 2003: 563). NR pledges do not share this interest in the minds of others. Despite their greater overall negativity, just a single negative token of mental capacity appears in the entire NR corpus, when Valle NR imagines his victim asleep and *not having the slightest clue of what we have planned*. The negative capacity used in the non-realized pledges is, instead, almost universally concerned with physical meanings.

In fact, this is the first instance of what will become a theme in these analyses, and so it is stated here explicitly in order to be taken up again at many points in the next few chapters: *the realized pledges often use stancetaking resources to contend with the minds and voices outside their textual borders; non-realized pledges focus instead on the physical world and how it may be changed or controlled.*

Beyond physical and mental, the most fundamental semantic question for threatening language lies, perhaps, in the division between violent and nonviolent meanings. Violent negative capacity is, of course, central to pledging as a genre, no matter the realization category, since judgements of future physical incapacity explicitly represent the 'harm' in a pledge to harm. Examples include statements like *I have to kill people* (Kinkel R) and *I can knock her out* (Valle NR). However, like negative capacity more generally, the relative strength of violent incapacity is distributed unevenly across

the two corpora, and in occasionally counterintuitive ways. This distribution is the focus of the next subsection.

5.2.3 VIOLENT VS NONVIOLENT NEGATIVE CAPACITY

From a folk linguistic perspective, authors who were emotionally invested enough in their ideations to proceed from language to action—i.e., the authors of the realized pledges—might be expected to focus more on the violence they would later attempt. As Figure 5.10 demonstrates, though, this is not the case.

Figure 5.10: Violent vs Nonviolent Lexemes of -Capacity by Realization Category

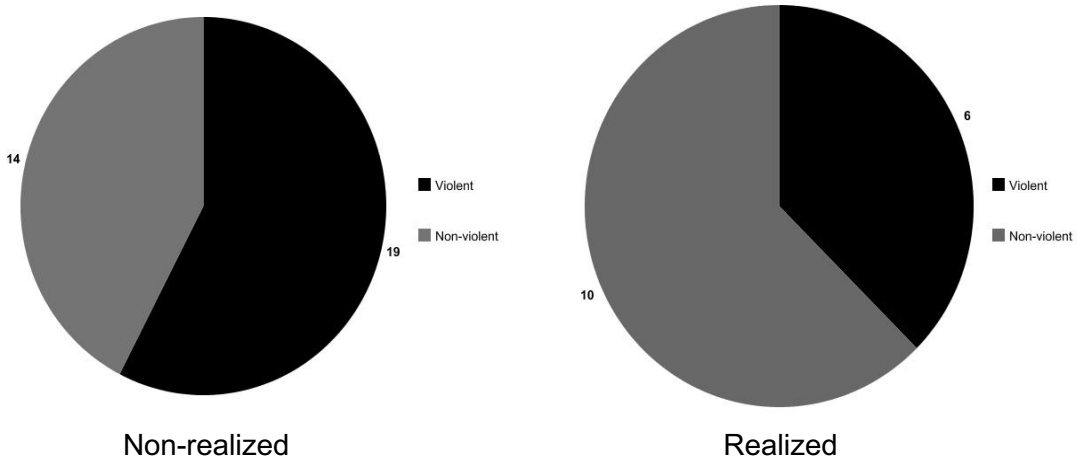

Instead, a greater amount of violent ideation is featured in the *non-realized* texts. This preoccupation with violence is found in two quantitative respects. First, the sheer number of tokens denoting violent actions (*kill*, *open fire*) or their results (*massacre*, *suffering*) is higher in the NR pledges, at a normed frequency of 19 tokens per 1000 words versus a frequency of just 6 per 1000 in the R corpus. This means that a token of violent incapacity is three times more likely to appear in a pledge written by someone who did not proceed to action. Second, the overall ratio of violent (*kill*) to nonviolent negative tokens (*faggots*, *asleep*) is higher in the NR corpus. Indeed, violent tokens constitute a healthy majority of the negative capacity present in the NR pledges, at a rate

of essentially 2:1. For the R pledges, by contrast, this ratio is inverted: nonviolent judgments of capacity appear at nearly double the rate of the violent tokens. Taken altogether, the NR pledges thus appear far more violent than their realized counterparts.

Qualitative differences are apparent between the two corpora as well. Table 5.2 below captures the lemmas, or lexical "base forms" (Biber, Conrad, & Reppen, 1998: 29) of violent content words which appear in each realization category. (For instance, *killing* and *killed* are counted as two instances of the base form KILL[17].) However, accurately schematizing sense relations between lexical items is a problematic task. One widely available tool is Princeton University's semantic dictionary, called WordNet. This lexical database groups content words into "sets of cognitive synonyms (synsets), each expressing a different concept" (Miller, 2013). Super-subordinate relations between synsets are then displayed. For the purposes of objectivity, WordNet is used here.

Despite all fourteen pledges being instances of the same illicit genre (Bojsen-Møller *et al.*, 2020), overlap between the two corpora occurs only in the first four lemmas listed: KILL, DIE, DESTROY, PAIN. According to WordNet, KILL and DIE are considered their own superordinate terms, while the noun form of PAIN is a hyponym of the very broad term 'symptom'. As superordinates, these three thus carry less marked semantic weight. Only DESTROY is identified by WordNet as a troponym—in the context of living beings, 'destroy' is subordinate to the verb 'kill'—and thus more marked. (To capture the full range of meanings, Rodger NR's text has been re-included. Lexemes unique to his pledge are marked in red.)

More semantic overlap between the pledges might have been expected. However, despite their generic similarities, this may reflect the varying aims and concerns of the individual authors. Interestingly, despite the fact that guns were used by

[17] In keeping with Biber *et al.*'s (1998) practice, lemmas will be distinguished from word forms by the use of small capital letters.

Table 5.2: Lemmas of Violent Ideation by Realization Category

Token of Violent Ideation	Frequency			
	Non-realized corpus		Realized corpus	
	Raw	Normed	Raw	Normed
KILL	9	6	2	1
DIE	5	3	1	<1
DESTROY	1	<1	3	1
PAIN	1	<1	1	<1
HIT	-	-	4	2
ATTACK	-	-	1	<1
SLAUGHTER	-	-	1	<1
PUNCH	-	-	1	<1
LOSS OF LIFE	-	-	1	<1
LIVES…TAKEN	-	-	1	<1
KNOCK…OUT	4	3	-	-
MASSACRE	2	1	-	-
SHOOT	2	1	-	-
TIE…UP	2	1	-	-
SUFFER	2	1	-	-
TORTURE	2	1	-	-
PUNISH	1	<1	-	-
CUT	1	<1	-	-
FLAY	1	<1	-	-
STRIP (FLESH)	1	<1	-	-
POUR BOILING WATER…	1	<1	-	-
BEHEAD	1	<1	-	-
SLIT (THROAT)	1	<1	-	-
OPEN FIRE	1	<1	-	-
COOK	1	<1	-	-
GAG	1	<1	-	-
STUFF INTO A BAG	1	<1	-	-

four of the six realized authors (Kinkel, Long, Rodger, and Shaw), gun-related terminology (*shoot*, *open fire*) only appears in NR pledges[18]. The violent ideation is thus more varied—and arguably more imaginative—in the writings of those authors who would not, in fact, attempt to realize the events and actions they were describing.

[18] This will be echoed in the analysis of the appreciation variable of composition (section 6.1.1) which notes that, for example, mentions of *bombs* only occur in NR pledges.

5.2.4 PARTS OF SPEECH

Recognizing that some difference exists in the semantics of violent lexemes in the two corpora raises an additional question about whether grammatical differences are also apparent. I.e., which linguistic resources are used to enact the stances identified in section 5.2.3 above? And is any pattern evident in how these resources are distributed across the pledge types? These questions require a momentary departure from the semantics of Appraisal into a more rudimentary consideration of lexical categories.

Interestingly, sorting the violent lexemes shown in Table 5.2 above according to their parts of speech reveals a supplemental, and somewhat counterintuitive, picture. With Rodger NR once again excluded, only two lexical categories are used to communicate violence in the dataset: nouns and verbs. (Rodger NR features the one violent adjective; violent adverbs do not appear). Of the two, verbs naturally spur the greatest analytical interest since they most directly address the question every assessor faces when analyzing a threatening communication: what does the author claim will *happen*?

Figure 5.11: Violent Lexemes of Capacity by Parts of Speech

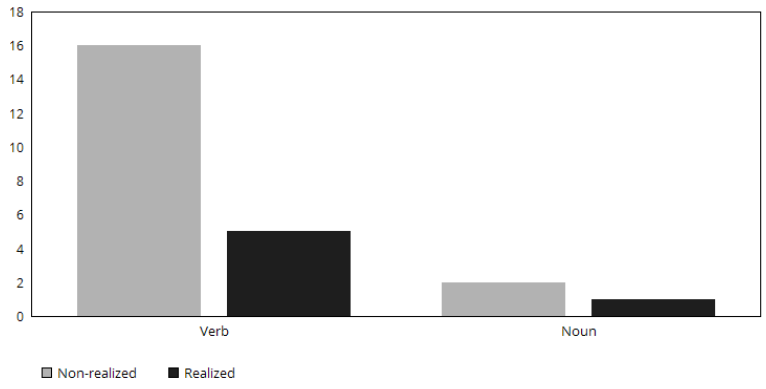

Frequency per 1000 words

Authors of the NR texts are over three times as likely than their R counterparts to include a verb of violence in their pledges, at a rate of 16 tokens per 1000 words in the NR corpus to just 5 in the R corpus. However, in Rodger NR's absence the semantic

diversity of the verbs does indeed suffer. Rodger NR is the source of more marked troponyms like *slit*, *flay*, and *behead* (Rodger NR). Without these lexemes, many of the differences between the two corpora disappear, leaving a few tokens of more marked verbs on either side of the realization divide (e.g., *open fire* in the NR corpus and *punch* in the R corpus). In general, only more superordinate lexemes like *kill*, *die*, and *destroy* remain—terms which are just as likely to appear in the realized texts as the non-realized. The heightened frequency of violent ideation appears to be a better indicator of a text's non-realized status than the markedness of the ideation. This too is a trend whose thread runs through other areas of this analysis.

5.2.5 TENSE

If violent verbs are responsible for much of the menace of a pledge to harm, then discovering what kinds of violence are placed in a grammatical future is one way of qualifying a threat. One question, then, is how many of the violent verbs in the dataset discuss actions in the future, and do the pledge types show differences in ideation in this regard? In English, this includes verbs projected forward in time via prediction modals (e.g., *will*, *shall*), verbs that appear with certain circumstantial adjuncts (e.g., *tomorrow*, *soon*), and verbs that appear with certain conjunctive adjuncts (e.g., *then*) (Lock, 1996).

As in the other areas of in/capacity discussed so far, the disparity between the two corpora in terms of tense is large, as Figure 5.12 demonstrates. And once again, these meanings are used predominantly by the NR pledges. A non-realized text is almost four times as likely to locate a violent event in a future time. For instance:

- *i <u>am going to</u> open fire* (Skyline NR)
- *i <u>will</u> kill all the blacks <u>tonight</u>* (McKelvey NR)
- *Might kill at least fifteen <u>tomorrow</u>* (Dickens NR)

Figure 5.12: Violent Verbs by Tense

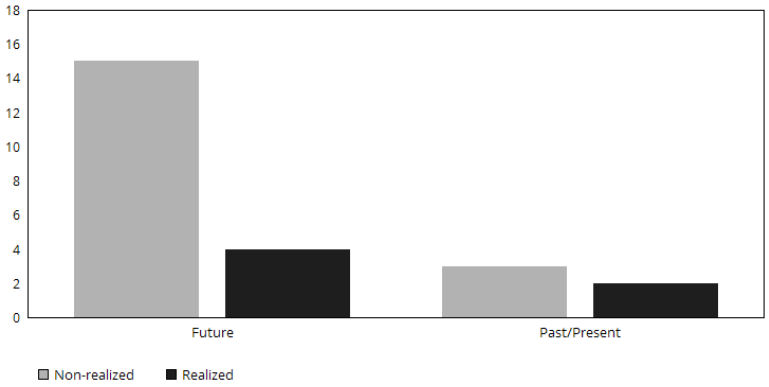

Frequency per 1000 words

By this metric, the non-realized pledges are far more threatening, despite being authored by people who would not proceed to real-world action. By themselves, then, violent verbs located in a future time would appear to be poor predictors of whether a pledge to harm will, in fact, be carried out.

5.2.6 TARGET SPECIFICITY

As section 2.3.1 notes, one way a pledge may satisfy its 'essential' condition, which requires the locution to be conventionally understood as pledging future harm (Harmon, 2008: 89), is by somehow "specific[ifying]…the victim's identity" (Harmon, 2008: 90). A concern for specificity can be found in the threat assessment literature as well, although there is disagreement about exactly which behaviors it may predict. Borum *et al.* (1999: 330), for example, call "evidence of ideas or plans to attack a specific target" an "attack-related behavior," i.e., a warning sign that a person's violent ideations should be taken seriously by the assessor. Conversely, Smith *et al.* (2014: 322) say that "if the victim's identity is named or implied, targeted violence or approach *is less likely* to occur." Thus, expressions which feature a higher degree of specificity about the target are generally considered predictive, although what they predict is disputed. Because incapacity covers the range of meanings associated with targeted physical violence, examining

who is being judged by a pledge author and whether or how these parties are specified is potentially useful.

In Table 5.3 below, the targets of violent incapacity appear in bold, while the phrasal or lexical tokens of incapacity themselves are underlined.

Table 5.3: Targets of Incapacity by Corpus

Corpus	Author	Instance of Incapacity
Non-realized	Brahm	*Due to the open air, the radiological fallout will destroy* ***those not killed*** *in the initial explosion.* (NFL fans)
	Dickens	*Death to* ***all white cops nationwide****.*
	LA Unified	*The* ***students at every school in the L.A. Unified district*** *will be massacred, mercilessly.*
	McKelvey	*i will shoot* ***any black person i see*** *at kean university*
	Rodger	*Silently killing as many* ***people*** *as I can around Isla Vista by luring them into my apartment through some form of trickery.*
	Skyline	*i am going to open fire on the* ***people in the commons*** *in the morning until i am either taken down by our schools police officer, or until i run out of mags.*
	Valle	*But I will really get off on knocking* ***her*** *out, tying up* ***her*** *hands and bare feet and gagging* ***her****.* (A woman personally known to Valle and referred to as [Victim-1] in the criminal complaint)
Realized	Hribal	*I can't wait to see the priceless and helpless looks on the faces of the* ***students of one of the "best schools in Pennsylvania"*** *realize their previous lives are going to be taken by the only one among them that isn't a plebeian.*
	Kinkel	*I have to kill* ***people***
	Long	*Therefore I must bring the same destruction that bad cops continue to inflict upon my people, upon* ***bad cops*** *as well as* ***good cops***
	Rodger	*I will attack the* ***very girls*** *who represent everything I hate in the female gender:* ***The hottest sorority of UCSB****.*
		I will sneak into their house at around 9:00 p.m. on the Day of Retribution, just before all of the partying starts, and slaughter ***every single one of them*** *with my guns and knives.*
	Shaw	***Every Asian Woman by herself*** *must be hit in the face.*

This list shows the imagined victim(s) in each text as they are somehow incapacitated through 1) a verb of physical violence in 2) a future time frame. Thus,

these utterances most explicitly satisfy the essential condition of each pledge in the dataset.

Despite the clear threat each of these tokens expresses, and an intuitive grasp that some of the targets mentioned seem more specific (*students of one of the "best schools in Pennsylvania"*) and less specific (*people*), quantifying specificity according to linguistic theory is actually a problematic exercise. This kind of referential question is not theorized within Appraisal because, in SFL, denotation is encoded at the ideational level (see, e.g., Halliday & Matthiessen, 2004: 63) rather than the interpersonal which Appraisal is built to measure. Thus, to address Harmon's (2008) observation and the concern for specificity found in the threat assessment literature, it is necessary to briefly step outside Appraisal. Unfortunately, no clear consensus exists among linguists or threat assessors about how to approach the question. If, for example, the "specificity as to the victim's identity" (Harmon, 2008: 90) is bound up with the noun phrases (NPs) themselves, then referring expressions which narrow their set of possible referents should equal a more identifiable target. One test for this is morphosyntactic in nature. Caraballo and Charniak (1999) argue that specificity may be detected through modification: more specific nouns will accept fewer modifiers. "It seems reasonable to suppose that very specific nouns are rarely modified, while very general nouns would usually be modified" (Caraballo & Charniak, 1999: 64). The authors use this metric to identify the less-to-more specific relationship between, e.g., *food → meat → ham*.

If, however, a NP's specificity is tied more to a "property of independent existence" (Bornkessel-Schlesewsky & Schlesewsky, 2009: 42), then this quality is best addressed through an analysis of theta roles. In threatening language, the imagined victims would be best categorized as Undergoers, because "they are the target of sentience, causally affected by an event and controlled" (Bornkessel-Schlesewsky & Schlesewsky, 2009: 41). But, as Bornkessel-Schlesewsky and Schlesewsky (2009)

point out, Undergoers “have no defining prototypical features of their own” (Bornkessel-Schlesewsky & Schlesewsky, 2009: 41) and are therefore not necessarily marked for features like +definite/specific in the same way as an Actor.

Elsewhere, a potentially useful scalar measure is the nominal hierarchy presented by Silverstein (1976) and employed by Bornkessel-Schlesewsky and Schlesewsky (2009). This is shown as Figure 5.13.

Figure 5.13: The Nominal Hierarchy (Bornkessel-Schlesewsky & Schlesewsky, 2009: 26)

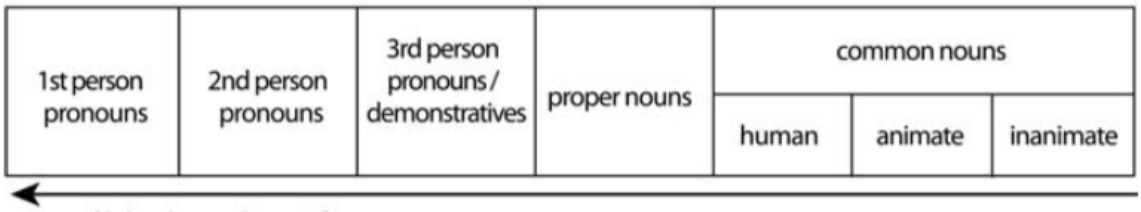

As glossed by Dixon (1994), the usefulness of the hierarchy is in how it arranges NPs:

> The further to the left an NP is, the more likely it is to be definite (at the extreme left it is always definite) and the further to the right it is the more likely it is to be indefinite. (Dixon, 1994: 91)

While definiteness and specificity are not interchangeable, the former may yet offer a quantifiable, grammatical approach to the latter. A customized version of this hierarchy is presented in Table 5.4 on the page below. In it, the NPs identified in Table 5.3 are categorized according to the hierarchy. Allowing for contextual anaphora, Rodger R’s use of *them* in the realized text as a pronoun for *Alpha Phi Sorority*, and Valle NR’s use of *her* to refer to a particular woman known as Victim-1 in the criminal complaint, might be better classified as proper nouns. If so, each would shift rightward, becoming (somewhat counterintuitively) less specific according to this system of classification. Similarly, if Brahm NR’s phrase *those not killed* operates as a stand-in for NFL attendees, then this might be better logged under ‘common noun’ as a ‘type of

human'—also losing the specificity afforded it by its demonstrative formulation. According to this test, then, there is actually very little specificity present in the noun phrases denoting the threatened parties in either corpora.

Table 5.4: Nominal Hierarchy of Stated Targets

Corpus	Author	← More specific				Less specific →
		1st Person Pronouns	2nd Person Pronouns	3rd Person Pronouns / Demonstratives	Proper Nouns	Common nouns[19]
						Human
Non-realized	Brahm			those not killed		
	Dickens					all white cops nationwide
	LA Unified					students at every school in the L.A. Unified district
	McKelvey					every black person i see
	Rodger					people
	Skyline					people in the commons
	Valle			her		
Realized	Hribal					students of one of the "best schools in Pennsylvania"
	Kinkel					people
	Long					bad/good cops
	Rodger					the very girls
				them		
	Shaw				every Asian Woman by herself	

[19] As shown in Figure 5.13, Silverstein's (1976) hierarchy includes two additional categories under the header of Common Nouns: Animate and Inanimate. Because the subsystem of judgement applies strictly to a speaker's views of other human beings (where the value of animals and objects is subsumed by the subsystem of appreciation) these two categories have been excluded.

The method proposed by Caraballo and Charniak (1999), whereby these nouns are classified according to whether and how the NPs are modified, yields an equally indistinct picture. The following Table 5.5 sorts the targets by their modifiers: either the scope of the NP is limited, thereby making it more specific, or the scope is expanded, making the NP more general. The Proper Noun category has been shifted to the left

Table 5.5: Scope of Proper and Common Noun Phrases of Stated Targets

Corpus	Author	← More specific			Less specific →
		Proper Noun	Limiting	No Modifier	Expanding
Non-realized	Brahm		those not killed (NFL fans)		
	Dickens				all white cops nationwide
	LA Unified				students at every school in the L.A. Unified district
	McKelvey				every black person i see
			black people at kean university		
	Rodger			people	
			some of the good looking people		
	Skyline		people in the commons		
	Valle			her (Victim-1)	
Realized	Hribal		students of one of the “best schools in Pennsylvania”		
	Kinkel			people	
	Long			bad cops/good cops	
	Rodger		the very girls		
					every single one of them (Alpha Phi Sorority)
	Shaw				every Asian Woman by herself

here to represent greater specificity, in accordance with the general consensus of the threat assessment literature, which argues that a threat is more urgent when a particular target has been specified by the threatener (Meloy, 2011).

Gray areas exist here as well. Shaw R's *every Asian Woman by herself* could be reclassified as a proper noun according to its right-most, more limiting prenominal descriptor (*Asian*) rather than by its left-most 'allness' term *every*. However, doing so would require McKelvey NR's *every black person i see* to be reclassified as well, erasing a potential cross-corpora difference. Similarly, reclassifying the specificity of Rodger R's anaphoric reference to *Alpha Phi Sorority* as a proper noun would be nullified by Valle NR's reference to Victim-1. And Hribal R's identification of the *students of one of the "best schools in Pennsylvania"* is roughly equivalent both to Skyline NR's focus on *people in the commons* and McKelvey NR's *black people at kean university*. All in all, the NR corpus seats four NPs in the left half of the table (as a noun with a limited scope) while the R corpus seats just two.

Viewed this way, the NPs in the NR corpus tend to be more specific, but only barely. This is essentially a null result, but one that sits at an interesting point in relation to other literature on the subject. For example, in their own examination of direct threats, Smith, Woyach, and O'Toole (2014: 325) found that "targeted violence or approach is less likely to occur…if the victim's identity is named or implied." However, as noted previously, the threat assessment literature views higher target specificity as indicating a greater risk of approach (e.g., Turner & Gelles, 2003). The fact that NP specificity is not a point of distinction in this dataset—neither risk-enhancing nor risk-reducing—is therefore noteworthy.

5.2.7 ACTION SPECIFICITY

For Mohandie (2014: 136) “evidence of intent” includes the “specificity of [the] plan” which the threatener hopes to enact. Complementary to the question of target specificity, then, is the level of detail of the imagined violence. Beginning from a grammatical rather than a pragmatic standpoint, violent physical actions are a type of material process in English, a *changing* or *doing to* of an element of the physical world (Halliday & Matthiessen, 2004: 172). In the case of pledges, this process changes the bodies of the imagined victims. How specific this process is potentially discoverable by examining the markedness of verbs denoting violence.

Helpfully, the super-/subordinate sense relationship of troponymy is one basis of Appraisal’s conception of ‘infusion’ in the system of graduation. Infusion occurs when “there is no separate lexical form conveying the sense of up-scaling or down-scaling” (Martin & White, 2005: 143). For example, the lexeme *frighten* is down-scaled compared with the term *terrify* (Martin & White, 2005: 144), although each are troponyms of the verb *scare*. *Terrify* is thus more semantically infused than *frighten*.

The violent verbs placed in a grammatical future time, captured quantitatively in Figure 5.12 above, are sorted according to the synset classifications found in WordNet (Miller, 2013). These are shown in Table 5.6 below. Unfortunately, WordNet does not supply a scoring or ranking system to determine how many levels of remove separate a subordinate term from its most unmarked, superordinate category header. This inhibits a cross-corpora comparison of the relative ‘depths’ of the violent verbs in each pledge type. However, the synsets are presented as a chain of sense meanings, running from the search term up through the broader and broader sense relations to the lexeme WordNet considers the most unmarked term of the category. In an *ad hoc* manner, then, a comparison is possible by assigning a score of (1) to each semantic echelon. An ultimate, superordinate term—the general node from which other more specific

meanings branch—would receive a score of 1, while its hyponyms would be scored as 2, 3, 4, etc., depending on how many levels separate them from the superordinate. The term *kill*, for example, receives a score of 1 because WordNet lists no direct hypernyms.

Table 5.6: Troponymy of Future Tense Violent Verbs

Corpus	Author	Violent Verb	Troponymic Level
Non-realized	Brahm	*destroy*	2
	Dickens	*Death to*[20]	2
		kill	1
	LA Unified	*massacred*	2
	McKelvey	*die*	3
		kill	1
		shoot	5
	Rodger	*killing*	1
		knock…out	2
		slit (throats)	5
		torture	5
		kill	1
		cut	3
		flay	4
		strip	2
		pour (boiling water)	2
		torture	5
		behead	2
		keep…heads	2
		suffer	2
	Skyline	*open fire*	5
	Valle	*knocking…out*	2
		tying	4
		tying up	4
		gagging	5
		stuffed	4
Realized	Hribal	*(lives are) going to be taken*	1
	Kinkel	*kill*	1
	Long	*bring…destruction*[21]	2
	Rodger	*attack*	1
		slaughter	2
	Shaw	*hit*	2

[20] Dickens NR employs no verb in her construction *Death to all white cops nationwide*. Contextually, however, *Death* is clearly presented as a process which *all white cops* will undergo.

[21] Similarly, d*estruction* is a noun, but the intent of Long R's verb phrase is only clear when considered in its entirety, i.e., the sense of the phrase depends on including the object *destruction* as the complement selected by transitive *bring*.

But *torture* receives a score of 5 since it is the fifth node of the following synset tree: *change – indispose – hurt – injure – torture*. By contrast, *destroy* is scored as 2 because it is second in the following tree: *kill – destroy.* These gradations are shown in Table 5.6 above. Rodger NR is briefly reintegrated for the purpose of completeness.

Totaling these counts, the mean specificity for verbs in the NR texts is higher relative to the R texts, with an average of 2.92 for NR and 1.5 for realized. Interestingly, when Rodger NR's heightened use of violent future tense verbs is accounted for, the mean score of the remaining six NR pledges actually increases to 3.0. Shifting from the mean to the median to better account for the occasional but unusual 4th and 5th level troponyms reveals a similar, though less stark contrast. The median of the R texts still calculates to 1.5—barely more marked than a source node lexeme. The NR texts are still more specific, now at 2.0. Notably, the median score for the NR corpus does not change with the removal or inclusion of Rodger R. Thus, outlier or no, the authors of the NR texts use slightly more precise descriptions of future actions than their realized counterparts.

A separate test of specificity is Martin and White's (2005) concept of isolated lexicalization in the system of graduation, wherein "the up-scaling/down-scaling is realised by an isolated, individual item which solely, or at least primarily, performs the function of setting the level of intensity" (p. 141). This is an essentially synonymous mechanism to the examination of noun phrase modifiers in section 5.2.6 above. An example of isolation in verbal processes given is the contrast between the phrases *this upset me <u>slightly</u>* and *this upset me <u>greatly</u>* (Martin & White, 2005: 142).

While Martin and White (2005) naturally concentrate on adverbials and adjectives in their framework—two content word classes whose function is to modify—the tokens which appear in this dataset are not limited to these categories only. As Table 5.7 on the next page shows, along with adverbs and adjectives, verbs are also made more

Table 5.7: Violent Verbs and Modifier Types

Corpus	Author	Violent Verb	Pre-/Postverbal Modifier	Word Class or Phrase Type
Non-realized	Brahm	*destroy*	-	-
	Dickens	*Death*	-	-
	LA Unified	*massacred*	*mercilessly*	Adv
	McKelvey	*die*	-	-
		kill	-	-
		shoot	-	-
	Rodger	*killing*	*Silently*	Adv
		knock them out	*with a hammer*	PP
		slit	*their throats*	NP
		torture	-	-
		kill	-	-
		pain	-	-
		suffering	-	-
		cut	-	-
		flay	-	-
		strip	*all the skin*	NP
			off their flesh	PP
		pour boiling water	*all over them*	PP
			while they are still alive	PP
		torture	*any other form of*	NP
		behead	-	-
		keep their heads	*in a bag*	PP
		suffer	-	-
	Skyline	*open fire*	-	-
	Valle	*knocking…out*	-	-
		tying	*her body*	NP
			onto some kind of apparatus	PP
		tying up	*her hands and bare feet*	NP
		gagging	-	-
		stuffed	*into a large piece of luggage*	PP
Realized	Hribal	*(lives are) going to be taken*	*previous (lives…)*	Adj
	Kinkel	*kill*	-	-
	Long	*bring the (same) destruction*	*same*	Adj
	Rodger	*attack*	-	-
		slaughter	-	-
	Shaw	*hit*	*in the face*	PP

specific through the use of prepositional phrases[22] (PP) and noun phrases (NPs), either separately (e.g., *their throats*) or in concert (e.g., *all the skin off their flesh*). (Rodger NR is once again reincluded for the purpose of completeness.)

Clearly, each word class and phrase type performs a different function in limiting the range of the described action. For instance, the two adverbs here—*mercilessly* (LA Unified NR) and *Silently* (Rodger NR)—both up-scale the quality of the violence by making the actions more menacing. PPs and NPs, on the other hand, provide more spatial specificity by isolating where the action will occur. In the case of PPs, this is achieved via locative (*in the face*) and instrumental (*with a hammer*) information. With NPs, the focus is on the body parts affected by the action (*throats*, *hands and bare feet*). Interestingly, the two adjectives (*previous lives*, *same destruction*) are both arguably neutral, serving neither to up- nor down-scale the stance of the nouns they modify. Meanwhile, the realized corpus has one non-specific temporal adjective (*previous*), one non-specific comparative adjective (*same*), and one specific locative PP (*in the face*).

However, when Rodger NR's pledge is removed, much of the distinction between the two threat categories disappears. In the non-realized corpus, what remains is an up-scaled adverb (*mercilessly*), a single more specific NP (*hands and bare feet*), one less specific NP (*her body*), one more specific locative PP (*into a large piece of luggage*), and one less specific locative PP (*onto some kind of cooking apparatus*). All in all, there appears to be no particular pattern of type distribution (lexical or phrasal) or of up-scaling regarding verbal modifiers across the two corpora.

Of course, such a narrow grammatical approach risks missing the forest for the trees. Evidence of research and planning can include verbal confirmation that an author has engaged in behaviors like "surveillance of the target, Internet searches, testing

[22] For current purposes, the NP internal to the PP constituent will not be counted separately since the PP would be rendered ungrammatical without it.

security around the target, and researching methods of attack" (Bulling & Scalora, 2013: 15). Such concerns are broader in scope that the specificity of the violence itself, and are therefore outside the discussion of capacity particularly. Planning is revisited from a discursive perspective at the close of Chapter 7, following the analysis of engagement.

5.2.8 LINKING IN THE RODGER NR AND VALLE NR TEXTS

Finally, a word is warranted on a stylistic choice present in Rodger NR's and Valle NR's pledges. As Figure 5.8 above shows, Valle NR's text is very concerned with meanings of incapacity, second only to Rodger NR. Rodger NR's pledge in particular is responsible not just for the heightened token counts which have been culled away as outliers at various points in this analysis, but also for the bulk of the semantic diversity and the specificity apparent in the imagined violent actions. Stylistically, all three of these divergences—from token counts, to semantic diversity, to specificity—are primarily the result of the same writerly choice on the part of these two authors. Both employ lists in their pledges. Rodger NR, for instance, writes (violent incapacity is underlined):

- *After that, I will start luring people into my apartment, knock them out with a hammer, and slit their throats.*
- *I will cut them, flay them, strip all the skin off their flesh, and pour boiling water all over them while they are still alive, as well as any other form of torture I could possibly think of.*
- *When they are dead, I will behead them and keep their heads in a bag, for their heads will play a major role in the final phase.*

While Valle NR states:

- *I was thinking of tying her body onto some kind of apparatus . . . cook her over a low heat, keep her alive as long as possible.*

- *But I will really get off on knocking her out, tying up her hands and bare feet and gagging her.*

Lock (1996: 249) refers to this kind of clausal relationship as 'linking,' and notes that often "the sequence of clauses represents the chronological sequence of the…actions." The fact that Valle NR would not be able to *cook her* before he had *tied her body* onto a cooking *apparatus*, and that Rodger NR cannot *keep their heads in a bag* until the men have been *beheaded* is evidence that this kind of grammatical relationship is indeed at work in the examples above.

From a folk linguistic perspective, lists like these can give the impression of a "plan that demonstrates considerable detail" (Turner & Gelles, 2003: 97). The presence of this stylistic choice would, thus, potentially increase the seriousness with which a text was viewed by a threat assessor. However, Gales (2010: 71) warns that folk linguistic impressions "oftentimes conflict with authentic language practices." This would seem to be the case here. Although linking cannot be generalized as a reliable difference between the two corpora, it is nevertheless worth commenting that structures like this, which are intuitively associated with pre-planning, only appear in two NR pledges.

5.2.9 SUMMARY OF CAPACITY

In summary, authors in both the realized and non-realized corpora exhibit similarities in the resources they employ to instantiate their judgements of others' in/capacity. Similarities include:

- nearly uniform negativity, with positive instances of others' capacity used as frames or underscores for the larger negative themes, rather than as antitheticals;
- relative equity in the semantic diversity and markedness of the base lexical forms denoting violence;

- relative equity in the specificity of who is targeted for violence.

However, the non-realized authors utilize these resources in ways that are measurably distinct from their realized counterparts. Adjusted for Rodger NR, a NR pledge is likely to:

- contain more than twice the lexical and phrasal tokens of negative capacity;
- contain three times as many violent as non-violent negative tokens of capacity;
- focus on physical incapacity to the exclusion of mental incapacity;
- contain three times the frequency of violent verbs;
- contain violent verbs in the future tense which are, on median, slightly more specific in the actions they describe, and;
- locate violent events in a future time frame at nearly four times the rate.

In short, the non-realized pledges forecast more future physical violence than the realized pledges, despite the fact that the realized texts are theoretically more predictive of the authors' future physical violence.

5.3 PROPRIETY

Where Martin and White (2005: 52) discuss capacity and tenacity in terms of 'social esteem,' an area that is "critical to the formation of social networks," they consider propriety a type of 'social sanction,' an area which encompasses ideas of "civic duty and religious observances." Moving from esteem to sanction, the authors analogize, is comparable to shifting from "venial to mortal sins" (Martin & White, 2005: 53). Transgressing against social expectations at the level of sanction, then, is thought to have more dire repercussions, i.e., "too much negative esteem, and we may need to visit a therapist; too much negative sanction, and a lawyer may need to be called in" (Martin & White, 2005: 53). Propriety thus serves as a gauge for how well or poorly an author

believes a third party has upheld a given ethical value, covering a range of meanings from ‘good,’ ‘kind,’ and ‘charitable,’ to ‘bad,’ ‘cruel,’ and ‘selfish’ (Martin & White, 2005: 53).

Fundamentally, ethical meanings are realized through a modulation of obligation (Martin & White, 2005: 54). This often occurs at the lexical level, through adjectives or adverbs like the examples just cited, or through such devices as the modal auxiliary *should*, verbal *have to*, etc. Propriety may also be realized through a meeting of grammatical choices and pragmatics, as in Halliday and Matthiessen’s (2004: 148) example, “Mary will help.” Despite the fact that modal *will* technically communicates likelihood rather than obligation, “if Mary is listening, she can hardly refuse” (Halliday & Matthiessen, 2004: 148).

Examples of positive and negative propriety from the two corpora include:

- *we are all selfish.* (Hribal R)
- *They were wonderful people.* (Kinkel R)
- *They are all spoiled, heartless, wicked bitches.* (Rodger R)
- *Hell they condone crimes against us.* (Dickens NR)
- *The bullying, the loneliness, the rejection… it is never-ending.* (LA Unified NR)

before

The dispersion of this resource across the dataset is shown below in Figures 5.14 and 5.15, first according to its positive polarity and then by its negative.

Before proceeding further, however, Martin and White (2005) share a word of warning: an analyst is not meant to approach ethical meanings through the lens of the *text’s* value system (whatever it may be). Rather, the question is how are “the feelings popularly construed by the culture” (Martin & White, 2005: 46)? With propriety, then, “[w]e are not concerned here with the value that a particular uncommon sense

Figure 5.14: Positive Propriety Dispersion by Corpus

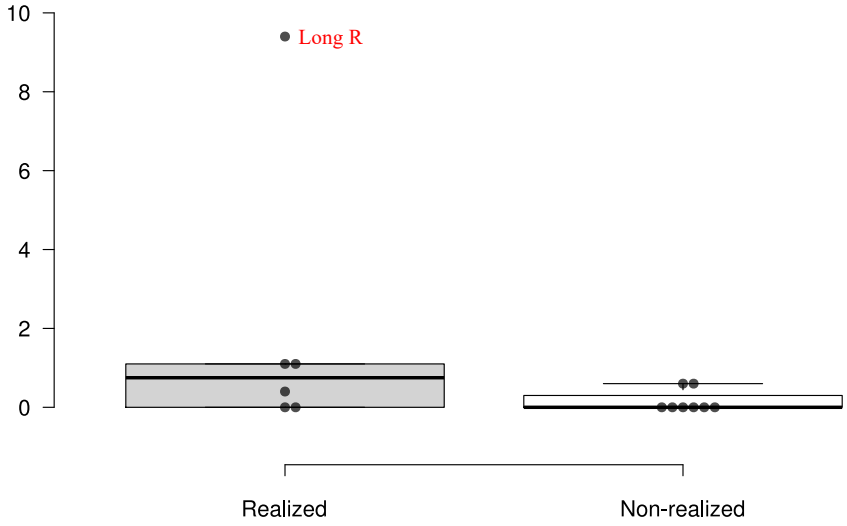

Frequency per 1000 words. One outlier (>1.5IQR) is found.

Figure 5.15: Negative Propriety Dispersion by Corpus

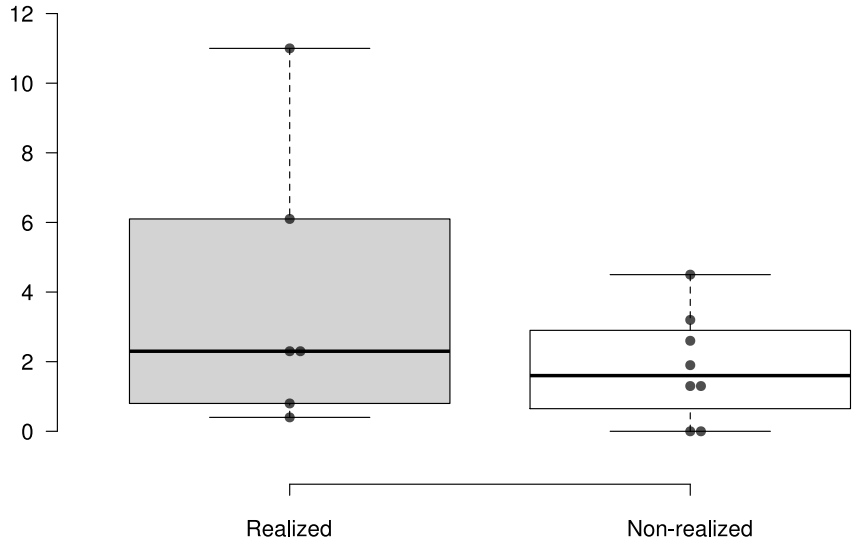

Frequency per 1000 words. No outliers (>1.5IQR) are found.

psychological framework might place on one or another emotion (cf. 'It's probably productive that you're feeling sad because it's a sign that ...')" (Martin & White, 2005: 46). In practice, this means that Roof R's obvious admiration for *skinheads* and the *KKK* is nevertheless coded as instantiating negative propriety, based on the mainstream American understanding that these two factions are, in fact, hate groups (Southern Poverty Law Center, 2015). Similarly, the *jihadist cell* of which LA Unified NR claims to

be a member is considered a token of impropriety in that both *jihadist* and *cell* ('a small group acting as a unit within a larger organization') culturally connote with 'terrorist'—a strongly negative lexeme in English—no matter that the author plainly presents the group's aims as synonymous with his or her own[23]. (In fact, in texts like LA Unified NR the fearsome connotation is arguably the point.) Propriety in the dataset has therefore been coded, first, from a sense perspective—asking if the polarity of a given lexicogrammatical item is positive or negative based on its place in the larger vocabulary system (Saeed, 2009: 12), e.g., *good* versus *bad*. Where the issue is instead one of connotation or invocation, tokens have been coded according to the cultural values predominant in the U.S. as of this writing. For example, Long R does not supply additional evaluative content indicating the positive or negative polarity of *holding them accountable*, but arguably this is because none is needed—holding authority figures to account for their failures is a positive value in modern U.S. society, of which Long R was a part.

Turning to the analysis, the dataset is suffused with ethical meanings. Of course, this is less than surprising in writings which contemplate such morally fraught topics as assault and murder. Nor is it surprising to see that a majority of the tokens present across the data, or roughly 74% of both corpora, are negative, as Figure 5.16 shows below. Of the 14 authors in the dataset, all but one employ judgements of propriety, and negative propriety in particular. However, even the one author, Archangel Michael NR, arguably engages with ethics as well, just through the matrix of social valuation within

[23] Additional precedent, beyond Martin and White's (2005) express advice, is found in Gales (2010) and her Appraisal analysis of the Eric Rudolph bomb threat. She observes that "through a repetition of lexical tokens infused with negativity, [Rudolph] portrays [the Army of God's] behaviors as improper"—i.e., as instantiating negative propriety—despite the fact that Rudolph personally approved of the group's violence as "biblically supported" (Gales, 2010: 247-8). There is thus no proscription against judgements of negative propriety being leveled at parties with whom an author claims solidarity.

the larger semantic region of appreciation—an Appraisal area with direct correlations to propriety[24] (O'Donnell, 2013). Thus, some level of preoccupation with morality is, for all intents and purposes, universal across the dataset.

Figure 5.16: Tokens of Propriety by Realization Category

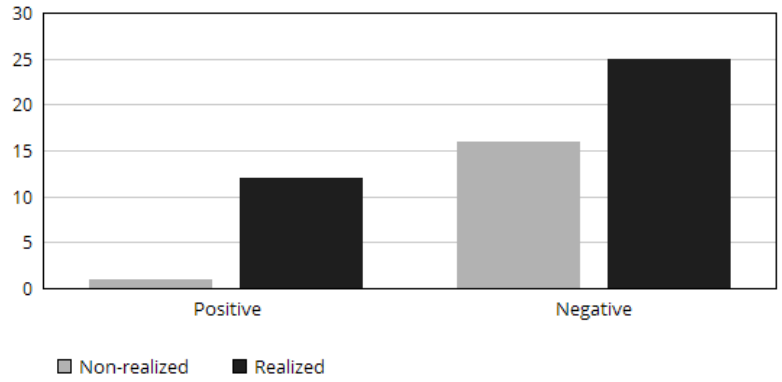

Frequency per 1000 words

However, the pledge of Long R is an outlier, constituting 58% of the propriety tokens in the realized corpus (55 of the 94 present). Unlike the outliers flagged in capacity and tenacity, though, Long R's text presents a pattern which may be systematically accounted for. His main topic is the political establishment's treatment of *good cops* and *bad cops* as equals. *Good* and *bad* are, obviously, ethical assessments

Figure 5.17: Tokens of Propriety by Realization Category (Adjusted)

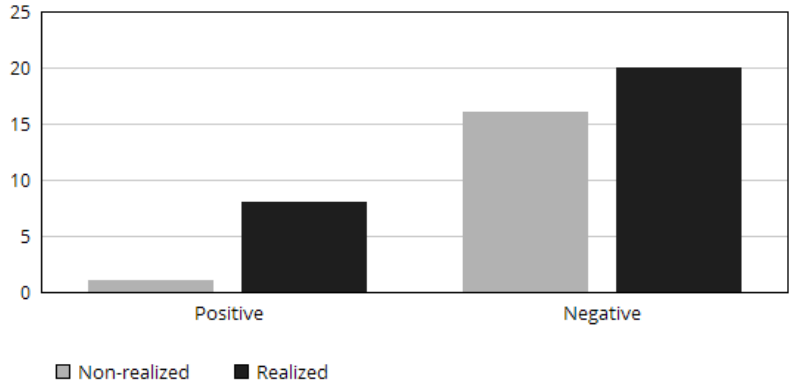

Frequency per 1000 words

[24] The ethical dimensions of valuation are analyzed within appreciation in Chapter 6.

which fall under the larger heading of propriety. Yet, here *good cop* and *bad cop* are used simply as category headers—functionally no different than Shaw R's repeated use of *Asian Women* or McKelvey NR's use of *black person*. Given this, a reasonable way to control for the artificial inflation is simply to remove these 25 adjectives (10 *goods* and 15 *bads*), so as not to lose outright the richness of the rest of Long R's writing.

Following this, the negativity in the R corpus drops from a high of 25 per 1000 words, as shown in Figure 5.16 above, to 20 per 1000 words, as the adjusted Figure 5.17 shows. However, the overall relationship between the four measures remains stable. The proportion of negative to positive propriety in the realized texts also remains roughly the same as before these select adjectives were removed, now at a slightly higher ratio of 5:2. Generally, adjusting for Long R's topic alters very little of the larger analytical picture, either internally or cross-corporally. These 25 adjectives will therefore be omitted from the analysis going forward.

Figure 5.18: Proportion of Propriety by Realization Category

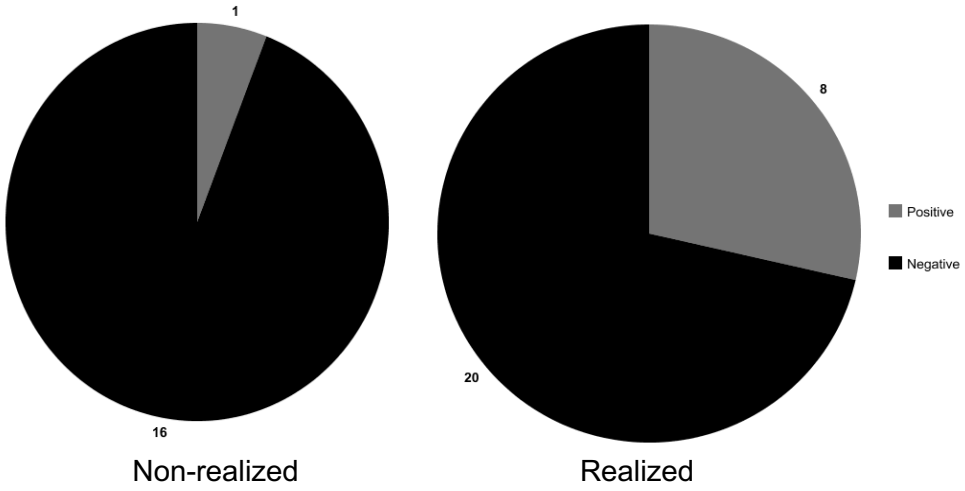

Frequency per 1000 words

However, though the two corpora share a preponderance of negativity, the proportional distribution of negative to positive tokens is not realized uniformly, as Figure 5.18 above shows. The NR authors judge outside parties as being somehow unethical

in 94% of the propriety tokens which appear in the corpus. For the R writers, a person's perceived moral failings are the content of this class of judgements at the lower rate of 71%. In other words, although the realized corpus is *quantitatively* more negative, the non-realized writers are 32% more likely to accuse people of failing their more solemn social obligations. The realized corpus may thus make greater use of propriety—with 28 tokens per 1000 compared to the NR corpus's 17—but these authors are also more evenhanded in their judgements.

