WHEN WORLDVIEWS COLLIDE: WHAT LINGUISTIC STYLE MATCHING AND DISTAL LANGUAGE REVEAL ABOUT DECEPTION IN POLITICAL DISCOURSE

by

Lucille M. Booker

A Dissertation

Submitted in Partial Fulfillment of the

Requirements for the Degree of

Doctor of Philosophy

Major: English

The University of Memphis

December 2012

Abstract

Booker, Lucille, M. PhD. The University of Memphis. December 2012.
When worldviews collide: What linguistic style matching and distal language reveal about deception in political discourse. Major Professor: Dr. Philip M. McCarthy.

Political discourse is an observable, measurable, and testable manifestation of political worldviews. However, when worldviews collide, notions of truth and of lies are put to the test. The challenge for researchers is how to establish confidence in their analysis. Despite the growing interest in deception research from a diversity of fields and industries, the trend is to focus on validating the assessment approach to the data without considering validity issues related to the data itself. Such a trend is concerned more with *how* to assess linguistic features and less with *what* is being assessed. By contrast, this dissertation is concerned with both the *how* and *what* of deception analysis. Of necessity, establishing validity of the data needs to be addressed first. To this end, I use the computational textual analysis tool of *the Gramulator* to facilitate validating a corpus of truthful texts and deceptive texts written by self-described liberals and conservatives. Specifically, I apply the *internal validity process (IVP)* to the corpus. The IVP comprises rigorous validity assessments of the *homogeneity* and the *markedness* of the data as well as the derived indices. This process aims to validate *what* is being assessed (i.e., truthful texts and deceptive texts). I then use the Gramulator to analyze the corpus. Specifically, I conduct a *linguistic style matching analysis* and a *distal language analysis* of the corpus. The linguistic style matching analysis assesses the degree to which deceivers and truth tellers from divergent political groups coordinate their language. By contrast, the distal language analysis assesses the degree to which those same groups differ in their pronoun and pronominal usage. These analyses reveal *how* computational approaches in combination with linguistic theory facilitate deception analysis. The results of my dissertation suggest that the IVP supplies compelling evidence for the validity of the corpus. These results also suggest that linguistic style matching and distal language usage may be deception strategies used to influence others into a false belief. Taken as a whole, my dissertation offers greater insight into how subtle differences in the framing of word choices facilitate the identification of prominent features of deception in political discourse.

Table of Contents

List of Tables

Chapter 1: Introduction

Deception research has advanced our understanding of a rather basic human behavior: lying in everyday life (DePaulo, Kashy, Kirkendol, Wyer, & Epstein, 1996). Intrinsic to understanding such a basic human behavior is *context*. That is, lying is a communicative norm and much can be understood about such a norm by assessing deception in context-specific social environments. This dissertation is concerned with such a task. More specifically, this dissertation is concerned with assessing deception within the social context of politics. To address such a concern, I conduct a deception analysis of political discourse.

Political discourse is a repository of communications expressed by politicians, political groups, and the public, many of whom represent contrasting worldviews (Lakoff, 2002; van Dijk, 1998). When such worldviews collide, notions of truth and of lies are put to the test. In other words, the same instantiations of political discourse may be interpreted differently by divergent political groups, and by different members of those divergent political groups. Consider the famous "You lie!" accusation expressed by Republican Representative Joe Wilson of South Carolina to Democratic President Barack Obama of the United States. During a joint session of Congress in 2009, President Obama stated that his administration's proposed health care reforms would not apply to immigrants who are in the United States illegally. On the cusp of that statement, Representative Wilson uttered "You lie!"to which President Obama retorted "It's not true." Clearly both politicians were accusing the other of perpetrating a lie. Such an accusation suggests that either one politician is lying while the other is communicating the truth, or perhaps that both politicians are lying. In either case, judgment about who is or is not lying has more to do with contrasting worldviews, and less to do with absolutism.