On one hand, this result is expected. Returning to Martin and White's (2005) analogy, if negative sanction is comparable to mortal sin—or some class of felony in secular terms—then ethical lapses are a potential rationale for the harshest punitive actions a society may deliver, including a death sentence. Meanings within propriety are therefore the strongest resource a threatener could call on to justify violent ideation, something like *X is immoral, therefore X deserves harm, therefore I will harm X*. And so, it makes a certain amount of sense that those authors who proceeded to action would include more of this kind of meaning than their non-realized counterparts.

That said, it is surprising to find a higher proportion of impropriety in the *non-realized* corpus. The NR pledges overwhelmingly devote their textual space to damning others, at a rate of 16:1, whereas the R writers indulge in blame rather than praise at a rate of just 2.5:1. The prosodic effect of this disparity is that the non-realized corpus appears far more incriminatory, or, to again indulge in the metaphor of mortal sin, more 'fire and brimstone.' Conversely, positive propriety is eight times more likely to appear in the realized texts. Somewhat counterintuitively then, positive ethical meanings are a potential indicator of violent intent.

5.3.1 NEGATIVE VS POSITIVE PROPRIETY

The substance of a great deal of the negative tokens present in the two corpora are predictable from the generic character of these pledges. For example, third parties are said to *exploit*, *hate*, and *steal* from others. They are portrayed as *evil*, *disgusting*, *heartless*, and so on. Virtually none of these pledges—either inter- or intra-corporally—share a concern with a single, homogeneous type of unethical person or behavior. Instead, as this small sampling shows, the ethical stances enacted by the authors in the dataset are diverse, drawing from social issues, political grievance, criminal activity, etc. Even so, certain patterns within this widespread negativity are indeed apparent, and may be understood by way of the following questions: Who is judged as unethical? What parts of speech are used to realize these judgements? And just how strong are these judgements, i.e., how many tokens are somehow graduated?

Before addressing each of these in turn, though, an additional comment is warranted on the subset of +propriety from a *qualitative* standpoint. This resource is used discursively in much the same way as +tenacity (section 5.1) and +capacity (section 5.2.1): although positive, the great majority are not instances of praise *per se*. Instead of mitigating the justifications for imagined violence, such tokens rather work to underscore them. For instance, Dickens NR says of the police force that (+propriety underlined) *they <u>condone</u> crimes against us*. Long R levels a charge, also against the police force, that *it <u>protects</u> all cops whether <u>good</u> or bad*. Dickens NR is actually decrying that the police excuse *crimes* against *Black ppl*, while Long R is constructing an argument against a *justice system* which shields *bad cops* from punishment. Although *condoning* and *protecting* have positive senses as standalone lexemes, in context they describe a state of the world which the authors view as unacceptable, and therefore something they wish to see changed.

Two other texts, LA Unified NR and Hribal R, are less overt in their use of +propriety, but the rhetorical effect is the same. LA Unified NR, for example, tells the pledge's addressees that *It is time to pray to allah, as this may be your last day*. In this pledge, the ongoing geopolitical conflicts between Islam and the West are a major frame. This exhortation to piety is, in context, actually an ironic taunt to the American school administrators whom LA Unified NR had emailed. For his part, Hribal R attributes a stance to his readers—*you preach that humans should be nice to each other*—then argues that a widespread failure to live up to this standard *makes you an exhibit of hypocrisy*. In his view, the state of the world is *harsh* and *evil*, and the change he hopes for is that his audience will recognize this reality the same way he has.

No matter if these epistemic stances are plainly stated or couched in taunts or attributions, tokens of +propriety like these help to communicate the author's *dis*satisfaction with the ethical condition of the world around them. And they do so by invoking the moral quality's photographic negative, e.g., the police force is *not* just (Dickens NR/Long R), the U.S. is *not* Islamic (LA Unified NR), and people are *not* altruistic (Hribal R).

Finally, the four tokens of +propriety which may be read contextually as genuinely complimentary are all found in realized texts. Having shot and killed his mother and father, the troubled Kinkel R praises them—sincerely, by all measures—as *wonderful people*, and argues that his violent actions are *not their fault*. And Roof R politely mitigates a threat to his audience's negative face, asking that they *Please forgive any typos*. The fact that only the realized texts allow for such earnest tokens of +propriety is interesting in itself.

5.3.2 WHO IS UN/ETHICAL?

According to Foster (2003: 36), ethics is tied inextricably to justice, and justice "is about receiving one's due or getting what one deserves." Or, as he says elsewhere, "[j]ustice served is ethics realized" (Foster, 2003: 35). Thus, a qualitative investigation of the various meanings found within propriety is discovering who, in the authors' view, is deserving of what. From a threat assessment perspective, a second question flows from this: are the perceived victims of future violence the objects of judgement in these texts, and, if so, how are they judged?

Addressing the initial question, Figure 5.19 shows the 'directedness' of each corpus, as a kind of snapshot of 'who' is being judged.

Figure 5.19: Directedness of Propriety by Realization Category

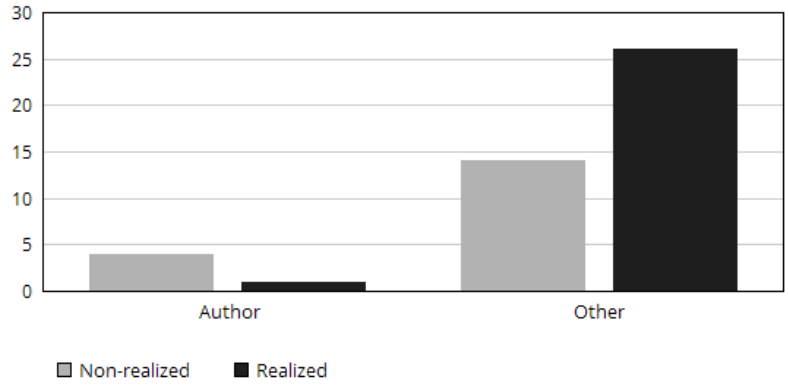

Frequency per 1000 words

While all judgements in the Appraisal framework are oriented towards a person other than the speaker, a certain subset of propriety tokens function similar to a "hybrid" (Martin & White, 2005: 68). A hybrid token inscribes one stance while managing to invoke a second type of attitudinal meaning—an operation very similar to lexemes which simultaneously denote one meaning while connoting an additional meaning. In this dataset, such hybrids are used by authors to inscribe a condemnation of others, but in a way which subtly reflects back on themselves. For instance, Valle NR says he aspires to be a *professional kidnapper*—literally, a paid criminal. Kidnapping is, of course,

considered a “serious felony offense” under U.S. federal criminal code 18 U.S.C. § 1201[25]. Despite his job as a policeman, Valle NR speaks admiringly of the person who engages in this kind of activity. However, the larger legal and social meanings of *kidnapper* implicate him in something inescapably negative.

A similar redounding of meaning occurs when LA Unified NR characterizes his or her group of compatriots as a *jihadist cell* (as discussed above), and when Rodger NR describes his own actions as *vengeance* (rather than, say, a related but more positive term like *justice*). Interestingly, the non-realized authors are four times more likely to permit this kind of stancetaking hybrid, thereby implicitly recognizing the immorality of their ideations.

The realized texts are far less ethically complicated in this regard. Just two such hybrid tokens appear in the R corpus, and both are used by the same author in the longest text of the dataset, Hribal R. First, he employs the 1st Person plural *we* to collectively denounce every person, including himself, as *selfish*, then frames his attack as *revenge* against his fellow students *for being so goddamn stupid* that they *value their lives*. Thus, despite the fact that tokens of negative propriety are nearly twice as likely to appear in the R corpus, such ethically dualistic lexemes are rare enough to be a stylistic choice by a single author.

Moving from author- to other-directed judgements addresses the next question: to what extent are the imagined victims of future violence the objects of negative propriety? Targets of violence were dealt with in capacity (section 5.2.6) from a grammatical standpoint—as NPs interacting with VPs and limiting or expanding elements like PPs and AdjPs—the current question calls for a different approach. For example, target NPs need not be governed by a future tense VP, since propriety is used

[25] http://criminal.findlaw.com/criminal-charges/kidnapping.html

not to make a threat but, potentially, to *justify* a threat by accusing the imagined victim(s) of past immoral acts. In other words, a target's prior misdeeds may be cited as warranting an author's future retaliation. A proper accounting of victims as stance objects requires a more discursive, contextual approach than simply focusing on NPs. A first step is to isolate them as a topic from the other human entities present in the text.

As van Dijk (1977: 16) defines it, "[t]he topic of discourse is a semantic structure which we take to be identical with the macro-structure of the discourse." Topics determine "the kind of possible events and actions which may take place in an episode" (van Dijk, 1977: 16). A topic, therefore, behaves as the organizational locus for various semantic threads, forming a kind of propositional umbrella. Those propositions which fall under the canopy are judged to be cohesive (per Eggins, 2004); those which fall beyond it are judged incohesive. Brahm NR's text provides perhaps the clearest example, in that every proposition relates directly to the bombing of the NFL stadiums, i.e., each clause is sensible and cohesive with its neighbors because it "originate[s] in the same range of semantic space" (van Dijk, 1977: 6), that of a terrorist attack. Propositions which do not originate in this space, then, would either be disallowed from entering the text by the author (either automatically or through later editing), or recognized as somehow irrelevant by the reader.

Thus, for instance, when Shaw R says that *people hurt my feelings*, this token is logged under 'Other', because 'people' in general are not the topical objects of Shaw R's ire. But when he says *I don't think Asian Women like me*, casting a moral aspersion on the preferences of Asian women, this token is catalogued as a judgement of his intended 'Victim(s).'

With these parameters in mind, the topical Victims of each text are isolated in Table 5.8.

Table 5.8: Primary Target by Text

Corpus	Author	Victim(s)
Non-realized	Brahm	NFL stadiums (metonym for the fans present)
	Dickens	all white cops nationwide
	LA Unified	students at every school in the L.A. Unified district
	McKelvey	every black person i see (at Kean University)
	Rodger	all of the men who have had pleasurable sex lives
	Skyline	people in the commons (area of Skyline High School)
	Valle	her (Victim-1)
Realized	Hribal	students of one of the “best schools in Pennsylvania”
	Kinkel	people
	Long	bad/good cops
	Rodger	Alpha Phi Sorority
	Shaw	every Asian Woman by herself

Per Figure 5.20 below, there is little difference between the pledge types in whether they consider their imagined victim(s) to be immoral. The NR authors are almost as likely to judge these people as somehow unethical or ‘bad’ as their R counterparts (14 tokens per 1000 words in the R corpus versus 11 tokens in the NR). Because propriety is perhaps the strongest resources a threatener may call on to advocate for retribution, it is interesting to find so little daylight separating the two realization categories.

Figure 5.20: Objects of Propriety by Realization Category

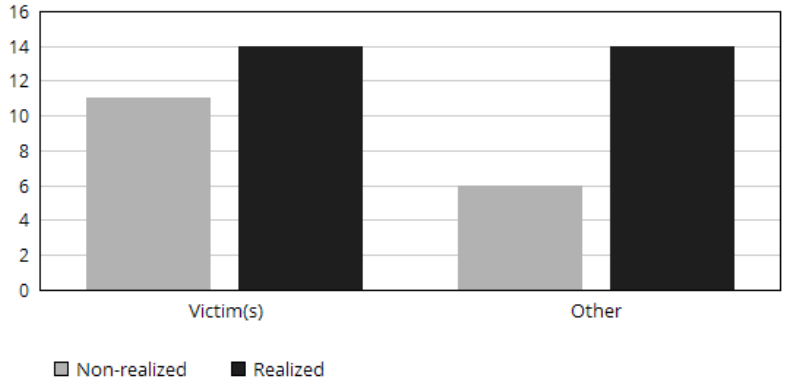

Frequency per 1000 words

However, the non-realized authors are more likely to focus on these imagined victims to the exclusion of other, outside entities. Realized authors, on the other hand,

split their focus evenly between the two. This category of Other may be further subdivided into separate, roughly delineated areas: Agent-less events (such as the nominalizations *attacks* and *civil wars*), and those non-victims named by the authors (e.g., *jihadist cell*, *professional kidnapper, cops*, *everyone who hates blacks people*). The realized writings utilize both of these general classifications as well (e.g., *the dehumanization of public school* for an Agent-less event, and *justice system leader's* as non-victims), as well as a third subcategory, here called simply 'generic'. This last category encompasses entities who are identified by the authors but whose scope is nevertheless so broad as to be essentially global, e.g., *humans should be nice* (Hribal R).

Figure 5.21: Objects of Propriety Subcategories

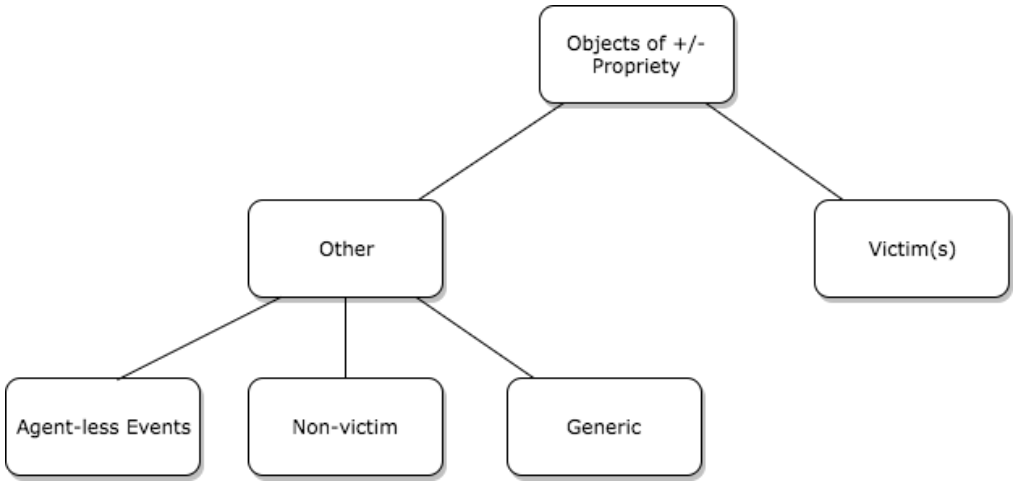

First, an interesting semantic difference is found in the kinds of judgements the realized authors levy against non-victims, in that all the instances of nominals marked for positive propriety are found in this category. By contrast, no NPs in the non-realized corpus communicate positive propriety. Second, the realized writers are twice as likely to address positive and negative ethical judgements towards non-victims. When these two trends are considered in combination, the realized texts read as less single-minded in their fixation on some 'bad' entity the author imagines harming, while also admitting

that ‘good’ entities do, in fact, exist. This is a more nuanced worldview than what is apparent in the non-realized texts, which are relentlessly negative and far more victim oriented.

Figure 5.22: ‘Other’ Objects of Propriety by Realization Category

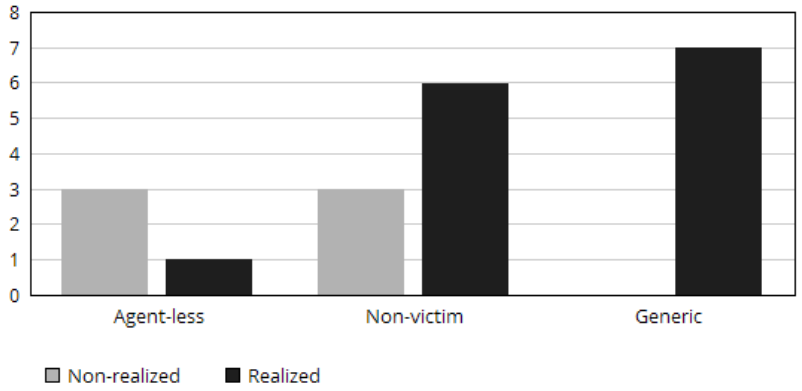

Frequency per 1000 words

An additional, noticeably different pattern of usage occurs with Agent-less events. The realized texts include just one token of propriety in such a semantically passive construction, again in Hribal R’s text when he blames his actions on the *dehumanization of public school*. In the NR texts, Agent-less events constitute fully half of the Other items present (three of the six tokens per 1000 words), as Figure 5.22 above shows. In short, the realized texts leave far less room for doubt concerning who is to blame for what.

Finally, and perhaps most significantly, over half the tokens of Other in the realized writings (8 per 1000 words out of a total of 14) fit into neither the Agent-less nor the Non-victim categories, but instead fall under what is here called Generic. Such tokens of propriety address society or the world at large. This indefiniteness is signaled via grammatical devices like the 1st Person plural (***we*** *are all* *selfish*), or the ‘generic you’ (*making* ***your*** *fellow man* *suffer*), as well as through non-delimited common nouns (***humans*** *should be* *nice*, ***people*** *stain the world with sins*), and ‘allness’ terms (***the***

whole World *Hates me*). Interestingly, the NR pledges employ none of these wide-angle devices to deliver judgements of propriety, limiting themselves rather to people or communities that are either more circumscribed or simply unnamed. The R texts are, thus, far more likely to let their ethical discussions range across the entire spectrum of possible targets, up to—at its widest end—a status quo which can include society in general or even humanity itself.

5.3.3 PARTS OF SPEECH

A similar exercise may be performed with these tokens of propriety as the one in section 5.2.4 above, which briefly shifted away from the semantics of Appraisal to consider the more rudimentary grammatical resources of stancetaking in the two corpora. Examining the parts of speech employed by these authors sheds light on how these attitudinal meanings are lexicalized. It also allows for examining whether one corpus is more fixated on im/proper entities (nouns), im/proper processes (verbs), or im/proper qualities of people or processes (adjectives, adverbs). The semantic class by which a judgement of propriety is realized can clarify where and in what ways the various authors believe society is either in harmony or out of joint.

Naturally, no one-to-one correspondence is required between the attitudinal load of a given adjective or adverb and the semantics of the item it is modifying. For instance, Skyline NR calls his or her imagined victims *smug and snooty cunts*. *Smug* and *snooty* come very close, respectively, to Martin and White's (2005: 53) examples of 'vain' and 'arrogant,' and so fall within the scope of -propriety. *Cunts*, however, connotes effeminacy and/or weakness, and so is coded instead under capacity. These two adjectives, *smug* and *snooty*, would be lost as tokens of impropriety, though, or their meanings muddied, if their stance functions were categorized solely by the phrase's constituent head, the noun *cunts*. For this reason, each lexeme is considered separately

in Table 5.9 below. Along with this, what van Leeuwen (1996) calls process nouns, and what Halliday (2004) calls nominalizations, are isolated from the conventional NPs present in the dataset. Process nouns "function as nominals, although they refer to activities" (van Leeuwen, 1996: 40), and so contribute more to understanding how concerned these authors are with in/offensive actions than with in/offensive entities. All the NPs marked for +/-propriety in the corpora appear in Table 5.9 below.

Table 5.9: Noun Phrases

Corpus	Noun Phrase	Process Nouns
Non-realized	jihadist cell	attacks
	hate	attacks
	offender	America's Hiroshima
	professional kidnapper	civil wars
		bullying
		discrimination
		vengeance
		vengeance
Realized	Selfishness	dehumanization
	not their fault	war
	fault	punishment
	wrong's	threats
	bitches	
	skinheads	
	KKK	

First, these may be grouped together for a quantitative snapshot of the two pledge types. For the sake of clarity, nominalizations are grouped together with VPs under the larger header of 'processes' in Figure 5.23 below. Both author types are equally as likely (or unlikely) to classify an entity under an attitudinally charged nominal, e.g., *bitches* or *offender*. The differences between the two corpora are more apparent in their relative focus on processes and qualities. While the non-realized authors are over three times as likely to nominalize actions, the realized authors are almost 60% more likely to speak to some kind of ethically fraught process, such as *protecting* (+propriety) or *dehumanization* (-propriety).

Figure 5.23: Parts of Speech in Propriety by Realization Category

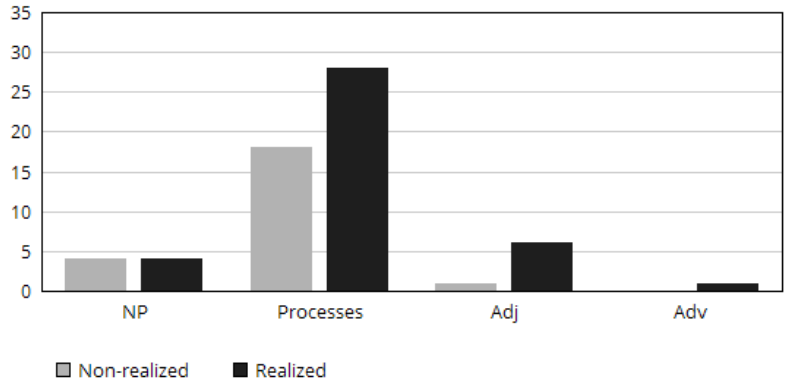

Frequency per 1000 words

One possible interpretation of this disparity is that the simple fact of a person's existence is not terribly troubling to either author type, but the realized authors are collectively more concerned with the ethics of other people's actions. Meaning, stance objects in either corpora are unlikely to be judged as simply 'bad' by virtue of, say, their social or ethnic groups. There are no blanket ethnic or religious slurs in the dataset, for instance, and only one sexist term (*bitches*, Rodger R). Instead, the objects of the authors' displeasure are far more likely to be judged as 'bad' because of their behavior (via verbs and process nouns) and the qualities these choices reveal about them (via adjectives and adverbs).

A qualitative difference in these processes exists alongside the numerical difference. While it is difficult to reliably sort these meanings based on their severity—no fully-fleshed out typology of ethical transgression exists outside of legal and religious codes, which are of limited help here—some broad strokes may be applied to the data using shared world knowledge and a native speaker's sense of proportion. For instance, *bullying* is more of a social crime, while *revenge*, *stealing*, and *harassment* arguably cluster nearer to issues of misused authority or social status[26]. Further, both corpora

[26] Even this sample distinction is admittedly murky. A person cannot bully or be bullied without some kind of extant status or power imbalance between the two parties. Yet,

Table 5.10: Process Lexemes of Negative Propriety

Corpus	Type of Meaning			
	Mental State	Social Conflict	Abuse of Power	World Historical
Non-realized	hates	act superior	(cops) wont save	attacks
		blamed	deserve	attacks
		bullying	deserve it	America's Hiroshima
		abuse	dont deserve	civil wars
			kidnap	
			discrimination	
			monger	
			vengeance	
			vengeance	
Realized	Hates	hurt my feelings	revenge	inflict (destruction)
	hate	don't think Asian Women like	holding them accountable	kill
		are only concerned about yourself	punished	dehumanization
		look down on me	reprimanded	war
		ignored	go unpunished	
		kicking (me)	punishing	
		accuse	punishment	
		blame	suffer	
		blamed	do this	
			exploit	
			threats	
			stain the world with sins	
			steal	
			harassed	
			blackballed	
			blacklisted	
			forces	

bullying connotes the schoolyard and the disciplinary measures that come with it, where the 'adult' version might be *harassment*, a recognized crime with known legal penalties. And so, the former is collected under the 'social' category while the latter is counted as an 'abuse of power'. Although the particulars of each token's classification are arguable, more would be lost by *not* addressing these meanings than by proceeding in the rough-and-tumble fashion followed here.

include some form of the lexeme HATE as an ongoing mental state, and discuss events that have wider societal implications, such as *war*, *killing*, *attacks*, etc. The full list of these process lexemes (verbs and nominalizations) is captured in Table 5.10 above.

The difference captured here is essentially one of scale. 'Smaller' ethical violations—e.g., *ignored*, *hurt my feelings*—more comfortably fall within interpersonal or social conflicts and are collected to the left of the graph below, while impersonal, societal conflicts are collected towards the right. An offensively negative predisposition (*hates*), unconnected with some kind of action, is thus placed at the far left. Meanwhile, larger scale transgressions—in terms of the number of people impacted as well as the severity of the aftereffects—are grouped further to the right. These include such concepts as war and racial discrimination. Figure 5.24 below provides a general outline of how much of these meanings is present in each corpus.

Figure 5.24: Unethical Processes by Realization Category

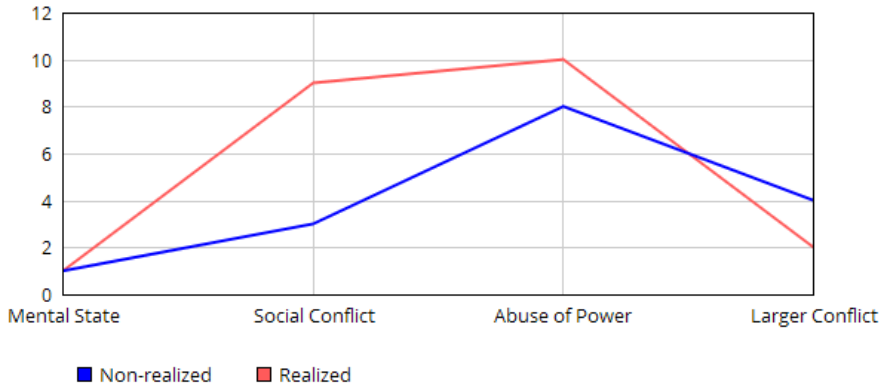

Frequency per 1000 words

When parsed this way, the R texts are three times more likely to concentrate on less serious 'sins' than the NR pledges, and only slightly more likely to point to the more serious misdeeds which fall under Abuse of Power as justifications for their grievances. This is despite the fact that the realized authors would later use some form of violence trying to rectify these perceived breaches. Meanwhile, the non-realized texts are twice

as likely to view conflict on a scale so macro as to be nearly abstract (*civil wars*, *attacks*). Metaphorically, then, the realized authors may be said to be more concerned with moral lapses that are 'closer to home,' i.e., more communal or personal, and not of world-historical significance.

Finally, there is an unmistakable difference in how the *qualities* of the various entities and actions are described in these writings. The R pledges are seven times more likely to speak to such qualities via modifiers like adjectives and adverbs, as Table 5.11 shows. The communicated meanings are also far richer in the realized texts.

Table 5.11: Modifiers Marked for Propriety

Corpus	Adjective		Adverb
Non-realized	smug		
	snooty		
Realized	selfish		cruelly
	spoiled		
	bad		
	Disgusting		
	disgusting		
	evil		
	evil		
	heartless		
	wicked		
	wrong		
	unpunished		
	nice	+prop	
	wonderful	+prop	
	right	+prop	
	Protected	+prop	

The meanings of the two adjectives in the NR corpus—*smug* and *snooty*—may both be conflated to a sense of snobbishness or arrogance. As moral transgressions go, this is far less serious than the meanings laid out in the R texts, where people are judged along essentially biblical lines as *evil* and *wicked*. Several of the tokens of +propriety also appear in this lexical category, adding to the variety of meaning employed by the realized authors. Thus, although both the NR and R authors are equally as unlikely to classify a person or group of people under an ethically charged nominal header (e.g.,

bitches, *cunts*), the realized authors are many times more likely to ascribe certain moral qualities to identified entities through these additional modifiers. Often, these are qualities of the utmost ethical deficiency. The larger prosodic effect in the R pledges of this combined with the 'closer to home' quality of the unethical processes is of a community more immediate to the author that is also more deeply broken.

5.3.4 TENSE

In section 5.2.5 above, capacity is examined according to the grammar of verb tense, to gauge how many of the violent actions in these texts are projected into a future time. Tense is also a useful tool for understanding the fantasized results of those actions. Does the author imagine that the moral universe will improve following the harm he or she envisions causing? Presumably, someone with a grievance strong enough to spur the writing of a pledge—much less a subsequent, real-world act of aggression—might be expected to depict a world that is better off following the violence. Such an expectation is prefaced on, e.g., the primacy of satisfaction in violent ideation (Harmon, 2008; Simons & Tunkel, 2014). However, as Figure 5.25 below shows, ethical judgements of future states do not figure especially prominently in either pledge type.

Figure 5.25: +/-Propriety Tense by Realization Category

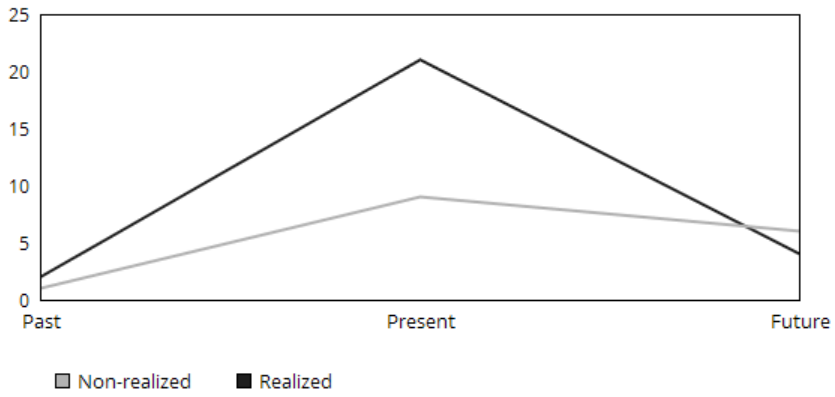

Frequency per 1000 words

The focus on the present in the realized corpus is perhaps the most striking result. Two folk linguistic expectations might be that realized authors believe the future will be better after the fantasized violence, and/or that past bad actions would be clear motivators for the authors' violent ideations. Neither is borne out. Instead, the realized authors more or less ignore both past and future in favor of positive and negative evaluations of people in the present tense, as Rodger R does when he says the sorority members *are all spoiled, heartless, wicked bitches*.

Figure 5.26: Present Tense +/-Propriety by Realization Category

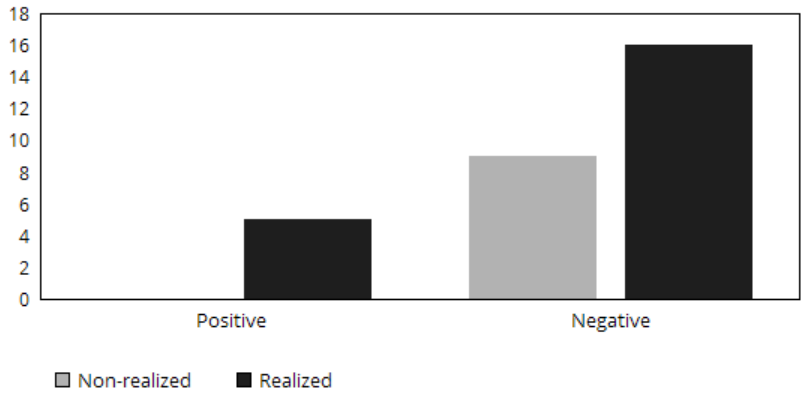

Frequency per 1000 words

Further, when these present tense tokens are separated by polarity, as shown in Figure 5.26, they are over three times more likely to be negative. Past bad deeds, therefore, do not motivate this anti-social ideation as much as current bad circumstances. Similarly, future imagined satisfaction—that things will be set right after the guilty parties are physically punished—does not figure prominently either. The here and now is by far the main concern in the R pledges, and the intolerableness of the here and now seems to be a much stronger prompt for violent ideation than any potential gains the authors imagine they will reap.

5.3.5 SUMMARY OF PROPRIETY

Propriety is a rich source of meaning for the authors in both corpora, and an area offering strong distinguishing features between the two realization types in this dataset. While some level of preoccupation with morality is, for all intents and purposes, universal across the dataset, the ethical preoccupations themselves, as well as how these are realized lexically and with what frequency, are not identical. A non-realized text is:

- more likely to blame than praise others, at a rate of 16:1, where a realized text offers blame rather than praise at a rate of just over 2:1;
- four times more likely to permit a stancetaking hybrid which reflects somehow negatively on the text's author;
- twice as likely to view conflict on a global or societal rather than a personal scale;
- and almost twice as likely to focus on imagined victims to the exclusion of other entities, where a realized text addresses both victims and non-victims with equal frequency.

By contrast, a realized text is:

- eight times more likely to utilize resources communicating +propriety, and to employ these tokens of praise sincerely (the +propriety which appears in the non-realized texts is all contextually negative);
- twice as likely to address ethical judgements towards non-victims;
- three times more likely to concentrate on less serious 'sins' (e.g., *selfishness* vs *discrimination*);
- seven times more likely to speak to ethical qualities, via modifiers like adjectives and adverbs, and these qualities tend to be more severe, even biblically, negative;

- three times more likely to feature ethically fraught processes (via verbs and process nouns) which are communal rather than global in scale, and of lesser ethical severity; and
- far more likely to situate their imagined victims within this communal scope, as an identifiable part of a larger, unacceptable status quo.

Having now discussed the significant variables within judgement, the final subsystem of attitude, appreciation, will be taken up in the next chapter.

CHAPTER 6 APPRECIATION

"Appreciation looks at resources for construing the value of things, including natural phenomena and semiosis" (Martin & White, 2005: 36). This final domain of attitude therefore addresses how writers evaluate objects, whether real or abstract. Broadly speaking, for texts which may contain evidence of premeditation, appreciation offers a way to understand whether the writers consider a proposed event or series of events to be feasible, whether they believe their reasons for desiring the event(s) are well thought out, and their stance towards the instruments by which the event(s) may be made to happen, e.g., via bombs.

Three variables are available for coding in appreciation: "our 'reactions' to things (do they catch our attention; do they please us?), their 'composition' (balance and complexity), and their 'value' (how innovative, authentic, timely, etc.)" (Martin & White, 2005: 56). Of the three, a statistically significant difference in usage is observed in two: composition and valuation.

Table 6.1: Statistical Significance in Appreciation by Realization Category

Appreciation Type	Tokens		Probability (p)
	Non-realized	Realized	
Reaction	5.15	4.39	> .05
Composition	26.35	5.02	< .001
Valuation	9.16	26.67	< .01

Frequency per 1000 words

Addressing the null result of reaction first, Martin and White (2005: 57) point out that "there are strong links between…reaction and affect" since both involve the expression of an author's feelings. So, for instance, the adjective 'weepy' would be coded as a token of affect when applied by a person to him- or herself, as in "I'm weeping," but as a token of reaction when used to evaluate a song, as in a "weepy rendition" (Martin & White, 2005: 58). The distinction in such a case is between

"construing the emotions someone feels (affect) and ascribing the power to trigger such feelings to things" (p. 57-58). The latter falls under reaction.

Given the strong correlation between these two semantic domains, it is merely interesting to note that both reaction and affect turn up statistically null results in the dataset. This result further reinforces the conclusion that personal feelings are of essentially equal value to both author types, in that descriptions of internal states are no source of distinction within the dataset. Once again, crucial differences in stance between the realized and non-realized pledges are oriented institutionally.

6.1 COMPOSITION

Composition captures an author's view of order—whether an item or idea is considered flawed or functional, simple or extravagant (Martin & White, 2005). Importantly, such evaluations—how well or poorly things "hang together" (Martin & White, 2005: 56)—exclude what the objects or their uses are worth. Questions of worth are instead addressed by valuation, which is discussed in section 6.2 below. Rather, composition is focused on issues of coherency[27]. If, as Malle and Knobe (2001: 55-56) say, "intentions serve to fulfill desires by identifying a course of action that is feasible to implement for the agent," then composition is one way to assess such feasibility. For instance, is a proposed course of action presented as logically sound or well-considered, as in *my strategic plan of using violence* (Shaw R)? How do the authors characterize the extent of the aftermath, e.g., *Global economies will screech to a halt* (Brahm NR)? Does shared world knowledge argue for or against the elements needed to enact the author's ideations? For instance, *bombs* are attainable (LA Unified NR) but a *WEAPON OF*

[27] Composition may be further subdivided into meanings related to the "balance" and the "complexity" of the evaluated item (Martin & White, 2005: 56); however, no statistical significance is detected at this finer level.

MASS DESTRUCTION is not (Archangel Michael NR). And finally, are the tactics implementable, e.g., *some form of trickery* (Rodger NR)? Examples of +/-composition in the two corpora include:

- *[The bombs] are strategically placed* (+composition, LA Unified NR)
- *my art has obviously been revealed* (+composition, Hirbal R)
- *civil wars will erupt across the world* (-composition, Brahm NR)
- *My head just doesn't work right* (-composition, Kinkel R)

Figure 6.1: Positive Composition Dispersion by Corpus

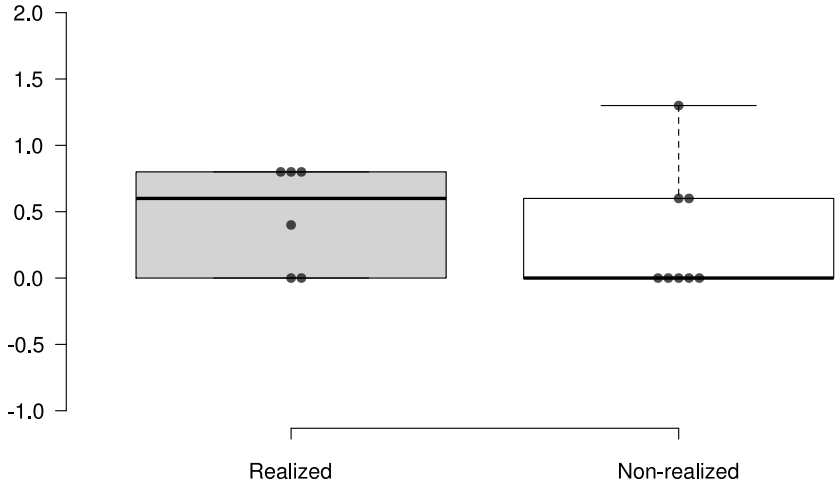

Frequency per 1000 words. No outliers (>1.5IQR) are found.

Figure 6.2: Negative Composition Dispersion by Corpus

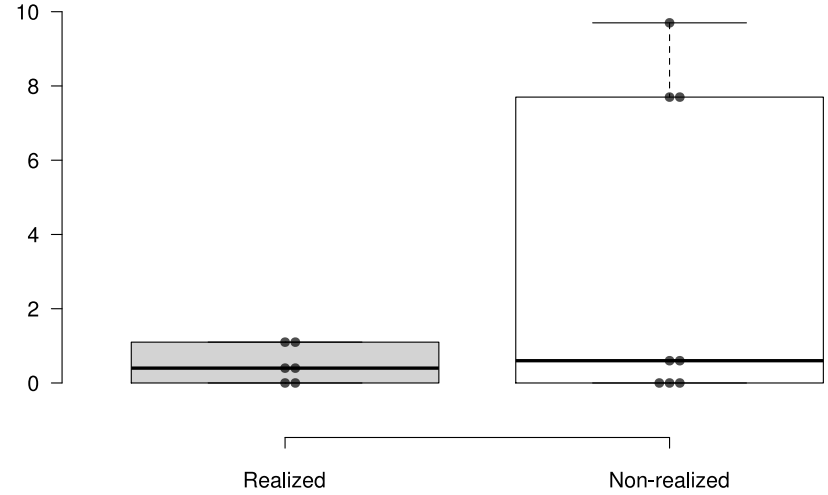

Frequency per 1000 words. No outliers (>1.5IQR) are found.

Figures 6.1 and 6.2 show the dispersion of composition across the dataset. Per text, evaluations of order are common: 12 of the 14 authors in the dataset utilize some form of composition, negatively or positively. Only Dickens NR and Rodger NR abstain. Interestingly, a bird's-eye view of this variable shows that its negative manifestations, i.e., tokens of disorder, are far more important to the authors of the *non-realized* texts.

Figure 6.3: Composition Polarity by Realization Category

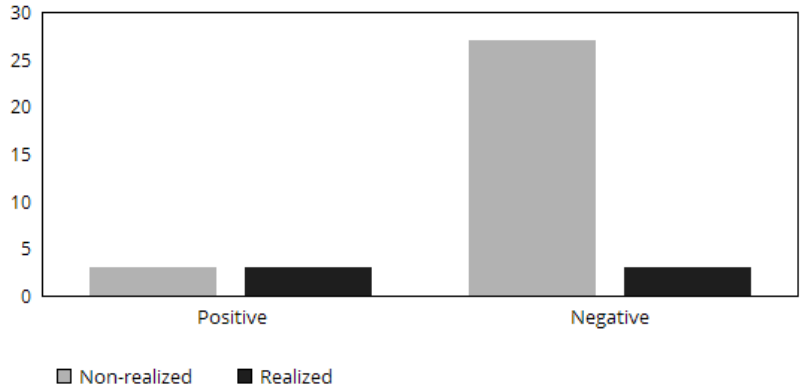

Frequency per 1000 words

Addressing the small scope of positive meanings first, one similarity and one difference between the corpora are apparent. As Figure 6.1 shows, tokens of +composition are rare, appearing at a rate of just 3 tokens per 1000 in both realization categories. Both R and NR authors employ this type of meaning in a similar fashion: to evaluate their proposed actions. For example, both Shaw R and LA Unified NR

Table 6.2: Tokens of Plan-related +Composition

Corpus	Author	Token
Non-realized	LA Unified	[The bombs] are *strategically* placed
	Skyline	[the jocks] are in the middle of the commons so it would be an *ideal* place to start [killing]
	Valle	The abduction will have to be *flawless*
Realized	Hribal	By now, my *art* has obviously been revealed to the world
	Shaw	I'm going to talk to a few more Asian Women, before I start my *strategic* plan of using violence
		I think its *brilliant* to give all Asian Women a legitimate reason to hate me.

characterize their plan for violence, in part or in whole, as being *strategic*. Table 6.2 shows all instances of these planning-related tokens in the dataset.

Where the author types differ is in the realized authors' use of +composition as a way of justifying their violent ideations. Such a use is not found in the non-realized writings. For instance, Long R says *the way the current system is set up* creates *fertile ground for bad cops to flourish*, and offers this metaphor as a pretext for killing police officers. Elsewhere, Roof R refers to his text as, simply, *An Explanation*, implying that his thoughts on committing hate crimes are reasonable enough to be both argued and understood. Although writers in both corpora employ tokens of +composition to praise the ingenuity of their designs, only authors of realized texts use this semantic resource to defend the rationality of the violent ideations themselves. That said, while this difference between the corpora is worth noting, more data would be needed to support (or contradict) the apparent variation in usage.

6.1.1 NEGATIVE COMPOSITION

According to Turner and Gelles (2003: 97), the mention of a "particular weapon or method that the author intends to use as an instrument of violence" signals that a threat is more credible. Gales (2010: 282) catalogs these under the simple header, "item[s] used to kill." Because weapons, whether man-made or natural, are fundamentally instruments of *dis*order, -composition offers a means of cataloging a text's reference to these objects. And in fact, the great majority of tokens in the dataset which are coded under this variable are either nouns denoting some class of 'Weapon' (*guns*, *knives*) or various parts-of-speech which are somehow 'Weapons-related', e.g., a verb like *detonated*. (An example of the third category shown in Figure 6.4 below, -composition categorized as 'Other', would be Brahm NR's assertion that *chaos will rule*.) However, the distribution of these three token types is highly uneven.

Figure 6.4: -Composition Senses by Realization Category

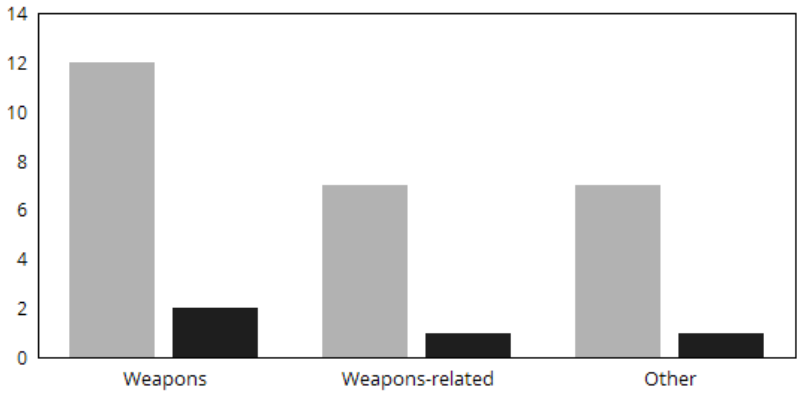

Frequency per 1000 words

A higher frequency of these types of tokens in any one text may be explained partly as the result of the semantic macro-structure provided by the topic (van Dijk, 1977): if a *bomb* is mentioned, chances are naturally higher that a term like *explosion* will collocate. However, the disparity between the corpora in weapons-related terminology cannot be wholly accounted for just through the idea that one member of a semantic field will attract more of its fellows. For one, the NR pledges not only contain more numerous tokens denoting weapons, these tokens are also more specific—i.e., they are hyponyms of more unmarked categories. Variations of GUN and BOMB appear alongside deeper subordinate terms like *rifles* and *pistols*, or *grenades* and *explosive devices*. Authors of the realized texts only cite superordinate terms such as *guns, knives*, and the most unmarked hypernym *weapons*.

Furthermore, only the authors of the NR texts somehow graduate these lexemes via an "isolated, individual item which solely, or at least primarily, performs the function of setting the level of intensity" (Martin & White, 2005: 141). In the Appraisal framework, isolating locutions are typically some kind of lexicalized modifier, such as "very miserable" (Martin & White, 2005: 141). However, rather than scaling for a more typical quality like 'size,' e.g., *big gun* or *large knife*, the nature of the weapons in these pledges is intensified instead through technical descriptors such as *Kalashnikov rifles*, or *PETN*

PLASTIC EXPLOSIVES. This is perhaps an effort to make the proposed violence more credible—and thus more menacing—since evidence that an author possesses the "technical expertise…to carry out the threat" (Gales, 2010: 25) is one measure of its seriousness. But again, of the 24 raw tokens explicitly denoting weapons in the dataset (19 in the non-realized texts and 5 in the realized), the 10 technical modifiers are appended only in NR pledges. Thus, not only is a non-realized author six times more likely to explicitly mention some kind of weapon (at a normed rate of 12 tokens per 1000 versus 2 per 1000 in the R corpus). The weapon cited is also likely to be more specific—through the resources of hyponymy, isolated technical modifiers, or both. If weapons-related lexemes are included in this tally, a NR pledge is over 6 times more likely to somehow discuss instruments employed for violence, the results of their use, etc. (at a normed rate of 19 tokens per 1000 versus 3 per 1000 in the R corpus).

This trend would seem to run counter to Turner and Gelles' (2003: 97) observation that an author who mentions a "specific make and model of a particular weapon" has a higher likelihood of acting. This finding also reinforces the patterns of usage uncovered within capacity—that NR writings tend to contain more violent ideation than their R counterparts, and that this ideation is realized textually through more varied lexicogrammatical choices.

Lastly, those few tokens of -composition which fall within the category of Other are nevertheless related to the central, violent event of the pledges in which they appear, if the events described by a pledge are viewed more as the temporal junctures within a story. However, as noted in section 3.3.3 above, the Appraisal framework is ill-suited to analyzing narrative structures. And so, to understand how these tokens of Other may relate to the central, violent event, they are, perhaps, best discussed in terms of Narrative Analysis. Labov (2003, 2010, 2013) categorizes the work performed by clauses in a narrative into six main functions:

- Abstract: briefly encapsulates the most reportable event (why the story is being told) as well as pragmatically signaling that a narrative is to follow;
- Orientation: answers questions of who, when, where, etc.;
- Complicating action: progression of events signaled by temporal junctures;
- Evaluation: accomplished by juxtaposing real and potential events;
- Resolution: often the aftermath of the most reportable event;
- Coda: returns the narrative to the present moment.

Following the conventions briefly laid out by Labov (2010), ABS = abstract, OR = orientation, CA = complicating action, EV = evaluation, and RES = resolution.

Table 6.3: Narrative Category of 'Other' Tokens of -Composition

Corpus	Author	Token	Category
Non-realized	Brahm	civil wars will **erupt** across the world	RES
		Global economies will **screech to a halt**	RES
		General **chaos** will rule	RES
	LA Unified	If you do end up…**canceling** classes for the day	EV
		If you **cancel** classes	EV
	Rodger	start luring them into my apartment through some form of **trickery**	CA
Realized	Hribal	**Ragnarok**	EV[28]
	Kinkel	My head just **doesn't work right**	EV

In these pledges, the most reportable event, or the event which is "the least common and has the greatest consequences for the life chances of the actors involved" (Labov, 2010: 7), is the primary act of violence described by the pledge. In Brahm NR, for example, the most reportable event is the simultaneous bombing of the NFL stadiums. In Table 6.3 above, these tokens of Other, i.e., -composition meanings which do not denote some kind of weapon and are not directly weapons-related, are

[28] Hribal R begins his text with this single, free-standing word, and while it is potentially classifiable as an orientation—answering questions about where (what kind of world) the textual voice inhabits—the mythical/metaphorical qualities of the term argue instead that its function is attitudinal, i.e., evaluative.

nevertheless shown to be instances of stancetaking related to events described by the pledges.

Despite the relative paucity of data, it is interesting that these eight tokens manage to conform to patterns discerned elsewhere in the corpora through capacity and propriety. The NR texts are far more process-oriented, offering stances which are temporally organized around the most reportable event—meaning they are presented as occurring either before or after the central violent ideation driving the creation of the pledge. For instance, Brahm NR closes his narrative chain by resolving the inchoative forms (*erupting*, *screeching*) into a final resultative state (*chaos*), while Rodger NR describes actions (*trickery*) which he believes will enable the most reportable event (the *torture* of those he has tricked). The use of evaluative clauses by LA Unified NR is similarly procedural. Labov (2013: 30) notes that evaluative clauses establish "[t]he human meaning of the events that did occur," and often perform this function through the grammatical resources of the irrealis mood. But while LA Unified NR places the *cancelation* of the classes in an irrealis space (via the conditional *if)*, these two constructions add no extra meaning beyond what might be called contingency thinking.

By contrast, the R texts are not concerned here with disordered processes. The clauses containing the two Other tokens are not subject to a temporal juncture, and so are not situated in relation to these pledges' most reportable event. Hribal R includes no tense information when he writes, simply, *Ragnarok*—a reference to the mythical Norse Armageddon when the world is at the apogee of disorder and decay. Kinkel R employs both the present tense and a negative—a device flagged by Labov (2013) as a common resource for evaluative phrases—thereby presenting his judgment as an ongoing or habitual state. Both Hribal R's and Kinkel R's tokens add an important layer of human meaning—of ongoing external (Hribal R) and internal (Kinkel R) disorder—to the other events described in their pledges. And they do so in ways which are absent in the NR

writings. This finding, as dependent as it is on the little data available here, nevertheless echoes a tendency uncovered in the analysis of propriety, where R pledges are more concerned with the circumstances surrounding the imagined violence, while NR pledges are instead more occupied with the imagined violence itself. Thus, within the variable of -composition, the micro can be seen recapitulating the macro.

6.1.2 SUMMARY OF COMPOSITION

The resources of composition offer a means of gauging an author's view of dis/order. These resources are found throughout the dataset: 12 of the 14 authors take stances related to the order of objects, processes, and/or ideas. Positive evaluations are used by authors of both realization categories to assess the cleverness of their violent ideations, though only realized authors use this resource to *justify* the ideations themselves.

Meanwhile, the great majority of -composition appears in the NR corpus. Quantitatively, a non-realized pledge is six times more likely to mention an "item used to kill" (Gales, 2010: 282), such an item's state (*a hidden bomb*), its purpose (*crumble the foundations*), and/or its effects (*damage*). Weapons cited by a NR author are also likely to be more specific, either through the resources of hyponymy, isolated technical modifiers, or both. Finally, instances of -composition which are not directly weapons-related nevertheless pertain to the violent ideation at the heart of a pledge to harm, although the NR authors use these tokens to evaluate other disordered processes, while R authors employ them to characterize conditions surrounding the imagined events. These findings—heightened violence in NR texts, and a focus on justification in R texts—mirror previous results uncovered in capacity and propriety, respectively.

6.2 VALUATION

Where composition covers an author's view of order, the final variable in the subsystem of appreciation captures estimations of an object's value, i.e., "was it worthwhile?" (Martin & White: 2005: 56). Examples of +/-valuation in the two corpora include:

- *This world would undoubtedly be better if we were all in heaven.* (+valuation, Hribal R)
- *Humans found racism popular.* (+valuation, Shaw R)
- *The death toll will approach 100,000* (-valuation, Brahm NR)
- *In such a terrible state of affairs* (-valuation, Hribal R)

Figures 6.5 and 6.6 show the dispersion of this resource across the dataset.

Figure 6.5: Positive Valuation Dispersion by Corpus

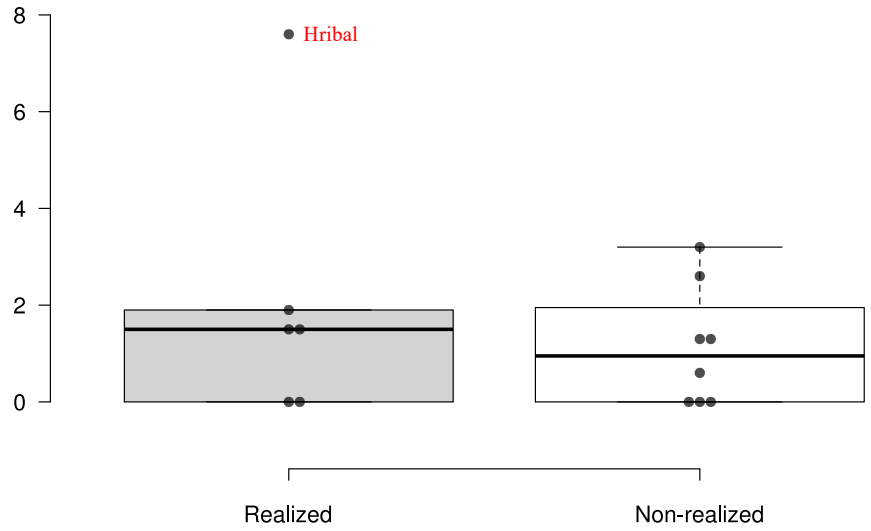

Frequency per 1000 words. One outlier (>1.5IQR) is found.

Figure 6.6: Negative Valuation Dispersion by Corpus

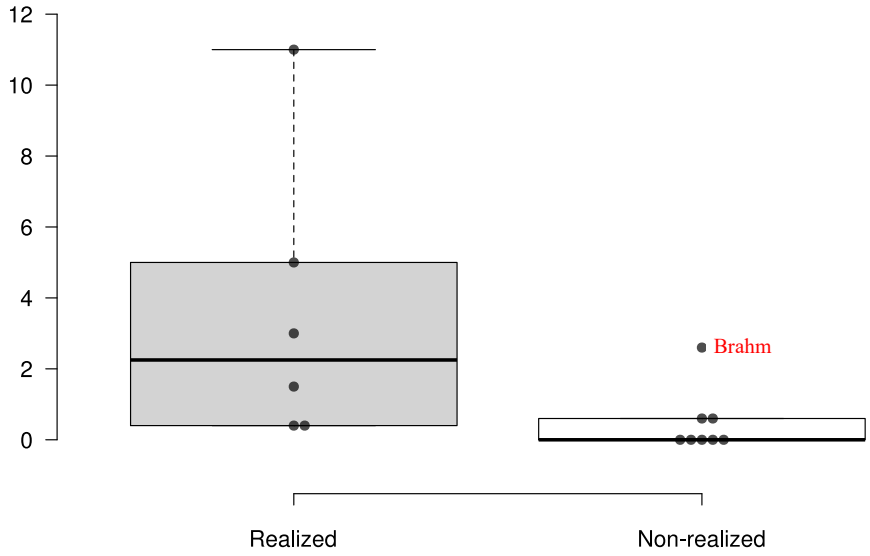

Frequency per 1000 words. One outlier (>1.5IQR) is found.

Both Hribal R and Brahm NR are flagged as outliers. Both will be discussed below. But first, an interesting theoretical opportunity presents itself in this area. Martin and White (2005: 27) note that valuation is very sensitive to the 'field' in which a text is operating. This systemic category is "concerned with the discourse patterns that realise the domestic or institutionalised activity that is going on." Or, in other words, the "value of things depends so much on our institutional focus" (Martin & White: 2005: 57). For example, with the genre of cooking recipes, where the institutionalized activity is instruction, the semiotic process performed by the text is one of 'enabling' (Matthiessen, 2015: 8).

As noted in the discussion of attitudinal null results (section 4.2), the two pledge types differ little in their usage of personal meanings (through affect) but significantly in their stances towards institutional meanings like "rules and regulations" and "criteria and assessment" (Martin & White, 2005: 45). However, the concept of field is not only useful for understanding the nature of the valuations in these texts, it is also potentially helpful from a theoretical standpoint. Field offers a means of directly connecting related stancetaking patterns spread across valuation and judgement. Illuminating these

connections, though, requires a brief return to SFL theory to expound on a suggestion from Martin and White (2005) concerning the behavior of valuation meanings.

In SFL, a text's institutional focus is a direct result of "the nature of the social action taking place" through its language (Halliday & Martin, 2003: 36). Because of this, field is said to perform a 'textual' metafunction, "organizing the discursive flow and creating cohesion and continuity as it moves along" (Halliday & Matthiessen, 2004: 30). Cohesion between kinds of meaning is common within Appraisal, as seen in the intra-system echoes between reaction ('does *it* please us?') and affect ('are *we* pleased?') discussed at the beginning of this chapter. Similarly, Martin and White (2005: 58) say that "positive and negative valuations of something imply positive and negative judgements." The authors point specifically to a link between valuation and judgement which runs through the meanings of capacity, e.g.:

judgement: capacity	appreciation: valuation
a brilliant scholar	*a penetrating analysis* (Martin & White, 2005: 58)

O'Donnell (2013) expands on the links between valuation and judgement when he observes that tokens of valuation often contain ethical meanings, and thus dovetail with the domain of propriety. And indeed, overlaps between valuation and the variables of capacity and propriety are readily apparent in the dataset. For instance, Shaw R uses valuation to instantiate +capacity when he writes *It was the greatest achievement of my life.* And Hribal R uses valuation to instantiate -propriety when he writes *they saw something wrong in the world*.

Importantly, neither Martin and White (2005) nor O'Donnell (2013) rule out the possibility that valuation may show qualities corresponding to the remaining three variables of judgement: *normality* (how special?); *tenacity* (how dependable?); and *veracity* (how honest?). And indeed, the full analysis of this variable reveals that

valuation meanings break along these already delineated semantic lines relatively easily. For instance:

- *It will make national headlines* (LA Unified NR, +normality via +valuation)
- *It opposes their hypocritical opinion* (Hribal R, -tenacity via -valuation)
- *their is a unseen & concealed war* (Long R, -veracity via -valuation)

Further, when all instances of valuation (positive and negative) are sorted according to the type of judgement with which they most closely align, and these are superimposed on the findings within judgement itself, correlations of the kind that Martin and White (2005) and O'Donnell (2013) postulate are immediately evident. The overlay of valuation on judgement is shown in Figure 6.7.

Figure 6.7: Valuation and Judgement Variables

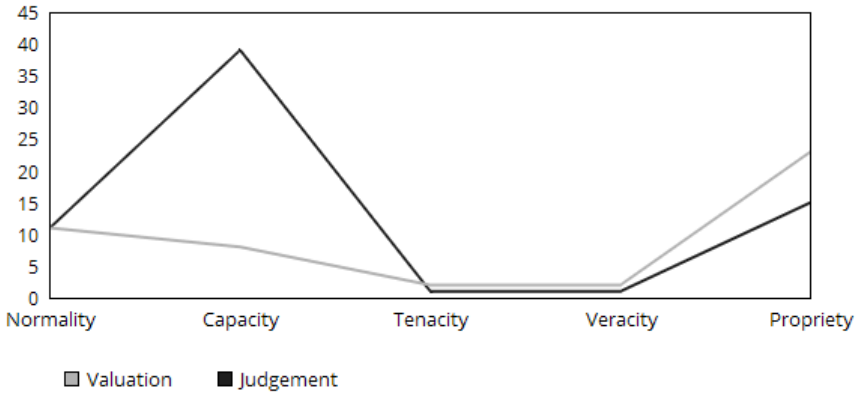

Frequency per 1000 words

In fact, the contours of each are nearly identical, with the notable exception of capacity. But this point of divergence is interesting in its own right: first, because this area of judgement is employed at a significantly different rate within the two corpora; and second, because Martin and White (2005) single out capacity specifically as being in direct relation with valuation. This disconnect between evaluations of capacity directed at other people (via judgement) and those directed at objects (via appreciation) will be explored in more detail in sections 6.2.1 and 6.2.2 below. Otherwise, the meanings

within valuation do indeed seem to echo and reinforce the kinds of evaluations uncovered in judgement, evidence that the cohesive metafunction of field described by Halliday and Matthiessen (2004) may indeed be in operation via this variable.

Figure 6.8 below shows the polarity of judgement as a whole side-by-side with the polarity of valuation. The contours reflect each other in their general use of positive meanings, but diverge sharply in their negativity.

Figure 6.8: Judgement Polarity by Realization Category

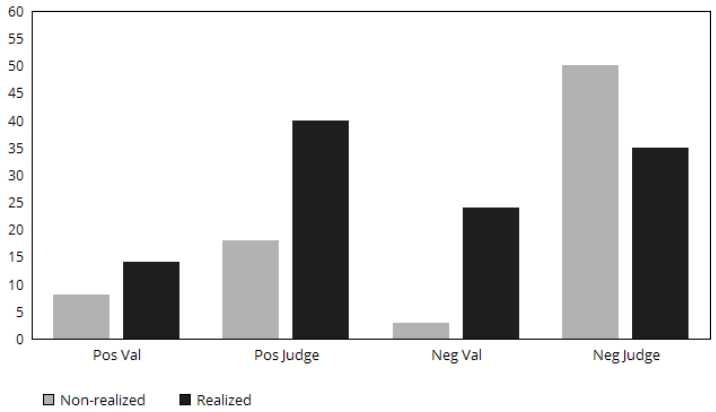

Frequency per 1000 words

However, while reflections between valuation and judgement are readily apparent (capacity excepted), valuation is not used merely as judgement's mirror. For instance, Figure 6.9 shows a bird's-eye view of the polarity of valuation in both corpora.

Figure 6.9: Valuation Polarity by Realization Category

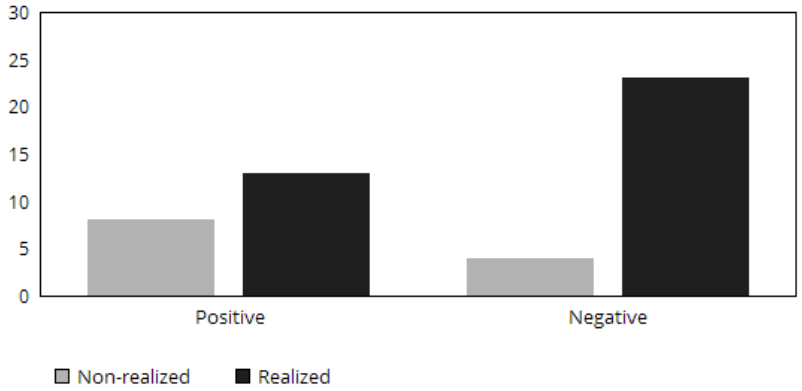

Frequency per 1000 words

In broad strokes, the R texts not only contain more of this kind of meaning than the NR writings, but this meaning is almost twice as likely to be somehow negative. The NR texts draw on this variable far less than their R counterparts and are more likely to use valuation in a (superficially) positive sense.

However, Hribal R's pledge is flagged as an outlier in its use of +valuation, and Brahm NR is similarly flagged for -valuation. With these two texts removed, the pledge types are brought almost to a parity in their use of +valuation, but the heightened use of -valuation in the R pledges is further exaggerated, as adjusted Figure 6.10 shows.

Figure 6.10: Valuation Polarity by Realization Category (Adjusted)

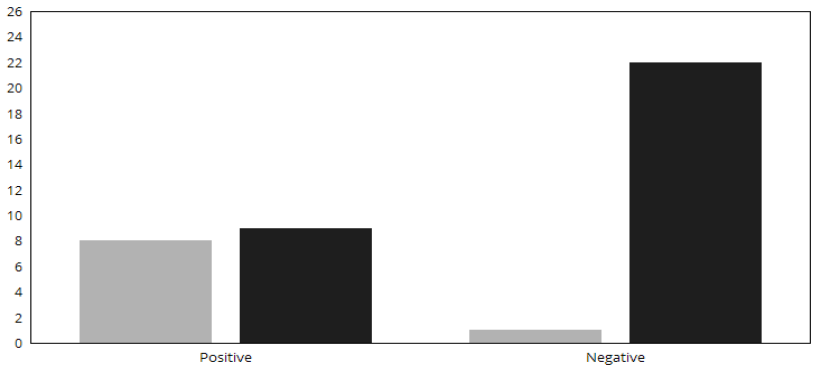

Frequency per 1000 words

All that said, valuation is its own domain, and cannot be assumed simply to conform to judgement at a broader level like polarity. Still, in certain key areas, the resources of valuation are employed by authors of each realization type in ways which complement the patterns uncovered in judgement, especially with respect to propriety, a connection which is discussed below.

6.2.1 POSITIVE VALUATION

As may be expected from a genre devoted to imagined violence, there are far fewer positively charged tokens of valuation in the dataset. Just half of the non-realized authors—four out of eight—employ this resource, at a frequency of 8 tokens per 1000.

The appearance of +valuation is slightly higher in the R writings. Positive valuation is found within six of the eight texts at a usage rate of 13 tokens per 1000. Interestingly, if somewhat unsurprisingly, 14 of the 32 total tokens of +valuation in the realized writings are found in Hribal R's pledge, by far the longest of the texts in the dataset. When the meanings of +valuation are analyzed according to the five variables of judgement, two pronounced trends are immediately apparent, as Figure 6.11 shows below. (For the moment, Hribal R will remain included, in order to get a full picture of the dataset.)

Figure 6.11: Judgement Meanings via +Valuation by Realization Category

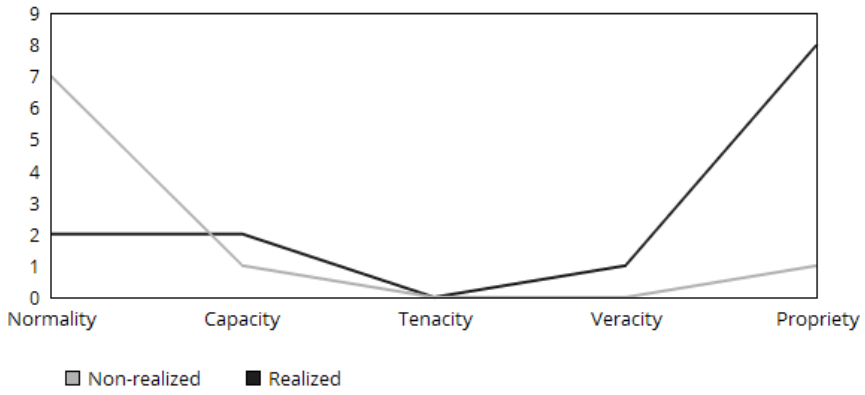

Frequency per 1000 words

Evaluations of objects and ideas which draw on resources similar to capacity, tenacity, and veracity occur at essentially the same rate in both corpora. Neither tenacity nor veracity were used at statistically different rates when these were examined within judgement directly. However, capacity appeared far more in the NR texts, and behaved as an expressive vehicle for much of the violent ideation found in those writings. Yet, tokens of +valuation which align with capacity occur with virtually the same frequency in both realization types here.

Instead of capacity, valuation meanings in the NR texts are directed at issues of normality ('how special?'). For instance, LA Unified NR writes that *Something big is going down*, while Archangel Michael NR says that bombs have been placed *ACROSS*

OFFICIAL BUILDINGS IN CHEYENNE. Essentially, these tokens and others are used by the non-realized writers to highlight how notable or exceptional some aspect of the imagined violent event will be. Indeed, all ten tokens of normality via +valuation found in the NR corpus function this way, serving to increase the overall tone of menace in the pledges. Archangel Michael NR, for example, plans to use not just *C4* as an explosive, but *U.S. ARMED FORCES STANDARD C4*, citing the substance's military origins as a way to reinforce the impression of its destructive power. Similarly, Valle NR imagines not just that he will abduct his victim, but that she will be abducted *out of her home*, a locative with connotations of safety and sanctuary. None of the six tokens in the R texts is employed this way. The closest is Roof R, who says he has chosen Charleston as the site of his attack because it is the *most historic city in my state*. This is arguably a justification of imagined action rather than a means of amplifying the menace of the threat. For Roof R, Charleston is *historic* due to its demographics. The number of African Americans living in the city means, for him, that it is where the *fight* must take place. The modifier *historic* is practical rather than theatrical: in order to kill his imagined victims, he must go to where they may be found. This trend—that similar stance resources justify violent ideation in the R pledges but increase the sense of danger in the NR pledges—was noted initially in the discussion of judgement. This distinction in discursive strategies thus appears to hold true in +valuation as well.

Other than normality, the other area of clear difference between the corpora is the heightened interest in ethical meanings by the realized writers. Of the 20 raw tokens of +valuation which correspond somehow with propriety, 14 are found in the Hribal R text. With all of Hribal R excluded, the frequency is diminished, but interestingly, the range of meanings is not. The adjusted counts are shown in Figure 6.12.

Figure 6.12: +Valuation Meanings by Realization Category (Adjusted)

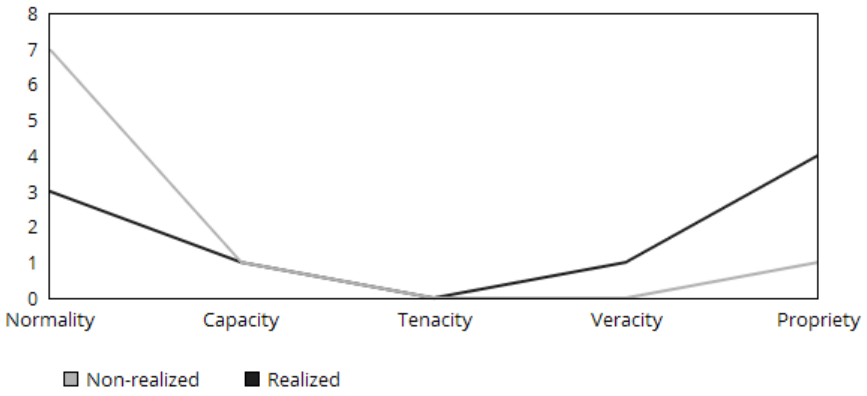

Frequency per 1000 words

Notably, every token of +valuation which denotes or connotes an ethical meaning is genuinely positive in the realized corpus. Long R offers—sincerely, by all measures—his *condolences to the good cops*. Roof R, for his part, states his intention to use *my life…for the good of society*. In fact, the uniform positivity across the realized writings is blunted only slightly when Hribal R's tokens are re-included. Hribal R consistently focuses on *good things*, and *the world of heaven*, allowing back-handed praise just twice: when he declares *the opinion of helping fellow human beings* to be *hypocritical*; and when he uses scare quotes to indicate that the application of the superlative *best* to his high school is intended ironically. By contrast, the NR texts contain just a single token of propriety via +valuation, when LA Unified NR writes that people will die *in the name of Allah*. As with the meanings of normality discussed immediately above, the *name of Allah* arguably draws on international religio-political tensions to increase this text's sense of terror to an American audience, and so may not be judged to be used sincerely in the same way as the Long R and Roof R tokens.

Both the increased awareness of ethical issues in the R pledges and their consistent admission that good exists make for a more nuanced moral attitude than simply declaring the world evil, a result which is fully in keeping with the analysis of propriety within the subsystem of judgement.

6.2.2 NEGATIVE VALUATION

Compared to +valuation, there is a far greater disparity between the corpora in how the meanings of -valuation are used. For instance, of the eight non-realized authors, tokens of -valuation appear in the writings of only three: Brahm NR, Rodger NR, and Valle NR. Indeed, the use of this variable is so infrequent in the NR corpus that Brahm NR is flagged as an outlier despite having a mere 4 tokens. On the other hand, all six realized authors make use of this resource, again in ways which echo—if not directly mirror—the findings in the domains of judgement. Before such echoes are examined, however, it is worth noting that the continued absence of capacity (as seen in Figure 6.13 below) as a distinguishing feature of the non-realized writings resists easy explanation.

Figure 6.13: -Valuation Meanings by Realization Category

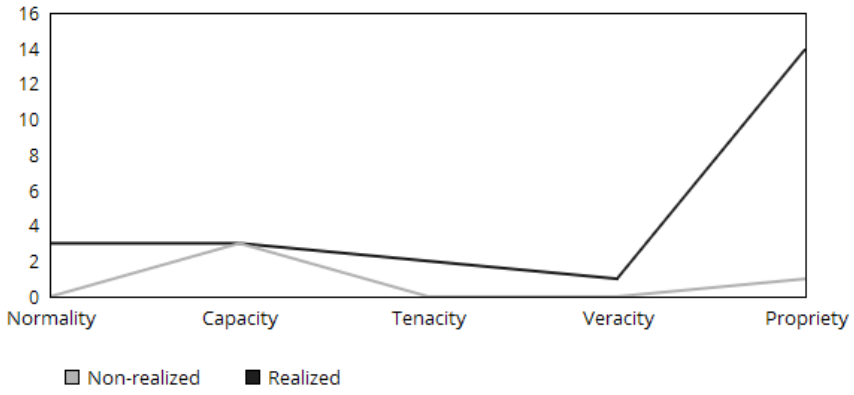

Frequency per 1000 words

One tentative hypothesis is that the more frequent mention of human targets by non-realized authors situates this violence squarely in the realm of judgement rather than appreciation. So, a realized author such as Hribal R discusses his proposed knife attack as *a violent action*, momentarily abstracting the event itself away from the people necessarily undergoing the attack. The trend for the non-realized texts is instead to tie mentions of violence directly to the humans being threatened. Thus, when a non-realized author mentions 'death,' it is far more likely to be in direct relation to a human

target, as when Dickens NR writes *Death to all white cops*. But the wider, prosodic effects of this tendency—if this interpretation is accurate—are difficult to pin down. Does this lend additional menace to the non-realized texts, in that violence is tied much more clearly to the imagined victims? And does this add a clinical touch to the realized writings, in that events which are inherently bloody are discussed with less direct reference to who will be bleeding?

The question of -capacity and its lack aside, tokens of -valuation which somehow speak to ethical concerns emerge once again as a hallmark of the R pledges, appearing at 13 times the rate versus the NR pledges. Further, while all six realized authors employ some variety of -valuation, only those tokens which appear in Roof R's text lack an obvious ethical component. -Propriety via -valuation is thus quite common across the corpus.

Interestingly, the range of meanings captured by this variable generally breaks across a fault line of 'moral' versus 'legal,' although moral meanings are more prevalent (as they are in judgemental propriety). Hribal R, for example, talks repeatedly of the *world* and of *life* as being *evil*; Long R judges his assault to be a *necessary evil*; Rodger, in both his R and NR pledges, refers to his *Day of Retribution*; etc. However, while tokens of morality appear in both corpora and in both positive and negative polarities, tokens of *legality* appear only in the R corpus and only within the frame of -valuation. And while this meaning is rarer across the dataset than moral concepts like 'evil' and 'wrong,' it nevertheless appears in the writings of three of the six realized authors. Returning to Hribal R, he characterizes his own impending attack as a *crime*. Kinkel R admits that the murder of his parents, which he committed before writing his note and before his murder of a classmate the next day, has earned him *two felonies*. And Long R discusses the behavior of America's police officers—his targets of choice—using terms like *illegal* and *unjust*.

The inclusion of these kinds of meanings is potentially indicative of a point of connection between the realized authors and what is called “consensus reality” (Gellerman & Suddath, 2005: 485), i.e., an apprehension of the world more or less as others agree it to be. The inclusion of legal meanings could indicate that R authors see the world as functioning not just according to what may be broadly described as Biblical values, e.g., good versus evil. They also see their own and other’s actions in terms of the more terrestrial rules codified by the state. The NR authors, when they engage with ethical meanings at all under valuation, recognize only questions of virtue and its lack, and nothing of the law.

6.2.3 VALUATION AND TENSE

Finally, this variable offers an additional way of assessing both the motivations and goals of the realized authors by examining *when* these ethical meanings are relevant. Was the world right and good at some point but now faces a moral crisis which the author finds intolerable? Will the world be a better place following the act of violence meant to set it right? Etc. Only two tokens of propriety via valuation appear in the NR corpus: *Me, and my 32 comrades, will die tomorrow in the name of Allah* (+prop via +val, LA Unified NR); *On the day before the Day of Retribution, I will start the First Phase of my vengeance* (-prop via -val, Rodger NR). Both are moral and both are in the future tense. The R pledges are far richer with these meanings, and so the brief discussion which follows will focus on this corpus alone.

As with capacity and propriety, when propriety appears in the R texts via +/- valuation, it may be broken along basic verb tense lines—past, present, and future—as seen in Figure 6.14 below. What emerges is intriguing.

Figure 6.14: +/-Propriety via Valuation Sorted by Tense in Realized Texts

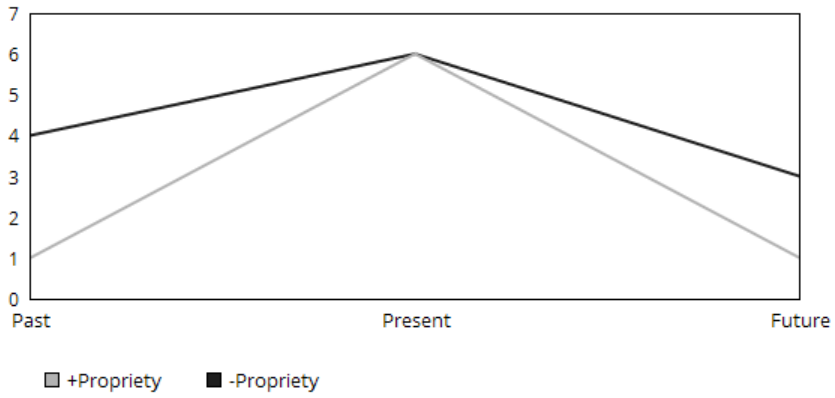

Frequency per 1000 words

That negative meanings continue to predominate in texts devoted to violent ideation is, by now, unsurprising. However, such predominance is less of a given when a corpus is compared with itself. Intra-corporally, there is no reason to assume the kind of pessimism captured in Figure 6.14. Despite the fact that fewer tokens of +propriety via valuation than negative appear in the realized corpus as a whole, tokens which are governed by verbs in the present tense are nevertheless split equally along negative and positive polarities. Interestingly, the R authors are also almost as cynical about the future as the past.

However, when shared world knowledge is applied to some of these tokens, and the discursive function of others is taken into consideration, a slightly different picture develops. For the former, several authors place their own violence in the past or the present rather than a time to come, despite not having committed the acts yet, e.g.:

- *what I did was <u>evil</u>* (Hribal R);
- *I know most of you who personally know me are in disbelief to hear that I am suspected of committing such <u>horrendous</u> acts of violence* (Long R).

The fact that these texts were authored *before* the attacks they describe argues that these tokens are best grouped with other instances of future impropriety. (The fact

that they are discussed via the past tense is extremely interesting in its own right, something which is examined in section 7.5.)

For the latter, discourse-based re-categorizations across polarities are possible as well, although such reversals are entirely the province of Hribal R, whose rhetorical style is often sardonic. For example, he is being ironic when claims he is a student at *one of the "best schools in Pennsylvania"*, and consistently contrasts *the world of heaven, where all is good*, with this world, plainly implying that this world is neither heavenly nor good. These three lexemes (*best, heaven, good*), therefore, are better understood as negatively charged comments on aspects of Hribal R's social experiences.

Figure 6.15 shows the adjusted tense categories when these discursive functions are taken into account.

Figure 6.15: +/-Propriety via Valuation by Tense and Semantics in Realized Texts

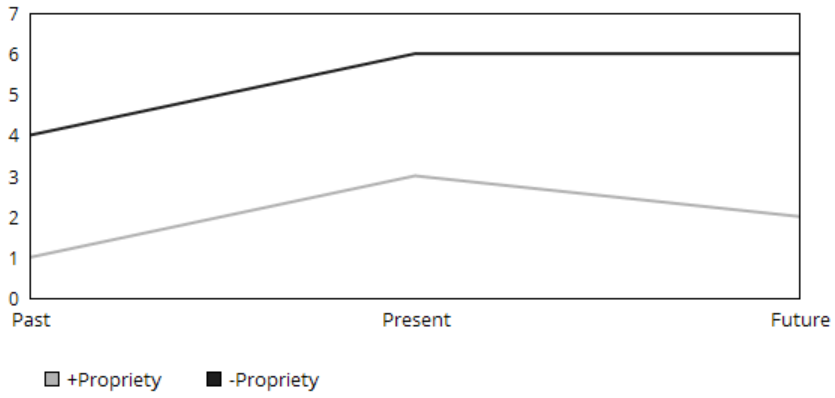

Frequency per 1000 words

Realized authors appreciate the past as a relatively unwelcoming place, marked by *unjust practices* (Long R), and requiring *malice* and *cruelty* if one is to thrive (Hribal R). Interestingly, the present and future are viewed even more cynically. While Long R sincerely offers his *condolences* (which, even here, is a potential hybrid realization, since the comfort of condolences are only given in times of mourning), and Roof R talks

earnestly of *the good of society*, ethical evaluations of the world as it is are nevertheless overwhelmed by references to *racism* (Shaw R), *felonies* (Kinkel R), *selfishness* (Hribal R), and *illegality* (Long R), among other things.

Perhaps the most surprising dynamic uncovered here is, again, the attitude that the future offers no relief. Ethical conditions are not portrayed as improving even after these authors have performed the violence they envision. Long R refers positively to his killing of police as a personal *sacrifice* (again, a potential hybrid) to his beloved community, a frame which indicates that what was offered is less valuable than what will be received in return. Yet, while he orders the police to *quit committing criminal acts against melanated people*, he never predicts that his sacrifice will result in police officers acceding to his wish.

Elsewhere, the future is a place where these authors understand they will be vilified for their planned actions: Hribal R knows his attack will be called a *"monstrosity"*, and Long R anticipates that the media reports will accuse him of *horrendous acts of violence*. Lastly, neither Shaw R nor Rodger R anticipates the triumph of what they may think of as good. Shaw R terms his attacks a *battle* but nowhere mentions winning this battle, and Rodger R calls his planned event the *Day of Retribution*, rather than anything having to do with justice. This cynicism is starker knowing every one of these authors would somehow attempt or successfully carry out the violent acts at the heart of their texts. And yet, none of them attests that this violence—which would risk the authors' personal freedom and even their lives—will change the world for the better. At worst, the results of the authors' actions may be characterized as futile, but the intolerable impropriety of the present appears to drive them on regardless.

6.2.4 SUMMARY OF VALUATION

The final variable in the subsystem of appreciation, called valuation, captures an author's estimation of an object's worth. Five of the eight non-realized authors employ this resource, at a frequency of 11 tokens per 1000. By contrast, valuation is found in all six of the realized writings, at a usage rate of 35 tokens per 1000—more than three times higher than the non-realized.

Qualitatively, Martin and White's (2005) assertion that valuation is sensitive to field seems borne out by the analysis, in that discourse patterns identified elsewhere in these texts are echoed and reinforced here—with the major exception of meanings related to capacity. Secondly, the suggested correspondence between valuation and judgement also appears valid, as tokens coded within valuation break down quite naturally into judgement's five semantic categories of *normality* (how special?), *capacity* (how capable?), *tenacity* (how dependable?), *veracity* (how honest?), and *propriety* (how far beyond reproach?). Analyzing this area of appreciation according to the meanings of judgement reveals where field is operating as a cohering force across the texts.

Evaluations of objects and ideas which draw on resources similar to tenacity and veracity occur at essentially the same rate in both corpora. This result mirrors the pattern uncovered in judgement, that both the R and NR authors are equally un/concerned with truth and dependability, no matter if it applies to people or things (objects, ideas, etc.). More crucially, ethical concerns emerge once again as a special theme of the R pledges. Tokens of -propriety via -valuation are 13 times more likely to appear in a realized pledge, and five of the six R texts employ -valuation to these ends. Only those tokens which appear in Roof R's text lack some kind of ethical component.

When this variable is viewed as prosodically interlinking with propriety, however, analysis reveals that a heightened concern for ethics is universal across the R corpus. Within the larger system of attitude, this is the only variable whose usage behaves so

consistently intra-corporally and so differently cross-corporally. The result, rhetorically, is a more nuanced worldview on the part of the realized authors, one where good *can* exist but often not in society as the authors currently find it. An example of this nuance is the realized texts' inclusion of ethical meanings which further divide along lines of morality and legality. The appearance of tokens with judicial denotations argues that the realized writers are not only focused on the moral state of the world around them, but that they are also aware of the real-world mechanisms by which this moral state is altered or preserved. This is, arguably, evidence that the violent ideations of the realized authors are in fact more connected to consensus reality than their non-realized counterparts.

However, contra field's cohering force, the NR texts evince a relative disinterest in meanings which correspond to capacity. In judgement, the heightened use of capacity by NR pledges to communicate imagined violence is a major way in which the two text types differ. In valuation, this difference is essentially erased. Again, there is no easy answer as to why this may be so, though one hypothesis is that the NR pledges' more frequent mention of human targets naturally situates this violence within judgement rather than appreciation.

Capacity aside, the non-realized texts are markedly more concerned with issues of normality via +valuation. This resource is mainly employed to answer questions of how notable or special some aspect of the imagined violent event will be, whether that aspect is the attention the author's actions will receive (e.g., *It will make national headlines*), the nature of the author's target (*OFFICIAL GOVERNMENT BUILDINGS IN WYOMING*), or the means by which the violence will be achieved (*COMPACT U.S. ARMED FORCES STANDARD C4; kidnap her right out of her home*).

One last point is uncovered by valuation in the R corpus. This is the attitude that the future will not be more just or moral following the authors' imagined actions, a finding

first noted in the discussion of the null results of dis/satisfaction. Importantly, the apparent futility of the imagined violence did not deter them from acting.

Chapter 6 closes the investigation of the system of attitude. Chapter 7 thus moves from this into the next system of Appraisal, that of engagement.

CHAPTER 7 ENGAGEMENT

The system of engagement seeks to uncover the "the linguistic resources by which speakers/writers adopt a stance towards to the value positions being referenced by the text" as well as how such resources are employed "with respect to those they address" (Martin & White, 2005: 92). The question of addressee is especially crucial, for reasons theorized by SAT (section 1.2.1) and ADT (section 1.2.2). The main difference between a pledge to harm and the kind of threatening language studied by Shuy (1993), Fraser (1998), Gales (2010, 2011), and others, is the pledge's intended audience. Direct threats presume that "the recipient…will not be in alignment with the threatened act" (Gales, 2010: 231), communicating hostility towards the 2nd Person of the text's addressee. A pledge to harm instead targets a grammatical 3rd Person. Thus, the threatened act in a pledge is not necessarily "to the hearer's detriment" (Shuy, 1993: 98)—arguably the most distinctive felicity condition of threatening as a speech act. And because the broader presumption of disalignment between addressor and addressee is not comparable between a direct threat and a pledge to harm, other conclusions about how threatening language operates must also be examined anew in the context of pledging. This is especially true of conclusions about how a pledge writer construes the text's intended audience and what this may say about the writer's underlying intent.

The system of engagement is expressly built to measure these dynamics. Exposing the double motion of an author's stancetaking towards both the content of a text as well as the text's 'imagined' or 'ideal' reader (Martin & White, 2005: 95) can shed light on other attitudinal patterns. Certain questions arise from the fact that pledge authors do not face an inherently combative audience. For example, any points in a pledge text where the writer anticipates disagreement, or feels the need to argue or to convince, are worth special interrogation. Do such moments comport with Smith's

(2006) finding, for example, that strategies of persuasion correlate significantly with real-world approach behaviors? Similarly, allowing outside voices into a threat is characterized as a weakening of authorial stance (Gales, 2010, 2011; Meyer, 1997), but might it also be evidence of higher conceptual complexity (Smith, 2006) on the part of the author? Conversely, points in a pledge where agreement is taken for granted are also of interest. What, exactly, is the audience being positioned to accept? Thus, both presumptive agreement *and* presumptive disagreement invite a closer look at the value position being advanced.

Generally, the system of engagement provides a means of assessing *patterns* of rhetorical moves which might otherwise be invisible to a regular reading. Engagement does so via a taxonomy of dialogistic positions aimed at "describing what is at stake when one meaning rather than another is employed" (Martin & White, 2005: 97). This system thereby allows for a greater understanding not just of the raw mechanics of stancetaking but also of the significance of the authorial choices responsible for the creation of the text itself. Such patterns may offer insights into how commitment is expressed to a pledge's imagined audience.

However, analyzing engagement requires something different than the primarily lexical semantic approach used in the previous three chapters. For its part, attitude captures the lexical 'road signs' which authors place at strategic moments in their discourse. These particles of stancetaking build on each other to create a prosody of interpersonal meaning. But the 'current' bearing these particles along is the larger syntactic structures which indicate how each utterance is meant to be understood, e.g., is the value position "presented as one which can be taken for granted for this particular audience, as one which is in some way novel, problematic or contentious, or as one which is likely to be questioned, resisted or rejected"? (Martin & White, 2005: 93).

Engagement values may indeed be marked lexically, like with attitudinal tokens. The conjunction ‘but’ may perform the rhetorical action of countering, negations like ‘no’ or ‘not’ may signal denials, etc. More often, though, engagement options are indicated syntactically. There is no overt lexical marker of engagement, for instance, in Hribal R’s pronouncement *I watch anime*. Instead, the dialogically contractive work of pronouncing is performed by the declarative mood of the clause itself. The larger implication of viewing engagement options syntactically versus lexically is that *some form of engagement is in operation at all times across the entire act of communication*, whether or not the rhetorical purpose is overtly flagged by standalone lexemes. As Stubbs (1996: 197) simply says, “whenever speakers (or writers) say anything, they encode their point of view towards it.” The most sensible approach to understanding engagement meanings is therefore not strictly lexical, but rather *clausal*.

In practice, this approach affects coding, and thus the statistical analysis that proceeds from it. The universality of engagement means that the kind of normed word count comparisons used to analyze judgement and appreciation are inapplicable here. The difference is, perhaps, best illustrated by two different instances of ‘proclaiming’ found in the dataset, an engagement option which represents a proposition as “highly warrantable (compelling, valid, plausible, well-founded, generally agreed, reliable, etc.)” (Martin & White, 2005: 98). Proclaiming may be accomplished using thirteen words, as Rodger NR does when he says *I will torture some of the good looking people before I kill them*, or using just two, as LA Unified NR does when s/he writes, simply, *No more*. Instead of being blinded by each statement’s respective word count, which would merely muddy the waters, the most sensible approach to coding engagement is to view each of these utterances as <u>one</u> token of proclamation. This, in turn, makes quantitative inter-corporal comparisons possible.

Taking the clause as the fundamental unit[29] of engagement, the non-realized corpus contains a raw count of 175 tokens, while the realized corpus contains 279. This leads to a final methodological note: for the purposes of norming ahead of a chi-square test, it became quickly apparent that calculating these token frequencies to the more typical baselines of either 100 or 500 distorted the data in suspect directions. Therefore, the admittedly unorthodox base of 300 has been chosen for norming the engagement tokens in this dataset, to better preserve the patterns as they occur.

Of the sixteen possible coding categories at each level of engagement's hierarchy, a statistically significant difference is found in five types of locutions. Table 7.1 shows these in order of p-value, with below significant findings shaded.

Table 7.1: Statistical Significance in Engagement

Engagement Type	Tokens per 300 instances		Probability (p)
	Non-realized	Realized	
Deny	10	44	< .001
Disclaim	29	76	< .001
Attribute	12	30	< .01
Proclaim	211	154	< .01
Acknowledge	9	22	< .05
Concur	62	30	> .05
Contract	240	230	> .05
Counter	19	32	> .05
Distance	3	9	> .05
Endorse	0	0	> .05
Entertain	38	28	> .05
Expand	50	58	> .05
Heteroglossic	290	288	> .05
Justify[30]	12	11	> .05
Monoglossic	10	12	> .05
Pronounce	137	113	> .05

[29] Smith (2006) notes that there is precedent for the clause as a unit of linguistic measure in the psychological literature as well: the Gottschalk-Gleser method "use[s] the grammatical clause as the unit of analysis, rather than single words" (Smith, 2006: 50).
[30] This category of meaning does not appear in Martin and White's (2005) original taxonomy but is included in O'Donnell's (2013) via UAM Corpus Tool. Needless to say, for analyzing texts whose potential purposes include the justification of violence, this addition has been welcomed here.

Like the other Appraisal systems, engagement is arranged as a pyramidal flow-chart of choices—but one with deeper embedding than the relatively superficial layering in the subsystems of judgement or appreciation. Because of this, daughter and mother variables in the system may both show significance. So, for instance, Table 7.1 shows that both deny and disclaim are of statistical interest, but because denial is a subtype of disclamation, the rhetorical action found in the one inherits to meanings found in the other. In the two instances where this is the case—attribute and acknowledge; disclaim and deny—the two engagement options are analyzed together, to better allow for these natural interdependencies of meaning to emerge.

The sections of analysis are arranged as follows. First, the broad contours of dialogic contraction versus expansion in the dataset are addressed. Next, the more significant, but more limited findings of dialogic expansion (attribute and its component variable, acknowledge) are analyzed, followed by the more prevalent instances of dialogic contraction (proclamation, disclamation, and denial). First, however, a brief note is warranted about the null result of justification.

7.1 JUSTIFICATION

Of the several engagement options which are not statistically significant, this is perhaps the most interesting from a threat assessment perspective. In the literature, motive is considered one of the key indications of the risk a threatener poses (Borum *et al.*, 1999; Calhoun & Weston, 2015). Justification would seem to be a prime resource for a threatener to express the motive for an imagined attack. Indeed, people surveyed about threatening language identified “the justification for the threat” as one of the few language functions they would expect to find in a threatening communication (Gales,

2010: 96)[31]. The lack of quantitative difference here between the two realization categories is therefore notable. If, in fact, the R pledges do encode commitment, then one might expect that justifications for the author's actions would loom larger than they would for their NR counterparts. Yet, this is not the case.

At least, this is not the case *numerically*. A qualitative analysis reveals a functional difference, however—one that is interesting enough to warrant a brief comment despite the lack of statistical significance in the area. Each of the five realized authors who employ this rhetorical resource do so in conjunction with a 1st Person pronoun either 1) to defend the choice of target or 2) to justify the attack itself.

Justifying the target:

- *Therefore **I** must bring the same destruction that bad cops continue to inflict upon my people, upon bad cops as well as good cops in hopes that the good cops (which are the majority) will be able to stand together to enact justice and punishment against bad cops b/c right now the police force & current judicial system is not doing so.* (Long R)
- *the kind of girls **I've** always desired but was never able to have because they all look down on me* (Rodger R)
- ***I** chose Charleston because it is most historic city in my state* (Roof R)
- *The reason is because **I** don't think Asian Women like me* (Shaw R)

Justifying the attack:

- ***I** do this because the world is bitter and painful.* (Hribal R)

[31] As of this writing, Gales's (2010, 2011) work is the only available literature which applies engagement to the genre of threatening language. Therefore, many of the results here will be framed, perhaps unavoidably, as either echoing or departing from her findings.