While the “You Lie!” example demonstrates both a breach in protocol as well as an exchange between contrasting worldviews, this example also demonstrates how political positioning (e.g., supporters and opposers) further complicates an understanding of such worldviews. In other words, President Obama’s claim suggests that he and his administration *oppose* health care benefits for illegal immigrants. However, Representative Wilson’s utterance suggests that President Obama and his administration *support* health care benefits for illegal immigrants. Thus, political discourse is a repository of collective views expressed by political groups and expressed by issue supporters and issue opposers across and within such groups.

Although political groups may trend towards a political position depending on the issue under consideration, political groups also comprise factions of members who stand for political positions that contrast with such trends. For example, *liberals* in the United States generally include more *supporters* than opposers of stricter gun control laws, whereas *conservatives* in the Unites States generally include more *opposers* than supporters of those same laws (Harris Interactive, 2010). By contrast, *conservatives* in the United States generally include more *supporters* than opposers of stricter immigration control laws, whereas *liberals* in the United States generally include more *opposers* than supporters of those same laws (The Pew Research Center, 2012). While political positioning may differ regionally (e.g., Southerners in the United States generally favor fewer restrictions on gun control)[1] or within political groups (e.g., conservatives are divided among supporters of migrant workers including access to citizenship for migrant workers, and supporters of the border enforcement-first approach)[2], it is reasonable to associate political groups with specific political positions based on their expressed political platforms.

In broad terms, political discourse functions as an integral resource for understanding political worldviews and how we construct and communicate such political worldviews. As such, although political discourse involves expressing our views on political issues, political discourse also involves a notion of the truth and lies that form the lens through which we express those views. Consequently, a deception analysis of political discourse offers an opportunity to assess the language used by divergent political groups to communicate deceptive discourse and how such language differs from the language used to communicate truthful discourse.

Theme and Focus of the Dissertation

This dissertation is a deception analysis that analyzes the prominent linguistic features of truthful discourse and deceptive discourse relative to each other. My concern is with instantiations of political discourse that are perceived to be representations of the truth or of a lie based on the worldview of the discourse participants. In other words, my focus is on assessing what discourse participants *perceive* or *believe* to be the truth or a lie even when existing factual information challenges and even contradicts those perceptions or beliefs (Nyhan & Reifler, forthcoming). To this end, I conduct an

[1] Brennan, P. G., Lizotte, A. J., & McDowall, D. (1993). Guns, Southerness, and gun control. *Journal of Quantitative Criminology*, 9, 3, 289-307.

[2] Hayworth, J. D. (2005). *Enforcement First Immigration Reform Act: H.R. 3938*. 109th Congress. http://www.govtrack.us/congress/bills/109/hr3938

experiment using the "Devil's Advocate Approach" (Leal, Vrij, Mann, & Fisher, 2012; McCarthy, Duran, & Booker 2012). With this approach, discourse participants communicate their political views under a truth condition and under a lie condition. However, under a lie condition, discourse participants suppress their personal beliefs to argue against their strongly held political views. Such an approach facilitates assessing differences in the language used by discourse participants under each condition (Leal et al., 2012).

My experiment involves a survey entitled, *Truth and Lies in Politics*. For this survey, the truth condition calls for anonymous participants to share their truthful views about *stricter gun control laws in the United States*, and *stricter immigration control laws in the United States*. These participants are instructed to share truthful views that represent what they believe to be "the truth, the whole truth, and nothing but the truth." The lie condition calls for anonymous participant to share their deceptive views about the same issues, but these participants are instructed to share deceptive views that *do not* represent what they believe to be "the truth, the whole truth, and nothing but the truth" (see Chapter 4: Methods). This task of sharing deceptive views about the issues under consideration is predicated on first knowing what one's truthful views are. In other words, such truthful views comprise *unmarked or default language*, which is simpler, or general language. By contrast, such deceptive views comprise *marked language*, which is more complex, or focused language (Battistella, 1990).