This personalized usage appears only once in the NR corpus, when Skyline NR says of the jocks that ***i*** *am going to start killing them first* <u>*because they deserve it the most*</u>. In the case of the R pledges, then, justification is used much more to perform the function which the respondents in Gales' (2010) survey expected it would, that of setting forth a personal motive for acting. Thus, despite the quantitative similarity between the corpora, a clear qualitative difference in usage is apparent.

7.2 CONTRACTION VS EXPANSION

Dialogic contraction has been singled out by Gales (2010, 2011) as a resource with direct ties to the question of intent: "through the use of contracted, heteroglossic forms, a deeper awareness of the author's underlying intent and assumed level of commitment can be gained" (Gales, 2011: 37). So, for instance, when Rodger NR writes *I will behead them*, his apparently high commitment to the proposition—that he will personally decapitate some group of people in the future—is communicated by the statement's exclusion of possible alternatives.

However, Gales's (2010) groundbreaking contribution to this area of study is that threateners regularly take stances which both strengthen *and* weaken their apparent commitment, often in ways which run counter to folk expectations of what kind of language 'should' appear in realized versus non-realized threats. (For instance, while Elliot Rodger did advance to assault and murder, he never beheaded anyone.) Figure 7.1 below may capture something like the dynamic uncovered by Gales (2010) in threatening more broadly. Mostly, the figure shows that the dataset is highly contractive, a finding that is, perhaps, unsurprising for a hostile genre like pledging. Yet, these texts are not *entirely* closed off to outside voices and possibilities. The fact that a genre so loaded with rage and imagined violence makes any room at all for alternate viewpoints is, perhaps, more surprising.

Figure 7.1: Dialogic Contraction and Expansion

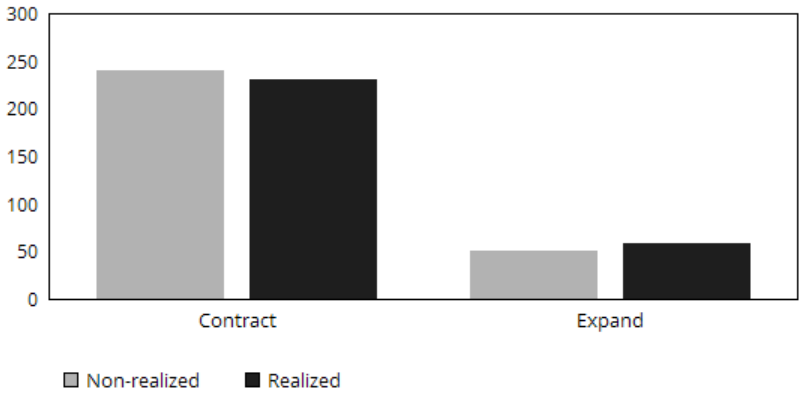

Frequency per 300 clauses

The broad contours of the two corpora are identical. As the following sections will demonstrate, though, there are notable functional differences in how particular engagement options are employed by each author type as they contract and expand their pledges in relation to outside voices, both in anticipation of them and in response to them.

7.3 DIALOGIC EXPANSION: ATTRIBUTE AND ACKNOWLEDGE

Through attribution, propositions are "disassociate[d]…from the text's internal authorial voice by attributing it to some external source" (Martin & White, 2005: 111). This area is broken into two sub-categories: acknowledge and distance. Through distancing—which shows no significant difference in usage between the corpora—the authorial voice "explicitly decline[s] to take responsibility for the proposition" (Martin & White, 2005: 114). This can imply skepticism of claims made by others, such as when Rodger R writes of the sorority members, *They think they are superior to me*. Acknowledgement, by contrast, is more neutral. No overt indication is provided "as to where the authorial voice stands in respect to the proposition" (Martin & White, 2005: 112). Such formulations appear, e.g., when Brahm NR says that *Al-Qaida will be blamed for the*

attacks, and when Long R writes *You are thinking to yourself that this is completely out of character of the MAN you knew.*

Figure 7.2: Acknowledge and Distance

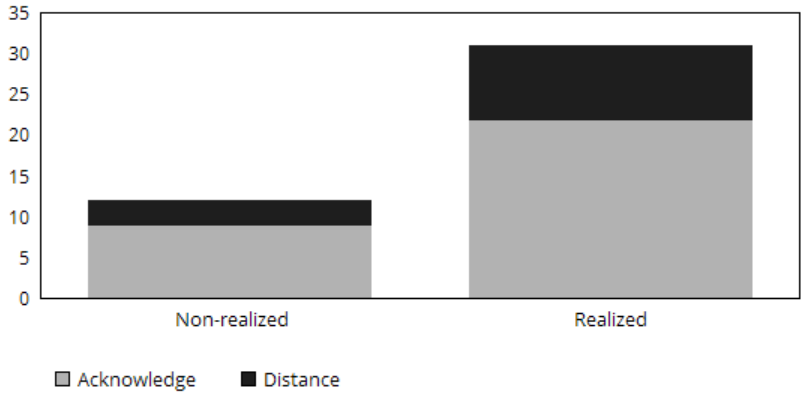

Frequency per 300 clauses

As Figure 7.2 above shows, outside voices are not primarily treated with skepticism or suspicion—the work performed by 'distancing'. Rather, they are handled more neutrally by both pledge types. Indeed, the bulk of attribution present in the dataset by far is acknowledgement, although realized authors are more likely to make use of this option, at a rate of nearly 2.5:1 (22 tokens per 300 versus 9 in the NR corpus). The dispersion of attribution across the dataset is shown below in Figure 7.3.

Figure 7.3: Attribute Dispersion by Corpus

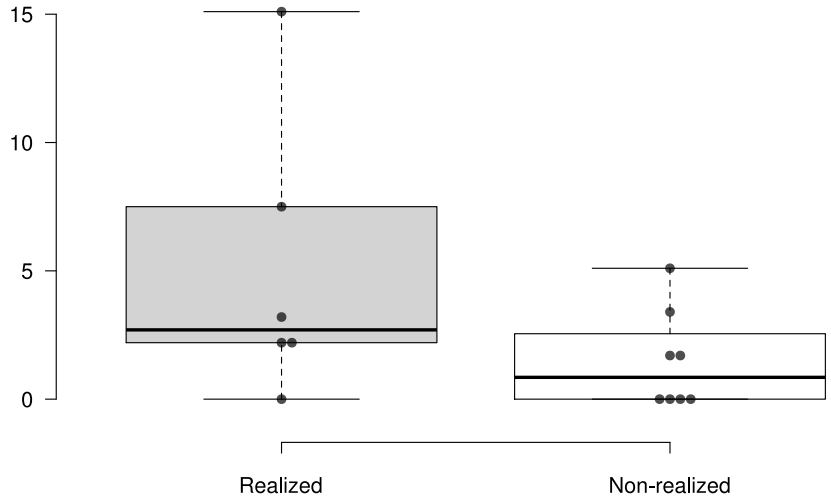

Frequency per 300 clauses. No outliers (>1.5IQR) are found.

Turning to acknowledgement, this is highlighted by Martin and White (2005: 113) as an anticipatory operation, one which "[makes] space in the ongoing dialog for those who might hold alternative views". The dispersion of acknowledgement across the dataset is shown below in Figure 7.4.

Figure 7.4: Acknowledge Dispersion by Corpus

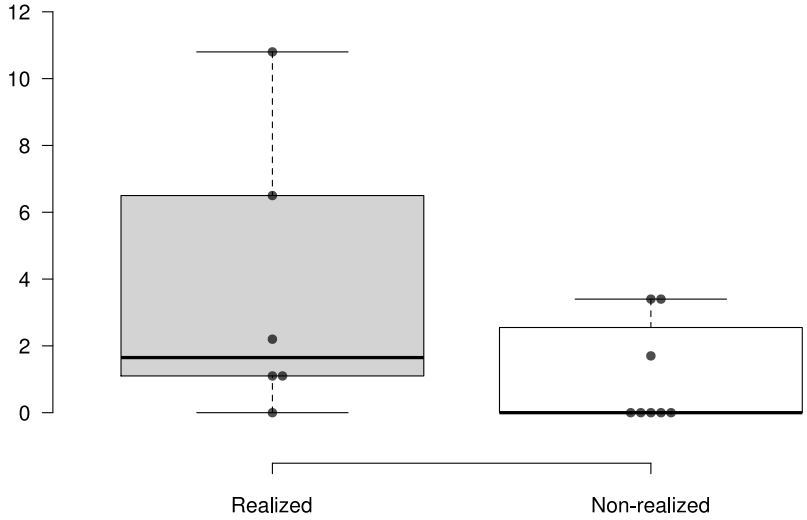

Frequency per 300 clauses. No outliers (>1.5IQR) are found.

That such space is more likely to appear in the pledges of people who proceeded to real-world action invites a further series of questions. First, whose voices are pledge authors allowing into their texts? And by opting for the more neutral rhetorical strategy of acknowledgement, are pledge authors anticipating that the cited propositions are somehow problematic for their intended audience? I.e., are the realized authors especially "making space in the ongoing dialog for those who might hold alternative views" (Martin & White, 2005: 113) in the interests of, say, minimizing face threats and maximizing writer-reader solidarity? Or are the propositions presented stance-neutrally because they require no further attitudinal signaling, i.e., they are *not* considered problematic? These questions are explored from several angles in the subsections which follow.

7.3.1 WHOSE VOICE IS ADMITTED?

Engagement's sensitivity to audience means that a pledge's intended readership is often detectable in the pledge itself. The voice of the reader being pulled into the discourse occurs, for example, in Dickens NR's text, when she challenges her audience's knowledge of the U.S. Constitution, writing *So when you can absolutely show me in the 1st amendment where it explicitly says you can't say "kill all cops", then I'll delete my status.* Similarly, Long R steps into the voice of his readers when he says *I know most of you who personally know me are in disbelief*.

The bifurcation in a pledge to harm of audience and target means that the people whom the threateners imagine harming are a second voice worth listening for. For example, Rodger R co-opts the voice of the women he would later attempt to kill when he claims *they all look down on me*.

Figure 7.5: Attributed Voices

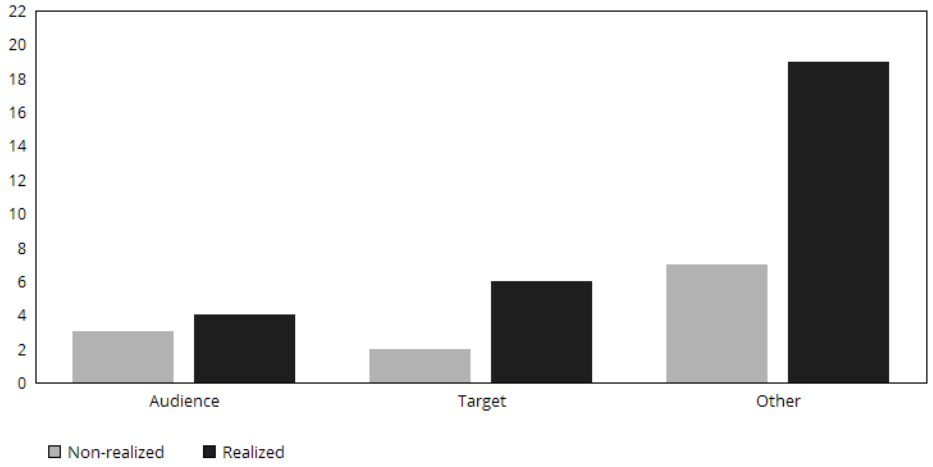

Frequency per 300 clauses

The analysis will begin with these two major categories: audience and target. For the moment, voices which fall outside these categories will be collected under the header of 'other'. This will itself be parsed more fully below. The frequency with which each voice type appears in the two corpora is shown in Figure 7.5 above.

Two things are immediately apparent. First, the realized texts engage more actively with all three categories than their non-realized counterparts. As noted previously, this kind of expansion "functions to weaken the [threatener's] stance by leaving room for other voices to vie for control" (Gales, 2010: 158). If the strength of an author's commitment to act correlates with the strength of their stance towards the propositions in their texts, then one might expect the realized texts to be more closed to outside voices—not more open—but this is clearly not the case. The realized texts are technically weaker in this regard despite discussing actions their authors are known to have later attempted.

Second, this expansion includes a higher level of engagement with the voices of the realized pledges' imagined victims. Rodger R, for instance, projects an opinion into the minds of the Alpha Phi sorority members he hopes to kill, saying *they all look down on me*, and Shaw R engages in the same process when he writes *I understand that Asian Women are afraid of African American Men*. Interestingly, the single instance of a non-realized author allowing in the voice of his or her imagined victims occurs in the Skyline NR pledge and is framed not as acknowledgement, but rather as distancing, an

Figure 7.6: Acknowledged Voices

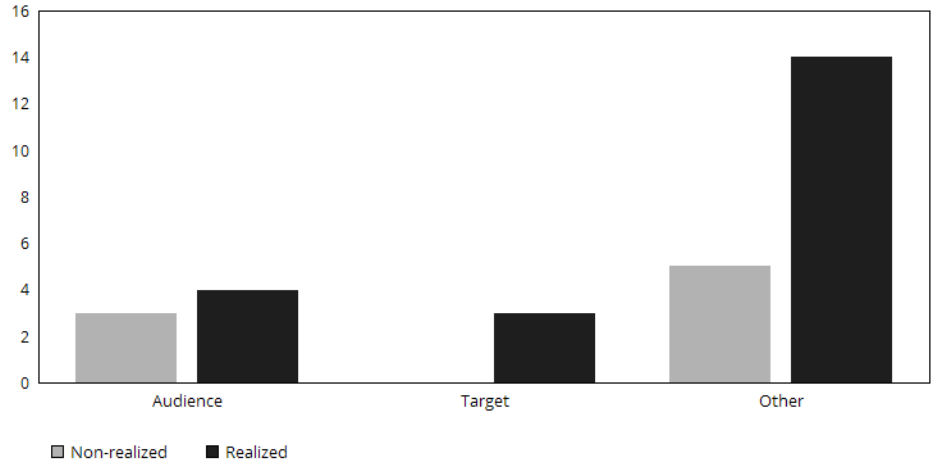

Frequency per 300 clauses

area which is not shown to be statistically significant: *they think it makes them better*. When only the statistically significant resource of acknowledgement is considered, victims' voices disappear from the NR pledges entirely, as shown in Figure 7.6 above. Each of the three areas (audience, target, other) will then be analyzed in turn.

Target: Examining how the victims' voices are included in these texts reveals two additional authorial tendencies. When viewed at the level of attribute, i.e., with tokens of distancing like Skyline NR's reintegrated, the inclusion of targets' voices often sits at a particular "projection nexus" (Halliday & Matthiessen, 2004: 443). Specifically, five of the seven occurrences in the dataset are realized at the level of 'ideas' and via the hypotactic mode of 'reporting.' In Halliday and Matthiessen's (2004: 443) formulation, ideas represent the content of a mental clause, i.e., what is thought by a person. These are contrasted with locutions, i.e., what is said. Reporting, meanwhile, is also held as distinct from what is called quotation. Grammatically, the difference is a matter of embedding—neither reported speech nor reported thought are able to stand alone from their projecting clause. The four quadrants of projection resources appear in Table 7.2 below.

Plainly, R pledges are given to describing—but not quoting—what their victims are thinking. Skyline NR is, again, the only text which engages with the target's voice in such a way. In other words, R pledges are particularly comfortable projecting content into the minds of the people they imagine harming but no pledges are comfortable enough to present these as the direct and literal thoughts of their victims.

The implications of such consistent choice-making are intriguing, especially in light of the second authorial tendency present here: while the victims are given space in these texts, in terms of polarity all seven tokens are meant to portray these people negatively. Further, the blanket negativity of these projections is used to bolster the

Table 7.2: Projection Nexuses of Target Voices

		Mode	
		paratactic ("direct, quoted")	hypotactic ("indirect, reported")
Level	idea ' mental		*they think it makes them better* (Skyline NR) *realize their previous lives are going to be taken by the only one among them that isn't a plebeian* (Hribel R) *they all look down on me* (Rodger R) *They think they are superior to me* (Rodger R) *I understand that Asian Women are afraid of African American Men* (Shaw R)
	locution " verbal	*Sorry I have a Boyfriend* (Shaw R) *Sorry I'm in a rush* (Shaw R)	

case that the imagined victims are somehow deserving of the anger they stir in the writer, e.g., because they are pretentious (Skyline NR, Hribal R, Rodger R) or because they are racist (Shaw R). Yet, no writer is willing to venture to the stronger level of evidentiality—that of direct quotation. The discursive effect of this reticence is subtle but potentially telling: these projected ideas have more of the character of the authors' own conclusions than of hard fact or observed behavior. The victims are imagined as thinking these things but, despite their incriminating quality, are not *known* to have thought them.

This disconnect between the presented evidence and a threatener's interpretation of the evidence is especially intriguing in the case of Shaw R, whose pledge is the only one which also employs the stronger, paratactic resource of quotation. Shaw R offers two excuses he presents as having been given verbatim by the Asian women he has spoken with: *Sorry I have a Boyfriend*; *Sorry I'm in a rush*. From the

standpoint of politeness theory (Brown & Levinson, 1987), both rejections are anodyne, each containing the positive face-threat minimizer *Sorry*. Moreover, neither sentiment could be credibly cited as evidence of these women being *afraid* of an African American man like himself, the emotion Shaw R projects onto them. And yet, both quotes are used to support exactly this conclusion. The direct evidence is misaligned with the idea Shaw R claims it represents.

Putting such fine parsing aside, in this regard the two corpora differ merely in quantity rather than quality. When the victims' ideas are reported, they are presented in a way intended to somehow condemn them as people. Such reporting is simply more prevalent in the R pledges.

Audience: Turning to the addressees, the ways in which these voices are incorporated is varied to the point that only trends but not patterns are discernible, as Table 7.3 below demonstrates. For instance, at the level of 'idea,' which contains only two instances from Long R, the thoughts projected onto the reader are meant to signal solidarity—that Long R literally knows what his audience is thinking. The statements at the level of locution, however, are universally combative: three of the four locutions include the reader's voice with the goal of disputing it (Dickens NR and Hribal R), while two of the four are bald on-record directives (Dickens NR and McKelvey NR), and thus highly face-threatening. But determining whether such patterns of function and distribution are true distinguishing characteristics between the realization types would, of course, require more data. Potentially, the preference by both author types for the less stringent mode of reportage may simply be a feature of the informal nature of these writings. That said, there are no known generic conventions dictating these authors' use of level—idea or locution—nor whether the projected voices must be presented in a positive or negative light. These, then, likely represent authorial choice, and are revealing of the author's stances for the same reason.

Table 7.3: Projection Nexuses of Audience Voices

		Mode	
		paratactic ("direct, quoted")	hypotactic ("indirect, reported")
Level	idea ' mental		*I know most of you who personally know me are in disbelief* (Long R) *You are thinking to yourself that this is completely out of character of the MAN you knew* (Long R)
	locution " verbal	*show me in the 1st amendment where it explicitly says you can't say "kill all cops"* (Dickens NR)	*tell every black person that you know they will die* (McKelvey NR) *you preach that humans should be nice to each other* (Hribal R) *It would be perfectly reasonable to say something is evil, so long as it opposes your point of view* (Hribal R)

Other: Here, the data itself suggests five roughly delineated categories of voices. The first three are:

1. the government and/or political sources (e.g., Long R quotes the U.S. Constitution when he writes *we the people*);
2. the media (Roof R directly quotes his *favorite film* twice); and
3. generic voices (identical here to the category presented in section 5.3.2 above, e.g., when Hribal R says *people value their lives*).

The fourth category is simply labeled 'unclear,' and it captures those instances where speech or thought is referenced without being tied to a speaker or thinker, as when Hribal R uses the passive voice to assert that *a dozen different things will be speculated to be at fault* but provides no semantic Agent responsible for the speculation. Finally, the fifth, somewhat idiosyncratic category of Individual Non-Victim captures two instances in Brahm NR's text where Osama bin Laden claims credit for the attacks which Brahm NR is describing, e.g., *Osama bin Laden will issue a video message claiming responsibility for what he dubs "America's Hiroshima"*. This is the only singular, 3rd

Person voice admitted to these writings who is not a victim. Figure 7.7 shows the frequencies with which these other voices appear.

Figure 7.7: Other Voices

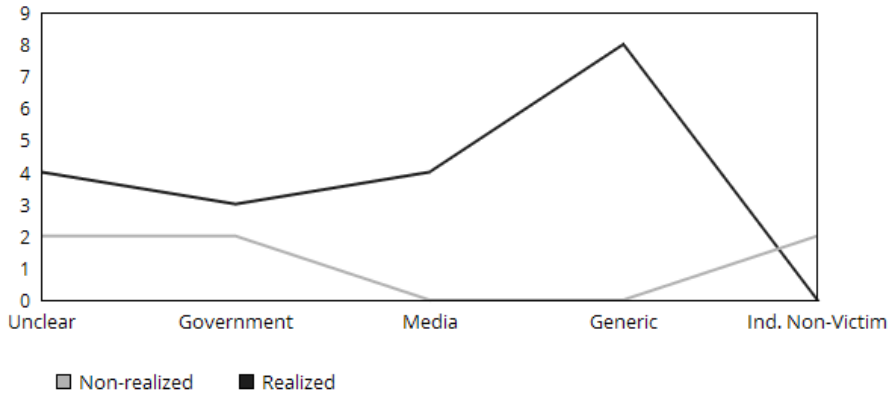

Frequency per 300 clauses

Setting aside Brahm NR's idiosyncratic category, a few things are apparent from the numbers alone. The first observation simply reinforces that the realized writers are much more willing to 'weaken' their stance by providing space for outside viewpoints. The second observation concerns the two null results in the NR writings, in regard to the media and to generic voices especially. While the disparity in the inclusion of generic voices is dramatic, every instance occurs in the text of a single author: Hribal R. (Again, this pledge also happens to be the longest in the dataset, likely allowing for the inclusion of comparably rarer meanings.) Similarly, while some form of media appears in three of the realized writings (Hribal R, Long R, and Roof R), the references are in no way uniform—*the news* and *film* are both cited, as is the less specific term *media*. Thus, while particular patterns of reference are not discernible, it appears that the simple inclusion of an outside voice marks a text as more likely to be realized rather than the *kind* of outside voice which is admitted.

7.4 DIALOGIC CONTRACTION: PROCLAIM

If the processes of attribution and acknowledgement are "dialogistic, making space in the ongoing dialog for those who might hold alternative views" (Martin & White, 2005: 113), then proclaiming is the active shrinking of such space within a text. It is a resource for confronting, challenging, and ultimately excluding alternate viewpoints. Authors do this by "representing the proposition as highly warrantable (compelling, valid, plausible, well-founded, generally agreed, reliable, etc.)" (Martin & White, 2005: 98) and therefore as something which a reader may not easily dispute. For instance, Skyline NR's claim that *i am going to start killing them first* is structured in a way which disallows all but the bluntest of rejoinders, e.g., "No, you aren't." Similarly, Shaw R's assertion that *Now the whole World Hates me* leaves no textual space for the possibility that this is not the case[32].

From a folk linguistic standpoint, the strength of commitment communicated by proclaiming seems like a natural component of threatening language. And both corpora are, in fact, heavy with clauses conveying this kind of meaning. For the sake of comparison, Figure 7.8 below shows proclaim along with the three other types of meaning at this level of the engagement hierarchy. Proclamations make up well over half of all engagement meanings found in the dataset—a finding which is likely a reflection of the genre of threatening. However, Figure 7.8 below also features a potentially counterintuitive trend, in that a significantly greater number of instances appear, not in the R texts where certainty of some variety might be more expected, but instead in the NR writings.

[32] Martin and White (2005) do not concern themselves with a distinction between objective and subjective assertions, since the propositional content put forward by a proclamation may be either. The basic factivity of an instance of proclaim—in this case whether the *whole World* actually hates Shaw R—is not at issue.

Figure 7.8: Portion of Proclaim by Realization Category

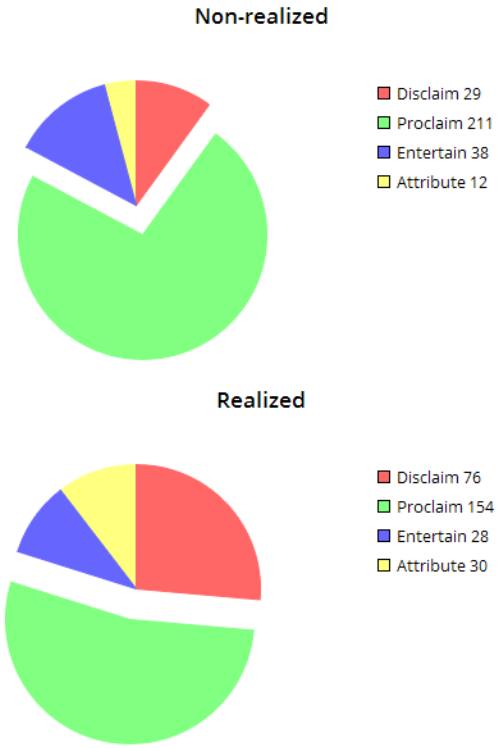

Frequency per 300 clauses

However, Hribal R—the longest pledge in the dataset by over 200 words—is again flagged as an outlier, as Figure 7.9 shows below.

Figure 7.9: Proclaim Dispersion by Corpus

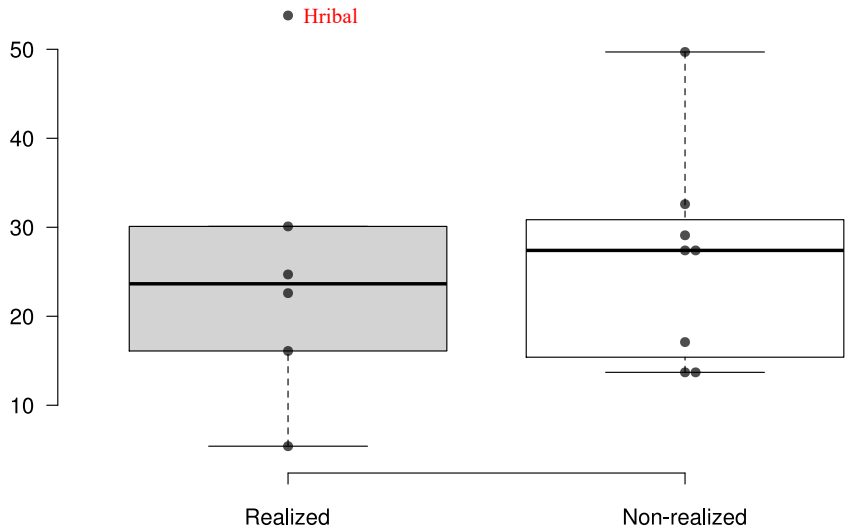

Frequency per 300 clauses. One outlier (>1.5IQR) is found.

With Hribal R's text removed, some—though not all—of the distributional difference of proclaim between the two corpora is erased, as adjusted Figure 7.10 now shows.

Figure 7.10: Portion of Proclaim by Realization Category (Adjusted)

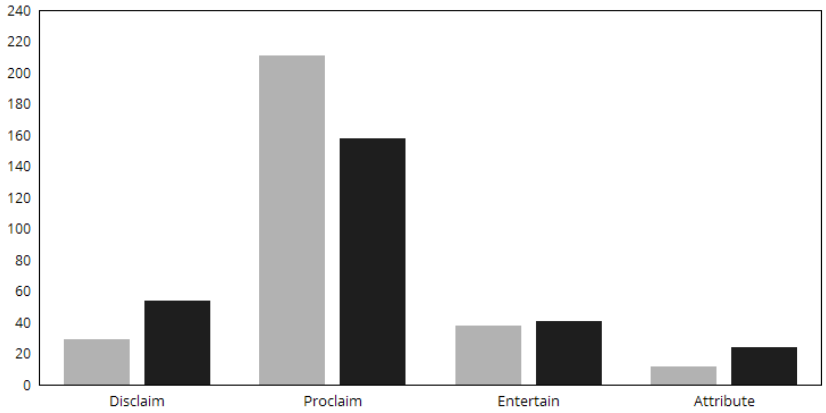

Frequency per 300 tokens

Even with the tightening of the numbers, a few observations are possible. At a general level, the NR pledges feature proposals which the authors would not later attempt. Nevertheless, these proposals tend to be presented in the rhetorical form most associated with certainty. In fact, this is of a piece with findings from several other areas of Appraisal examined thus far, e.g., that the non-realized texts take stances which are in many ways more violent, more threatening, more black-and-white in their moral thinking, more closed to outside voices, etc. To all this, the preponderance of proclamatory clauses adds a further prosodic feeling of conviction that is more pronounced in the NR corpus. Martin and White (2005) mention a potentially interesting theoretical possibility, which is that "[s]uch insistings or emphasisings imply the presence of some resistance, some contrary pressure of doubt or challenge against which the authorial voice asserts itself" (Martin & White, 2005: 128). An extremely interesting conjecture sits at this intersection of data and theory: perhaps the NR authors make greater use of this—and other strengthening resources—not to signal a true commitment

to act, but rather as a means of coping with a greater “pressure of doubt” (Martin & White, 2005: 128) about whether their pledges will be viewed as credible by their audience, a doubt which their historical lack of action shows is warranted.

7.5 DIALOGIC CONTRACTION: DISCLAIM AND DENY

Like acknowledgement’s subordinate relationship to attribution, Martin and White (2005) theorize denial as a type of disclamation. (Deny’s sister under disclaim, called ‘counter,’ was not identified by chi-square as warranting further investigation.) Disclamatory meanings are evident in the dataset whenever a counter-expectational is used to frame a proposition, e.g.: *However, to get one thing, another must be left behind* (Hribal R); *the bombings will take place regardless* (LA Unified NR). They are also in use when a proposition is explicitly negated[33] (i.e., denied), such as in the utterances *I have never been happy* (Kinkel R) and the *cops wont save you* (McKelvey NR).

Figure 7.11: Disclaim and Deny by Realization Category

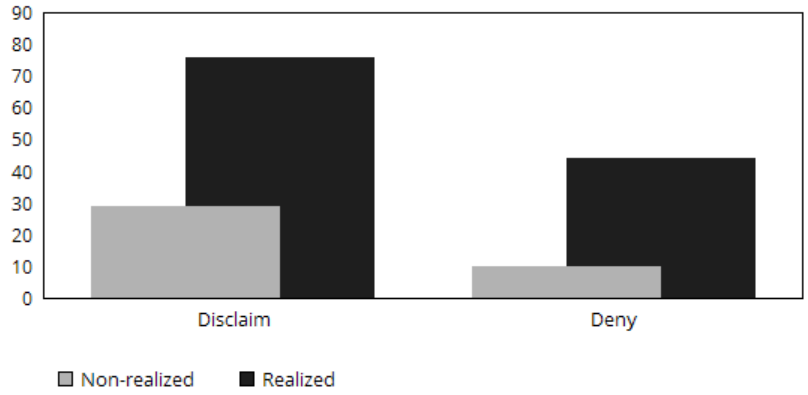

Frequency per 300 tokens

[33] Halliday and Matthiessen (2004: 46) note that in every language some meanings will be *lexicalized* (e.g., negative polarity is lexicalized in the verb ‘fail’) and some will be *grammaticalized* (e.g., the same negative polarity is construed grammatically in the expression ‘not succeed’). Martin and White (2005) do not say so explicitly, but coding for denial in engagement is a function of identifying grammaticalized forms of negation (e.g., *not*, *never*, *none*). In Appraisal, lexicalized negation (e.g., *fail*) is properly dealt with via the subsystems of attitude.

A wide-angle view of the two meanings, seen in Figure 7.11, shows their distribution in the dataset. As a category, disclaim covers "those formulations by which some prior utterance or some alternative position is invoked so as to be directly rejected, replaced or held to be unsustainable" (Martin & White, 2005: 118). In other words, constructions of this type serve to identify ideas with which an author in some way disagrees. Disclaim—and denial especially—is thus a tool of *dis*alignment between a writer and a second voice. The dispersion of disclamation across the dataset is shown below in Figure 7.12.

Figure 7.12: Disclaim Dispersion by Corpus

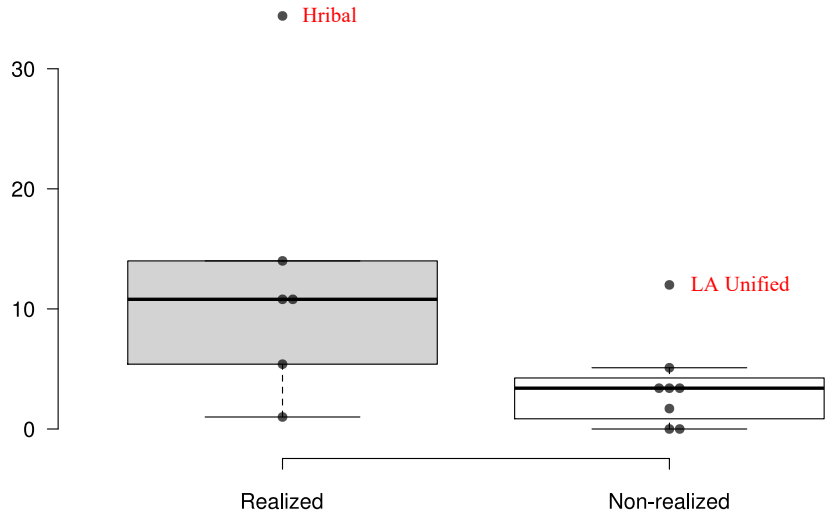

Frequency per 300 clauses. Two outliers (>1.5IQR) are found.

However, because disclaim is classified as a 'broad' category (Martin & White, 2005: 117), and because the disclaim subtype of countering shows no significant difference in usage by the two author types, the primary focus of this section is the area of denial. The dispersion of denial across the dataset is shown below in Figure 7.13. Interestingly, both Hribal R and Brahm NR are flagged as outliers at the level of disclaim but not at the level of deny, meaning that each text's aberrant usage is at the level of counter. Thus, their status as outliers does not affect the proceeding analysis of denying.

Figure 7.13: Deny Dispersion by Corpus

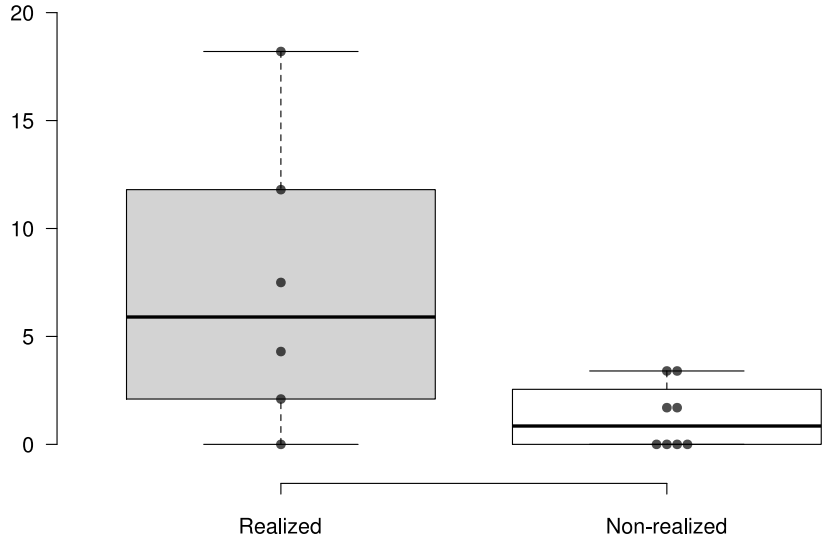

Frequency per 300 clauses. No outliers (>1.5IQR) are found.

As shown in Figure 7.13, the R corpus makes greater use of this resource. Not only that, a glance at the lexical mechanics of denying (seen in Figure 7.14 below) shows that R pledges are also more varied in the types of negation they employ. While this variation is likely a simple byproduct of denial's higher frequency in the R writings, it is still interesting to discover that the NR texts rely on just one negative marker, adverbial *not*. According to the Corpus of Contemporary American English (COCA)[34],

Figure 7.14: Grammatical Markers of Denial by Realization Category

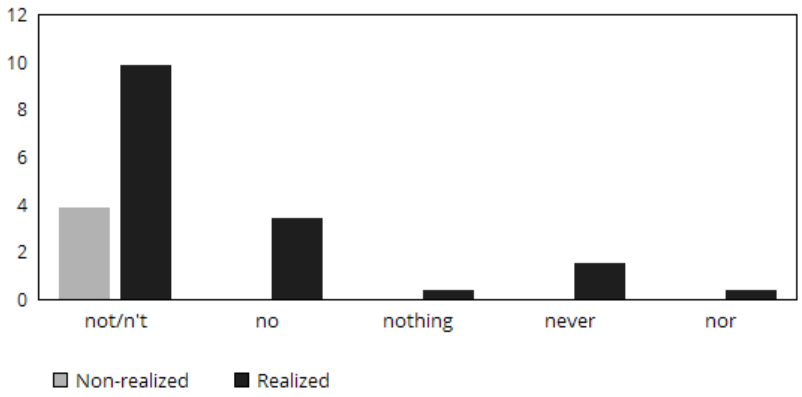

Frequency per 1000 words

[34] https://www.wordandphrase.info/frequencyList.asp

this preference for *not* over *no* is common: *not* is ranked the 28th most common word in American English versus 93rd for *no*. (The other three terms follow the same general AmE distribution patterns, according to COCA, with *never* ranked 138th, *nothing* ranked 316th, and *nor* ranked 1,066th.) One possible interpretation for the heightened frequency and variety of denial in the R pledges is as an additional piece of evidence that these texts are more connected to voices outside their textual borders than the NR writings. This is because denial does not operate as the rhetorical equivalent of psychological repression—it is not a means "by which we avoid becoming conscious of unpleasant or dangerous truths" (Maslow, 2014: 58). Instead, the act of denying signals that others' ideas have been both anticipated by the author and digested enough for the author to determine that he or she disagrees. Or, put more bluntly, denial is an admission that consensus reality exists. As such, it is a potential marker of higher conceptual complexity (Smith, 2006). But more than this, it is an assertion that the author sees consensus reality *more* clearly than the outside voice the pledge writer is contradicting.

Turning from mechanics to meaning, denial is characterized as a strong rhetorical action, since "to deny or reject a position is maximally contractive" (Martin & White, 2005: 118). From the viewpoint of folk linguistics, it seems unsurprising that the R texts would make greater use of this resource: the realized authors are, through denying, signaling the kind of epistemic certainty which goes hand-in-hand with contractive formulations, all of which serve to strengthen the author's apparent commitment (Gales, 2010). Previous parts of this analysis, however, argue that such contractive processes are mainly an activity of the NR pledges. Yet, Martin and White (2005) complicate the idea that denial is merely contractive. The authors note that the grammar of negation is at the heart of denial, and "the negative necessarily carries with it the positive, while the positive does not reciprocally carry the negative" (Martin &

White, 2005: 118). In other words, a proposition which is summoned up in order to be denied nevertheless leaves its positive imprint on the content of the text.

Thus, the specific work of rejecting an alternative position includes a revealing rhetorical double action. So, for instance, when Hribal R claims that he is *not trying to justify my actions whatsoever here*, the reader is reflexively faced with the inverse prospect, i.e., the notion that Hribal R is indeed using the text to justify his actions. As noted above, the trend so far catalogued in this analysis is that the NR texts are more closed to outside voices than their R counterparts. That the realized authors make greater use of denial—a means of contracting their discourse—would seem to upset this pattern. Yet, there is an element of expansion at work here as well, as Hribal R's example demonstrates, even if such discursive space is created for the sole purpose of being collapsed. Thus, the one area in which the R corpus appears more rhetorically closed is nevertheless one whose operation is contingent on textual expansion. From a theoretical perspective, then, the previously identified pattern seems to hold.

For the purposes of analysis, the idea that the positive travels arm-in-arm with the negative offers an interesting opportunity. By examining what an author feels the need to deny, a clearer picture is possible of just which topics have captured that author's attention, either as a past statement the writer seeks to rebut or a possible objection the writer foresees the need to intercept. Such topics can offer a glimpse into what Shuy (1981: 115) calls the "cognitive thrust" of a discourse. This is explored in the next subsection.

7.5.1 TOPICS OF DENIAL

Hribal R writes of his friends and family that *I don't belong with them*. No great analytical leap is necessary to see that the issues of social inclusion and solidarity were of some concern in the creation of this pledge. Even those denials which are ostensibly aimed at

strengthening the overall feeling of menace and speaker control, such as when LA Unified NR claims that *Your security will not be able to stop us*, read as the outcropping of a particular anxiety when considered alongside the proposition's positively framed double, in this case that the writer and his or her compatriots can indeed be thwarted by campus security. In other words, *whether or not LA Unified NR is in control* is the topic which he or she has opted to address via this denial. When each token of denial in the dataset is interrogated this way—by asking 'what underlying question is the denial addressing'—the data itself suggests three overarching areas of concern which both author types engage with through negation.

The first might best be covered by a term like 'autonomy.' As the example from LA Unified NR shows, many denials in the dataset assert some kind of authorial control, or self-control, over various events and contingencies. Hribal R, as noted previously, uses denial to say *I am not trying to justify my actions whatsoever here*. In doing so, he not only passively calls to mind the positive inversion—that he is, in fact, justifying his actions—he also actively claims authority over the interpretation of his own text. The denial essentially stresses Hribal R's will to be free from any reading which is different than what he intends[35].

The second area of concern has to do with claims to special knowledge or truth. Of the three identified topics here, this is perhaps most in keeping with Martin and White's (2005: 120) conclusion that denials are 'corrective,' i.e., that they "present the addresser as having greater expertise in some area than the addressee and [are] acting to correct some misunderstanding or misconception." Arguments made from this position of 'superior knowledge' are common in the pledges, such as when Skyline NR determines that his imagined victims *dont deserve the lives they have been given*, and

[35] Whether his protestations are sincere, and the pledge is indeed something other than a justification, is beside the current point.

when McKelvey NR predicts that *the cops <u>wont</u> save you*. In such formulations, the pledge writer positions him- or herself as an arbiter of consensus reality (Gellerman & Suddath, 2005) by contradicting an anticipated or imagined assumption on the part of an outside voice. The denial essentially 'sets the record straight'.

The last broadly defined topic could be called 'social standing.' Denials in this category may be used to assert in-group status, such as when LA Unified NR writes that *I would <u>not</u> be able to do it alone*, or when Hribal R credits Eric Harris and Dylan Klebold with inspiring his ideations, saying *if it <u>weren't</u> for [them] this whole event would <u>never</u> occur*. Secondarily, an author may also seek to stand apart from a group, claiming status as an outsider, as Hribal R does when he says of anime and its followers that he does *<u>not</u> consider myself a fan*. These denials thus address questions of solidarity, social distance vs. social acceptance, etc.

Just how concerned each pledge type is with these three topics in general is captured in Figure 7.15 below.

Figure 7.15: Denial Topics by Realization Category

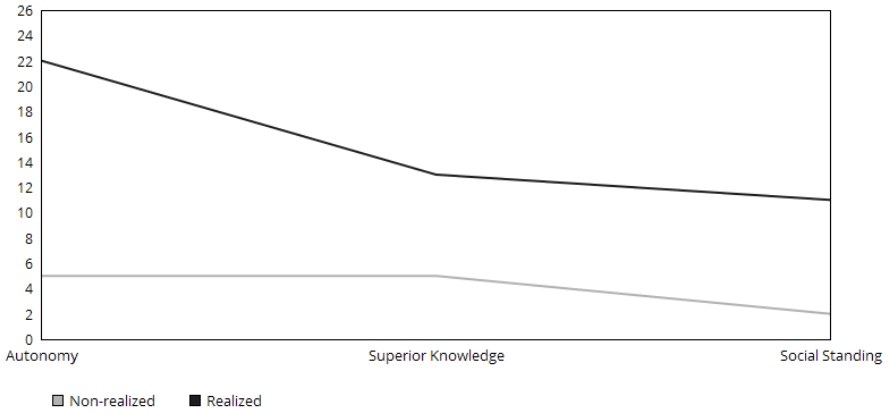

Frequency per 300 tokens

Besides the clear difference in frequency already flagged by the chi-square calculations, the main difference in the contours of the distribution lies in autonomy. Questions of an author's self-determination are addressed substantially more often by

denials in the R texts, at a rate of over 4:1. Of course, any question leaves open whether a response will be negative or affirmative. Looking at the three data-driven areas of concern from this perspective—is the 'answer' to the issue being denied positive or negative—reveals a litany of interesting stances, especially on the part of the R authors.

For example, denials may be employed to communicate an absence of one of the three qualities, e.g., an absence of autonomy, superior knowledge, or some kind of social standing. Addressing each in turn, authors of both realization categories routinely abdicate their autonomy by denying that they can manage a particular situation, e.g., *My head just doesn't work right* (Kinkel R). Second, while access to knowledge and information is its own type of power (van Dijk, 2001), several authors admit to a lack of understanding about the issues they raise, as Shaw R does when he says *I never understood how they even made it on a date*. Finally, this dynamic is perhaps most interesting when at work in questions of social standing. Denials in this area are not only employed to move closer to or farther from a particular social group. They are also used to express a *longing* for inclusion. This is apparent when Roof R writes that *I am not in the position to, alone, go into the ghetto and fight*, a denial which serves as a tacit admission that his preferred social circumstances are beyond his control and works in tandem with his proclamatory lament that *We have no skinheads, no real KKK, no one doing anything but talking on the internet*.

Such protestations of weakness are interesting primarily because they run counter to folk ideologies about the role of control in threatening language: control of the parameters of the conflict; the outcome of the threat; the recipient's role in that outcome; etc. (Gales, 2010). Instead, grammatical negation in the dataset is frequently used to *remit or relinquish* authorial control. Moreover, such admissions of powerlessness are

actually more prevalent in the R texts, pledges which the authors would later attempt. Figure 7.16 illustrates the difference in usage between the two realization categories.

Figure 7.16: Denial Topic 'Polarity' by Realization Category

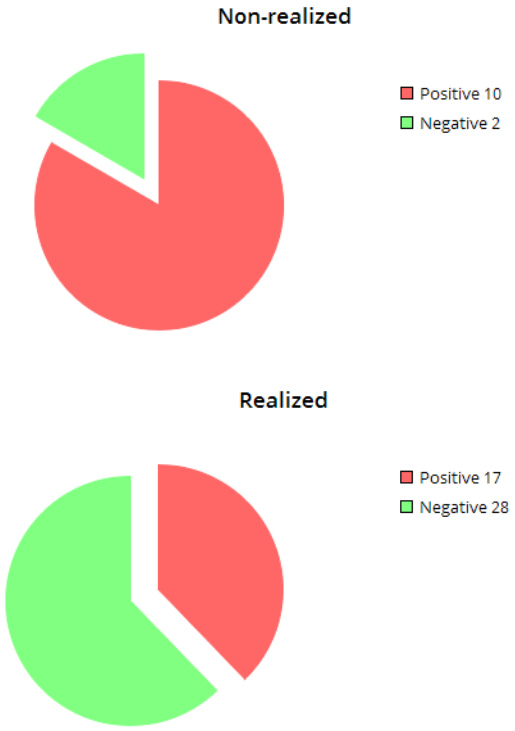

Frequency per 300 clauses

The R pledges are over 16 times more likely to use grammatical negation to announce that they suffer some type of impotence or vulnerability. Such constructions account for 62% of all denials in the R corpus, versus just 17% in the NR writings. In fact, a denial used to concede some kind of weakness occurs just once in the non-realized corpus, when LA Unified NR writes that *I would not be able to do it alone*. And even this token must be judged a kind of hybrid, since it admits to a lack of autonomy by also making a positive claim to in-group solidarity. Compared to this, the realized texts are far blunter in their admissions. Hribal R, for example, bemoans his lack of control over the side effects of his ensuing knife attack (answering a question of autonomy) when he writes that *The only thing I wish didn't come out of this was the pain it would*

inevitably cause my closest friends and family. Elsewhere, Roof R implies his desire for an absent community (topically: attachment) via a string of denials when he says that *We have no skinheads, no real KKK, no one doing anything but talking on the internet*. And Kinkel R professes a lack of understanding when he says, simply, *I don't know what is happening* (topically: superior knowledge).