Political discourse comprises both default language and marked language expressed by political groups. However, such a combination of types of language poses challenges for political groups. That is, default language (e.g., truthful views) and marked language (e.g., deceptive views) are not created equal, and such inequality complicates the ease with which political groups refer to and exploit instances of such language types. More specifically, default language and marked language require different levels of cognitive complexity (Givón, 1991; Sweller, 1994). In other words, default language is easier to access and tends to require less working memory to process. By contrast, marked language is relatively harder to access and tends to require more working memory to process. This dissertation is concerned with the language of TRUTH (default language), and the language of LIES (marked language) across divergent political groups. More specifically, I analyze the distinguishing linguistic features of truthful views and deceptive views expressed by liberals and conservatives.

Research Question

My research sets out to answer the following question: "What does *linguistic style matching* and *distal language* reveal about deception in political discourse?" More specifically, I investigate *which features of language uniquely characterize the views of deceivers and truth tellers relative to each other; and of self-described liberals and self-described conservatives, relative to each other?*" To this end, my dissertation is concerned with two areas of interest: *linguistic style matching* and *distal language*. I provide an explanation of my usage of each term, and follow such explanations with the corresponding hypotheses.

Linguistic Style Matching. My use of the term *linguistic style matching* refers to the degree to which discourse participants match their word use (Niederhoffer & Pennebaker, 2002). This use of *linguistic style matching* is highly related to, or overlaps with *communication accommodation theory* (Giles, Coupland, & Coupland, 1991; Giles

& Smith, 1979), *interaction synchrony* (Bernieri, Davis, Rosenthal, & Knee, 1994; McDowall, 1978), and *interpersonal coordination* (Bernieri & Rosenthal, 1991; Fowler, Richardson, Marsh, & Shockley, 2008). What distinguishes linguistic style matching from the other references is that linguistic style matching is primarily concerned with coordination between discourse participants at *the word level* (e.g., pronoun usage, cohesion, negation). By contrast, the other references are concerned with coordination between discourse participants at the word level and at *the behavioral level* (e.g., facial expressions, gestures, pitch, pauses). Such a distinction has important implications in terms of how analyses of linguistic style matching are conducted compared to analyses of the other related references. More specifically, a focus on words (as opposed to gestures) allows for the use of NLP (natural language processing) algorithms and technologies. For example, Allen and Bryant (1996) used an adaptive parser to measure linguistic style matching at the phrase level. More specifically, the authors tested whether the adaptive parsing algorithm was appropriate for interpreting phrases that have undergone an incremental deviation process. With such a process, the authors deleted, transposed, inserted, and substituted words of a given phrase and analyzed the algorithm's ability to parse the deviations. For a more relevant example to this dissertation, Niederhoffer and Pennebaker (2002) used the computational textual analysis tool of *LIWC (Linguistic Inquiry Word Count)* to measure linguistic style matching between former President Nixon and his aides from transcribed conversations connected to the Watergate scandal of 1972. More specifically, the authors tested the efficacy of LIWC to identify the coordination of *prepositions, adverbs,* and *tentative words* between the discourse participants. Similarly, Danescu-Niculescu-Mizil, Gamon, and Dumais (2011) used

LIWC to measure linguistic style matching between discourse participants of Twitter conversations, formally known as *tweets*. More specifically, the authors test the efficacy of LIWC to identify coordination of style dimensions in the short tweets that are restricted by the Twitter platform (maximum of 140 characters). These style dimensions include *first person pronoun plural usage, prepositions*, and *words of certainty*. Taken as a whole, these studies demonstrate that linguistic style matching is a measurable phenomenon, and that NLP tools such as parsers and LIWC are appropriate for measuring such a phenomenon.