Along with the greater tendency for a realized author to highlight various kinds of personal and/or social deficiencies using denial, there are three tokens in the area of autonomy where the realized authors directly echo a single structural choice to communicate a lack of control:

- *I have no other choice* (Kinkel R)
- *I have no choice* (Roof R)
- *I have no choice* (Shaw R)

Of course, the simple fact that half of the realized authors in the dataset—each with very different underlying concerns and motivations—reach for a practically identical utterance is fascinating, especially in light of the stancetaking work these denials perform. Grammatically, these simple declaratives are strongly stated. Functionally, however, the denials actually serve to *weaken* the authors' stance, by framing their commitment to the act as something which is beyond their control. Taken together, claims that future violence is compulsory may, in fact, be evidence that an author's commitment to performing the act is quite strong. Or, in other words, authors denying that they have the strength to resist what they see as an obligation may mean, somewhat paradoxically, that chances are higher the authors will attempt the 'required' action. This counterpoint of volition versus requirement will be explored again from a different grammatical perspective in the analysis of modal auxiliaries in Chapter 8.

7.6 DIALOGIC CONTRACTION: TEMPORAL SHIFTING

The final area of analysis examines an oddity in the dataset. For writings whose primary purpose is a description of a future event, it is difficult to imagine a way in which such a text could be more closed to other eventualities than to frame the event as if it has already happened. For instance:

- *People say what I did was evil.* (Hribal R)
- *I know most of you who personally know me are in disbelief to hear from media reports that I am suspected of committing such horrendous acts of violence.* (Long R)

This phenomenon is so specific that, quite understandably, no mention of it as a stance marker is found in Appraisal (Martin & White, 2005), but the construction occurs often enough in this dataset to warrant comment. These are peculiar specimens from the standpoint of linguistic theory. In context, however, it is quite possibly the strongest stancetaking device which an author may employ to communicate epistemic certainty about how a future event will unfold—which is to say, *that the event will indeed unfold*. The incongruity hinges on the choice of a realis over an irrealis mode. Because future events are not fact but merely potential, they are most naturally located in the 'virtual' grammatical world of irrealis (Halliday and Matthiessen, 2004: 425). Instead, pledge authors opt for the language of the "real, or actual" in these utterances (Halliday and Matthiessen, 2004: 425).

Interestingly, this choice treats the reality of the event as a presupposition, i.e., a "proposition [that] is assumed to be common knowledge shared by speaker and addressee" (Cruse, 2011: 42). As noted in Chapter 2, future actions that are presented as a given are a potential indicator that the threatener views them as highly feasible (Geurts *et al.*, 2016). However, there is more at work here than the semantic slyness inherent to presuppositions. When framed contextually, temporal shifting is an

especially subtle linguistic trick: if all goes as the author hopes, then the pledge will be read quite attentively (by friends, family, authorities, etc.) *after* the described attack has occurred. At that point, the realis mode will seem perfectly natural, almost as a kind of chronicle. So much so that the incongruity of its use in a text written *before* the event in question might seem unremarkable, almost grammatically 'invisible' to whomever happens to read it.

This is especially important because temporal shifting does not appear at irregular or difficult to explain moments in the dataset. Authors who employ this kind of presupposition inducer do so only when discussing the central bad act of these pledges. No other, more minor events are framed as if they have already happened. Additionally, and perhaps critically, *these tokens occur only in the R pledges*. The phenomenon is relatively widespread in this corpus: of the six realized pledges, only Rodger R and Kinkel R's texts do not feature some version of this formulation. Every instance of temporal shifting in the dataset is italicized in Table 7.6.

Table 7.6: Tokens of Temporal Shifting

Realized Author	Token
Hribal	*By now*, my art has obviously *been revealed*
	People say what I *did* was evil.
Long	I know most of you who personally know me are in disbelief to hear from media reports that I *am suspected* of committing such horrendous acts of violence.
Roof	I *chose* Charleston because it is most historic city in my state
Shaw	Truthfully, I *feel* so much better *after* hitting an asian Woman in the face with a steel rod.
	It *was* the greatest achievement of my life.

Of course, the violent event is not shifted into the grammatical past by these authors on a consistent basis. In each of the four, violence is also somewhere discussed in the future tense (a necessary criterion for these writings to qualify as a

pledge to harm rather than as, say, confessions). However, in the six utterances listed, each has in common a projecting of the authorial voice forward in time to the moment when an imagined reader has discovered the piece of writing—at which point the author's words will have taken on the added weight of history. By contrast, no NR writer steps into the grammatical time machine this way. Discussing their imagined actions as a *fait accompli* is a choice made only by those authors who would, in fact, attempt the described violence.

7.7 A NOTE ON PLANNING

Before leaving engagement entirely, a note on planning in threatening language is warranted. As mentioned on section 5.2.7, "evidence of intent" includes the "specificity of [the] plan" which the threatener hopes to enact (Mohandie, 2014: 136). Evidence of research and planning can include verbal confirmation that an author has engaged in behaviors like "surveillance of the target, Internet searches, testing security around the target, and researching methods of attack" (Bulling & Scalora, 2013: 15). Such behavioral concerns are broader in scope than any single linguistic variable or rhetorical type. However, because such a question is essentially discursive—stretching across clausal boundaries—it is a fit final question to address as an addendum to the analysis of engagement.

First, though, no single variable or bundle of stancetaking devices is associated with planning by the various literatures. Furthermore, just what kind of expressed ideation qualifies as research or planning is also opaque. For example, does Archangel Michael NR's series of explosions qualify?

> Text 7.1: Archangel Michael NR Pledge (Excerpt)
> THERE ARE 50 40MM ALUMINUM THERMITE PIPE BOMBS ACROSS OFFICIAL BUILDINGS IN CHEYENNE. THERE ARE 600 40MM THERMITE PIPE BOMBS HIDDEN ACROSS MULTIPLE BIG CITIES AND OFFICIAL

GOVERNMENT BUILDINGS IN WYOMING. THERE ARE 50 SETS OF COMPACT U.S. ARMED FORCES STANDARD C4 AT THE CYS AIRPORT.

Or perhaps Rodger NR's statement is more prototypical: *I will start luring people into my apartment, knock them out with a hammer, and slit their throats*. If the question is approached as a loose kind of linking (Lock, 1996), or a sequence of events joined by explicit or implied temporal junctures (Labov, 2013), as in *I will do X, Y, and Z*, then both sequences would seem to qualify. Yet, both seem more fleshed out—though less feasible—than Shaw R's self-described *strategic plan of using violence* which was, in its entirety: *Every Asian Woman by herself must be hit in the face*. However, intuition is unreliable for quantifying the level of detail in expressions like these. Without further, more precise guidance from the relevant literatures, then, the "specificity of [the] plan" (Mohandie, 2014: 136) is not something which can be pursued further from a linguistic perspective.

7.8 SUMMARY OF ENGAGEMENT

The system of engagement is structured with two main analytical ends in mind. The first is to uncover the "the linguistic resources by which speakers/writers adopt a stance towards to the value positions being referenced by the text" (Martin & White, 2005: 92). The second is how those resources are applied "to the relationship which the text construes as holding between speaker/writer and the text's putative addressee" (Martin & White, 2005: 95). Viewing threatening language through the combination of these two lenses, as Gales (2010, 2011) has done previously, potentially "reveals the threatener's apparent level of commitment and investment in the threat" (Gales, 2011: 37). This analytical function gets to the very heart of the current study. Engagement is therefore a critical framework for understanding the rhetorical ways in which R and NR pledges differ from each other.

Martin and White (2005) offer fifteen types of engagement options which a writer may call on to take a stance towards their subject matter and their audience, of which nine are surface forms, or leaf nodes of the engagement hierarchy (e.g., attributing, distancing, etc.). Of the fifteen, five meaning types in the dataset were identified by chi-square as warranting further analysis: attribute, acknowledge, proclaim, disclaim, and deny. Interestingly, two of these—attribute and deny—are surface categories, and prove to be the richest areas of investigation, showing stark patterns of divergence between the two realization categories.

Justification: First, however, a word on one null result was warranted. Justification is of paramount interest from a threat assessment perspective, since understanding a threatener's motive is key to determining the level of risk he or she poses to a target (Borum *et al.*, 1999; Calhoun & Weston, 2015). Furthermore, ideologies about threatening language indicate that "the justification for the threat" is one of the expected language functions of the genre (Gales, 2010: 96). Because of these considerations, justification was examined further, potentially uncovering an important distinction between the corpora, e.g.: five of the eight R pledges employ justifications in conjunction with a 1st Person pronoun to defend their choice of target or to justify the attack itself. This kind of personalized usage appears just once in the NR corpus, when Skyline NR says of the jocks that ***i*** *am going to start killing them first because they deserve it the most*.

Attribute / Acknowledge: Of the meanings whose frequencies are flagged as statistically significant, only two are dialogically expansive, and these two are related as mother and daughter nodes: attribute and acknowledge. Attribution is the means by which authors may externalize a proposition, a process which is primarily accomplished "through the grammar of directly and indirectly reported speech and thought" (Martin & White, 2005: 111). That writings so loaded with rage and imagined violence make any

room at all for outside voices is interesting by itself. Further, the fact that the R texts are more than twice as likely to open themselves in this way invites interpretation—an invitation that will be taken up in the conclusory discussion of chapter 11. Generally, both author types opt for indirectly reported thought over directly or indirectly reported speech. This disinterest in, or hesitancy about, employing the stronger evidentiality of direct quotation may simply be a product of the informal register of these texts. However, only reporting others' ideas and not their words lends this reportage the flavor of subjectivity instead of empirical fact, subtly weakening the writers' arguments. In other words, the victims are *imagined* as having thought a host of damning things but are not *known* to have actually thought them.

Beyond this similarity, further findings in attribute and acknowledge indicate that a realized text:

- is over twice as likely to make space for an outside voice;
- is more likely to include the voice of the pledge's imagined victim(s); and
- when the statistically non-significant area of distancing is removed from consideration, victim's voices disappear from the NR corpus entirely.

In general, the simple inclusion of outside voices is likely to mark a text as realized. However, a pattern of usage is potentially discernible with the few instances of acknowledgement present in the NR texts. When the pledges are viewed as a sort of narrative—telling the story of the imagined attack—each of the seven tokens appears immediately after what is arguably the "most reportable event" (Labov, 2013: 23) of the text, adding evaluative weight to this central moment. Furthermore, this is the *only* point at which outside voices are admitted into the NR writings. Thus, despite the paucity of acknowledgement in the NR corpus, this particular meaning does indeed seem to be deployed strategically by these authors.

Moving on from dialogic expansion, three areas identified by chi-square fall within the realm of dialogic contraction: proclaim; disclaim; and deny.

Proclaim: Of the three, the meanings of proclaim are perhaps the simplest to parse in the present dataset. Proclaiming, either by concurring, pronouncing, or endorsing a particular viewpoint, is used to present a proposition as "highly warrantable" (Martin & White, 2005: 98). Given the strengthening effect this option lends to an author's stance, it is less than surprising that this particular rhetorical formulation constitutes well over half of all engagement meanings found in the dataset. In all likelihood, the preponderance of this type of meaning reflects the importance of control in the wider genre of threatening language, as discussed by Gales (2010).

The discovery, then, that a 37% greater chance exists for a proclamation to appear, not in a R pledge where heightened commitment might be expected, but instead in a NR pledge, is more surprising. However, "[s]uch insistings or emphasisings imply the presence of some resistance, some contrary pressure of doubt or challenge against which the authorial voice asserts itself" (Martin & White, 2005: 128). This counterpoint invites an interesting conjecture: perhaps the non-realized authors make greater use of this resource as a means of coping with a greater 'pressure of doubt' about whether their pledges will seem credible, a doubt which their historical lack of action shows was warranted. This point will be taken up again in the broader discussion of chapter 11.

Deny: Disclaim is theoretically broad, covering "those formulations by which some prior utterance or some alternative position is invoked so as to be directly rejected, replaced or held to be unsustainable" (Martin & White, 2005: 118). Because of this breadth, analysis focused on disclaim's subtype, deny, which was also identified as an area of statistical divergence between the corpora. Denials are an option for disalignment through the resources of grammatical negation (*no*, *not*, *never*, etc.), introducing a value position into the text with the sole purpose of rejecting it (Martin &

White, 2005: 118). That the R texts make use of this type of meaning at a rate of over 4:1 is in keeping with other findings such as, e.g., acknowledgement, which show that realized authors are far more comfortable providing space in their writings to outside voices, even if this space is offered up only to be closed down, as is the case with denying.

From a theoretical standpoint, this interdependence of expansion and contraction leaves what may best be described as a positive propositional residue, "since the negative necessarily carries with it the positive" (Martin &White, 2005: 118). First, each token of denial in the dataset was approached with the intent of uncovering the positive proposition being rebutted. The results were then able to be collected under three larger, data-driven topical headings: denials which made some assertion about the author's *autonomy* or self-control; those which laid claim to some kind of *superior knowledge* or special truth; and finally, those which addressed *social standing* by placing an author within or without a certain social group.

This analysis discovered that a realized pledge is:

- four times more likely use denial to address concerns of autonomy; and
- over 16 times more likely to use grammatical negation to announce that they suffer some type of impotence or vulnerability—such constructions account for 62% of all denials in the realized corpus, versus just 17% in the non-realized which concede some kind of weakness.

Such protestations of weakness are interesting primarily because they run counter to folk ideologies about threatening language, which instead focus on threatener control (Gales, 2010). In reality, grammatical negation in the R pledges is often used to *remit or relinquish* authorial control.

Lastly, three tokens show an interesting convergence on a single realization of denial, whose form may be glossed as 'personal pronoun + a lack of volition', e.g., *I have*

no choice (Roof R). Not only do three of the six realized authors (Roof R, Shaw, R, Kinkel R) make use of this nearly identical utterance, but there is no equivalent anywhere in the NR corpus. Hypothetically, then, authors who deny that they are strong enough to resist an obligation may, paradoxically, be signaling a stronger commitment to performing the 'required' action. This interplay of inclination versus requirement will be investigated more fully in the following chapter analyzing modal auxiliaries.

Temporal Shifting: The analysis of engagement closes with a somewhat unorthodox finding, one which is not anticipated or addressed by Martin and White (2005) under engagement, but whose discursive value as a stance marker of commitment is remarkable. Four of the six realized authors—and only the realized authors—engage in what is herein called 'temporal shifting', where a future event is discussed as if it has already occurred, e.g., *People say what I did was evil* (Hribal R). For writings whose primary purpose is a description of a future event, it is difficult to imagine a way in which an author's stance could be stronger or more closed to other eventualities than to frame the violence as if it has already occurred. The systemic peculiarity hinges on the choice of a realis over an irrealis mode, or for the language of the "real, or actual" instead of "the potential, or virtual" (Halliday and Matthiessen, 2004: 425). This choice communicates a certainty that the event will occur, by semantically presupposing that it has *already* occurred. However, the final red flag raised by this odd phenomenon is the fact that such shifts only occur in clauses discussing the primary violent act of these pledges, thereby treating them as a *fait accompli*.

Because modal auxiliaries are one resource by which authors may realize various engagement meanings, several of the threads uncovered in this chapter will be taken up again and expanded from different perspectives in the analysis of Chapter 8, below.

CHAPTER 8 MODAL AUXILIARIES

Broadly defined, modality[36] provides resources for speakers to "construe the region of uncertainty that lies between 'yes' and 'no'" (Halliday & Matthiessen, 2004: 147). These gradations of indeterminacy serve to construct interpersonal meaning (Halliday & Matthiessen, 2004; Lock, 1996) by "enacting roles and relations between speaker and addressee" (Matthiessen *et al.*, 2010: 126). Essentially, modality is the means by which language users package an epistemic or affective stance towards particular content along with the content itself, simultaneously providing a listener with the intended information as well as the speaker's opinion of the information's validity (Halliday & Matthiessen, 2004: 116).

Because modality is a resource of "intersubjective positioning" (Martin & White, 2005: 95), this area is at least implicitly addressed in any analysis of engagement. And indeed, modality influenced the feature coding process which underlies the efforts of Chapter 7. For example, the high likelihood communicated by the helping verb *will* in *It will make national headlines* (LA Unified NR) is largely responsible for that token's classification as a proclamation. Likewise, the statement *I could say life is evil* (Hribal R) is considered a token of entertainment due to the bundling of the unreal condition and the semantics of ability comprising the modal *could*.

However, modal auxiliaries are a tangible lexical element held up by threat assessors and linguists alike as signals of commitment. For linguists, auxiliary verbs like

[36] According to Halliday and Matthiessen (2004: 147), the term *modality* "strictly belongs" to the realm of propositions, which exchange information through assertions and denials. The exchange of goods-&-services through offers and commands are instead considered proposals, meanings which are not modalized but rather modulated. Differences between the two categories will be considered at appropriate moments in the analysis, but, in keeping with Halliday and Matthiessen's (2004: 128) practice elsewhere, *modality* will be used as the blanket term covering both modalization and modulation, and *modals* will be used to describe the lexical units responsible for conveying this information.

will and *shall* are regularly identified as a grammatical resource by which authors encode their commitment to the validity of a proposition or proposal (Biber *et al.*, 1999). In her discussion of the modal area of 'requirement,' for example, Gales (2010: 97) makes the connection between modality and speaker intent explicit, commenting that "a *commitment* to the *intended action* [emphasis added]" is one expected use of "modals of obligation" in threatening communications.

These linguistic intuitions are reflected in the threat assessment literature as well. If a practitioner singles out a subject's language as a behavioral cue worth watching, some mention of modal auxiliaries will often follow. Mardigian (via Gales, 2010: 26), for instance, refers to *must*, *have to*, and *will* as "modals of intent" to which an analyst should be sensitive. Elsewhere, in her survey of the folk linguistic ideology surrounding threats, 27% of respondents identified "forceful modals" such as *will*, *must*, *shall*, and *have got to*, as a type of language they would expect to find in threatening texts (Gales, 2010: 96). These are circled in Figure 8.1 below.

Figure 8.1: Student Ideologies about Threatening Language (adapted from Gales, 2010: 96)

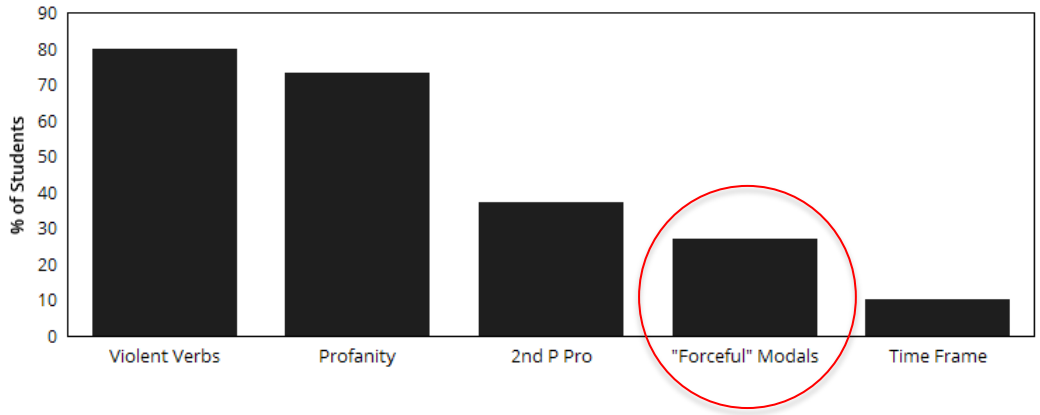

More typically for authors discussing threat assessment, Turner and Gelles (2003) make no mention of lexicogrammatical features worth watching for. Yet, by way

of illustrating an argument, the authors invent for their readers an archetypal threat which prominently features up-scaled modal auxiliaries (underlined):

> Company XYZ has continued its manufacturing process that places the environment at risk. Since Mr. Jones is clearly in charge and refuses to make changes, he will have to suffer. I have tried to get others to recognize the terrible impact to no avail. Now that the time of your stockholders meeting is approaching, Mr. Jones must be removed by force. Specifically, he will be shot to death—an execution, not a murder. (Turner & Gelles, 2003: 99)

No matter if these lexemes are singled out by practitioners as a marker of intent or if their value as a diagnostic tool is merely implied, modal auxiliaries thus constitute a lexical class which the literatures indicate is key to understanding how intent and commitment may be expressed linguistically. Thus, despite their implicit consideration in the analysis of engagement presented in the previous chapter, this special relationship with commitment earns modal auxiliaries their own focused analysis here. The current chapter will therefore zero in on this smaller class of words[37], to see if their behavior in this dataset comports with the previous literature on threatening.

Before presenting the analysis itself, however, a brief word on coding the auxiliaries is offered here, since the process differs slightly from that used in the preceding chapters. As noted elsewhere, all three systems of Appraisal were manually coded into specialized software, UAM Corpus Tool (O'Donnell, 2013). As a concordancer, Corpus Tool also offers an automatic tagging function. However, modality is a rich area of English grammar (Lock, 1996), with a correspondingly rich linguistic literature. As a result, linguists commonly differ in their taxonomies of these meanings. Indeed, Corpus Tool slightly diverges from other systemic functional taxonomies, e.g., Halliday and Matthiessen (2004) and Lock (1996). While the software lists 'intention' and 'willingness' as two branches of the modal meaning of 'volition,' it

[37] Such a limited focus, it should be noted, is also in keeping with Lock's (1996: 193) "narrow definition of modality" which "encompasses only the modal auxiliaries."

opts for ‘ability’ as the third meaning where Lock (1996) theorizes meanings of ‘determination’. Small departures like this, combined with the relative lack of transparency in the word list which Corpus Tool references for its automatic tagging, complicates any *post hoc* quality control. For these reasons, the auxiliaries in the pledges were both coded *and counted* manually. Lock’s (1996) taxonomy, reproduced in Table 9.1 below, is both the most transparent and the most thorough of the known options on offer, while also in keeping with the systemic functional theory underlying the rest of Appraisal, so this was the guide of choice.

Table 8.1: Modal Auxiliaries and their Meanings (Lock, 1996: 213)

Modal Meaning	*Modal Auxiliaries*
Likelihood	
Positive	
High (certainly)	must, have (got) to, will, shall, would
Mid (probably)	should, ought to
Low (possibly)	may, might, could
Negative	
High (certainly not)	cannot (can’t), could not (couldn’t), will not (won’t) (shan’t), wouldn’t
Mid (probably not)	should not (shouldn’t)
Low (possibly not)	may not, might not
Requirement	
Positive	
High (obliged/necessary)	must, have (got) to, will, shall, need*
Mid (advised)	should, ought to
Low (permitted)	can, may
Negative	
High (obliged/necessary not to)	must not (mustn’t), may not, cannot
Mid (advised not to)	should not (shouldn’t), ought not to (oughtn’t)
Low (permitted not to)	need not, don’t have to, haven’t got to
Frequency	will, can, could, may
Inclination	will, will not (won’t), shall, shall not (shan’t)
Potentiality/ability	can, could

*Strictly speaking, not formally a modal auxiliary

Finally, forms of the verbal construction *be going to* are included for the purpose of analysis since 1) these forms may signal “self-volitional control” on the part of the

author (Gales, 2010: 171); and 2) “there is no reason why [the *going to* forms] should not be included in an account of the tenses of English” (Lock, 1996: 150).

8.1 OVERVIEW

The distribution of all modal auxiliaries across the dataset is shown in Figure 8.2.

Figure 8.2: Modal Auxiliaries

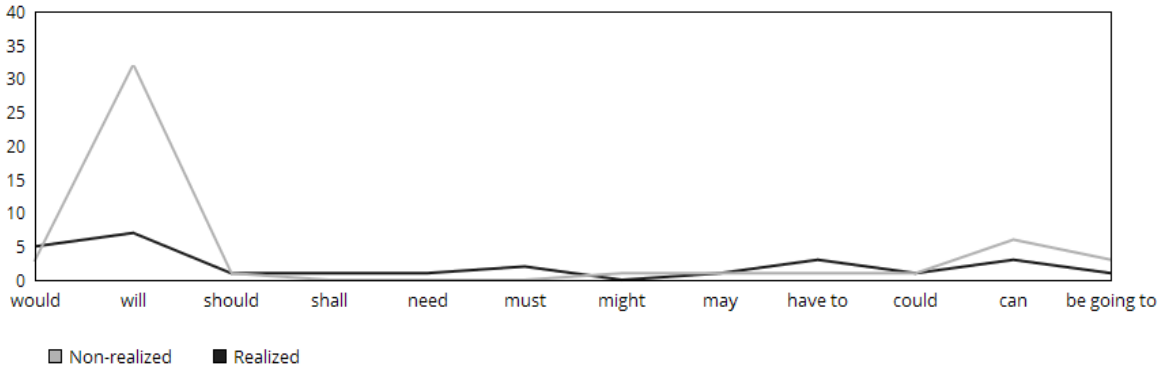

Frequency per 1000 words

Of the many potential points of interest in Figure 8.2, the peak of *will* in the non-realized corpus is plainly conspicuous. From a systemic point of view, this is only half as striking as it may initially appear: four of Lock’s (1996) five modal meanings employ *will* as a lexical resource; only the area of ‘ability’ manages without. However, neither the ubiquity of *will* nor its protean nature helps to explain its outsized presence in the NR corpus particularly. The possible reasons for this, and the further questions it raises, are discussed in the four areas of modality which include *will*, each examined in turn.

Interestingly, when chi-square calculations were performed on the output of Corpus Tool’s automatic tagging, four types of meaning were identified as having a statistically significant difference in their usage across the two corpora: the modality of *likelihood* ($p < .01$) and its daughter node of *certainty* ($p < .05$); and *volition* ($p < .01$) plus its daughter node of *intention* ($p < .05$). Using Lock’s (1996) categories, however, the distribution falls just below significance in all areas. Still, a broad, quantitative look at

these five meanings (Figure 8.3) offers a useful impression of which particular modal spaces occupy the authors' collective attention. Qualitative analyses then follow.

Figure 8.3: Modal Meanings

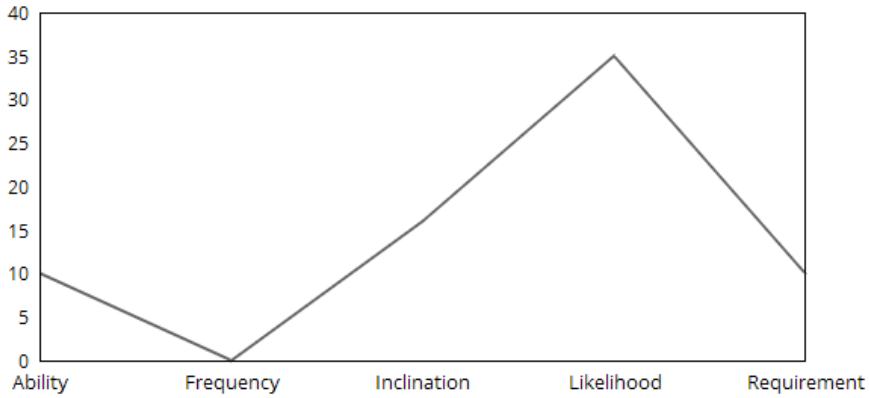

Frequency per 1000 words

Having now come to the tail end of a full Appraisal analysis on these texts, trends identified elsewhere in the data are clearly echoed here. Pledges across the dataset are very interested in the likelihood of certain happenings, for example. This is at least partly a reflection of the future-facing nature of pledges as a genre. Yet, when the data is broken apart by corpus (as in Figure 8.4 below), the preoccupation with likelihood

Figure 8.4: Modal Meanings by Realization Category

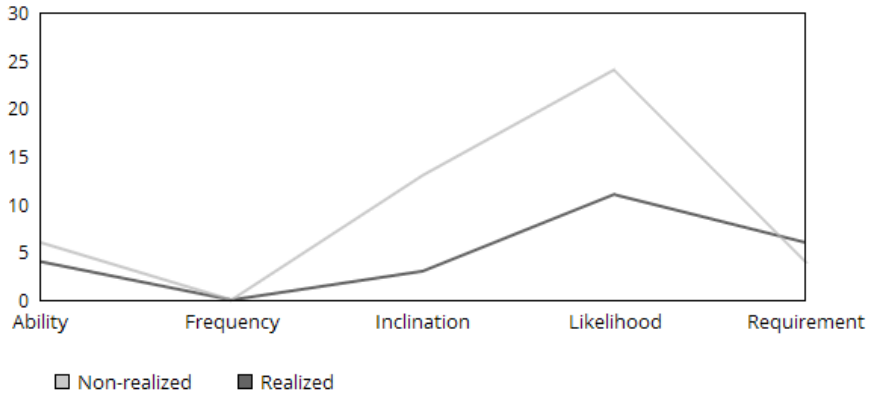

Frequency per 1000 words

more plainly mirrors the trend uncovered in the examination of imagined future violence in capacity (section 6.2.4). Once again, concern with events of the future is simply more

prevalent in the writings of the NR authors. This dovetailing of likelihood with capacity is not entirely coincidental, of course—judgments of likelihood shade into the grammatical future tense through the lexeme *will* (Lock, 1996: 196), which occupies a difficult to parse space as both a predictor of high likelihood and as a future tense marker in English. Thus, while the future tense was being *explicitly* considered in conjunction with violent incapacity, an element of modal likelihood was being *implicitly* considered as well. At the quantitative level, then, this particular pattern has been explored. Qualitative findings on likelihood will be presented in section 8.2 below.

More surprising, perhaps, is the heightened interest in inclination on the part of the NR authors, simply because no evidence exists that these writers were, in fact, inclined to action. This finding has a corollary in a separate but intimately related region of Appraisal: the variable of dis/inclination in the attitudinal subsystem of affect (Martin & White, 2005: 51). As noted in Chapter 4, distribution in this area of attitude falls beneath the level of significance. Nevertheless, a similar tendency is apparent in both resources: NR pledges are far more likely to include both modals of inclination (13 words per 1000 words versus just 3 for R pledges) and attitudinal dis/inclination (14 per 1000 words versus 10). The higher frequency of these auxiliaries is just one more way in which the NR pledges appear more menacing than their realized counterparts.

This interpretation is supported by the results shown in Figure 8.5 below, which offers a snapshot of how much of each kind of meaning is 1) placed in a grammatical future that is 2) under the control of the 1st Person of the author. (Exactly what kind of events and actions the authors profess an inclination towards will be unpacked in section 8.4 below.)

Figure 8.5: 1st Person + Future Tense Modal Meanings by Corpus

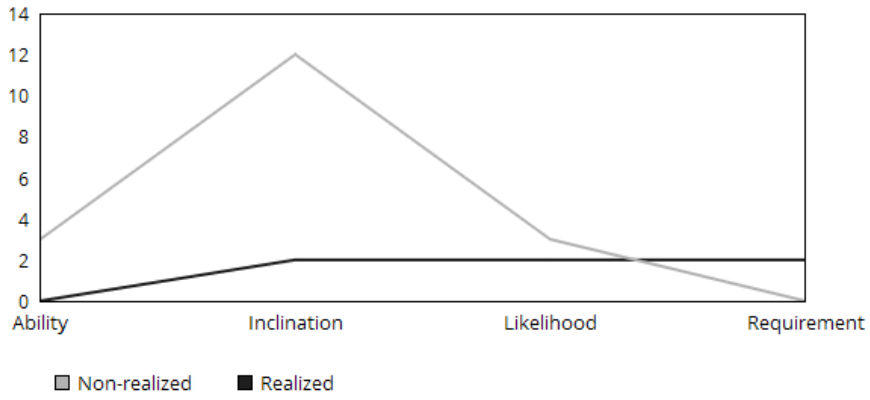

Frequency per 1000 words

Finally, the finding that modals of frequency do not occur at all in either corpora (as shown in Figure 8.4 above) is less than surprising. The lack of occurrence of frequency is actually in keeping with Lock's (1996: 209) observation that English speakers rarely rely on verbal auxiliaries to express these meanings, opting instead for circumstantial adjuncts of time like *always* and *usually*. Because of this, frequency will be omitted from the analysis moving forward.

8.2 LIKELIHOOD

Likelihood is typically considered an epistemic assessment, one which qualifies a speaker's knowledge of, or commitment to, the truth of a proposition (Huddleston & Pullum, 2005). While expressions of past likelihood are certainly possible, e.g., "Someone must have taken the message" (Lock, 1996: 198), judgments of how likely an event is to occur in the future—and a violent event especially—are of primary interest to a threat assessor. A pledge may appear to contain more commitment, for instance, if it expresses the kind of high likelihood found in statements like *the bombings will take place regardless* (LA Unified NR) and *TODAY WYOMING WILL BE TURNED TO DUST* (Archangel Michael NR). The dispersion of likelihood across the dataset is shown below in Figure 8.6.

Figure 8.6: Likelihood Dispersion by Corpus

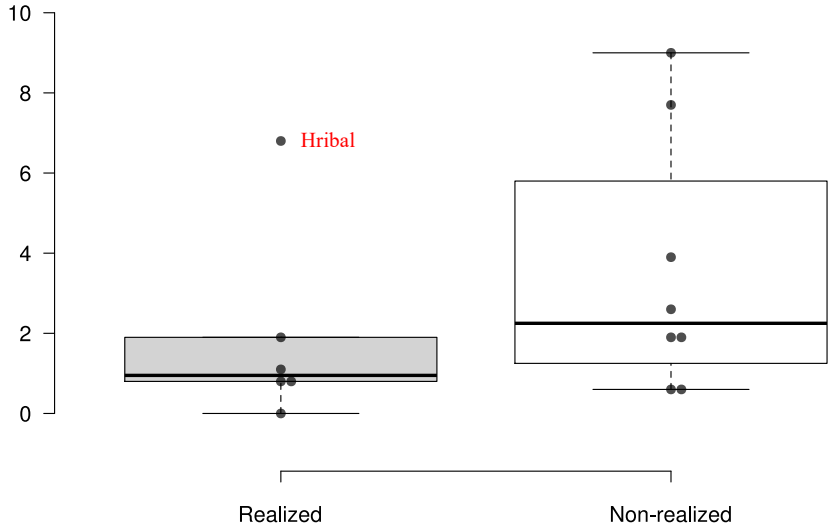

Frequency per 1000 words. One outlier (>1.5IQR) is found.

As mentioned above, predictive functions link tightly with the future tense in English through the lexeme *will*—a systemic overlap which contributes to this word's status as, among other things, the most frequently used modal in the register of conversation (Biber *et al.*, 1999). Thus, in a genre devoted to the likelihood of future violence, *will* would naturally be expected to appear at a higher frequency relative to the other auxiliaries, something which is certainly the case in this dataset. However, the distribution of its usage is interestingly lopsided, as Figure 8.6 shows.

Figure 8.6: Likelihood Modals by Realization Category

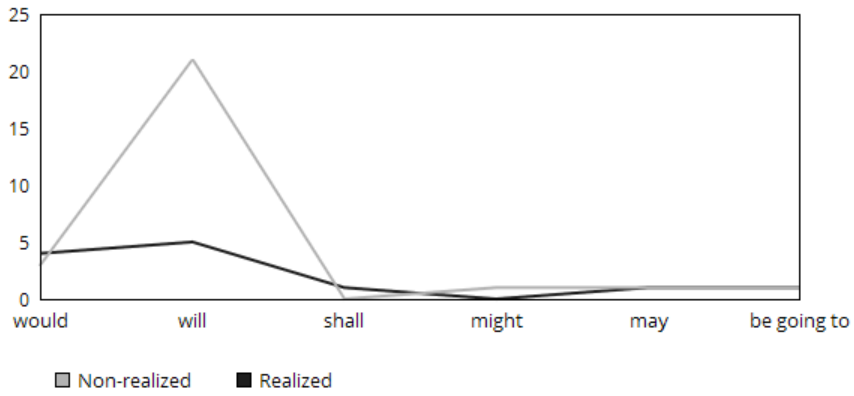

Frequency per 1000 words

This pattern is only accentuated with the removal of Hribal R's text—the longest pledge in the dataset and again flagged an outlier—as we see in the adjusted Figure 8.7.

Figure 8.7: Likelihood Modals by Realization Category (Adjusted)

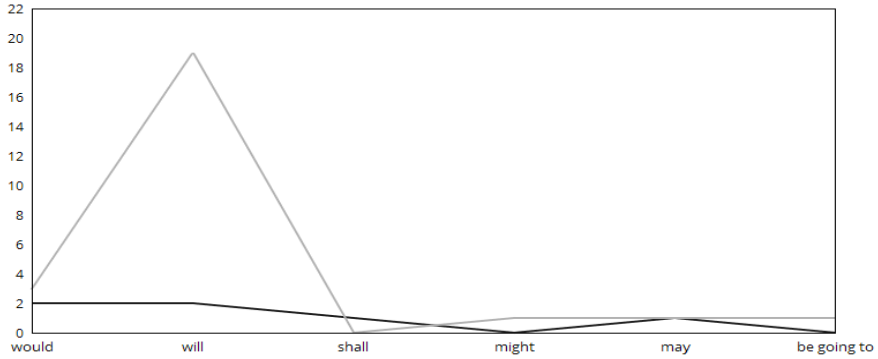

Frequency per 1000 words

Gales (2010, 2011, 2017) has also paid special attention to modality in her investigations of threatening language. She notes that "modals of prediction are the most frequent class of modals and are significant to the category of *non-realized threats* [emphasis added]" (Gales, 2017: 10). While the results of this study do not rise to the level of significance, her finding is clearly echoed here with *will*. Where other lexical choices like *may* and *be going to* occur at roughly equal rates across the corpora, an NR author is almost four times more likely than an R author to reach for modal *will* in expressions of likelihood.

Figure 8.7 also illustrates a tendency found elsewhere in the dataset. In terms of engagement, *will* is highly contractive, closing down the text to outside voices and possibilities. Modal *would*, on the other hand, "is used only in predictions based on unreal conditions" (Lock, 1996: 200), thereby expanding the text to encompass other potential realities. Given this, it is interesting to see that the R texts employ each resource almost equally. The contrast between *will* and *would* in the two realization categories appears to be a microcosm of tendencies discovered elsewhere, in that the more expansive communicative strategy—the one which serves to weaken the author's

apparent commitment—correlates more consistently with historically realized texts, while resources which strengthen an author's stance are most often found in texts whose authors did not proceed to action.

Of course, from a threat assessment perspective a highly relevant question is bound up with likelihood: what is the nature of the harm (Meloy *et al.*, 2014: 3)? The data itself suggests two ways of parsing attack-related modality. The first asks, of all the modal auxiliaries used to predict the likelihood of future events, how many are explicitly attack-related? This would include the usage of modals like *would* in the phrase *I realized that I <u>would</u> have to use violence* (Shaw R) and *will* in *They <u>will</u> be detonated via Cell Phone* (Archangel Michael NR), while excluding these same auxiliaries in statements like *i just thought <u>id</u> share my story* (Skyline NR) and *The 22nd of October <u>will</u> mark the final day of Ramadan as it would fall in Mecca* (Brahm NR). Admittedly, classifying modals in this way requires an interpretive, and thus potentially inconclusive, reading. For instance, McKelvey NR's claim that *the cops <u>wont</u> save you* is considered attack-related here because *saving* means 'saving from' the physical harm being threatened against the Kean University students. Alternately, Brahm NR's conclusion that *General chaos <u>will</u> rule* is not marked as attack-related because this is presented as the aftermath of the truck bombings, and not a state of the world during the bombings themselves.

Even proceeding conservatively, dividing the data this way reveals a stark difference between the corpora, as shown in Figure 8.8 below. The NR writings are eight times as likely to make attack-related predictions. These predictions are also largely up-scaled according to Lock's (1996) tiers of intensity. Of the 29 raw tokens across the entire dataset which are interpreted as being attack-related, 27 communicate high likelihood, or "certainty," in Lock's (1996) taxonomy. These are modals like *will*, *shall*, *would*, etc. Interestingly, the two which fall beneath this uppermost urgency both

Figure 8.8: Ratio of Attack-related Likelihood Modals

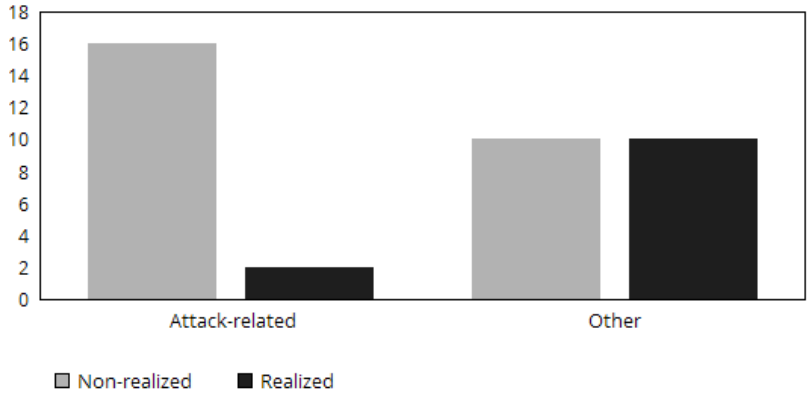

Frequency per 1000 words

appear in the NR corpus: when LA Unified NR says *this may be your last day*, and Dickens NR writes that she *Might kill at least fifteen tomorrow*. Both tokens are rated "low (possibly)" according to Lock (1996). This is, perhaps, just by dint of the NR authors making greater use of this resource. Generally, the non-realized authors are much more likely to present an attack as being highly probable using the semantics of likelihood.

The second angle of analysis involves using the attack itself as a temporal fulcrum, to examine how focused the authors are on events leading up to and including the violence, as well as the states, conditions, or events which are imagined as coming afterwards. This is a similar process to the one pursued under propriety (section 5.3.4), which examined whether the prospective assaults described by the writers were also imagined as setting right whatever injustice gave rise to the ideations. With modality, drawing a bright line between temporal frames is often impossible since the lexis overlaps so widely with tense and aspect. This is especially true of auxiliaries used for the unreal condition, e.g., *the biggest offender of that would definitely be the jocks* (Skyline NR). Attempts to demarcate actions immediately preceding the start of violence from the violence itself are also difficult, even in close discursive readings. For instance, Brahm NR's proclamation that *The explosions will be near simultaneous* is plainly

related to the onset of violence, but LA Unified NR's claim that *Your security <u>will</u> not be able to stop us* is more difficult to push into one frame or another, since it is left unspecified whether the author imagines the security personnel failing in their duties as the *army of Allah* is preparing to attack or as the attack itself is underway. For these reasons, the onset of violence (or the initial imagined event, if a text features a series of violent occurrences) is used as the pivot point, with modals covering the lead-up and the attack grouped together. Meanwhile, modals dealing with the aftermath of the attack are set apart, e.g., *Access to weapons, including guns, <u>will</u> be debated* (Hribal R). Figure 8.9 offers a picture of the data divided this way.

Figure 8.9: Likelihood Modal Auxiliaries by Time Frame

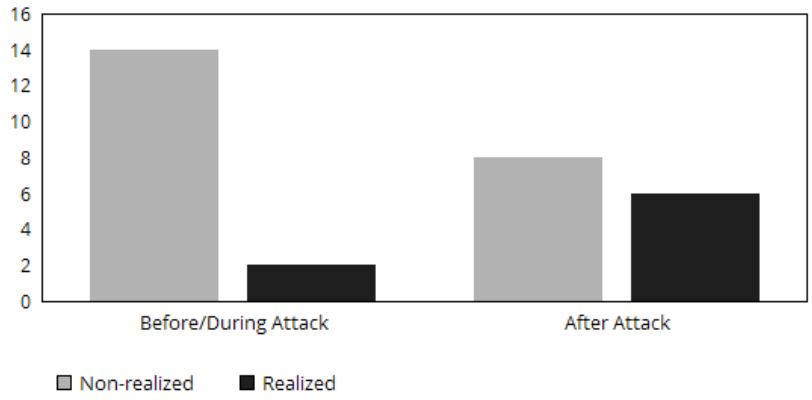

Frequency per 1000 words

First, there is a clear disparity in how focused each realization category is on the lead-up to and start of the attack: the NR writers are far more interested in sharing violent ideations located in this time frame. Interestingly, the texts themselves suggest five topics of concern which appear to be common to all of the NR writings. (These topics appear in the R corpus as well, though at a far reduced rate given the paucity of likelihood modals which address events before and during the imagined attack.) These topics are:

1) Targeted Area/Time: in terms of thematic roles (Saeed, 2009), these modals of likelihood are used essentially to communicate Locative information, of either a physical or temporal nature, e.g.:
    a. *These trucks will pull up to stadiums hosting NFL games* (Brahm NR)
    b. *they are in the middle of the commons so it would be an ideal place to start* (Skyline NR)
2) Means/Method: similarly, in terms of thematic roles, these modals convey the likelihood of certain Instruments being employed in the attack, e.g.:
    a. *their heads will play a major role in the final phase* (Rodger NR)
    b. *The bombs themselves will be delivered via trucks* (Brahm NR)
3) Inevitability: there is no obvious corollary to thematic roles here. Nevertheless, this category covers modals of the highest likelihood which the authors combine with additional semantic elements to characterize the attack as a foregone conclusion, e.g., the combination of *will* and *regardless* in the following statement:
    a. *If you cancel classes, the bombings will take place regardless* (LA Unified NR)
    b. *the cops wont save you…you're black* (McKelvey NR)
4) Victim Injury/Death: these modals predict some manner of physical harm occurring to the imagined victims, e.g.:
    a. *TODAY WYOMING WILL BE TURNED TO DUST!* (Archangel Michael NR)
    b. *But I will really get off on knocking her out* (Valle NR)
5) Attacker Death: only one token addresses the threatener's imagined death during the attack, but the predictive purpose of the modal is clear enough to stand alone:

a. *cant say the fact that im leaving this world tomorrow saddens me at all because ill be doing it with a good cause* (Skyline NR)

As Figure 8.10 shows below, the NR authors make more predictions about events and actions in this pre- and mid-attack frame. However, clear preference is given to predicting some kind of harm to the imagined victims.

Figure 8.10: Pre-attack and Attack Modals of Likelihood

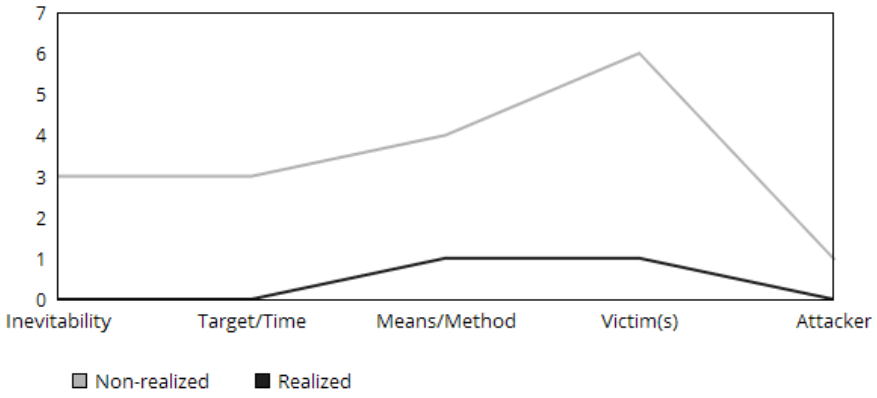

Frequency per 1000 words

Other than forecasting injury to their imagined victims—a trend explored more fully under negative capacity—it is somewhat counterintuitive to find that only NR authors present the attack as inevitable. While the frequency with which this occurs is obviously low, it is intriguing that none of the R pledges engages in this kind of chest-pounding rhetoric, given that the realized authors would later attempt the described attacks.

Additionally, the relatively more frequent mentions of Locatives and Instruments in the NR corpus reignites the question of specificity as a metric of authorial commitment. Turner and Gelles (2003: 99) advise that "identifiable targets…and a well-articulated action plan and time frame suggests considerably more pressure…in regard to the potential for action." And yet, only the NR texts employ likelihood to make predictions about where (*stadiums hosting NFL games*; *the middle of the commons*) and

when (*The 22nd of October*) the attacks are imagined to occur, while also providing moderately more information about how (*detonated via Cell Phone*).

Interestingly, the two tokens of likelihood which appear in the R corpus fall comfortably within the category distinctions driven by the NR data: Shaw R muses on the low possibility of a Means/Method when he writes *I may even take a photo before hitting them*, and Hribal R makes a stronger prediction of his victims' deaths when he claims that *their previous lives are going to be taken by the only one among them that isn't a plebeian*. Other than these two statements, though, the R authors do not concern themselves with the likelihood of events or actions in the lead-up or execution of an attack.

As Figure 8.10 above shows, the R corpus is comparatively more interested in making predictions about what follows the imagined violence. This post-attack time frame may also be broken down into topical categories suggested by the data:

1) Aftermath: likelihood modals which cover states and events following the attack, e.g.:
    a. *By starting an independent civil war where I will hit over a million Asian Women in the face with a stick will change history* (Shaw R)
    b. *Global economies will screech to a halt* (Brahm NR)
2) Satisfaction: this topic has an affective element indicating that the author will feel a measure of gratification from completing of the attack, e.g.:
    a. *Things will be fair once I make them suffer* (Rodger NR)
    b. *tomorrow ill make sure that they wont ever abuse them again* (Skyline NR)
3) Attention: these modals foresee some level of recognition of the exceptional nature of the attacks, from the media, the government, or society in general, e.g.:
    a. *It will make national headlines* (LA Unified NR)

b. *a dozen different things <u>will</u> be speculated to be at fault* (Hribal R)

The frequencies of these are shown in Figure 8.11 below.

Figure 8.11: Post-attack Modals of Likelihood

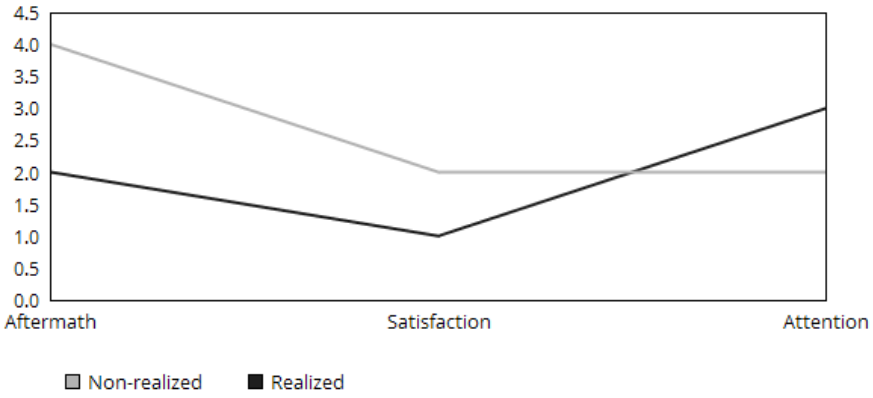

Frequency per 1000 words

Before addressing the differences presented here, however, two texts are potentially problematic. Brahm NR's text contains all six of the tokens speaking to the attack's aftermath, and two of the three related to attention. Meanwhile, seven of the eight tokens of attention are found in Hribal R's pledge. With these two texts removed, the numbers are less dramatic, as seen in the adjusted Figure 8.12.

Figure 8.12: Post-attack Modals of Likelihood (Adjusted)

Frequency per 1000 words

While the frequency of the remaining tokens is obviously quite low, the distribution of the topics in the corpora is worth some small comment. Particularly, the

NR texts are ever more slightly focused on the imagined satisfaction of hurting or killing the victim(s), and the attention the author's actions will receive. Meanwhile, a concern for the ramifications of the violence is now only present in the R writings, but these tokens have no semantic tinge of vindication. These results mirror those uncovered by propriety in section 5.3.4—that the realized authors are relatively disinterested in whatever peace of mind they may find after the fact—and are likely bound up with the same analytical considerations, e.g., the future tense.

Taken by itself, the varying use of the modality of likelihood by the two author types could argue for or against any number of hypotheses. For instance, the heightened penchant for predicting harm to the imagined victims in the non-realized writings could work in support of a "safety valve" theory (Gellerman & Suddath, 2005: 485), where socially unacceptable feelings are expelled via the act of their expression (Jackson, 1981). However, the fact that the contrast between the modulation of likelihood and the modalization of inclination (discussed below) turns in large part on the absence of authorial Agency could just as well argue *against* the safety valve effect here. Meaning, likelihood is not the area where an author would discuss *personally* inflicting the imagined death and destruction, only that death and destruction is likely. All that may be comfortably posited from this investigation of likelihood is that, here again, the NR pledges predict more future violence, thereby constructing a prosodic profile which is more threatening than their realized counterparts.

8.3 ABILITY

Modals of ability assess "internal abilities and skills that make it possible for a certain action to be performed or situation to come about" (Lock, 1996: 211). As with likelihood, the semantic content of this area is largely propositional, in that such expressions assert,

deny, or question the validity of a claim (Halliday & Matthiessen, 2004). With ability modals, this claim is that some entity has the wherewithal to accomplish a task.

The modal resources for encoding these meanings in English is actually quite small, essentially limited to the lexemes *can* and *could*. A speaker's range of expression is further limited by the observation that these two terms offer no gradations of intensity of the kind found in likelihood (Lock, 1996: 213). Although, of course, other mitigating language is available, such as how the verb *think* is used to down-scale Dickens NR's statement *I think I <u>can</u> pull it off*. Still, the ability of a threatener to carry out the described actions is one of the primary metrics used to assess a threat's credibility (Bulling & Scalora, 2013; Calhoun & Weston, 2015). And so, when claims like *I <u>can</u> knock her out* (Valle NR) appear in a pledge to harm, they may carry additional weight in determining whether a threat is considered "real" by a practitioner or merely "hypothetical" (Bulling & Scalora, 2013: 4). The dispersion of ability across the dataset is shown below in Figure 8.13.

Figure 8.13: Ability Dispersion by Corpus

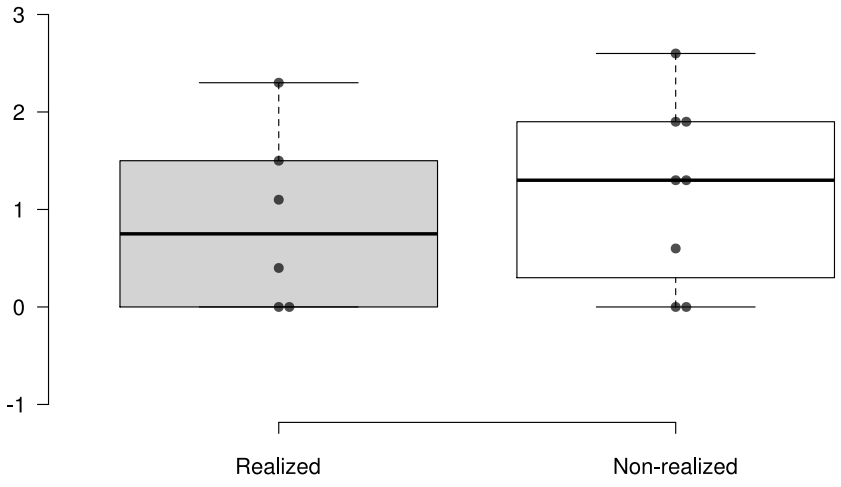

Frequency per 1000 words. No outliers (>1.5IQR) are found.

The frequency across the dataset of the two auxiliaries and their negations is shown in Figure 8.14.

Figure 8.14: Ability Modals by Realization Category

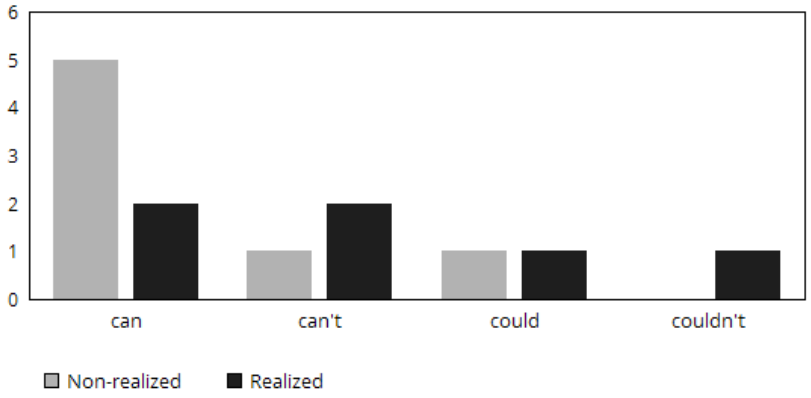

Frequency per 1000 words

At the lexical level, the NR texts include more mentions of ability, in whatever context, including the positive polarity of *can*. Once again, however slightly, these authors work to make their pledges potentially more menacing than their R counterparts. This question is explored more fully through Table 8.2 immediately below. Of more minor interest, each token of *could/n't* in the dataset, across both the R and NR corpora, is used to communicate an unreal condition rather than past ability, e.g. *I will [perform on them] any other form of torture I could possibly think of* (Rodger NR), *[It] would be too much…They couldn't live with themselves* (Kinkel R), and *I could say life is evil* (Hribal R). The normed count for these is exceedingly low but worth mentioning, since this is one of the only areas in which the NR and R texts stand at a relative parity in their employment of the irrealis mood.

Of course, an assertion of ability is not menacing by itself. A natural, subsequent question for a threatening communication is whether the threatener is "capable of carrying out an act of violence" (O'Toole & Smith, 2014: 273). When the tokens are flagged for the grammatical 1st Person—thus addressing the 'who'—and a discursive reading is applied to which tokens are attack-related (A-R)—thereby getting at the 'violence'—a picture begins to emerge. This sorting is shown in table 8.2.

Table 8.2: Concordance Lines of Ability Modals

Corpus	Author	Pre-text	Aux.	Post-text	1st P	A-R
Non-realized	Dickens	I think **I**	can	pull it off. Might kill at least fifteen tomorrow	✓	✓
		So when **you**	can	absolutely show me in the 1st amendment where it explicitly says		
	LA Unified	And there is nothing **you**	can	do to stop it.		✓
	Rodger	Silently killing as many people as **I**	can	around Isla Vista	✓	✓
		as well as any other form of torture **I**	could	possibly think of.	✓	✓
	Skyline	i will use to take my life if **i**	can	make it up to that point. the people at	✓	✓
		my story before **i** head to bed tonight.	cant	say the fact that im leaving this world tomorrow saddens	✓	✓
	Valle	["Victim-1"], **I**	can	just show up at her home unannounced, it will	✓	✓
		it will not alert her, and **I**	can	knock her out, wait until dark and kidnap her	✓	✓
				Total	7	8
Realized	Hribal	So **life**	can	be more enjoyable,		
		A world where most **people**	can	only find happiness or enjoyment in doing drugs, drinking		
		such a terrible state of affairs, **it** actually	CAN	get worse.		
		me Siddhartha, a wandering student of many that	can	only find himself when he meets a god almost no	✓	
		Absolutely Disgusting. **I**	can't	wait to see the priceless and helpless looks on the	✓	✓
		helping fellow human beings. **I**	could	say life is evil, because it blocks everyone in	✓	
	Kinkel	My **parents**	can't	take that! It would destroy them.		

		I destroy everything I touch. **I**	can't	eat. I can't sleep.	✓	
		I can't eat. **I**	can't	sleep. I didn't deserve them. They were wonderful	✓	
		embarrassment would be too much for them. **They**	couldn't	live with themselves. I'm so sorry. I am		
	Long	**we** the people	cannot	differentiate the the good from the bad.	✓	
				Total	6	1

Looking at the distribution of the raw tokens, five of the eight NR authors employ a modal of ability, while three of the six R authors do the same. The occurrences in the realized writings, however, tend to coalesce around two writers in particular—Hribal R and Kinkel R—making this potentially a question of style. The other four realized authors make use of these modals either sparingly (once by Long R) or not at all. However, the disparity in how modals of ability are used is remarkable. When the above tokens are normed and grouped together according to 1st P + A-R, a stark contrast between the two pledge types is immediately evident. This updated distribution is shown in Figure 8.15.

Figure 8.15: 1st Person Attack-related (A-R) Modals of Ability by Corpus

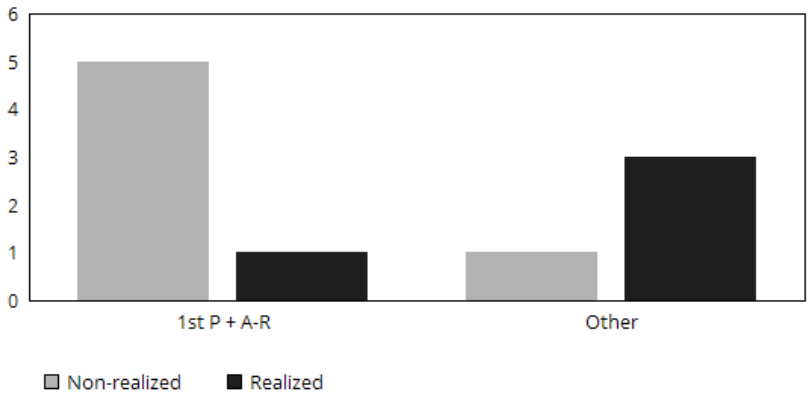

Frequency per 1000 words

NR authors are five times more likely than the realized authors to present themselves as *personally* able to carry out the pledged violence, and five times more likely to focus on this use to the exclusion of Other judgments of ability. This imbalance helps create a prosody of confidence, an effect which may rightly be interpreted by a threat assessor as raising the perceived danger a pledge writer poses to his or her imagined target.

Interestingly, only one of the 11 tokens present in the R texts is coordinated with a 1st Person pronoun while addressing a violent event, when Hribal R writes *I can't wait to see the priceless and helpless looks on [their faces]*. Semantically, this is similar to Skyline NR's statement *[i] cant say the fact that im leaving this world tomorrow saddens me at all*, in that both are less about an actual capacity for violence and more a statement of the author's affective stance towards the imagined attack. The use of ability modals in the R pledges to communicate a capability for violence is thus rare to the point of vanishing.

Looking at the category of Other, a qualitative difference is apparent between the corpora in the modals which somehow fall outside the scope of a 1st Person subject. The two tokens like this in the NR writings are both bald, on-record challenges to the abilities of the pledge's intended audience: *So when you can absolutely show me in the 1st amendment where it explicitly says you can't say "kill all cops", then I'll delete my status* (Dickens NR); *And there is nothing you can do to stop it* (LA Unified NR). Thus, when not employed for communicating violent ideation, modals of ability nevertheless serve a general sense of combativeness on the part of the NR writers.

In the R corpus, by contrast, these auxiliaries appear primarily in observations of either a philosophical bent, e.g.:

- *In such a terrible state of affairs, it actually CAN get worse* (Hribal R)
- *So life can be more enjoyable* (Hribal R)

- *we the people cannot differentiate the good from the bad (cops)* (Long R)

Or as a comment on a more personal state-of-being, like when Kinkel R laments that *I can't eat. I can't sleep* and worries about what his parents will think of his actions (despite the fact that he has already shot and killed both at the time of the pledge's writing). Not only is a professed ability for violence lacking in the R texts, then, the antagonistic edge is also missing. Taken altogether, this area of modality amplifies a prosody of menace in the NR texts uncovered elsewhere in the broader analysis, a prosody that is decidedly absent in the R corpus.

8.4 INCLINATION

Modal meanings of inclination are theoretically crucial to understanding how (or whether) psychological intent is encoded linguistically. Lock (1996), for example, includes the quality of *intention* as the middle gradation within a larger semantic space that includes *determination* at the high end and *willingness* at the low. Key to this modulation of meaning is that it is inherently 'intrinsic', i.e., it involves some kind of human control over events (Quirk, 1985). It is therefore a more specialized resource compared to likelihood and ability, modals which are just as comfortably used with animate and inanimate subjects alike. When the events in question are violent and the control rests with the human person of the writer, the resulting proposal is highly likely to be construed as a threat. And so, more than authors saying they *can* do violence to a third party, or that violence is *likely* to befall a third party, authors claiming that they *will* do violence to a third party is very close to the prototypical expression of intent, and the core definition of a pledge to harm. The dispersion of inclination across the dataset is shown below in Figure 8.16.

Figure 8.16: Inclination Dispersion by Corpus

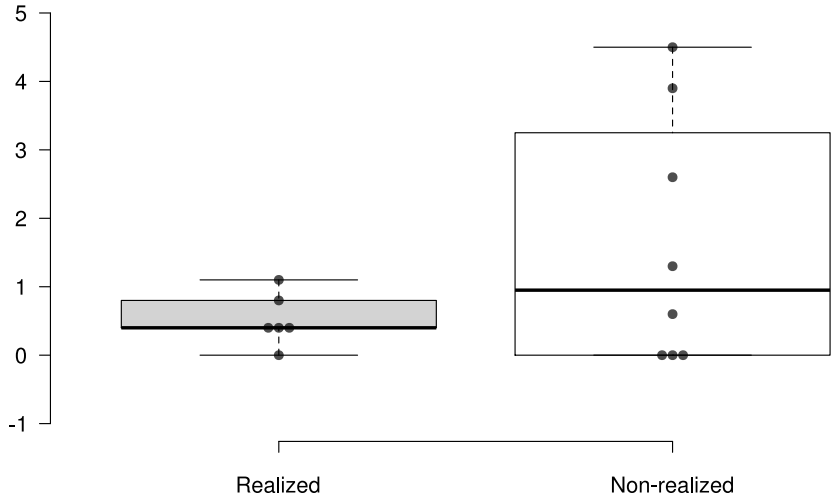

Frequency per 1000 words. No outliers (>1.5IQR) are found.

As with ability, an English speaker's lexical resources are somewhat constrained in this area, as "only a few modal auxiliaries are used for inclination" (Lock, 1996: 210): *will*, *would*, *shall*, and their respective negations. As noted above, forms of *be going to* are added to this tally, since they are commonly used to express intention in English (Bybee, 2014), as in *i am going to open fire on the people in the commons* (Skyline NR). The distribution of these auxiliaries is shown in Figure 8.17.

Figure 8.17: Modal Auxiliaries of Inclination

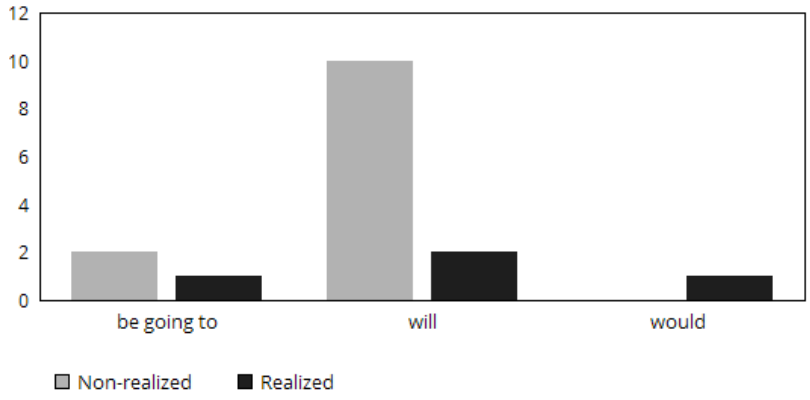

Frequency per 1000 words

The main lexical resource is clearly modal *will*, as in McKelvey NR's declaration *i will kill all the blacks tonight*. *Would* as inclination makes just one appearance in the

dataset, when Hribal R writes *I <u>would</u> just rather say*. *Shall* does not appear as a marker of inclination in any text. The fact that modal *will* has a heavy presence here is not surprising, since it is the lexeme of choice for this kind of meaning. What is more surprising is its heightened frequency in the NR writings: a non-realized author is five times as likely as a realized author to reach for *will* to communicate inclination, despite the theoretical *lack* of inclination on the part of these writers.

Asking of these tokens what a threat assessor would—is the author announcing a personal inclination toward future violence—reveals sharper distinctions between the two realization types, as seen in Figure 8.18 below. As shown, the NR writings are not only more likely to engage with these modal meanings, but to do so in service of communicating violent intent. In fact, just two of the six realized writers use modal *will* this way—Rodger R and Shaw R—compared to four of the eight non-realized—LA Unified NR, McKelvey NR, Rodger NR, and Skyline NR. Simply put, there is more inclination present in the NR corpus, and far more violent inclination.

Figure 8.18: Attack-related Inclination

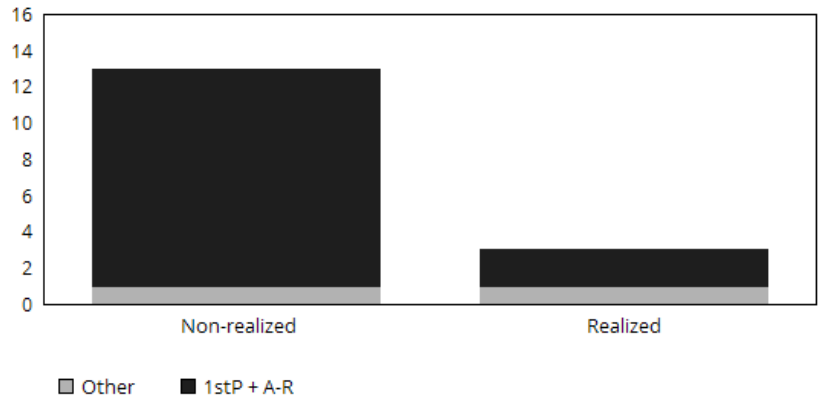

Frequency per 1000 words

Be going to provides an even more interesting qualitative image. Lock (1996) argues that the choice between *will* and *be going to* betrays a subtle difference in a speaker's decision-making process:

> *Going to* usually implies that the intention already exists, that is, that the speaker has already made up his or her mind to carry out the action. The use of *will*, on the other hand, often implies that the decision is made more or less at the moment of speaking. (Lock, 1996: 210)

Whether this holds for written utterances is unclear. But if so, the implications are provocative. In her wider investigation of threatening language, Gales (2010: 184) found that forms of *be going to* were used predictively at roughly four times the rate in her sub-corpus of non-realized threats. In NR pledges, *be going to* occurs just twice as often. But the picture is not as simple in this area as it is in other areas concerning the heightened menace of NR writings. While all three raw tokens in the NR corpus are violent, they are all found in the pledge of a single author:

- *tomorrow i am going to be taking my fathers erma smg* (Skyline NR)
- *i am going to open fire on the people in the commons* (Skyline NR)
- *i am going to start killing them first because they deserve it the most* (Skyline NR)

The two times a R pledge uses *be going to* in order to express inclination, neither conveys violent intent:

- *But am I honestly going to tell you I am a victim?* (Hribal R)
- *I'm going to to talk to a few more Asian Women* (Shaw R)

If Skyline NR is sidelined as a stylistic outlier, then this form disappears from the corpus entirely as an expression of personal inclination. Modal *will* is universally the resource of choice for NR pledges. If Lock (1996) is correct, that "implies that the decision is made more or less at the moment of speaking" (Lock, 1996: 210), then the choice of *will* over *be going to* may betray an impulsivity underlying the creation of these texts. This provides an interesting, albeit minor, data point arguing in favor of the "safety valve" hypothesis (Gellerman & Suddath, 2005: 485) for NR writings.

What may be said with certainty is that the NR pledges are six times more likely to propose future violent actions, and to invest the validity of these proposals squarely in

the Subject of the author. The prosodic effect of this, once again, is to bolster the appearance of actual intent through the language of inclination, despite a theoretical absence of psychological intent on the part of the non-realized authors. If the proposed disconnect between linguistic and psychological intent is true, then these results imply that modal auxiliaries of inclination are *not* reliable indicators of real-world intent in practice, despite being a prime systemic resource for expressing intent in the grammar of English.

8.5 REQUIREMENT

The deontic area of modality which Lock (1996) calls *requirement*, and is elsewhere referred to as *obligation* (Gales, 2010; Halliday & Matthiessen, 2004) or *necessity* (Biber *et al.*, 1999), covers the broad area of meaning between 'do it' and 'don't do it' (Lock, 1996: 204). In pledges, whose core generic distinction is a proposal to personally harm a third party, the value of understanding what an author feels obligated to do or not do is obvious. This is especially true given the previous finding, discussed in propriety, that feelings of obligation may signal a higher risk of dangerousness on the part of a pledge author.

The predictive value of requirement has been recognized in the analysis of threatening language more generally. As Gales (2010: 38) observes, "a commitment to the intended action through modals of obligation" is a linguistic feature commonly associated with threatening language by both scholars and practitioners. More typically, this association is left unstated in the threat assessment literature but still implicitly recognized at relevant moments. In the 73-word threat mocked up by Turner and Gelles (2003: 99) and quoted at the beginning this chapter, for instance, half of the modals the authors choose to include express requirement, e.g.: *have to suffer*; *Mr. Jones must be removed by force*. Their invention is obviously anecdotal, i.e., not an authentic forensic

text, yet this text points to a prevailing notion in the literature that the presence of obligation in violent communications is somehow important. The dispersion of requirement across the dataset is shown below in Figure 8.19.

Figure 8.19: Requirement Dispersion by Corpus

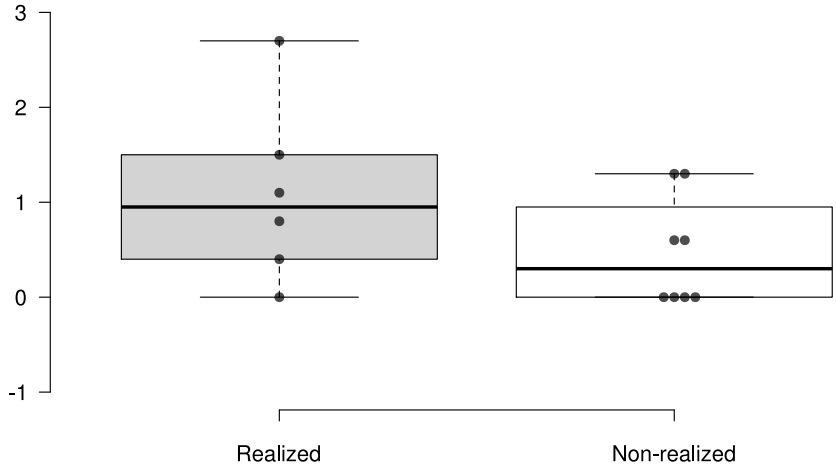

Frequency per 1000 words. No outliers (>1.5IQR) are found.

The ways in which this area of meaning is employed by pledge authors are various. For instance, Hribal R often frames his pontificating in terms of necessity, as in *All good things must end*, and *You don't have to endure this harsh and evil world*. Elsewhere, Dickens NR places the burden of obligation on her readers with her challenge to *show me in the 1st amendment where it explicitly says you can't say 'kill all cops', then I'll delete my status*.

From a numeric perspective, the overall presence of requirement in these texts is low compared to, say, judgments of likelihood or inclination: requirement of some stripe appears at a frequency of 3 words per 1000 in the non-realized pledges versus 6 words in the realized. The particular lexical items used by the two author types are shown in Figure 8.20 below. While the disparity in lexical frequency is small, at a difference of just over 3 words per 1000, certain key differences are apparent in how these modals are put to work in the two pledge types. The first difference is perhaps the simplest. Of all five of Lock's (1996) categories, this is the *only* modal resource which is more common to

Figure 8.20: Modal Auxiliaries of Requirement

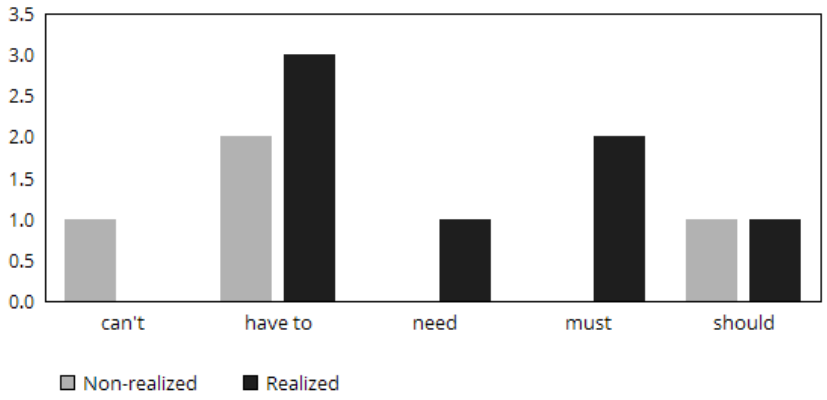

Frequency per 1000 words

the realized texts than the non-realized, even if the margin is slim. Second, when the lexemes in Figure 8.20 are grouped according to their intensity—i.e., whether the term is graduated towards obligation (high), advice (mid), or permission (low) (Lock, 1996: 213)—the realized authors are seen to employ stronger senses of obligation at twice the rate of their counterparts (Figure 8.21).

Figure 8.21: Levels of Requirement

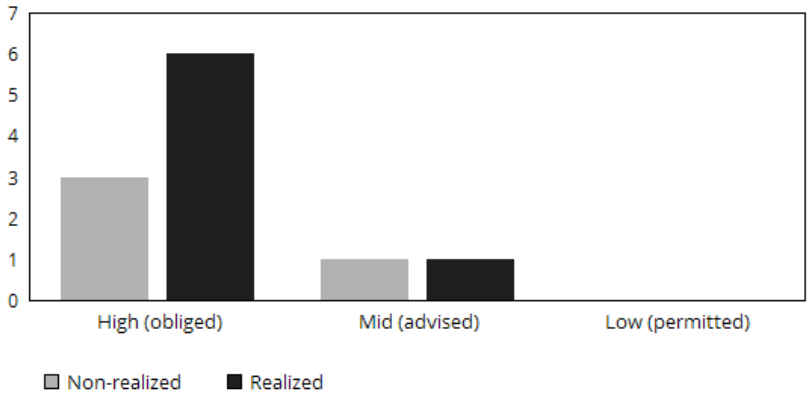

Frequency per 1000 words

Beyond this, the most crucial differences between the two realization categories are discursive. When the question turns to violence, the realized authors are three times as likely to represent this violence as somehow obligatory, as Shaw R does, for

example, when he claims *I had to punch a White dude in the mouth for kicking me*. The difference between the corpora is shown in Figure 8.22.

Figure 8.22: Required Violence by Corpus

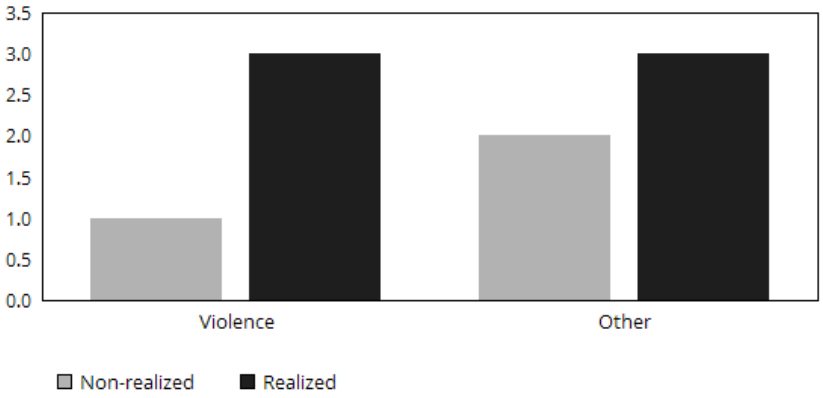

Frequency per 1000 words

However, Hribal R's text—again, the longest in the dataset—is an outlier here, although only regarding the Other uses of requirement modals. Seven of the eight tokens which fall in this non-violent category appear in Hribal R's writing. Figure 8.23 adjusts for this.

Figure 8.23: Required Violence by Corpus (Adjusted)

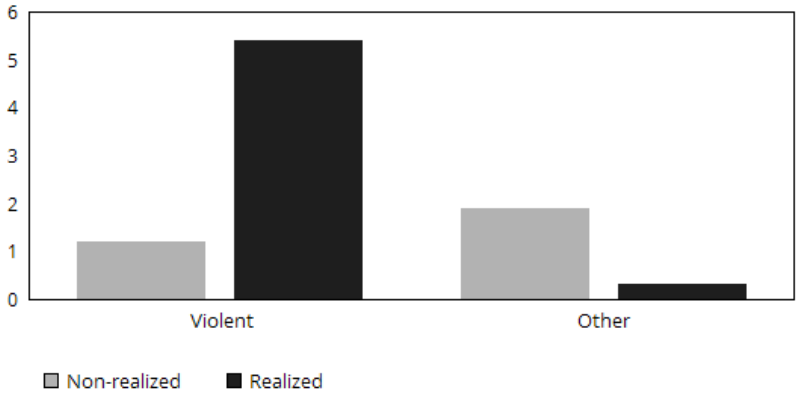

Frequency per 1000 words

Stylistically, Hribal R most often uses this area of meaning to philosophize, e.g., *You don't have to live*, as well as for textual metafunctions, e.g., *It should be noted*. With his text removed, though, the contrast between the two corpora in this regard actually

grows starker, as seen in Figure 8.23 above. Absent Hribal R, the realized writings use requirement almost exclusively to present violence as being somehow necessary or obligatory. The single remaining token of Other in the realized corpus appears in the pledge of Roof R, and is not used ideationally but textually, and even here a typo[38] leaves open to interpretation whether the statement qualifies as a use of requirement at all: *Unfortunately at the time of writing I am in a great hurry and some of my best thoughts, actually many of them have been to be left out and lost forever*. Potentially, then, with Hribal R removed and Roof R's typo re-interpreted, the realized pledges employ modals of requirement *only* for the purpose of communicating how necessary the authors judge various violent actions. Compared to the other areas of modality, this single-mindedness on the part of the realized authors is unique.

Finally, a close contextual reading of the remaining tokens reveals two uses of this modality which appear to be exclusive to the respective corpora. First, and more minorly, only in the NR pledges are tokens of Other used to complain about the author's life conditions in a way which justifies the imagined attack:

- *This First Phase will represent my vengeance against all of the men who have had pleasurable sex lives while I've had to suffer* (Rodger NR)
- *im tired of the pretentious faggots (and life in general) that i have to deal with on a daily basis at my high school* (Skyline NR)

The second pattern of usage is potentially more important, despite the limited data. When the 1st Person of the author is added as context, four of the six R texts are

[38] Roof R's statement shows evidence of misediting, leaving two competing auxiliary constructions—*have to be left out* and *have been left out*—grafted within his single clause *have been to be*. Based on the presence of infinitival *to* and the doubling of the auxiliary *be*, this token has been interpreted as *have to be*, and thus counted as a modal of requirement.

shown to use modals of requirement to mark the author's stance toward his or her own future violent actions:

- *But I have to kill people.* (Kinkel R)
- *Therefore I must bring the same destruction that bad cops continue to inflict upon my people, upon bad cops as well as good cops* (Long R)
- *Well someone has to have the bravery to take it to the real world, and I guess that has to be me.* (Roof R)
- *… I realized that I would have to use violence in order get the response that I desire. … Every Asian Woman by herself must be hit in the face.* (Shaw R)

This particular combination of 1st Person subject + requirement modal + proposal of violence appears nowhere in the NR pledges. The Dickens NR and Valle NR tokens come closest:

- *All Black ppl should rise up and shoot at every white cop in the nation starting NOW.* (Dickens NR)
- *The abduction will have to be flawless . . .* (Valle NR)

However, what is being required and by whom is unlike the bluntness of the R texts. Although Dickens, as an African American woman, certainly falls within the scope of her address (*All Black ppl*), her intended audience is essentially global, diffusing the responsibility for violent action well beyond her own person. In the case of Valle NR, the sense of requirement does not apply to the act itself (*The abduction*) but rather to how the act should be executed. Furthermore, the abduction is merely proposed as preliminary to the larger bad act of cooking and eating the abductee, and is therefore not the central violent event of Valle NR's text. In comparison, the actions described by the four realized authors quoted above are central to each pledge—arguably, the reason for communicating the pledge at all. That 2/3rds of the R authors characterize the core

proposals of their pledges as obligatory, while the NR authors never do, would seem to support both the intuitions of the student respondents as well as the conclusions of the practitioners and scholars cited by Gales (2010), not to mention the implicit attestations of threat assessors like Turner and Gelles (2003).

8.6 SUMMARY

In a few ways, the prosodic picture painted by the use of modal auxiliaries across the two pledge types is counterintuitive. A non-realized text, written by an author who would not later attempt the described ideations, is nevertheless:

- eight times more likely to predict that violence will occur;
- five times more likely to depict the author as personally able to carry out this violence; and
- six times more likely to express an inclination or intention to enact the imagined violence.

However, only realized pledges use modal auxiliaries to characterize the core threat of the text as somehow required. Grouping the four modal types into the three broader semantic categories put forward by Huddleston and Pullum (2005)—epistemic, dynamic, and deontic—helps bring these results into some kind of interpretive focus. Through likelihood, the non-realized texts make greater use of epistemic modality, meanings which "qualif[y] the speaker's commitment to the truth" of a proposition (Huddleston & Pullum, 2005: 52). These texts also use a higher number of dynamic meanings, which "generally concern[] the properties and dispositions of persons" (Huddleston & Pullum, 2005: 52) and encompass notions of ability and willingness (e.g., inclination). The only area in which the realized texts outstrip their counterparts—requirement—is instead deontic, having to do with permission, obligation, and prohibition. And while deontic meanings do appear in the non-realized corpus, only the

realized authors characterize themselves as personally duty-bound to carry out the central violent proposal of their pledges.

It is not immediately clear why, in either pledge type, there appears to be a mismatch between the use of linguistic resources of intent and the subsequent behavior of the authors. However, Halliday and Matthiessen (2004) offer one base for hypothesizing. They posit that a contradiction lies at the heart of modal expressions: speakers only reach for the most confident-sounding language when they *lack* full confidence in their judgment. "The importance of modal features in the grammar of interpersonal exchanges lies in an apparent paradox on which the entire system rests—the fact that we only say we are certain when we are not" (Halliday & Matthiessen, 2004: 624-625). The authors elaborate on this point in a way which, interestingly, speaks to both the linguistic and psychological questions at hand:

> If unconsciously I consider it certain that Mary has left, I say, simply, *Mary's left*. If I add a high value probability, of whatever orientation, such as *Mary's certainly left, I'm certain Mary's left, Mary must have left*, this means that I am admitting an element of doubt—which I may then try to conceal by objectifying the expression of certainty. (Halliday & Matthiessen, 2004: 625)

In other words, each up-scaled modal is an inadvertent admission of doubt, however niggling. If this assessment is correct, then modality supporting an author's personal willingness and ability for violence may appear more frequently in the non-realized corpus precisely because the authors are *less willing* and *less able* to realize their ideations. In other words, "the lady protests too much, methinks[39]." Just who these authors are working to convince of the truth of their claims—whether it is their audience, or perhaps even themselves—is unknowable and likely irrelevant in the final tally.

This paradox—certainty signaling doubt—may also be applied to the heightened appearance of requirement in the R texts without undermining the current hypothesis.

[39] *Hamlet*, Act III, Scene ii, by William Shakespeare

On a variety of fronts, the R writings have been weaker than the NR writings (e.g., their penchant for dialogic expansion, analyzed in Chapter 7), and that pattern arguably holds here. Halliday and Matthiessen's (2004) observation concerns the strongest of modal meanings, the ones most closely approaching a full 'yes/no,' 'do it/don't.' Although the modals used in the R pledges to express deontic qualities are of a high semantic intensity (*must*; *have to*), the stance they communicate is actually pragmatically weaker than the epistemic and dynamic formulations found in the NR writings. Why this is so lies in the speaker's relationship to epistemic and dynamic meanings on the one hand, and deontic meanings on the other.

Strong statements of likelihood, inclination, and ability, all carry a burden for the author making them. For epistemic predictions of high certainty, the risk lies beyond the text, in "the happenings and conditions of the world" (Halliday & Matthiessen, 2004: 24). When Archangel Michael NR asserts that *TODAY WYOMING WILL BE TURNED TO DUST!* the author is inviting the possibility of being proven wrong by the real-world events which follow the publication of the pledge. This bald, on-record statement of modal likelihood is therefore highly threatening to the author's negative face. For an author who asserts the ability and inclination to personally perform an action, the burden is instead one of agency. When McKelvey NR makes the dynamic claim that *i will kill all the blacks tonight*, she invests herself (as *i*) with the responsibility of ensuring that the threat is valid[40]. Thus, when likelihood, inclination, and ability are used to up-scale a proposition or proposal whose validity is somehow invested in the person of the author,

[40] Whether the author is writing anonymously—as both Archangel Michael NR and McKelvey NR were at the time—is functionally irrelevant. The authorial responsibility for the validity of the claim remains, by virtue of the modals themselves. Similarly, that the NR authors historically failed to follow through does not negate the fact that, linguistically, each willingly picked up this yoke via their pledge.

these three areas carry with them the damaging possibility of blame or accountability if the claims are later proved invalid.

By contrast, authors who contend that they are personally required to do something are *abdicating* this responsibility. Someone may fail in their duty, of course, and be held accountable for it, but duty, by its nature, is an expectation handed down by a higher authority. (Whether any real duty exists in the R pledges in a way which society would recognize as legitimate, or is merely imagined or concocted by these writers, is also irrelevant—the fact that the claim has been made is all.) Thus, the assertion *I <u>have to</u> kill people* (Kinkel R) is strong semantically but paradoxically weak pragmatically. Kinkel R is placing the onus for his behavior somewhere beyond his own person. Similarly, when Long R says *I <u>must</u> bring the same destruction that bad cops continue to inflict upon my people, upon bad cops as well as good cops*, he is placing the blame for his own actions implicitly on the circumstances created by the *bad cops* and not explicitly on a personal inclination to harm anyone. In a roundabout way, this is not terribly different than Martin and White's (2005) conception of dialogic expansion (though, of course, strong modals of requirement are technically contractive). I.e., the realized authors co-opt an outside, often unnamed, social force whose dictates they have no choice but to follow.

Taken altogether, then, modal auxiliaries show a similar pattern of stancetaking uncovered elsewhere in this analysis. The non-realized writers again present themselves as more capable of and more inclined to violence, as well as more certain violence will occur. Meanwhile, the realized writers' use of requirement is interestingly ambiguous: simultaneously strong, in that the authors present themselves as highly compelled to act; and weak, in that this compulsion is placed on them by something beyond themselves, a moral pressure which they are powerless to resist.

Finally, modals may indeed be one measure of an author's 'action imperative,' i.e., "the need on the part of the person to take personal action" (Turner & Gelles, 2003: 97). However, if these results are at all indicative of larger trends of usage in threatening language, then 'forceful modals' (Gales, 2010: 96), taken broadly, may actually communicate very little about the presence of psychological intent. Instead, what Mardigian (via Gales, 2010: 26) identified as "modals of intent"—particularly *must* and *have to*—may, in fact, be more reliable markers.

CHAPTER 9 GRADUATION

The third and final system of Appraisal, graduation, attends to the ways that speakers "present themselves as more strongly aligned or less strongly aligned with the value position being advanced by the text" (Martin & White, 2005: 94). The central idea of graduation is that stancetaking is not merely a binary choice between polarities—e.g., 0 or 1, A or B—but that stances are also *scalable*. Speakers initially choose between positive and negative evaluative resources, either aligning or disaligning themselves with certain values, in order to create solidarity with or distance from their intended audience. But speakers may then adjust how strongly these stances are presented, whether "subtly or boldly" (Gales, 2010: 90). In the attitudinal area of appreciation, for instance, the aesthetic dimension of *good* may be up-scaled to *very good* or *beautiful*, or down-scaled to *sort of good* or just *okay* (Martin & White, 2005: 56). Choices like these "construe greater or lesser degrees of positivity or negativity" (Martin & White, 2005: 135), which has "the effect of turning up or down the 'volume' of an utterance" (Gales, 2010: 90). In Appraisal, both "attitude and engagement are domains of graduation which differ according to the nature of the meanings being scaled" (Martin & White, 2005: 136). The semantics of graduation is therefore central to the Appraisal method.

This system is divided into two broad regions. The first looks at how "feelings are amplified" (Martin & White, 2005: 35). Amplification is accomplished through the various resources of 'force,' which "covers assessments as to degree of intensity and as to amount" (Martin & White, 2005: 140). Intensification may be applied both to qualities (e.g., *slightly* versus *extremely foolish*) and processes (e.g., *slightly* versus *greatly hindered*) (Martin & White, 2005: 140). Force also addresses assessments of amounts, in terms of number (*few*, *many*), size (*small*, *large*), proximity (*near*, *far*), etc. The quality common to the many meanings gathered under the header of 'force' is their 'inherent'

scalability (Martin & White, 2005: 137). There is no boundary, for example, on judgements of quantity—there may be *one* or *two* of a thing, *less* or *more* of a thing, etc.

However, the system also attends to complementary categories, i.e., semantically binary, either/or propositions. A person either is or is not a *father*, for example. And yet, speakers manage to scale these meanings as well, by referencing "the degree to which they match some supposed core or exemplary instance of a semantic category" (Martin & White, 2005: 137), i.e., by referencing the prototypical

Figure 9.1: The System of Graduation (Martin & White, 2005: 154)

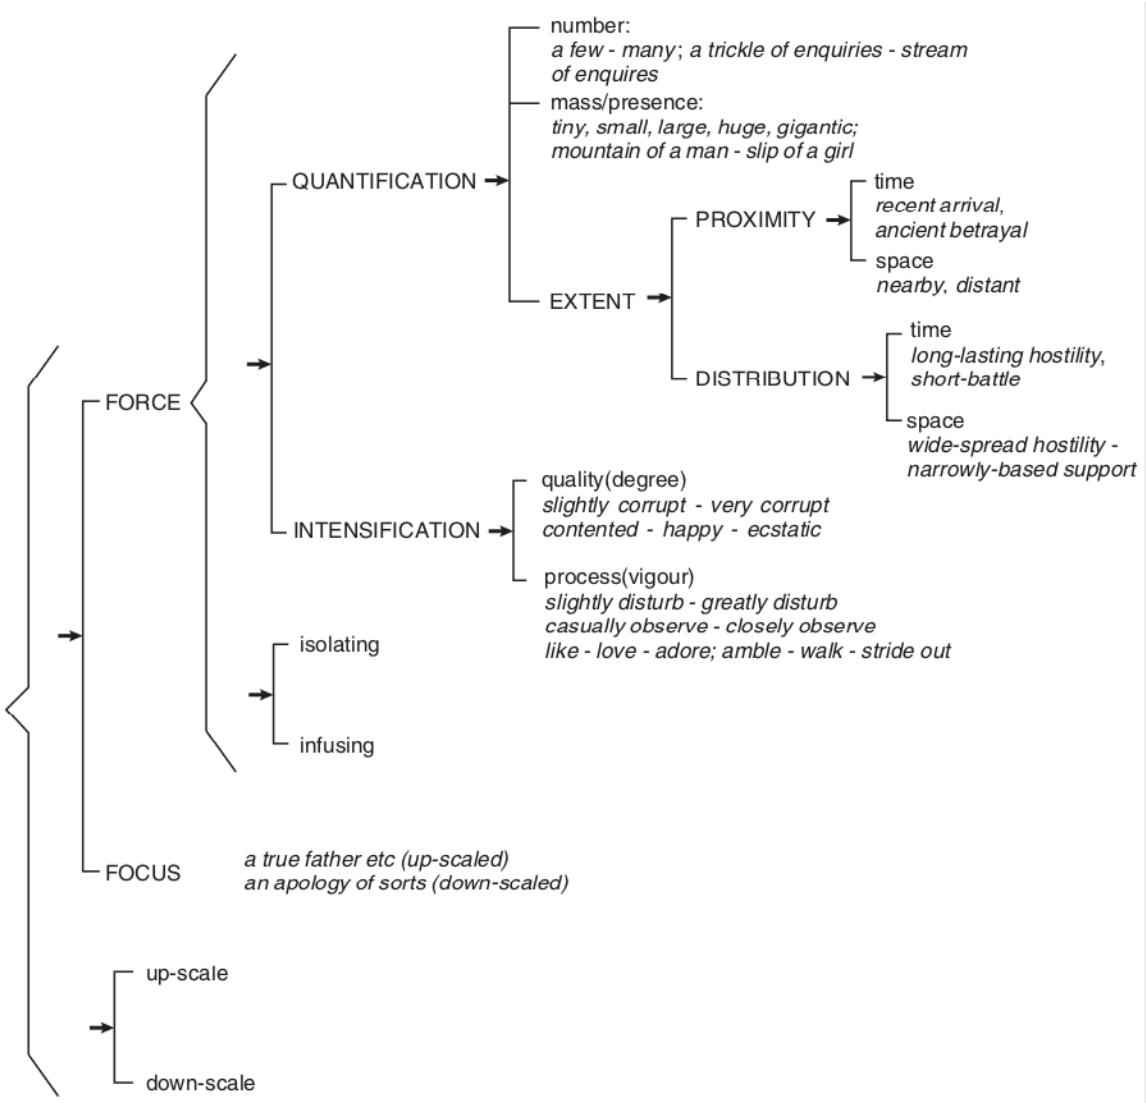

member of the category. Thus, a male with offspring—definitionally a member of the category *father*—may nevertheless be positively evaluated as *a real father* (Martin & White, 2005: 138), which is to say, a true exemplar of the category. In this way, graduation looks not just at amplification through force, but also at how "categories [are] blurred" (Martin & White, 2005: 35). This simpler collection of meanings is called 'focus,' and involves either upscaling by 'sharpening' (*a true friend*) or downscaling by 'softening' (*sort of*, *-ish*) an entity's relationship to the prototype (Martin & White, 2005: 138). The full system of graduation is shown in Figure 9.1 above.