Linguistic style matching assumes that discourse participants share linguistic features at *the turn-by-turn level*, as well as at *the broader level of discourse* (Niederhoffer & Pennebaker, 2002). Such an assumption suggests that at the turn-by-turn level, linguistic style matching may occur between discourse participants expressing their views to each other (i.e., A and B are talking, or writing about a topic to each other). Such an assumption also suggests that at the broader level of discourse, linguistic style matching may occur between discourse participants expressing their views independent of each other (i.e., A and B are talking, or writing about a topic, but not necessarily to each other). For example, Niederhoffer and Pennebaker (2002) assess linguistic style matching between discourse participants engaged in direct, topic-centered conversations (i.e., Watergate scandal) at a turn-by-turn level, in real-time. By contrast, Danescu-Niculescu-Mizil and colleagues (2011) assess linguistic style matching between discourse participants engaged in public conversations, or tweets, at the broader level of discourse. Specifically, Danescu-Niculescu-Mizil and colleagues analyzed tweets that were not on the same topic, and not in real time. This dissertation is concerned with linguistic style matching at the broader level of discourse between deceivers and truth tellers. Specifically, I analyze the degree to which deceivers ascribed to a particular political group (e.g., liberal deceivers) match the linguistic style of truth tellers ascribed to another divergent political group (e.g., conservative truth tellers). Independent of each other, deceivers and truth tellers express their views about the topics of stricter gun control laws in the United States (a liberal issue), and stricter immigration laws in the United States (a conservative issue).

Linguistic style matching studies that use NLP tools such as LIWC focus on individual words in isolation (Danescu-Niculescu-Mizil et al., 2011; Gonzales, Hancock, & Pennebaker, 2010, Neiderhoffer & Pennebaker, 2002). Although I analyze words in isolation (e.g., pronouns) to compare my results with the results reported in such studies, I primarily focus on *words in context*. That is, I use the computational textual analysis tool of the *Gramulator* to analyze *bigrams*, or the consecutive sequences of *any two words* that have an above average frequency. The Gramulator is appropriate for identifying and processing contextualized linguistic features (e.g., bigrams) used to measure linguistic style matching.

My hypotheses on linguistic style matching are divided into two parts: Analysis 1 and Analysis 2. Analysis 1 tests four hypotheses, two of which pertain to the liberal issue of stricter gun control laws, and two of which pertain to the conservative issue of stricter immigration control laws:

Hypothesis 1a. On the liberal issue of stricter gun control laws, liberal liars will match the linguistic style of conservative truth tellers *to a greater extent than* liberal truth tellers will match the linguistic style of conservative truth tellers because under a liberal lie condition, liberal liars and conservative truth tellers *have a shared rhetorical goal*, whereas under a liberal truth condition, liberal truth tellers and conservative truth tellers *have contrasting rhetorical goals*.

Hypothesis 1b. On the liberal issue of stricter gun control laws, conservative liars will match the linguistic style of liberal truth tellers *to a greater extent than* conservative truth tellers will match the linguistic style of liberal truth tellers because under a conservative lie condition, conservative liars and liberal truth tellers *have a shared rhetorical goal*, whereas under a conservative truth condition, conservative truth tellers and liberal truth tellers *have contrasting rhetorical goals*.

Hypothesis 2a. On the conservative issue of stricter immigration control laws, conservative liars will match the linguistic style of liberal truth tellers *to a greater extent than* conservative truth tellers will match the linguistic style of liberal truth tellers because under a conservative lie condition, conservative liars and liberal truth tellers *have a shared rhetorical goal*, whereas under a conservative truth condition, conservative truth tellers and liberal truth tellers *have contrasting rhetorical goals*.

Hypothesis 2b. On the conservative issue of stricter immigration control laws, liberal liars will match the linguistic style of conservative truth tellers *to a greater extent than* liberal truth tellers will match the linguistic style of conservative truth tellers because under a liberal lie condition, liberal liars and conservative truth tellers *have a shared rhetorical goal*, whereas under a liberal truth condition, liberal truth tellers and conservative truth tellers *have contrasting rhetorical goals*.

To allow for processing such complex hypotheses, I show the logical form of each argument under individual representations. The hypotheses from Analysis 1 are formalized in Table 1. Liberals are represented by *L*, and conservatives are represented by *C*. I use the negation symbol ($^-$) to represent *liars*, the addition ($^+$) symbol to represent *truth tellers*, and the pound (#) symbol to represent linguistic style matching. The greater than sign (>) is used to hypothesize the greater degree of a given phenomenon under compared conditions. For example, the representation of (L^- # C^+) > (L^+ # C^+) reads, "liberal liars match the linguistic style of conservative truth tellers to a greater extent than liberal truth tellers match the linguistic style of conservative truth tellers." All representations similar to the given example are to be read in the same way as the given format.

Table 1

Representations of Analysis 1 (H1a, H1b, H2a, and H2b)

Linguistic Style Matching	
H1a: Stricter Gun Control Laws	(L^- # C^+) > (L^+ # C^+)
H1b: Stricter Gun Control Laws	(C^- # L^+) > (C^+ # L^+)
H2a: Stricter Immigration Control Laws	(C^- # L^+) > (C^+ # L^+)
H2b: Stricter Immigration Control Laws	(L^- # C^+) > (L^+ # C^+)

The four hypotheses of Analysis 1 predict that a greater degree of linguistic style matching occurs between discourse participants who share the same rhetorical goal. For example, H1a predicts that on the liberal issue of stricter gun control laws, there will be a greater degree of linguistic style matching between liberal liars and conservative truth tellers (L^- # C^+) than between liberal truth tellers and conservative truth tellers (L^+ # C^+)

because liberal liars and conservative truth tellers ($L^- \# C^+$) have a shared rhetorical goal: *to oppose stricter gun control laws*. However, liberal truth tellers and conservative truth tellers ($L^+ \# C^+$) have contrasting rhetorical goals: *liberal truth tellers support stricter gun control laws,* whereas *conservative truth tellers oppose such laws*. Thus, Analysis 1 assumes that evidence of a greater degree of linguistic style matching between two groups is explained by shared rhetorical goals.

Analysis 2 builds from Analysis I. That is, Analysis two assesses whether conservative liars or liberals liars are better at matching the linguistic style of their corresponding truth tellers. In other words, when deceivers and truth tellers from divergent political groups share the same rhetorical goal, is there a greater degree of linguistic style matching between conservative liars and liberal truth tellers, or between liberal liars and conservative truth tellers. The hypotheses of Analysis 2 are as follows:

Hypothesis 3a. On the liberal issue of stricter gun control laws, conservative liars will match the linguistic style of liberal truth tellers to a greater extent than liberal liars will match the linguistic style of conservative truth tellers because conservatives are better than liberals at framing deceptive discourse.

Hypothesis 3b. On the conservative issue of stricter immigration control laws, conservative liars will match the linguistic style of liberal truth tellers to a greater extent than liberal liars will match the linguistic style of conservative truth tellers because conservatives are better than liberals at framing deceptive discourse.

Hypotheses 3a and 3b of Analysis 2 are formalized in Table 2. The representations of both hypotheses are the same. However, these representations correspond to different issues under consideration.

Table 2

Representations of Analysis 2 (H3a and H3b)

Linguistic Style Matching	
H3a: Stricter Gun Control Laws	$(C^- \# L^+) > (L^- \# C^+)$
H3b: Stricter Immigration Control Laws	$(C^- \# L^+) > (L^- \# C^+)$

The two hypotheses of Analysis 2 predict that there will be a greater degree of linguistic style matching between conservative liars and liberal truth tellers who share the same rhetorical goal than between liberal liars and conservative truth tellers who share the same rhetorical goal. In the case of H3a, conservative liars and liberal truth tellers *support* stricter gun control laws. In the case of H3b, conservative liars and liberal truth tellers *oppose* stricter immigration control laws.

In sum. Analysis 1 on linguistic style matching tests whether a lie condition has been satisfied for the experiment. By extension, Analysis 2 on linguistic style matching tests whether conservatives are more convincing than liberals under such a condition. In other words, Analyses 1 addresses the following question: *Are liars lying?* And Analysis 2 addresses the following question: *Are conservative liars more convincing than liberal liars?* Taken as a whole, the linguistic style matching analyses is concerned with the performance of discourse participants on two levels: the performance of liars at the level of the group, and the performance of liars across groups.

Distal Language. My use of the term *distal language* is grounded in the concept of *deixis*. Deixis refers to the interrelation between lexical items and the context in which those items are expressed (Levinson, 2006). That is, distal language is determined by a context, and that context minimally involves someone who communicates a message and someone who receives that communication. According to Rauh (1983), the communicator "*encodes*" the message and the receiver "*decodes*" the message. As such, lexical items that comprise any message are encoded relative to the communicator, or the *agent* of the discourse. In other words, the perspective of the agent represents the center of orientation for the message (Rauh, 1983). Thus, distal language is based on a notion of relativity that depends on the perspective of the agent.

Such an understanding of distal language as a relative concept implies that distal language is diametrically opposed to *proximal language*. For example, the pronoun *my* (e.g., *my values*) generally suggests that a given noun or abstraction belongs to, or is closely related to the agent (proximal), whereas the pronoun *their* (e.g., *their values*) generally suggests that a given noun or abstraction belongs to, or is closely related to some other person(s) independent of the agent (distal). Pronouns are examples of *deictic expressions* because the meaning (i.e., distal or proximal) of pronouns depend on or relate to the agent of the discourse.

Distal language in my dissertation is used as a general category of deictic expressions in which distal personal, social, spatial, and temporal information is encoded. I focus on the deictic expressions of pronouns and pronominals. More specifically, I use the Gramulator to measure deceivers' use of deictic pronoun and pronominal expressions relative to truth tellers use of such expressions. I also measure the significance of those expressions in the LIE corpus relative to the TRUTH corpus using the statistical procedures of the Fisher's Exact test for frequencies and independent and paired *t*-tests for derived values.

My hypothesis on distal language is as follows:

Hypothesis 4. More *distal language* will be used to communicate deceptive views because deceptive discourse is characterized as being distanced from the self. That is, deceivers are more likely to refer to primary figures other than themselves and to use more distal pronouns and pronominals in their deceptive discourse.

Hypothesis 4 is based on the results observed in deception studies including Duran, Hall, McCarthy, and McNamara (2010), Hancock, Curry, Goorha, and Woodworth (2008), and Newman, Pennebaker, Berry, and Richards (2003). These studies suggest that deceivers use fewer self-references and more other-directed pronouns than truth tellers as a way to distance themselves from their deceptive views. Hypothesis 4 is also based on the results observed by Duran, Hall, and McCarthy (2009) on the language used by pro-life supporters in contrast to pro-choice supporters. Their analysis suggests that truth tellers (pro-life and pro-choice) are more likely to use referents that are relationally close to them (e.g., *my mom, my dad,*), but deceivers avoid such easily verifiable information. Instead, deceivers are more likely to use referents that are relationally distant from them, and harder to identify (e.g., *high school friend*).

To test Hypothesis 4, I conduct two analyses. Analysis 1 assesses discourse participants (i.e., liberals and conservatives) explicit choice of primary figures in their deceptive discourse. That is, participants were instructed to provide one example of an event or a story to support their truthful views and their deceptive views, respectively. The supporting example could primarily involve the agent of the discourse or some other person(s). Such an explicit choice involves some degree of conscious processes (Hagège, 2005) that may reveal distinguishing features of deceptive discourse. By contrast

Analysis 2 assesses discourse participants (i.e., liberals and conservatives) implicit choice of pronouns and pronominals in their deceptive discourse, relative to their truthful discourse. Such an implicit choice involves some degree of unconscious processes (Arrivé, 1992). Thus, Analysis 1 and Analysis 2 assess distinguishing features of political discourse involving explicit and implicit lexical choices.

In sum. Analysis 1 on distal language addresses deceivers' explicit choice of distal language usage under a lie condition. By contrast, Analysis 2 on distal language addresses deceivers' implicit choice of distal language usage under a lie condition. Taken as a whole, the distal language analyses are concerned with the manifestation of conscious and unconscious cognitive processes.

Overview of the Chapters

In this section, I provide an overview of Chapters 2 – 8 of the dissertation. In addition to describing the purpose and thematic focus of each chapter, I also address the main concerns or issues of each chapter. Combined, these chapters provide evidence that new technologies, rigorous validation procedures, and corpus-based contrastive analysis contribute to a greater understanding of the prominent linguistic features of deception in political discourse.

The second chapter of this dissertation is a literature review of deception research. The purpose of this review is to situate my deception analysis of political discourse into the context of existing methods and approaches to deception detection studies. Accordingly, the deception detection studies summarized in this review address three themes: *behaviorism*[3], *natural language processing (NLP)*, and *applied natural language processing (ANLP)*. The chapter is concerned with the evolution of contemporary computational approaches to deception analysis.

The third chapter is a review of the Gramulator, which is the computational textual analysis tool used in this deception analysis. This review introduces the Gramulator and summarizes studies that feature the Gramulator. Ultimately, the chapter is concerned with providing confidence in the efficacy of the Gramulator to analyze deceptive discourse.

The fourth chapter is a detailed description of the methods used in this deception analysis. That is, I describe the mixed methods approach adopted to design and conduct the experiment of this dissertation and to prepare and assess the data collected as a result of the experiment. Such an approach is concerned with the use of new technologies to experimental research.

The fifth chapter is a validation study. The purpose of this study is to establish internal validity of the corpus under consideration. To this end, I introduce the internal validation process (IVP) and apply such a process to the data comprising the truthful and deceptive views of liberals and conservatives (referred to as the TRUTH and LIE corpora) and the data comprising the political views of issue supporters and issue opposers (referred to as FOR_IT and AGAINST_IT corpora). Taken as a whole, the chapter is concerned with providing confidence in the assessments of this deception analysis.

---

[3] The use of the term "behaviorism" in this dissertation should be interpreted broadly to refer to *behaviorial studies* and *behavioral measures*. That is, "behaviorism" here is not necessarily referring to the particular school of thought associated with such names as B. F. Skinner and John Watson.

The sixth chapter is a linguistic style matching analysis. The chapter assesses the degree to which deceivers and truth tellers from divergent political groups coordinate their language. To this end, I test three hypotheses on linguistic style matching across the two political issues under consideration: *stricter gun control laws in the United States* and *stricter immigration control laws in the United States*. In addition, I discuss the results generated from testing these hypotheses (H1a, H1b, H2a, H2b, H3a, and H3b) and address limitations and future analyses. Specific to this dissertation is the use of linguistic style matching to test that a lie condition has been satisfied. As such, the linguistic style matching analysis addresses the pervasive issue in deception analysis, which is the need to provide greater confidence that participants assigned to a lie condition are indeed lying.

The seventh chapter is a distal language analysis. The chapter assesses distal language usage by deceivers relative to truth tellers. To this end, I test the fourth hypothesis on distal language, which predicts that deceivers use more distal language than truth tellers. In addition, I discuss the results generated from testing the fourth hypothesis (H4) and address limitations and future analyses. Overall, the distal language analysis is concerned with the explicit and implicit lexical choices of deceivers.

The dissertation ends with the eighth chapter, which is a general discussion. The chapter synthesizes