This system has interesting potential applications to texts as charged with hostility as pledging. In the system of attitude, graduation resources are the means by which pledge authors demonstrate greater or lesser degrees of antipathy towards their targets, such as when Rodger R aims a string of up-scaled invective at the sorority members whom he would later attack, appraising them as *spoiled, heartless, wicked bitches*. In the system of engagement, graduated meanings can also demonstrate greater or lesser commitment to the violent proposal at the heart of the pledge itself, which Rodger R does, for example, when he uses a modal of high inclination to strengthen his pronouncement that *I will attack the very girls who represent everything I hate in the female gender*. The strength of his engagement would have been measurably diminished if he had opted for an auxiliary from a lower tier of inclination, e.g., *I might attack*.

Folk linguistic ideologies expect that threatening language will be somehow up-scaled, an attitude which is reflected in Gales's (2010) student survey of threatening language. Respondents imagined threats to include attitudinally charged features like a "cold, angry, distraught tone" and up-scaled elements like profanity (*shit*, *fuck*) (Gales, 2010: 96-97). Students also shared folk linguistic impressions of graduated meanings in engagement, such as the expectation that threat authors would work to dialogically

contract their texts “to disallow for pleading or other voices to be heard,” as well as showing “a commitment to the intended action through modals of obligation” like *will*, *must*, and *shall* (Gales, 2010: 96-97).

Graduation is thus widely viewed as central to performing the work of threatening, and two competing hypotheses are available for how these meanings might behave in the pledge dataset. On the one hand, if targeted violence begins with a grievance, then up-scaled language might be more likely to appear in texts by authors who felt the grievance keenly enough to act, i.e., in realized pledges. On the other hand, if targeted violence is more typical of authors with higher conceptual complexity and less typical of authors with higher ambivalent hostility (Smith, 2006), then realized texts would be more likely to feature *down-scaled* forms—since, in theory, these would evince “more deliberate and less emotional thinking” (Smith, 2006: 98). Interestingly, neither hypothesis is borne out by a statistical comparison of the corpora, as Table 9.1 shows below. Simply put, there is no quantitative difference between the pledge types in terms of graduation.

Table 9.1: Statistical Significance in Graduation

Graduation Category		Tokens		Probability (p)
		Non-realized	Realized	
Force	Degree	35.51	45.50	> .05
	Vigor	25.77	28.87	> .05
	Number	32.07	32.01	> .05
	Mass	4.01	0.63	> .05
	Extent	30.93	27.61	> .05
	Proximity	14.89	7.84	> .05
	Distribution	16.04	19.77	> .05
	Upscale	120.85	129.90	> .05
	Downscale	7.45	4.71	> .05
	Isolating	75.60	80.64	> .05
	Infusing	52.69	53.97	> .05
Focus	Soften	3.40	1.26	> .05
	Sharpen	6.30	4.08	> .05

Frequency per 1000 words

Of course, elements of graduation have been discussed at relevant points in the preceding analyses, when graduated meanings intersected with other areas of meaning (e.g., the up-scaling of likelihood modals in section 8.2). Still, of the statistical null results identified elsewhere between the two corpora, the lack of a global distinction across the system of graduation is perhaps the most unexpected.

This is because, in FCI, a "full-fledged intention" requires the agent to have answered "(a) whether she is capable of performing the action and (b) whether she has other desires that outweigh her desire to perform the action" (Malle & Knobe, 2001: 55). In the threat assessment literature, consideration (a) is equivalent to 'feasibility,' a variable which "refers to the ease or difficulty of reaching the end state" desired by the thinker (Geurts *et al.*, 2016: 55), while consideration (b) corresponds to 'desirability,' or how positively or negatively a person views the imagined end state (Geurts *et al.*, 2016: 55). The difference between people who intend to act and people who are threatening for effect is that the two "value the *desirability* and *feasibility* of their threat differently" (Geurts *et al.*, 2016: 55). Desirability—how positively or negatively a person evaluates an end state—is tied directly to questions of scalability and answered by 'more' or 'less'. Differences in desirability between the realization categories should theoretically appear in one, or across some collection, of graduation variables. But this is not the case.

Similarly, conceptual complexity is "significantly associated with outcome" in Smith's (2006: 79) dataset, in that writers with lower scores on this psychological characteristic were less likely to act. The sample words which Smith (2006) presents for coding low conceptual complexity—and a correlated failure to act—are allness terms (e.g., *without a doubt*, *absolutely*), which have an obvious dimension of up-scaling. Hypothetically, then, up-scaled meanings should appear more in the NR corpus. But, if anything, it is the R texts which make slightly more use of this resource (130 instances of

up-scaling per 1000 words in the R pledges versus 121 instances for the NR pledges, as seen in Table 9.1 above).

One possible interpretation of the null results across the entire system of graduation—compared to the more particular instances where graduated meanings were relevant to variables in other systems (e.g., capacity, in section 5.1.6)—is that graduation in this dataset cannot be fruitfully separated from the secondary system it is inflecting. In other words, graduation may add analytical depth to discussions of attitude and engagement, but, in this dataset, it says nothing as a system on its own. To return to specificity, precise descriptions of both the threateners' plans and their targets are considered important indicators of commitment (Borum *et al.*, 1999; Mohandie, 2014). As Chapter 5 shows, one approach to specificity is through hyponymic and troponymic relations. In terms of graduation, super- and subordinate semantic relations are best captured by lexical 'infusion,' a process by which up- or down-scaling is "conveyed as but one aspect of the meaning of a single term" rather than through a separate lexical item (*somewhat*, *extremely*) (Martin & White, 2005: 143). This is the difference between, say, the 'infused' term *joyous* and its 'isolated' semi-equivalent *very happy* (Martin & White, 2005). The word *joyous* is, among other things, a hyponym of *happy*. The up-scaled infusion of the word thus communicates its hyponymic markedness—being *joyous* is a more specific kind of *happy*. And yet, while aspects of graduation can help to illuminate attitudinal analyses like this, the quantitative frequency of infused tokens is essentially identical in both corpora.

One final observation: the various literatures concerned with threatening typically focus simply on whether certain qualities and processes are *present* in a text. Whether such qualities or processes have more predictive power if they are *intensified* in some way is not something these literatures tend to address. For instance, Gales (2010) discovers that the use of modals as a grammatical class are significant to non-realized

threats. This finding is unconcerned with graduation: up-scaled modals of likelihood such as *will* carry the same statistical weight in her dataset as down-scaled modals like *might* and *may*. The simple presence of modal auxiliaries in a threat thus sends a larger signal than their relative graduation. (Indeed, this trend has played out across several of the preceding analyses. The heightened presence of violent incapacity, for example, is a better indicator of a NR pledge than the semantics or specificity of the violence itself.) Meanwhile, in the psychological literature Malle and Knobe (2001) contend that an intention is only formed at the conclusion of a reasoning process. "Before making a decision about how to act, the person needs to consider various desires, balancing them against each other and asking which of them can potentially be fulfilled" (Malle & Knobe, 2001: 46). The authors, however, argue only that reasoning needs to be present in the process, not that it needs to be in any way intensified or drawn out (or even fully conscious!).

One example may serve to illustrate the limits of graduation when it is isolated from the attitudinal or engagement meanings being scaled. In Dickens NR's pledge, she weighs her desire to shoot police officers against the consequences of being *caught by the police or killed by them*. As argued in Chapter 2, Dickens NR does not emphasize or up-scale either *caught* or *killed* because she arguably has no need to—both consequences are so severe that shared world knowledge performs this work instead. Thus, at key textual moments such as this, the need for graduated meanings may be secondary to the experiential meanings themselves.

9.1 SUMMARY

Statistical null results are found across the entire system of graduation. This finding is potentially counterintuitive because pledging to harm as a speech act is shot through with hostile feelings, and graduated attitudinal meanings offer one way for an author to

verbalize such hostility. Similarly, graduating engagement meanings could be a way to signal a heightened commitment to the proposal of a pledge to harm. Nevertheless, the varied resources of graduation are not employed differently by either author type in a significant way. One possible reason for this may be that graduation in this dataset depends heavily on the secondary system it is modifying, whether attitude or engagement. Another may be that the presence of certain language features (e.g., “all modals” in the non-realized threats of Gales’s (2010: 182) dataset) carries more predictive power than how intensified these features are. Or, finally, authors may choose to rely instead on contextual information or shared world knowledge (per the Dickens NR example) to add heft to important moments in their pledges. Of course, none of these three potential explanations is mutually exclusive. Each, alone or in tandem, could affect how graduation is employed across this dataset. Whatever the case may be, graduation by itself explains little to nothing about the stancetaking of either realization category.

CHAPTER 10 A BLIND TEST

An overarching question is whether the findings of the preceding six chapters have any predictive power. How useful are these language features for classifying previously unexamined pledges as either realized or non-realized *through the assessment of their language alone*? The answer has real-world implications. For example, it is common for defendants to argue—disingenuously or no—that they did not mean their language as a threat. Could the kind of Appraisal analysis performed in this thesis shed a scientific light on such after-the-fact claims? Meanwhile, the purview of threat assessors is before-the-fact, the span between the issuance of a threat and the violence it threatens. In this window of time, assessors must determine the strength of the threatener's commitment to act so that authorities may respond accordingly. "Simply put, [threat assessors] must identify the doers from the non-doers" (Smith, 2006: 9). But to perform their task as accurately as possible, "professionals who analyze threat cases need valid correlates to predict which threateners are most likely to move from violent words to violent deeds" (Smith, 2006: 9). Do the correlates identified by this analysis qualify? Just how 'valid' are the Appraisal variables in discerning a pledge writer's true psychological intent?

To better understand which stance markers may extrapolate more broadly, this chapter tests the findings of the previous analyses on a new set of six authentic pledges to harm. Hereafter, the fourteen pledges which served as the basis of the thesis up to this point will be referred to as the 'working set' of texts. The six additional pledges will henceforth be referred to as the 'validation set.'

10.1 PROCESS AND ANALYSIS

The parameters of this test are as follows:

1) Six new pledge texts were supplied by a retired law enforcement officer;
2) The officer also supplied the known realization status for each text (i.e., whether each is considered realized or non-realized by law enforcement);
3) However, *these texts and their realization statuses were not supplied to the analyst directly* but rather to the thesis supervisor, who withheld the realization statuses and anonymized elements in the validation set as needed;
4) After the anonymized versions were provided to the analyst by the thesis supervisor, an Appraisal analysis was performed on each text to determine whether it showed markers more similar to the R or NR pledges of the working set.

The analysis itself was a two-part process:

5) First, the likely realization status was determined through a discursive, 'human' reading of the text based on the results of an Appraisal analysis;
6) Next, the text's normed feature counts were compared to the normed counts from the R and NR corpora across 147 features identified in the preceding analyses (e.g., attitudinal polarity);
7) For each feature, the text was given a score of 1 for the corpus it most closely resembled, and a score of 0 for the corpus it least resembled (these *unweighted* count totals are shown as Appraisal Features Score in Table 10.1 below);
8) Each text was marked as 'realized' or 'non-realized' depending on which corpus it shared the majority of its language features with;
9) *After* these two initial determinations (qualitative and quantitative), the thesis supervisor revealed the realization status of each text, as provided to him by the retired officer; and finally,

10) Accuracies and inaccuracies between the analysis and the realization status provided by the officer were evaluated.

Thus, like many of the live assessment situations which this research hopes to assist, the process of this test was 'blind.' *Only* the language of the threat was available as data for determining whether the author harbored real-world violent intent. The individual analyses are discussed text-by-text below. A summary of the analyst determinations and the unweighted Appraisal feature scores for each text is shown in Table 10.1 below, along with the realization statuses provided by the retired law enforcement officer familiar with the cases. Note that the short summaries which begin with the next paragraph were written *before* the statuses of each text were revealed, although the known realization status is noted for easy reference.

Text1: Analyst determination: **Non-realized**. Unweighted Appraisal feature score: **Non-realized** (85 NR; 62 R). Known status: **Non-realized**. At 81 words, Text1 is the shortest in the validation set. The initial determination is that this is a non-realized pledge, primarily due to a preponderance of appreciation meanings, e.g., *beautiful mountainous state*. Additionally, weapons and weapons-related terminology (in this case, *four bomb blasts*) appear much more frequently in NR texts in the working set, as do explicit mentions of the time and place of the attack (*12:00 noon on Tuesday July 7*; *State Capitol Complex*). And only NR pledges in the working set threaten bombings. Finally, despite its short span, this text opens once to acknowledge an outside voice, presumably that of the reader (anonymized here as *Gov. Jones*). However, this instance of acknowledgement is framed as a rhetorical question: *What ever happened to "Take only memories, leave only footprints"?* The closed nature of this question, along with the uniformly contractive formulations in the rest of this text, communicates a 'sureness' that is more a hallmark of non-realized pledges.

Text2: Analyst determination: **Non-realized**. Unweighted Appraisal feature score: **Non-realized** (85 NR; 62 R). Known status: **Realized**. Text2 has been initially categorized as non-realized because this pledge appears designed more to instill fear than to communicate a true intent to act. First, the grammatical 1st Person (*I*) appears with the modal auxiliary *will* in attack-related constructions four times (*I will kill*…)—and no other modal auxiliary is used. This is more typical of NR pledges. Second, Text2 features a linking of multiple events (three murders and a bombing). Series of violent events like this only appear in the NR pledges in the working set (Rodger NR and Valle NR). Third, every ethical meaning in the text is negative. Realized pledges in the working set tend to praise as often as they blame; a lack of moral nuance, i.e., black-and-white thinking, is much more a feature of non-realized pledges. Finally, this text also features a bomb threat (*I will blow up MD Hospital* [anonymized]). The mention of a weapon—and a bomb in particular—is much more common in NR pledges.

Text3: Analyst determination: **Realized**. Unweighted Appraisal feature score: **Realized** (69 NR; 78 R). Known status: **Non-realized**. This is the only text in the validation set that reads as a realized pledge. This initial assessment is based as much on what is *not* in the text as what is. For instance: no weapons are specified; no time and place are given; the lexeme *die* is repeated four times, but no more marked verb of physical violence is offered (e.g., *shoot*, *stab*). Along with this, the ethics of the text are relatively sophisticated. Four of the six imagined victims are condemned for moral reasons (e.g., *stuck up*, *Bitch*). However, the author also shows empathy, however unusually. The fifth victim listed is a loved one who is threatened because *pain from this event would be unbearable*. (This formulation has direct echoes in Kinkel R and Hribal R.) Additionally, the author lists him- or herself as the sixth victim who *must DIE*, and does so explicitly because the imagined murders are acknowledged as immoral (even as they are framed as being necessary). Finally, Text3 is shot through with modals of

requirement (four of the seven modal auxiliaries in the text are *must*) and moral reasons used to justify the imagined violence. This relative focus on the *why* of a threat over *what* the author imagines doing is much more typical of R pledges in the working set.

Text4: Analyst determination: **Non-realized**. Unweighted Appraisal feature score: **Non-realized** (74 NR; 73 R). Known status: **Non-realized**. Of all the texts in the validation set, this one was the most difficult to categorize, an ambiguity which is interestingly captured in the near parity of the unweighted feature count. The core bad acts of this pledge are a threat of *Systematically assassinating United States Senators* and *Having all Christians shot*. These are presented, somewhat whimsically, in the form of an unrhymed song lyric, which also fantasizes about unreal events like *sleeping on freight cars full / Of comfortable pillows*. Like Text3, this pledge is notable mostly for what it lacks. Text4 mentions no weapons, no effects of weapons, no time or place for the imagined assassinations or pogrom, etc., all of which argues for a status as a realized pledge. However, it also has no tokens of propriety whatsoever. And the two instances of requirement—*It was imperative for me* and *I had to*—are not connected to either assault, but rather to the author's argument that he or she had to depart from whatever space (perhaps a home) that the author shared with the intended reader. While meanings of requirement are a hallmark of realized pledges in the working set, they most often govern the central violent act of the pledges. In sum, there is very little *what* offered in this threat, and no *why* at all, i.e., why senators? why Christians? The author does not say. This overall lack of compulsion to act led, finally, to the classification of this pledge as non-realized, a determination barely borne out by the unweighted Appraisal feature count.

Text5: Analyst determination: **Non-realized**. Unweighted Appraisal feature score: **Non-realized** (76 NR; 70 R). Known status: **Non-realized**. While Text5 qualifies as a pledge to harm, in that the author threatens the reader's family, the validity of using

this text as a test case is questionable. The primary purpose of this text is extortion, an illegal act in its own right separate from the pledge it contains. Arguably, the threat of extortion is aimed at the 2nd Person of the reader, making this more of a direct threat than anything in the working set, and the transactional nature of this communication is unlike the pledges examined to this point. In fact, the profit motive runs directly counter to Bulling and Scalora's (2013: 9) definition of 'intended violence' as "[v]iolent acts that meet the following criteria: intent to commit the act; selecting an attack mode that ensures injury, death, or property damage; and *a motive that does not profit the attacker* [emphasis added]." Nevertheless, this text was included in the bundle of pledges supplied by the retired law enforcement officer, and so it has been analyzed here.

At 473 words, Text5 is the longest of the validation set (and is longer than everything except Hribal R and Long R in the working set). Interestingly, the length itself argues for a classification as realized, since realized authors tend to write more. In terms of Appraisal, this text has other features of a realized pledge. For instance, no violent hyponyms are used (e.g., *shoot*, *stab*). Instead, the *health and well being* of the reader's family will be put at *risk* if the recipient fails to pay the $500,000 demanded. The reader is told to *pay us or wonder when…* but no specific violent actions are described. Along with this, tokens of propriety are slightly more likely to be positively than negatively charged. Both a relative lack of explicit violence and an evenhanded approach to ethical meanings argue for classification as a realized text.

And yet, through appreciation, Text5 shows a much greater interest in the world of things than of people. Admittedly, this could be a result of the generic nature of an extortion note. Nevertheless, the most apt comparison on this front is with the non-realized pledges of the working set. This text also lacks any mention of why the reader has been targeted, focusing much more on what the reader is required to do to deliver the money. More minorly, the author characterizes him- or herself as part of a *deadly*

and serious team of professionals—a dubious claim, and one that has echoes of the LA Unified NR text. Perhaps most importantly, nearly the entire system of engagement in this text fits the contours of a NR pledge: from the low use of denial to the relative lack of dialogic space for outside voices and a heavy reliance on proclamations over other, less assertive forms of engagement. Finally, all twelve modal auxiliaries are some form of the lexeme *will*, used to communicate either high likelihood or high inclination. No other modals are used by the author. This preponderance of *will* comports more with the behavior of NR pledges in the working set. Despite the potentially imperfect comparison of an extortion note with the pledges analyzed previously, this analysis argues for classifying this text as non-realized. In this case, that means that while the author may or may not have followed through on the extortion attempt, the chances are low that he or she ever approached the reader's family with intentions to do them harm.

Text6: Analyst determination: **Non-realized**. Unweighted Appraisal feature score: **Non-realized** (81 NR; 66 R). Known status: **Non-realized**. Lastly, Text6 showcases a collection of language features which tend to appear in non-realized pledges. The text is a litany of violent anti-Semitic ideation aimed at the Jewish members of *Rhoades High School*. The author imagines *killing* Jews, Jews being *burned*, *branding a swastika on their forehead*, etc. This use of negative capacity to describe injuries to imagined victims is a hallmark of NR pledges in the working set, as is the linking of the several violent events in sequence. Additionally, the discursive use of propriety in the text is entirely negative. Jews are described as *bastards*, as *selfish*, and as *evil*. This focus on extreme, negative ethical meanings is also much more common to the NR pledges. Finally, the choice of modal auxiliaries is also more typical of non-realized writings. Of the eight modals present, four are a form of *will* and two are forms of *be going to*, both used with the 1st Person to communicate a threat (e.g., *I'll kill him*; *I'm going to get him*). There is a single modal of requirement, a resource more common

to realized texts. However, it appears in a passive construction—*Jews should be burned*—which places the onus of action elsewhere, a device that is instead found in NR pledges. There is very little that this text shares with the R pledges of the working set.

10.2 RESULTS

Table 10.1 below shows the initial 'projected' realization statuses of the texts in the validation set, as well as the known realization statuses provided by the officer and only revealed by the thesis supervisor after the full analyses were performed. There were two incorrect predictions: Texts 2 and 3. These two are highlighted in red.

Table 10.1: Blind Test Results

Text	Word Count	Projected Realization Status			Known Realization Status
		Analyst Determination	Unweighted Appraisal Features Score (of 147)		
			NR	R	
Text1	81	Non-realized	85	62	Non-realized
Text2	103	Non-realized	85	62	Realized
Text3	136	Realized	69	78	Non-realized
Text4	130	Non-realized	74	73	Non-realized
Text5	473	Non-realized	76	70	Non-realized[41]
Text6	152	Non-realized	81	66	Non-realized

10.3 DISCUSSION

It is superficially interesting to see that the analyst determinations (which produced the short analyses in section 10.1 above) agreed with the unweighted Appraisal feature scores in all cases. The five texts which read to the human eye as non-realized (Texts 1-2; 4-6), and the one which read as realized (Text3), were subsequently flagged as

[41] According to the case notes, the extortion itself was indeed attempted by the author, but no friends or family of the recipient were ever approached, and the author was shown to be working alone, not as part of a team.

such by the Appraisal feature count which followed. In practice, of course, these are not fully independent measures. Both have the same point of origin: the patterns of language usage in the working set uncovered in the preceding chapters. That the results of both analyses were in sync is thus of limited significance. This particular aspect of the process will therefore not be explored further.

Of more interest is the accuracy of these analyses in comparison to the known realization statuses supplied by the retired officer. As shown in Table 10.1 above, the patterns established by the working set were largely effective at identifying the *non-realized* pledges in the validation set. Of the five actual NR cases, only Text3 was misidentified as a realized pledge. However, when the numeric counts are limited to the 104 metrics in Appraisal which were shown to be statistically significant—rather than the full 147 features which yield the unweighted scores shown in Table 10.1—then Text3 is, in fact, correctly re-classified as non-realized (now 54 NR to 50 R). In fact, using only statistically significant features correctly classifies all five of the non-realized validation texts.

The issue, it seems, is with correctly identifying a realized pledge. Both the discursive, human analysis and the unweighted feature counts misclassified the single R pledge in the validation set: Text2 (85 NR; 62 R). But unlike with Text3, considering only statistically significant areas of Appraisal makes no difference—Text2 remains incorrectly flagged as non-realized (now 65 NR to 39). Indeed, if statistical significance is used as the measure of choice, then *all six* of the validation set texts would have been classified as non-realized by this quantitative measure. Categorizing every text as NR while misclassifying the one R pledge in the test is problematic for obvious reasons. There is an undeniable value in knowing which texts are composed for purposes other than to express true violent intent (as, e.g., Justin Carter's case demonstrates), but the urgency of threat assessment lies in preventing violence that has been sincerely

promised. Ideally, any analysis would pinpoint this metaphorical needle in the haystack, not just correctly catalog the hay.

Two hypotheses are worth entertaining if the consistent misclassification of Text2 is to be explained:

- Hypothesis One: Text2 encodes a true violent intent but there is a mismatch between the patterns of usage identified in the realized texts of the working set and how they overlay onto this particular pledge;
- Hypothesis Two: Text2 itself *lacks* true violent intent, and both the qualitative and quantitative analyses are, in fact, correct, despite the evident case history that the author would later violently assault one of his imagined victims.

Hypothesis Two flies in the face of the guideline followed up to this point—that authorial behavior is the best metric of intent. Further, it is also practically untestable. The best course of action would be to interview the author, assuming such a thing is even possible, and even then, the answer would rely on the frailty of human memory (to say nothing of the honesty of a violent felon). Hypothesis Two can be summarily dispensed with, then, in favor of exploring Hypothesis One: that Text2 encodes a true violent intent but does so in a way which draws from resources typically associated with the *non-realized* pledges in the working set. To better understand this text's status in relation to the analytical techniques applied during this blind test, however, a few additional steps were taken. Eventually, an exploration of Text2 in relation both to prototype theory (Lakoff, 1987) and, interestingly, Elliot Rodger's realized pledge will help illuminate the analytical problems presented by Text2.

The first step in this direction, though, requires recognizing that the feature counts used to test the validation set are an amalgamation of many *unique* texts. Thus, one follow-on question is whether the unweighted feature score would actually misclassify individual pledges from the *working set*. The same unweighted counts were

therefore applied to all 14 texts previously analyzed as two full corpora—the eight non-realized and the six realized pledges.

In fact, a misclassification does occur. The unweighted features correctly mark 13 of these 14 texts in line with their known realization statuses, i.e., as a member of either the realized or non-realized corpus. The single misclassified text is Rodger R's (89 NR; 58 R). To see if realized texts more generally are misidentified by these metrics, an additional pledge from outside either the working set or the validation set was analyzed. The chosen text is a piece of realized writing produced by a man named Jerry Varnell. In 2017, Varnell "attempted to detonate what he believed to be an explosives-laden van he had parked in an alley next to [a local bank]" in Oklahoma City (United States Department of Justice, 2019). Before this, Varnell had composed a message meant to be published to social media by an associate of his after the bombing was complete. (His pledge to harm is available in Appendix C.)

Discursively, Varnell's text reads like a realized pledge. To note just one example, despite its brevity, it contains four different instances of temporal shifting, where the bombing is discussed as having already occurred—a phenomenon quite particular to realized texts in the working set. Nevertheless, the quantitative results are mixed. Initially, Varnell's text is flagged incorrectly as a non-realized pledge by the full, unweighted feature count (76 NR; 71 R). But, as with Text3, considering only statistically significant features proves to be corrective—accurately re-classifying the pledge as realized (now 48 NR; 56 R). With Varnell's text included, this means that the count of statistically significant features correctly identifies 19 of the 21 texts to which it has been applied. Yet, the two consistently misclassified pledges are both realized.

The adjusted results are shown in Table 10.2 below, with misclassifications highlighted in red. (The reason for Rodger NR and Hribal R's pledges appearing in blue will be explained below.)

Table 10.2: Statistically Significant Appraisal Features Counts

Text	Known Status	Statistically Significant Appraisal Features Count		Projected Status
		NR	R	
Text1	Non-realized	61	43	Non-realized
Text2	Realized	65	39	Non-realized
Text3	Non-realized	54	50	Non-realized
Text4	Non-realized	56	48	Non-realized
Text5	Non-realized	57	47	Non-realized
Text6	Non-realized	57	47	Non-realized
Archangel	Non-realized	67	37	Non-realized
Brahm	Non-realized	66	38	Non-realized
Dickens	Non-realized	72	32	Non-realized
LA Unified	Non-realized	73	31	Non-realized
McKelvey	Non-realized	70	34	Non-realized
Rodger	Non-realized	78	26	Non-realized
Skyline	Non-realized	66	38	Non-realized
Valle	Non-realized	68	36	Realized
Hribal	Realized	28	76	Realized
Kinkel	Realized	50	54	Realized
Long	Realized	31	73	Realized
Rodger	Realized	69	35	Non-realized
Roof	Realized	36	68	Realized
Shaw	Realized	37	67	Realized
Varnell	Realized	48	56	Realized

Given this, it seems clear that the numeric tool as currently constituted is insensitive to aspects of certain—but by no means all—of the resources an author may draw on to express a true violent intent. (And because this tool and the discursive, human reading have been trained on the same dataset, the analyst obviously suffers from the same prejudices, at least in the case of Text2.) However, a look at this issue from a *qualitative* standpoint reveals that this error is not as random as it may first appear.

Because this effort is one of classification, a prototype effect is potentially detectable here. In particular, the misclassification of Rodger's realized pledge may, in fact, be bound up with the compositional character of Rodger's *non-realized* pledge. And this could, in turn, shed light on the analysis of Text2. Highlighted in blue in Table 10.2 above is the pledge from each realization type which contains the most features

associated with its realization category. For the realized pledges, Hribal R's text is the most marked, at 28 NR to 76 R. What these numbers arguably represent is that this pledge is the most prototypical *realized* member of all 21 texts listed. Meaning, on a goodness-of-example scale of the kind discussed by Lakoff (1987) and other cognitive linguists, Hribal R's writing may be judged as the best example of the category 'realized' by virtue of having the most features of a realized pledge of any in the extended dataset. Or, put more inelegantly, it is the *realizedest* of all 21 pledges analyzed here.

When the non-realized texts are viewed from this perspective, the picture is potentially illuminative. In this case, Rodger NR's pledge is the best example of the category, at 78 NR to 26 R. This is informative because of how it may interact with authorial style. Johnstone (2009: 1) observes that "linguistic styles emerge out of stancetaking strategies that prove repeatedly relevant and useful for particular speakers in particular kinds of interactions." Both of Rodger's pledges have been pulled from the same longer stretch of discourse, which means both were authored by one particular speaker in one particular kind of interaction—in this case, the catalog of violent fantasy which Rodger indulges toward the end of his autobiography. Discovering not just that his non-realized writing is the most *non-realizedest*, but also that his realized text is the only working set pledge consistently misclassified by the features count raises the question of whether both are somehow bound up with the style of his larger stretch of discourse, i.e., if there are "repeated stancetaking choices" (Johnstone, 2009: 2) that are suffusing his NR and R writings alike. Rodger, it would seem, leans heavily toward stance resources more associated with non-realized writings, and strongly enough that the center of gravity of his realized pledge appears to be pulled in the same direction. Or, put differently, the numbers show not only that Rodger's non-realized pledge is the best example of its category, but that his realized pledge is the *worst example* of its own.

The implications of this for the persistent misclassification of Text2 are only speculative. However, Text2 has definite echoes of Rodger's style—enough so that Rodger NR was cited as an influence in the blind analysis shown in section 10.1 above. (A reminder that those analyses were written *before* the known statuses were revealed.) This similarity is reflected in the numbers of the individual Appraisal variables as well. As Table 10.3 shows below, not only were Rodger R and Text2 misclassified as non-realized in the *exact same four areas* of statistical significance (and no others), these four areas were misclassified at almost the *exact same* ratios. (The high count is shown in red.) Indeed, the two only diverge in the system of engagement, and there only by a single feature count.

Table 10.3: Rodger R and Text2 Feature Count Errors

Text	Appraisal Variable							
	Negative Capacity		Propriety		Composition/ Valuation		Engagement	
	NR	R	NR	R	NR	R	NR	R
Rodger R	4	3	10	7	18	3	19	4
Text2	4	3	10	7	18	4	18	5

Gales (2010) has demonstrated, in part through her own use of the Appraisal method, that "a variety of form-based functional patterns were found to be salient to each category of threat (realized vs. not realized)" (p. 264). Nevertheless, "this dichotomy of interpersonal functions does not divide cleanly along threat realization lines" (Gales, 2010: 263). A similar interplay and overlap of linguistic resources could be at work in the pledges of Rodger R and Text2. Each of the two realized pledges makes use of resources more common to non-realized texts—and in practically identical ways—placing both well within the fuzzy boundary separating a realized from a non-realized pledge. However, their many systematic similarities imply that this

misclassification is not random, which further implies that it can somehow be accounted for and corrected.

Thus, while the current accuracy rate of 19 out of 21 for this extended dataset (or just over 90%) is a promising start for identifying violent intent in a pledge to harm using only its language, the nature of the error means that this rate of accuracy is potentially improvable. Clearly, correctly capturing *less* typical realized texts such as Rodger R and Text2 will require weighting the Appraisal features in as-yet undetermined ways, something which is beyond the scope of the current thesis. This therefore represents one possible avenue of future research.

CHAPTER 11 DISCERNING INTENT

"Considerations of intentions and intentionality permeate human social life" (Malle *et al.*, 2001: 1). Indeed, how we evaluate the words and actions of those around us depends in large part on discerning the *purpose* of their words and actions. This applies equally to legal questions, since "[t]he law relies on this concept as well" (Malle & Knobe, 1997: 102). In the U.S., as in many other countries, the court system routinely considers not just the nature of the injury but also whether the injury was intended; evidence of *mens rea* alters the perceived severity of the *actus reus*. Thus, when questions of intent arise in forensic contexts—where weighty matters of justice, life, and freedom hang in the balance—a clear-eyed view of what does and does not constitute intent is crucial.

This is no less true with threats. Any attempt to predict future violence based on a threatening communication must minimally address two questions, as formulated by Gales (2010: ii): "Is the intent real?" and "Is the threatener likely to act?" How threat assessors answer these questions is consequential, since "their decisions may involve injury or even death and may require extensive personnel resources and large expenditures of money" (Smith, 2006: 2). The stakes for discerning intent can therefore be quite high. But "[n]obody can know for sure, of course, what another person's intentions are" (Shuy, 2005: 15). For this reason, "[i]t remains a formidable task, in daily life as in the court room, to make judgments about other people's mental states with a high degree of accuracy" (Malle & Nelson, 2003: 564). The current research seeks to ease the burden faced by threat assessors on the front end and the judicial system on the back end by investigating whether violent intent is detectable in a 'pledge to harm' (Harmon, 2008), which is not only the most common type of threatening language (Meloy & O'Toole, 2011) but also one which is likely to become *more* common as social

media penetrates further into world societies (Bojsen-Møller *et al.*, 2020; Lidsky, 2012). The novel aspects of the study's methodology are summarized below.

11.1 ORIGINAL CONTRIBUTIONS

While the ability to detect violent psychological intent is one with applications in several different forensic areas, only Smith (2006) is known to have tried using linguistic science to uncover its presence in the language of threats. However, her use of linguistics is more as a helpmate to what is primarily a psychological investigation. This thesis inverts Smith's (2006) model by using psychological theories to illuminate a linguistic analysis. To date, then, this is the only study which focuses on the question of violent psychological intent from a linguistic perspective. More uniquely, this research explores a kind of authentic linguistic production—a pledge to harm—that has been expressly delimited by linguistic theory.

Improving the degree of accuracy for detecting violent intent has required more than a single theory. Indeed, a novel coordination of several different theoretical models from both the psychological and linguistic sciences has been employed. Linguistically, Speech Act Theory (Austin, 1975; Searle, 1976) provided the means for classifying a threatening text as a pledge to harm rather than as a direct threat, and Audience Design Theory (Bell, 1984) argued for why the distinction matters. Forensically, pledging was shown to be a direct correlate of what is known as 'leakage' in the threat assessment literature (e.g., O'Toole, 2000). Psychologically, Biological Naturalism (Searle, 2004) opened a window onto the *internal* cognitive make-up of intent—its components and processes—while the Folk Concept of Intentionality (Malle, 1999) offered a framework for understanding how *external* social actors (including judges and jurors) perceive these psychological processes in others. The two together supplied a working inventory of

psychological elements worth being alert to in the subsequent analyses, e.g., desires, calculations about feasibility, etc.

Finally, and most fundamentally, the investigation was conducted within the tradition of Systemic Functional Linguistics (e.g., Halliday & Matthiessen, 2004), which posits that language is, first and foremost, a resource for meaning-making. But SFL does not view meaning as simply ideational, it is also *interpersonal*. This makes SFL a natural fit for assessing intent as it may be transmitted through language since the "social role of intentionality" (Malle & Knobe, 1997: 102) means that it too functions interpersonally. Finally, the tools for analyzing these interpersonal meanings were provided by Appraisal (Martin & White, 2005), a discourse semantic method born of SFL. Appraisal facilitated an extremely close reading of the fourteen pledges in the working set, as well as a comparison of these findings against the seven texts of the validation set.

As helpful as this varied framework proved, certain limitations were inescapable. These are laid out in the next section, which are then followed by the study's primary findings.

11.2 DIFFICULTIES AND LIMITATIONS

The simplicity of the main question under examination—is violent intent discernible in pledges to harm—perhaps belies the challenges inherent to answering it. Social actors, whether they are regular people or legal authorities, assume that intent is something which can be both detected and measured in the behavior of others, including their language behavior. Because the aim of this study is to uncover a convergence point between psychology and linguistics, it was important that each pledge analyzed herein be an *authentic* production—a text authored with a legitimately personal communicative purpose in mind. The reasons for this requirement are several. Only authentic writings

could potentially contain true violent intent, and therefore be used to populate a corpus of realized threats. Populating a non-realized corpus might be possible in a more controlled setting, e.g., in a laboratory and at the instigation of a researcher, in that such texts would feasibly lack real-world intent. However, these would also be missing the genuine communicative purpose that often makes a non-realized pledge *look like* a genuine threat. Furthermore, any artificial approach would be hamstrung by the fact that a laboratory setting is inappropriate for curating realized pledges: manufacturing the conditions needed to instigate true violent intent in human subjects would be both impractical and unethical. Therefore, both the realized and the non-realized pledges needed to be gathered, not created. Yet, this crucial stipulation—that texts be authentic—carried several difficulties and limitations in its wake.

Data: The main challenge was finding data. As noted in Chapter 3, the authentic nature of the texts means that they were often the subject of live legal disputes. Populating the dataset was therefore challenging, and the dataset itself remained relatively small. Two limitations follow from this. The first is that the small number of texts, while manageable for a single analyst, also reduces the likelihood that the study's findings can be extended to a wider population of pledges. Indeed, this was part of the reasoning for performing the blind test in Chapter 10. Along with assessing the predictive power of the various possible stance markers of intent—an interesting task in its own right—the test helped identify which linguistic forms may indeed extrapolate more broadly. The second limitation has to do with the non-realized texts in particular, and is likely insurmountable no matter the size of the study or its analytical approach.

Non-realized Data: Quite simply, there is no way to ever be sure that an authentic non-realized text—including the eight in the dataset—has been authored *without* true violent intent. As discussed in Chapter 3, the only point of comparison is with an author's behavior after he or she has communicated the pledge. But this relation

is correlative at best. Despite best evidence, the link between a lack of real-world action and the mental state which gave rise to a pledge can only be assumed, never proven, leaving a non-zero chance that any comparison between a 'realized' and a 'non-realized' pledge is inherently faulty.

This lack of clarity is admittedly frustrating. It is not, however, unusual. Applied linguists are forced to work with likelihoods rather than certainties as a matter of course. (As is the legal system, it should be noted.) The best that can be done is to acknowledge the ever-present ambiguity while gathering the clearest evidence possible supporting the classification of a pledge as 'non-realized.'

Language Variety: Finally, this study was limited to American English (AmE). It did not consider other axes of dialectal variation (e.g., age, gender, ethnicity, etc.), nor did it examine other world Englishes (e.g., British English, Australian English, etc.). Instead, preference was given to an informal register of the standard variety of AmE, no matter the demographic information available about the author. Although detecting intent in language is a question which legal systems wrestle with the world over, the findings of this study are thus limited to American English only.

11.3 PRIMARY FINDINGS

The preceding analyses uncovered several trends in language use which might be considered primary. These are presented as they relate to the two realization types.

Non-realized pledges: The NR pledges are more negatively charged. This is evident across several metrics. Looking at general attitudinal polarity, which captures the prosodic tilt of stance markers across the subsystems of affect, judgement, and appreciation, this corpus is 64% negative. (For comparison, the R corpus is almost evenly split, at 52% negativity.) The NR pledges are also more violent, and place more of this violence in the future, as something which is yet to occur. Additionally, this future

violence is more likely to be framed as something the pledge's author will undertake personally, e.g., via 1st Person pronouns combined with modal auxiliaries of inclination like *will*. Similarly, weapons-related terminology (including the bad effects of using weapons) is much more prevalent in the NR pledges. All of this results in texts which appear more threatening than their realized counterparts. In terms of their expressed morality, the NR pledges are more black-and-white in their judgements, with an emphasis on the 'black': opinions of others' ethics is almost universally negative in these texts.

These attitudinal trends are buttressed by the non-realized authors' tendency to close their pledges to outside voices. Although both realization types are roughly equal in how much they expand and contract their writings, the NR pledges are 88% likely to do so using the grammatical resources of proclamation, "meanings by which, through some authorial interpolation, emphasis or intervention, dialogic alternatives are confronted, challenged, overwhelmed or otherwise excluded" (Martin & White, 2005: 117-118). (By contrast, the R pledges employ this resource for 67% of their contractiveness, opting instead for disclamation at a higher rate, something which is discussed again below.)

Prosodically, then, a non-realized pledge is likely to be: more violent; more menacing; less morally nuanced; and more confident in the rightness of its own conclusions.

Realized pledges: The R pledges are more attitudinally balanced. While tokens of violence certainly appear in this corpus, the focus is instead on moral meanings. These writings are 64% more likely than their non-realized counterparts to judge others against ethical standards. And *only* the R pledges allow for earnest or non-ironic judgements of positive propriety. For these authors, good people and good behavior can indeed be found in the world. Unfortunately, such goodness is not found in the

present moment, i.e., at the time the pledge is authored. The intolerableness of current circumstances is a recurrent theme across several Appraisal resources, e.g., propriety as well as the variable of composition in appreciation, which is used to present things as disordered or out of balance. These assessments are also closer-to-home than in the NR texts. For instance, meanings related to the mundane legal world only appear in the R pledges, which contrasts with the occasionally biblical depictions of conflict in both realization types.

This even-handedness extends to voices other than the authors' own. An outside voice is over two times as likely to appear in a R text, even if it is admitted in order to be refuted. Along with resources like attribution and acknowledgment, the R authors' employment of denial—at over four times the rate than in the NR corpus—is of a piece with this. As a form of negation, a denial "necessarily carries with it the positive" position that is being denied (Martin & White, 2005: 118). In other words, even when R authors are explicitly closing their texts to alternative views through dialogic contraction, they prefer to do so in a way which introduces these voices *implicitly*.

An additional finding in the R corpus is their authors' lack of personal agency. This is manifest in several areas. For instance, R authors are 14 times more likely to use the resources of denial to present themselves as somehow powerless. The most striking examples of this, perhaps, are the nearly identical versions of "I have no choice" which appear in half of the R pledges, and nowhere in the NR corpus. This stance of helplessness is echoed in the modality of requirement, which is used particularly by R authors to argue that they feel compelled to perform the imagined violence by forces beyond themselves, e.g., *I must kill people* (Kinkel R). This particular usage also appears nowhere in the NR corpus. Rhetorically, this functions hand-in-hand with the several instances where imagined assaults are expressed in the past tense, presupposing for the reader that the violence is a *fait accompli*. In these ways, the R

pledges tend to present the attack as inevitable and both as something which the authors are obligated to perform and powerless to resist.

Prosodically, then, a realized pledge is likely to be: more ethically nuanced; more aware of the wider world, particularly human society and its workings; and more likely to view future violence as a moral imperative which the authors are obligated to perform, whether or not they consider themselves willing.

11.4 DISCUSSION

Having now come to the end of the analysis, the overriding question is whether Appraisal has uncovered patterns of stancetaking that are consistent with the psychological difference theoretically underlying the creation of the two text types. In other words, to what extent do the preceding chapters show that the decision-making process leading *toward* violent intent is, in fact, distinguishable from the process leading *away* from it (and toward a different intent) *using only the language of the pledges themselves as indicators*? And to what extent are the patterns of stancetaking either expected or surprising from the standpoint of the linguistic and psychological theories used throughout this thesis? More generally, what interpretive value can theory bring to these empirical findings?

Addressing the potential correspondence between linguistic forms and psychological states first, the results are promising, as Chapter 10 shows. While only correlative, the correlation between certain collections of linguistic forms and the subsequent behavior of the authors nevertheless appears quite strong. Any attempt to interpret why this might be, however, must begin with the foundational psychological mechanisms which gave rise to the texts themselves. According to the Folk Concept of Intentionality (e.g., Malle, 1999), *any commitment to act, including the intent to commit real-world violence, is formed at the conclusion of a reasoning process*. The process

takes as inputs both a belief about the world and a desire for the world to be somehow different. Such a desire is then considered in light of any conflicting desires on scales of both intensity and feasibility. For instance, a desire for violence which risks detainment or death might be weighed against the competing desires to remain alive and free. Finally, the output of this process is an intention: a commitment to pursue whichever goal has been calculated as both the most desirable *and* the most achievable. In sum, "intentions serve to fulfill desires by identifying a course of action that is feasible to implement for the agent and is compatible with the agent's other desires" (Malle & Knobe, 2001: 55-56). It is therefore possible to view intention formation as a process of *stancetaking*, whereby agents crystallize their attitude towards their own wants and needs.

This research has theorized that realized and non-realized pledges are linguistic products of a reasoning process with similar inputs but divergent outputs. In all cases, there seems to be a belief that something is wrong in the world and that a desire for violence is a valid response. In the case of R pledges, the output of this process is the authors' expressed intent to pursue real-world violence. In the case of NR pledges, the intention is different, e.g., to intimidate, to challenge authority, to vent anger, to get attention, to be humorous, etc. (Fraser, 1998). And yet, recognizing that the two pledge types take divergent stances toward their violent content still fails to explain *why* the linguistic forms appear as they do—i.e., why these forms and not others?—and whether these forms can be linked back to different kinds of authorial intent.

However, positing that the communicative purpose of each text type is different appears to be the strongest hypothesis. Arguably, *the empirical patterns uncovered by Appraisal do reflect two divergent communicative aims*. At this point, it is worth revisiting the hypotheses presented at the beginning of this thesis, that: 1) different systemic resources are employed in the communication of true violent intent because the

functional aim of such a communication is different than that of a text created with a different intent; and 2) because these systemic resources are manifested in the tangible realm of language, the difference between violent intent and other intents is visible to linguistic tools, and therefore detectable in the texts themselves. The argument to be put forward in the discussion which follows supports the validity of both hypotheses. However, a full accounting of the data will require a series of speculative—yet logical—steps. These are pursued in the next two subsections.

11.4.1 Linked Reasoning Processes

The initial and most important step is the premise that Appraisal is actually measuring the results of, not one, but *two* reasoning processes. The first process determines whether violent intent is formed at all, e.g., by an agent weighing the desirability versus the feasibility of future, real-world violence. On its own, this first process is invisible, being confined entirely to the mind of the person in question. Thus, a second intention-forming process is needed if the agent is to become a threatener: one that determines whether the results of the first process are somehow verbalized. In other words, *the decision to communicate a pledge to harm is the result of its own subsequent reasoning process*, one which involves its own unique set of calculations about the desired social effects versus the potential personal costs of speaking out.

In theory, dividing the processes this way allows space not only for realized authors (who might be characterized as +violence/+communicate) and for non-realized authors (who might be characterized as -violence/+communicate). It also accounts for those people who harbor violent intent but do not engage in leakage before acting out—the 'hunters' who do not 'howl,' to use Calhoun and Weston's (2015) parlance. This third group could be characterized as +violence/-communicate. And finally, of course, separate reasoning processes allow for the possibility that any number of people might

decide, first, not to commit violence and, second, not to tell anyone that they ever entertained such a socially problematic desire (this last group could be characterized as -violence/-communicate)[42].

Along with making room for the kind of typological observations discussed in the threat assessment literature, a second premise follows from this division. This is that, despite the separation, these two reasoning processes are nevertheless linked, i.e., *the decision for or against committing real-world violence becomes a belief input in the reasoning process of whether to communicate the desire for violence*[43]. Further, if the agent does decide to speak then this recursion necessarily influences how the imagined actions are discussed. Someone who plans to attack a target would carry a different set of affective and epistemic stances towards the imagined event than someone who has concluded that violence is morally or practically infeasible.

Perhaps the 'hunter' decides it is in his or her best interest to explain their motivations? Perhaps the 'howler,' nursing a frustrated desire, has decided to use language as an emotional "safety valve" (Gellerman & Suddath, 2005: 485) to expel feelings that are considered socially unacceptable (Jackson, 1981). (And, in fact, it is exactly scenarios like these that are taken up again in more depth below.) In these cases, the stance resulting from the first reasoning process becomes, in turn, its own "object of stance" (Du Bois, 2007: 151) against which the hunters or howlers position themselves in the course of the second process. Or, to put the premise more succinctly,

---

[42] The first three typological combinations are addressed again below. Because the fourth (-violence/-communicate) conceivably creates no observable data in terms of either language or non-verbal behavior, this type will not be considered further.

[43] Helpfully, the case can be made that this is indeed a simple, binary distinction, with no real grayscale to worry over. For an agent vacillating between the poles of *yes* and *no* or between *I will* and *I won't*—i.e., someone who has not yet made up their mind—a reasonable argument is that such a person has so far failed to form a violent intention. Or, put differently, any decision that is not *yes, I will* is, in fact, merely a form of *no, I won't*, leading to non-action on the part of the agent.

the language which results from the second (whether to communicate) inevitably reflects the conclusions of the first (whether to commit violence).

The logical implication of two 'chained' or 'linked' reasoning processes, if accurate, is that the presence of true violent intent can indeed be detected in threatening language. However, before delving more fully into the explanatory power of these premises, a linked decision-making process also offers an interesting perspective on why detecting and measuring intent in language has proven so difficult, and why linguistic attempts to do so are still in their infancy.

Essentially, any effort to discover *mens rea* in language (such as this thesis) is limited to looking for evidence of the first reasoning process in the results of the second—a kind of 'Perseus's shield' effect. This is compounded by the contention, voiced in section 2.2.3, that language is an unfit device for indicating two of the three measures of commitment put forward by Malle and Knobe (2001). To briefly revisit this dilemma, utterances by which an agent could show 1) an early investment and 2) the acceptance of opportunity costs actually only work as 3) public announcements. When the data is linguistic, then, the typical signals an outside observer relies on to classify a behavior as intentional are simply unavailable. Because of this refractive nature of psychological intent, it is little surprise to see that social scientists of various disciplines have "questioned the use of linguistic form as an indicator of behavior" (Gales, 2010: 264)[44]. And, of course, more research is needed before this hesitancy can be comfortably shed, a topic taken up again in section 11.5 below.

An additional difficulty is that, if the decision to communicate a pledge to harm is the result of its own reasoning process, then this decision involves motivations that are

[44] Despite pushing further into psychology than other linguistic studies, this research unequivocally supports the view that a "a one-to-one correspondence between a particular feature of language" with any internal mental state is indeed a "hopeful myth" (Lord, Davis, & Mason, 2008: 375).

separate from those driving a desire for violence. Thus, while two threateners may differ in their willingness to act, there is no reason to believe they might not share an identical motivation for speaking, i.e., that their *communicative intent* could be the same regardless of their psychological intent. For instance, a hunter and a howler might both decide to share their ideations through social media with the goal of causing fear or social disruption. It is easy to imagine how this similarity in linguistic intent could obscure the differences in their commitment to act, further complicating a forensic analysis.

All that said, "[a]ny use of language is motivated by a purpose" (Eggins, 2004: 5). Systemic functional theory recognizes—indeed, is built upon—the idea that speakers choose linguistic forms which maximize the realization of their communicative purpose (Eggins, 2004). The final proposal of this thesis is that the communicative purpose of each text type is discoverable from the patterns of stancetaking present in each, and, furthermore, that these communicative purposes are *different*. This argument hinges on the systemic functional understanding of register and the related concept of solidarity between author and audience—in particular, the way the realized pledges appear to cultivate solidarity while the non-realized pledges appear to reject it.

11.4.2 Register Variation and Communicative Intent

In SFL theory, register is a way of approaching "functional variation according to language use" (Martin & White, 2005: 24). More specifically, because a speaker's choice of register "is determined by what the speaker is doing socially" (Matthiessen *et al.*, 2010: 176), the concept allows for the potential reverse engineering of the *purpose* of this doing, i.e., the *what* of the communication may be used to better understand the *why* of the communication. Three contextual variables are theorized as 'controlling' for register variation: mode, field, and tenor. Mode deals with language's place in the

activity in question (Martin & White, 2005), e.g., the work performed by dialogue in an image-heavy medium like film. Field is the "social and/or semiotic process that the interactants in the context are engaged in" (Matthiessen *et al.*, 2010: 95), i.e., the domestic or institutional activity itself, such as threatening. (The concept of field was taken up previously to parse valuation meanings in Chapter 6.) Finally, tenor captures "the role relationships entered into by the interactants" (Matthiessen *et al.*, 2010: 217), e.g., employer/employee. As triplet axes within the single "sphere of meaning" that is context (Matthiessen, 2015: 6), any change in the value of one variable affects the other two. However, because all of the pledges in the dataset were created and transmitted in roughly the same way, language's place in a pledge to harm is more or less identical across the dataset. Mode will therefore not be discussed further. The focus instead will be on apparent differences in tenor between the realization categories and the information this offers about the activity or field each type seems to be engaged in.

Values of tenor are influenced by interpersonal considerations like status, formality, and politeness. "What they have in common is a very general sense of the social distance between the speaker and the addressee" (Halliday & Matthiessen, 2004: 631). The fact that a pledge to harm does not automatically conflate its audience and its imagined victim means that a different kind of interpersonal space is theoretically possible than that found in direct threats[45]. Broadly speaking, the R and NR pledges appear to take differing approaches in how or whether they attempt to establish solidarity with the 2nd Person of the addressee.

Interestingly, it is the realized authors who appear most interested in bolstering stances of alignment with their putative readers. And, in fact, viewing the primary

[45] Indeed, it makes a certain amount of sense that tenor would be a crucial pivot point between a realized and a non-realized pledge since this is the register category which directly correlates with the 'interpersonal' metafunction of language (see Martin & White, 2005: 27)—i.e., the metafunction to which Appraisal is most sensitive.

findings of section 11.3 above through the lens of solidarity shows that the patterns of stancetaking are far from random. Appraisal reveals that these authors are far more likely to open their texts to outside voices through resources like attribution, acknowledgement, and even denial. These authors are also morally more even-handed, allowing that good and evil exist side-by-side in the world. Generally, these texts focus on the *why* of the violent ideation, particularly on why an author feels the need to commit violence, and they expend tangible rhetorical resources to convince their audience of the rightness of these ideations. Indeed, such openness and evenhandedness could be just as much at home in more persuasive genres, e.g., editorial writing, though it is found here in threats of physical harm. Despite the negative content, attempts at solidarity in these writings are comparably frequent.

Non-realized authors, on the other hand, generally reject such stances in favor of disalignment. But neither are the primary findings of this corpus random. NR pledges are more negative, more violent, and almost uniformly closed to outside voices. These pledges do not engage in the same kind of shared meaning-making with their audience as the R pledges. Instead, they use their textual space like a bully pulpit to impose their preferred meanings on their putative readers, e.g., through the much-used resource of proclaiming. Rather than arguing, convincing, or persuading, a non-realized text is more likely to simply assert the value positions expected of its addressees. Thus, where R pledges tend to answer the deontic question of *why* an author must act, NR pledges revolve instead around epistemic and dynamic questions of *what* the authors want to do and *why* they want to do it. The attempted level of solidarity in these writings is measurably lower in comparison to the R pledges.

This disparity in tenor leads, in turn, to a difference in the two realization categories' perceived field or "focus of the activity" (Eggins, 2004: 9)—i.e., in the purpose driving the communicative effort. The proposed distinction is slight but crucial

and may be glossed thusly: realized pledges perform the task of *explaining*; non-realized pledges perform the task of *announcing*. This means, roughly, that where R pledges work *to interpret* imagined violent events for their putative addressees, NR pledges instead *give notice* of these events[46]. This rather intuitive division in the "field of activity" (Matthiessen, 2015) of the realization types has two important correspondences with the linguistic research done by Gales (2010) and Smith (2006). First, this interpretation of the data echoes Smith's (2006: 88) finding that "[t]hreateners were significantly more likely to approach/stalk or harm when they used the language strategy of persuasion in their threat communications." Second, conceiving of NR pledges typologically as announcements also comports with Gales' (2010) findings that modals of prediction were statistically more likely to appear in the non-realized threats of her dataset. And third, drawing such a distinction is premised on Muschalik's (2018) own use of function to categorize direct threats.

Hypothesizing this sliver of daylight between the communicative purposes of one realization type and the other leads to a final question: what social effect(s) are these slightly different registers being used to achieve? If texts created to *explain* hope to achieve a state of understanding in their audience, and texts created to *announce* hope to achieve a state of alertness or perhaps even alarm, then why do authors with different psychological intentions typically choose one field of activity over the other?

The answer must in some way boil down to the perceived social benefits of communicating, and the conclusion on the part of the texts' authors that these benefits outweigh the social sanction they might face for voicing such unacceptable desires (Jackson, 1981). For example, if realized authors have concluded that violence is feasible, then explaining their motives would be one way to control the narrative that will

[46] More minorly, the motivation to *explain* might also be the reason realized texts tend to be longer than their non-realized counterparts.

result from the bad act. This could be attractive to authors who feel there is a personal meaning in their imagined actions (e.g., a duty to create a more just world). Such authors might view pledging as an understandably valuable investment, one which will secure their personal meaning a place in the discussion after the fact. For realized authors, to commit violence but remain silent risks ceding the interpretive authority to others—such as the media—who almost certainly would not have the author's personal interests in mind. In these cases, silence is a vacuum[47]. Indeed, concerns about how the author's actions will be perceived are cited explicitly in some realized pledges in the dataset (and only in realized pledges). For example, Hribal R writes that *a dozen different things will be speculated to be at fault* before explaining that *all this was caused by was dehumanization of public school*. Similarly, Long R begins his pledge by noting that his shooting of the police officers *does seem to be out of character* but then immediately pleads for his readers to consider his rationale, writing *I ask that you finish reading before you make that decision*.

Authors of non-realized pledges, on the other hand, may feel powerless to act but not to cause fear or social disruption with their language, goals arguably best served by the most attitudinally charged language they can muster. For some non-realized authors, a pledge may be a response to a "heightened state of emotional arousal"

[47] A potentially instructive example of this risk is the difference between Elliot Rodger and Stephen Paddock. Both died in the commission of their crimes, so neither had the opportunity to interpret their actions for the world after the fact. However, the two are treated very differently by the media and others. Rodger's writings have managed to turn him into a veritable martyr for the 'incel' community (Branson-Potts & Winton, 2018). By contrast, Paddock—the perpetrator of what is currently the deadliest mass shooting in U.S. history (Rosenblatt, 2017)—left no writing behind and no overt clues as to his motive. This has caused even close friends and family to feel "puzzled by the attack" (Subramanian, 2017). In other words, rather than arguing over the *validity* of the attack's meaning, as in the case of Rodger, with Paddock onlookers have been forced to debate whether there was any meaning at all. Paddock's actions have not been embraced by any particular community or cause because he did not offer any ready-made interpretative frame for them.

(Smith, 2006: 90). If the emotion is strongly negative, this arousal could be indicative of both ambivalent hostility (Gottschalk *et al.,* 1979) and lower conceptual complexity (Hermann, 2005)—two interdependent psychological variables flagged by Smith (2006) as risk-*reducing*. In these cases, the "act of writing the threatening communications may assist these threateners in defusing their anger" (Smith, 2006: 90), a notion consistent with the 'safety valve' effect discussed by the psychiatrists Gellerman and Suddath (2005: 485). For other non-realized authors, causing strong emotion in their readers seems more the point (e.g., Brahm NR, Skyline NR, etc.), a goal that is *also* best served by highly provocative language.

Interestingly, why someone might choose to issue a pledge to harm at all could be very broadly the same in all cases: language is a means of controlling what others believe. In the case of realized texts, this could be control over the meaning people see in the author's violent act. For non-realized texts, language offers potential control not just over others' emotions (e.g., through fear) but also, perhaps, control over the author's own (e.g., through venting anger).

Table 11.1: Communicative Intent by Realization Category

		First Reasoning Process (Commit Violence)	
		+Violence	-Violence
Second Reasoning Process (Communicate Ideation)	+Communicate	Realized **Explain**: Asserts control over any *post factum* interpretation of the violent act	Non-realized **Announce**: Causes fear or social disruption, vents negative emotion, etc.
	-Communicate	N/A **N/A:** Action justifies itself; Jeopardizes control over any *post factum* interpretation of the violent act	

Table 11.1 above features a simplified view of the two reasoning processes—revisited as +/-violence and +/-communicate—and the potential social benefits which pledging might serve at the culmination of the different decision-making pathways.

As noted in section 7.1, motive is considered one of the key areas indicating the degree of risk a threatener poses (Borum *et al.*, 1999; Calhoun & Weston, 2015). Indeed, people surveyed about threatening language identified "the justification for the threat" as one of the few language functions they would expect to find in a threatening communication (Gales, 2010: 96). These assessors and survey respondents are, rightly, focused on the results of the first reasoning process—what is the motive for attacking? But having now split the *ideation* apart from the *decision to discuss the ideation*, there is no reason to believe that the motive for speaking is any less crucial in determining a threatener's commitment to acting, and that these motives are revealed by the language.

11.5 FUTURE RESEARCH

Whether or not "work on threatening communications is still in its infancy" (Gales, 2010: 269), understanding how—or even *if*—psychological intent is encoded linguistically in threatening language is certainly still beginning. The assumption built into legal codes in the U.S. and elsewhere is that intentions exist and that they are measurable through the medium of behavior, including language behavior. The results of the current study support this view. However, there is a great deal further to go before anyone can claim that intent is being detected and measured *reliably* in linguistic data. Efforts to address this question are only a few steps into what is, undoubtedly, a thousand-mile journey.

To that end, there are several potential avenues of future research suggested by these results. These are roughly divisible into practical concerns—areas that might benefit from exploring these empirical findings further—and theoretical concerns—areas

that could be better illuminated by the combination of linguistic and psychological frameworks used to arrive at the empirical findings.

Practically, this study has been aimed at reducing the number of unknowns confronting threat assessors and jurists tasked with identifying *mens rea* in threatening language. While building a diagnostic tool was not the original goal of this research, the possibility of creating one now exists. An ongoing project, of course, would be to further extend the dataset with additional pledges, and continue testing to see which Appraisal features remain reliable indicators of the different realization categories. Beyond this basic endeavor lie several important possible improvements. For instance, because the compilation of features used in the blind test in Chapter 10 is unweighted, a next step could be a statistically-motivated refinement of the tool based on an extended dataset, one that better represents which features more heavily correlate with the different realization types (e.g., temporal shifting with realized pledges). However, applying Appraisal as a method of analysis requires a level of technical knowledge that even skilled assessors and jurists are unlikely to acquire. A final goal, then, could be to take the weighted Appraisal features and produce a faithful translation into terms more centered on the lay distinction between explaining and announcing. (And indeed, further refining this distinction should run concurrently with any and all future efforts.)

In terms of theory, it would be interesting to see if the mixture of linguistic and psychological frameworks employed here could shed light on other areas of forensic linguistics. The combination of SFL, SAT, ADT, BN and FCI might produce fruitful commentary on any situation where a misalignment between linguistic and psychological intents is hypothesized to occur, e.g., deception detection, false confessions, etc. These theories may also help understand the presence or absence of violent intent in 'fictionalized' threats, where arguably real people and events are represented in artistic media like stories and music. The line between fiction and personal fantasy is often thin

and legally precarious; these theories could potentially demarcate it more clearly. Similarly, this framework could help linguists better understand phenomena like abusive language and the intent behind communications sent by stalkers. Even language crimes whose commission is the act of communication itself—including direct threats—could benefit from a deeper understanding of the intentions spurring their creation and potentially encoded in their language.

If anything, this thesis reaffirms a basic premise underlying the field of threat assessment, that even though leakage and pledges to harm are cause for serious legal concern, *statements of violent intent should not be taken at face value*. The potential psychological motivations for issuing a pledge to harm are as numerous as the potential contextual purposes a pledge might serve. Continued efforts to understand the interdependence of the two can only aid the cause of justice in dealing with this fear-inducing but all-too-common speech act.

REFERENCES

Abramson, B. (1978). Referees and umpires under fire. *Rome News-Tribune*. Retrieved from: https://news.google.com/newspapers?nid=348&dat=19780305&id=Dk4jAAAAIBAJ&sjid=1zIDAAAAIBAJ&pg=3832,4186666&hl=en

The Associated Press. (2008, June 5). Former clerk gets 6-month term for stadium attack hoax. *ESPN*. Retrieved from: http://www.espn.com/espn/wire/_/section/nfl/id/3427307

Austin, J. L. (1975). *How to do things with words*. Oxford University Press.

Barker, K., Wright, C., & Goodman, J.D. (2015, June 28). How a life that showed promise veered into violence against Asian women. *The New York Times*. Retrieved from: https://www.nytimes.com/2015/06/29/nyregion/how-a-life-that-showed-promise-veered-into-violence-against-asian-women.html

Bell, A. (1984). Language style as audience design. *Language in society, 13*(2), 145-204.

Berlinger, J. (2016, August 4). Gavin Long: Who is Baton Rouge cop killer? *CNN.* Retrieved from: https://www.cnn.com/2016/07/18/us/who-is-gavin-long/index.html

Biber, D. (2006). Stance in spoken and written university registers. *Journal of English for Academic Purposes*, *5*(2), 97-116.

Biber, D., Conrad, S., & Reppen, R. (1998). *Corpus linguistics: Investigating language structure and use.* Cambridge University Press.

Biber, D., Johansson, S., Leech, G., Conrad, S., & Finegan, E. (1999). *Longman grammar of spoken and written English.* London: Longman.

Black, H. C. (1968). *Black's Law Dictionary (Revised Fourth Edition)*. St. Paul, Minn.: West Publishing Co.

Black Lives Matter (2013). Herstory. *BlackLivesMatter.com*. Retrieved from: https://blacklivesmatter.com/about/herstory/

Bojsen-Møller, M., Auken, S., Devitt, A. J., & Christensen, T. K. (2020). Illicit genres: The case of threatening communications. *Sakprosa, 12*(1), 1-53.

Bornkessel-Schlesewsky, I., & Schlesewsky, M. (2009). The Role of Prominence Information in the Real-Time Comprehension of Transitive Constructions: A Cross-Linguistic Approach. *Language and Linguistics Compass, 3*(1), 19-58.

Borum, R., Fein, R., Vossekuil, B., & Berglund, J. (1999). Threat assessment: Defining an approach to assessing risk for targeted violence. *Behavioral Sciences & the Law*, *17*(3), 323.

Borum, R., & Reddy, M. (2001). Assessing violence risk in Tarasoff situations: A fact-based model of inquiry. *Behavioral sciences & the law, 19*(3), 375-385.

Branson-Potts, H., Ceasar, S., & Blume, H. (2015, December 15). L.A. schools to reopen Wednesday; threat against schools was 'not credible,' officials say. *Los Angeles Times.* Retrieved from: http://www.latimes.com/local/lanow/la-me-ln-all-lausd-schools-closed-threat-20151215-story.html

Branson-Potts, H., & Winton, R. (2018, April 26). How Elliot Rodger went from misfit mass murderer to 'saint' for group of misogynists—and suspected Toronto killer. *Los Angeles Times.* Retrieved from: https://www.latimes.com/local/lanow/la-me-ln-elliot-rodger-incel-20180426-story.html

Brown, R., & Gilman, A. (1960). The pronouns of solidarity and power. *Style in language*, 253-276.

Brown, P., & Levinson, S. C. (1987). *Politeness: Some universals in language usage* (Vol. 4). Cambridge university press.

Bucqueroux, B., & Seymour, A. (2009). A guide for journalists who report on crime and crime victims. *Justice solutions*. Retrieved from: http://www.mediacrimevictimguide.com/journalistguide.pdf

Bulling, D., & Scalora, M. (2013). *Threat assessment glossary.* Public Policy Center, University of Nebraska.

Bybee, J. (2014). Cognitive processes in grammaticalization. In *The new psychology of language* (pp. 159-182). Psychology Press.

Calhoun, F. S., & Weston, S. W. (2015). Perspectives on threat management. *Journal of Threat Assessment and Management, 2*(3-4), 258.

Caraballo, S. A., & Charniak, E. (1999). Determining the specificity of nouns from text. In *Proceedings of the 1999 Joint SIGDAT Conference on Empirical Methods in Natural Language Processing and Very Large Corpora* (pp. 63-70).

Center for Constitutional Rights. (2017, March). Solitary confinement: Torture in U.S. prisons. *Center for Constitutional Rights*. Retrieved from: https://ccrjustice.org/home/get-involved/tools-resources/fact-sheets-and-faqs/torture-use-solitary-confinement-us-prisons

Clement, J. (2017, May 15). Internet usage worldwide – Statistics & facts. *Statista*. Retrieved from: https://www.statista.com/topics/1145/internet-usage-worldwide/

Cruse, A. (2011). *Meaning in language: An introduction to semantics and pragmatics.* Oxford University Press.

Davidson, D. (2001). *Essays on actions and events: Philosophical essays* (Vol. 1). Oxford: Clarendon Press.

Davis, D. A. (1997). *Threats pending fuses burning: Managing workplace violence*. Palo Alto, CA: Davies-Black Publishing.

De Costa, P. I. (Ed.). (2015). *Ethics in applied linguistics research: Language researcher narratives.* Routledge.

Dietz, P. E., Matthews, D. B., Martell, D. A., Stewart, T. M., Hrouda, D. R., & Warren, J. (1991). Threatening and otherwise inappropriate letters to members of the United States Congress. *Journal of Forensic Science, 36*(5), 1445-1468.

Dixon, R. M. (1994). *Ergativity*. Cambridge University Press.

Dormond, K. (2014). Domestic violence threat assessment: Putting knowledge and skills into practice. *International handbook of threat assessment*, 330-339.

Du Bois, J. W. (2007). The stance triangle. *Stancetaking in discourse: Subjectivity, evaluation, interaction*, 139-182.

Durant, A. (2010). *Meaning in the media: Discourse, controversy and debate.* Cambridge University Press.

Durant, A., & Leung, J. (2016). *Language and law: A resource book for students.* Routledge.

Eggins, S. (2004). *Introduction to systemic functional linguistics.* A&C Black.

Eynon, R., Fry, J., & Schroeder, R. (2008). The ethics of internet research. *Sage internet research methods*, 23-41.

False Alarm or Report, Texas Penal Code § 42.06. (1973).

Fein, R. A., & Vossekuil, B. (1998). *Protective intelligence and threat assessment investigations: A guide for state and local law enforcement officials*. U.S Department of Justice, Office of Justice Programs, National Institute of Justice.

Foster, G. D. (2003). Ethics: Time to revisit the basics. *The Humanist, 63*(2), 30.

Fraser, B. (1998). Threatening revisited. *Forensic Linguistics, 5*(2), 159-173.

Gales, T. (2010) 'Ideologies of Violence: A Corpus and Discourse Analytic Approach to Stance in Threatening Communications', unpublished dissertation, University of California, Davis.

Gales, T. (2011). Identifying interpersonal stance in threatening discourse: An appraisal analysis. *Discourse Studies, 13*(1), 27-46.

Gales, T. (2017). Threatening Stances: a corpus analysis of realized vs. non-realized threats. *Language and Law= Linguagem e Direito, 2*(2).

Gales, T. (2020). ‘Prison has been a proper punishment’: Investigating stance in forensic and legal contexts. *The Routledge handbook of forensic linguistics*, 675-693.

Gellerman, D. M., & Suddath, R. (2005). Violent fantasy, dangerousness, and the duty to warn and protect. *Journal of the American Academy of Psychiatry and the Law Online*, *33*(4), 484-495.

Geurts, R., Granhag, P. A., Ask, K., & Vrij, A. (2016). Taking threats to the lab: Introducing an experimental paradigm for studying verbal threats. *Journal of Threat Assessment and Management, 3*(1), 53.

Gibbs, R. W. (2001). Intentions as emergent products of social interactions. *Intentions and intentionality: Foundations of social cognition*, 105-122.

Glasgow, K., & Schouten, R. (2014). Assessing violence risk in threatening communications. *In Proceedings of the Workshop on Computational Linguistics and Clinical Psychology: From Linguistic Signal to Clinical Reality,* 38-45.

Goffman, E. (1981). *Forms of talk*. University of Pennsylvania Press.

Goldman, A. I. (2001). Desire, intention, and the simulation theory. *Intentions and intentionality: Foundations of social cognition*, 207-225.

Gottschalk, L. A., Winget, C. N., & Gleser, G. C. (1979). *Manual of instructions for using the Gottschalk-Gleser Content Analysis Scales: Anxiety, hostility, social alienation-personal disorganization*. Univ of California Press.

Grant, T., Clark, U., Reershemius, G., Pollard, D., Hayes, S., & Plappert, G. (2017). *Quantitative research methods for linguists: A questions and answers approach for students*. Taylor & Francis.

Grant, T., & MacLeod, N. (2018). Resources and constraints in linguistic identity performance–a theory of authorship. *Language and Law/Linguagem e Direito, 5*(1), 80-96.

Grice, H. P. (1957). Meaning. *The philosophical review, 66*(3), 377-388.

Halliday, M. A. K., & Martin, J. R. (2003). *Writing science: Literacy and discursive power.* Taylor & Francis.

Halliday, M., & Matthiessen, C. (2004). *An introduction to functional grammar.* Routledge.

Harmon, A. G. (2008). Back from wonderland: A linguistic approach to duties arising from threats of physical violence. *Cap. UL Rev., 37*, 27.

Haydon, T. (2016). Kean U. grad gets 90 days in jail over threatening tweets. *NJ Advance Media*. Retrieved from: https://www.nj.com/union/index.ssf/2016/06/kean_twitter_suspect_gets_3_months_in_jail_judge_c.html

Hersher, R. (2017, January 10). What happened when Dylann Roof asked google for information about race? *NPR*. Retrieved from: https://www.npr.org/sections/thetwo-way/2017/01/10/508363607/what-happened-when-dylann-roof-asked-google-for-information-about-race

Hermann, M. G. (2005). Assessing leadership style: Trait analysis. *The psychological assessment of political leaders with profiles of Saddam Hussein and Bill Clinton,* 178-212. The University of Michigan Press.

Huddleston, R., & Pullum, G. K. (2005). *A student's introduction to English grammar*. Cambridge University Press.

Jackson, R. (1981). *Fantasy: The literature of subversion.* London and New York: Methuen.

Jaffe, A. (2009). Introduction. The sociolinguistics of stance. *Stance. Sociolinguistic perspectives 3*(28).

Johnstone, B. (2009). Stance, style, and the linguistic individual. *Stance: sociolinguistic perspectives*, 29-52.

Kess, J. F. (1992). *Psycholinguistics: Psychology, Linguistics, and the Study of Natural Language*. Amsterdam: John Benjamins Publishing Company.

KOMO Staff. (2012). Former student arrested in Skyline High School shooting threat. *KOMO News*. Retrieved from: https://komonews.com/archive/former-student-arrested-in-skyline-high-school-shooting-threat-11-21-2015

Kreuz, R. J., & Roberts, R. M. (1995). Two cues for verbal irony: Hyperbole and the ironic tone of voice. *Metaphor and symbol, 10*(1), 21-31.

Kropp, P. R., & Cook, A. N. (2014). Intimate partner violence, stalking, and femicide. *International handbook of threat assessment*, 178-194.

Labov, W. (2003). Uncovering the event structure of narrative. *Round table on language and linguistics*, 63-83.

Labov, W. (2010). Where should I begin. *Telling stories: Language, narrative, and social life*, 7-22. Georgetown University Press.

Labov, W. (2013). *The language of life and death: The transformation of experience in oral narrative.* Cambridge University Press.

Lakoff, G. (1987). *Women, fire, and dangerous things: What categories tell us about the nature of thought*. Chicago: University of Chicago Press.

Lidsky, L. B. (2012). Incendiary speech and social media. *Texas Tech Law Review, 44*, 147-164.

Lock, G. (1996). *Functional English grammar: An introduction for second language teachers.* Cambridge University Press.

Malle, B. F. (1999). How people explain behavior: A new theoretical framework. *Personality and social psychology review, 3*(1), 23-48.

Malle, B. F., & Knobe, J. (1997). The folk concept of intentionality. *Journal of experimental social psychology, 33*(2), 101-121.

Malle, B. F., & Knobe, J. (2001). The distinction between desire and intention: A folk-conceptual analysis. *Intentions and intentionality: Foundations of social cognition, 45*(67).

Malle, B. F., Moses, L. J., & Baldwin, D. A. (Eds.). (2001). *Intentions and intentionality: Foundations of social cognition*. MIT press.

Malle, B. F., & Nelson, S. E. (2003). Judging mens rea: The tension between folk concepts and legal concepts of intentionality. *Behavioral sciences & the law, 21*(5), 563-580.

Manslaughter, 18 U.S.C. § 1112 (1948)

Mann, J. J., Apter, A., Bertolote, J., Beautrais, A., Currier, D., Haas, A., ... & Mehlum, L. (2005). Suicide prevention strategies: a systematic review. *Jama*, 294(16), 2064-2074.

Mao, L. R. (1994). Beyond politeness theory: 'Face' revisited and renewed. *Journal of pragmatics, 21*(5), 451-486.

Martin, J. R., & White, P. R. (2005). *The language of evaluation*. Palgrave Macmillan.

Matthiessen, C. (2015). Register in the round: Registerial cartography. *Functional Linguistics, 2*(1), 9.

Matthiessen, C., Teruya, K., & Lam, M. (2010). *Key terms in systemic functional linguistics.* A&C Black.

Maslow, A. (2014). *Toward a psychology of being*. Sublime Books.

Meloy, J. R. (2015). Threat assessment: Scholars, operators, our past, our future. *Journal of Threat Assessment and Management, 2*(3-4), 231.

Meloy, J. R., Hart, S. D., & Hoffmann, J. (2014). Threat assessment and threat management. *International handbook of threat assessment*, 3-17.

Meloy, J. R., & Hoffmann, J. (2014). *International Handbook of Threat Assessment.* Oxford: Oxford University Press.

Meloy, J. R., James, D. V., Farnham, F. R., Mullen, P. E., Pathé, M., Darnley, B., & Preston, L. (2004). A research review of public figure threats, approaches, attacks, and assassinations in the United States. *Journal of Forensic Sciences, 49*, 1086-1093.

Meloy, J. R., & O'Toole, M. E. (2011). The Concept of Leakage in Threat Assessment. *Behavioral Sciences & The Law, 29*(4), 513-527.

Meyer, P. G. (1997). Hedging strategies in written academic discourse: Strengthening the argument by weakening the claim. In *Hedging and Discourse: Approaches to the Analysis of a Pragmatic Phenomenon in Academic Texts*, R. Markkannen and H. Schröder (eds.), 21–42. Berlin/New York: Walter de Gruyter.

Miller, G. A. (2013). Wordnet-about us. *WordNet. Princeton University*.

Mohandie, K. (2014). Threat assessment in schools. *International handbook of threat assessment*, 126-141.

Morton, T. (2016, September 22). Statewide bomb threat provokes confusion, caution, closings. *Townsquare Media*. Retrieved from: https://k2radio.com/statewide-bomb-threat-provokes-confusion-caution-closings/

Moyer, J. (2015). 'Cannibal cop' wins in court again. *The Washington Post.* Retrieved from: https://www.washingtonpost.com/news/morning-mix/wp/2015/12/04/cannibal-cop-cleared-in-court-again/

Murder, 18 U.S.C. § 1111. (1948).

Muschalik, J. (2018). *Threatening in English: A mixed method approach* (Vol. 284). John Benjamins Publishing Company.

Napier, M., & Mardigian, S. (2003). Threatening messages: The essence of analyzing communicated threats. *Public Venue Security*, 16-19.

Naranjo, C. (2015, February 20). Elliot Rodger rehearsed killing roommates before UCSB rampage. *Nexstar Broadcasting*. Retrieved from:

https://www.kron4.com/news/elliot-rodger-rehearsed-killing-roommates-before-ucsb-rampage/

Nini, A. (2017). Register variation in malicious forensic texts. *International Journal of Speech, Language & the Law, 24*(1), 99-126.

O'Donnell, M. (2013). UAM Corpus Tool (Version 3.3) [Computer software]. Available from www.wagsoft.com/CorpusTool.

O'Donnell, M. (2017). Interactions between natural–language processing and systemic functional linguistics. *The Routledge Handbook of Systemic Functional Linguistics,* 585-598.

O'Toole, M. E. (2000). *The school shooter: A threat assessment perspective.* Critical Incident Response Group, National Center for the Analysis of Violent Crime, FBI Academy.

O'Toole, M. E., & Smith, S. S. (2014). Fundamentals of threat assessment for beginners. *International handbook of threat assessment*, 272-282.

PBS. (2014). Chronology. *WGBH Educational Foundation*. Retrieved from: https://www.pbs.org/wgbh/pages/frontline/shows/kinkel/kip/cron.html

Pérez-Peña, R. (2015, December 26). Schools evaluate threats, questioning when to shut down. *The New York Times*. Retrieved from: https://www.nytimes.com/2015/12/27/us/schools-evaluate-threats-questioning-when-to-shut-down.html

Pew Research Center. (2019, July 12). Social media fact sheet. *Pew Research Center*. Retrieved from: https://www.pewinternet.org/fact-sheet/social-media/

Pinsof, A. (2013, July 12). Teen terrorist or Facebook prankster? *The Austin Chronicle*. Retrieved from: https://www.austinchronicle.com/news/2013-07-12/teen-terrorist-or-facebook-prankster/

Precht, K. (2003). Stance moods in spoken English: Evidentiality and affect in British and American conversation. *Text - Interdisciplinary Journal for the Study of Discourse*, *23*(2), 239-258.

Preston, D. (2007). The uses of folk linguistics. *International Journal of Applied Linguistics, 3*(2), 181-259.

Quirk, R. (1985). *A comprehensive grammar of the English language.* Longman.

Rahman, K. (2016, June 17). Black woman who tweeted anonymous racist threats to fellow black students is jailed. *Daily Mail*. Retrieved from: https://www.dailymail.co.uk/news/article-3647516/Black-woman-tweeted-anonymous-racist-threats-fellow-BLACK-students-jailed.html

Rodger, E. (2014). *My Twisted World*. Self-published.

Rodham, K., & Gavin, J. (2006). The ethics of using the internet to collect qualitative research data. *Research Ethics, 2*(3), 92-97.

Rogers, A., Wall, B., & Murphy, J.P. (1977). Introduction. *Proceedings of the Texas conference on performatives, presuppositions, and implicatures* (pp.6-18). Center for Applied Linguistics.

Rosenblatt, K. (2017, October 2). Las Vegas shooting is deadliest in modern U.S. history. *NBC News.* Retrieved from: https://www.nbcnews.com/storyline/las-vegas-shooting/las-vegas-shooting-deadliest-modern-u-s-history-n806486

Routh, J. (2018, January 22). Alex Hribal sentenced to 23 1/2 to 60 years in Franklin Regional stabbings. *Pittsburgh Post-Gazette*. Retrieved from: https://www.post-gazette.com/local/westmoreland/2018/01/22/Alex-Hribal-sentenced-in-Franklin-Regional-stabbings-pittsburgh-school-attack-westmoreland-county/stories/201801220093

Saeed, J. I. (2009). *Semantics, 3rd Edition*. Wiley-Blackwell.

Salinger, T. (2016, May 6). Georgia woman who called for ‘death to all white cops’ apologizes publicly, charges dropped. *New York Daily News*. Retrieved from: https://www.nydailynews.com/news/crime/georgia-woman-called-death-white-cops-apologizes-article-1.2628073

Sanders, A. (2018, April 6). Felony charges dropped in “Facebook threat” case. *The Austin Chronicle*. Retrieved from: https://www.austinchronicle.com/daily/news/2018-04-06/felony-charges-dropped-in-facebook-threat-case/

Scalora, M. J., Baumgartner, J. V., Zimmerman, W., Callaway, D., Maillette, M. A. J., Covell, C. N., Palarea, R. E., Krebs, J. A., & Washington, D. O. (2002). Risk factors for approach behavior toward the U.S. Congress. *Journal of Threat Assessment, 2*(2), 35-55.

Scheff, T. J. (1979). *Catharsis in healing, ritual, and drama.* Univ of California Press.

Searle John, R. (1965). What is a speech act? *Philosophy in America*, ed. Max Black. London: Allen and Unwin, 221-239.

Searle, J. R. (1969). *Speech acts: An essay in the philosophy of language*. Cambridge University Press.

Searle, J. R. (1975). Indirect speech acts. *Syntax and semantics, 3*(1), 59-82.

Searle, J. R. (1976). A classification of illocutionary acts. *Language in society, 5*(1), 1-23.

Searle, J. R. (2004). *Mind: A brief introduction*. Oxford University Press.

Shaw, T. (2015). Rejected too many times. *Mr. Talented.* Retrieved from: https://mrtalented.wordpress.com/2015/06/18/rejected-too-many-times/

Shuy, R. W. (1981). Topic as the unit of analysis in a criminal law case. *Analyzing discourse: Text and talk*, 113-126.

Shuy, R. W. (1993). *Language crimes: The use and abuse of language evidence in the courtroom.* Cambridge, MA: Blackwell.

Shuy, R. W. (2005). *Creating language crimes: How law enforcement uses (and misuses) language.* Oxford University Press.

Shuy, R. W. (2014). *The language of murder cases: Intentionality, predisposition, and voluntariness.* Oxford University Press.

Silva, R. S. (2013) 'Detecting plagiarism in the forensic linguistics turn', unpublished dissertation, Aston University.

Silverstein, M. (1976). *Hierarchy of features and ergativity* (pp. 112-171). Humanities Press.

Simons, A., & Tunkel, R. (2014). The assessment of anonymous threatening communications. *International handbook of threat assessment*, 195-213.

Singer, J. L. (2014). *Daydreaming and Fantasy (Psychology Revivals)*. Routledge.

Smith, S. (2006). *From violent words to violent deeds? Assessing risk from threatening communications.* Ph.D. dissertation. Georgetown University.

Smith, S., Woyach, R., & O'Toole, M. E. (2014). Threat triage: Recognizing the needle in the haystack. *International Handbook of Threat Assessment*, 321-329.

Smothers, R. (2006). Man, 20, arrested in stadium threat hoax. *The New York Times*. Retrieved from: http://www.nytimes.com/2006/10/21/us/21stadium.html?_r=0

Solan, L. M., & Tiersma, P. M. (2005). *Speaking of crime: The language of criminal justice.* University of Chicago Press.

Sornig, K. (1989). Some remarks on linguistic strategies of persuasion. *Language, power and ideology: Studies in political discourse, 7*, 95-113.

Southern Poverty Law Center. (2015). Active hate groups in the United States in 2015. SPL Center. Retrieved from: https://www.splcenter.org/fighting-hate/intelligence-report/2016/active-hate-groups-united-states-2015

Storey, K. (1995). The language of threats. *Forensic Linguistics 2*(1), 74-80.

Stubbs, M. (1996). Towards a modal grammar of English: a matter of prolonged fieldwork. *Text and Corpus Analysis*. Blackwell, Oxford.

Subramanian, C. (2017, October 6). Las Vegas shooting: What was Stephen Paddock's motive? *BBC News*. Retrieved from: https://www.bbc.com/news/world-us-canada-41518852

Tarasoff v. Regents of the University of California, 131 Cal. Rptr. 14. (Cal. 1976).

Terroristic Threat, Texas Penal Code § 22.07. (1973).

Texas Court of Appeals (2015). Carter, Ex Parte Justin River, PD-1291-15.

Thimm, C., Koch, S. C., & Schey, S. (2003). Communicating gendered professional identity: Competence, cooperation, and conflict in the workplace. *The handbook of language and gender*, 528-549.

Thompson, G. (2008). Appraising glances: evaluating Martin's model of APPRAISAL. *Word, 59*(1-2), 169-187.

Thompson, G. (2014). Affect and emotion, target-value mismatches, and Russian dolls. *Evaluation in context, 47*(66).

Torpy, B. (2016, May 8). Bill Torpy at Large: Ironic or not, her Facebook foolishness backfired. *The Atlanta Journal-Constitution*. Retrieved from: https://www.ajc.com/news/local/bill-torpy-large-ironic-not-her-facebook-foolishness-backfired/vrwClG1p7GnXzYfOuhk7vI/

Turner, J. T., & Gelles, M. (2003). *Threat assessment: A risk management approach*. The Haworth Press.

U.S. v. Massino, 356 F. Supp. 2d 227, 230 (E.D.N.Y. 2005)

United States Department of Justice. (2019, February 25). Federal Jury Convicts Man Who Attempted to Bomb Downtown Oklahoma City Bank. *Western District of Oklahoma.* Retrieved from: https://www.justice.gov/usao-wdok/pr/federal-jury-convicts-man-who-attempted-bomb-downtown-oklahoma-city-bank

Van Dijk, T. A. (1977). Semantic macro-structures and knowledge frames in discourse comprehension. *Cognitive processes in comprehension*, 3-32.

Van Leeuwen, T. (1996). The representation of social actors. *Texts and practices: Readings in critical discourse analysis, 1*, 32-70.

Williams, S. G. (2012). The ethics of internet research. *Online Journal of Nursing Informatics, 16*(2).

Yamanaka, N. (1995). On indirect threats. *International Journal for the Semiotics of Law, 8*(1), 37-52.

APPENDIX A

Justin Carter Facebook pledge.

APPENDIX B

Ethical decision-making tree (Williams, 2012).

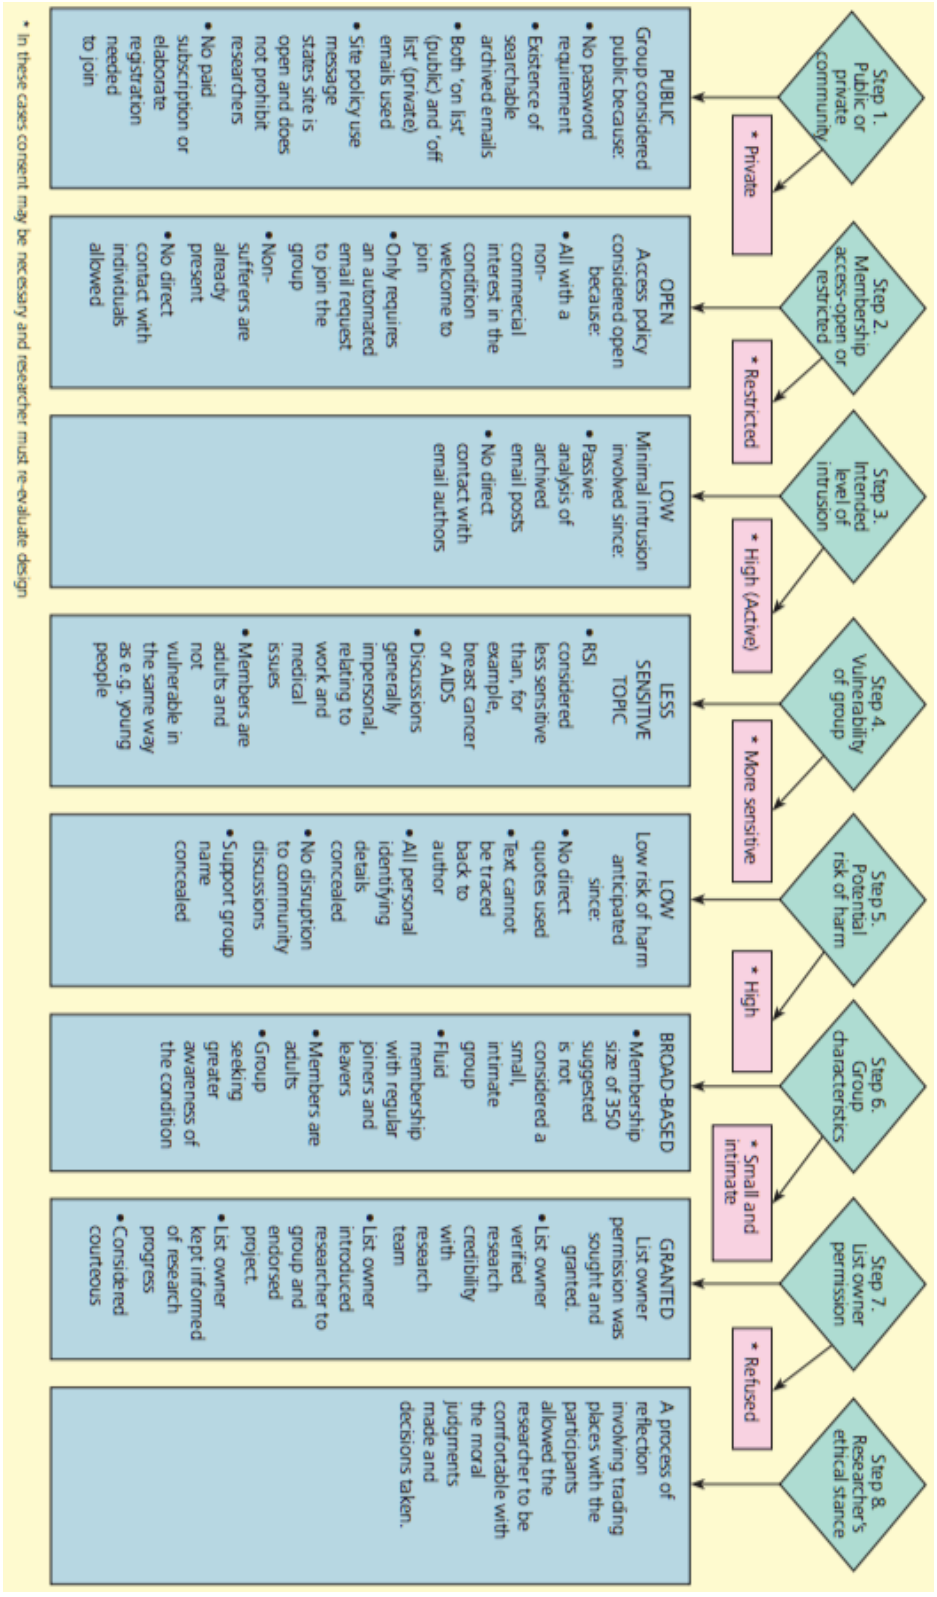

APPENDIX C

Jerry Varnell pledge.

What happened in Oklahoma city was not an attack on America, it was retaliation. Retaliation against the freedoms that have been taken away from the American people. It was a wake up call to both the government and the people. An act done to show the government what the people thinks of its actions. It is also a call to arms, to show people that there are still fighters among the American people. The time for revolution is now.

A final word on desirability and feasibility: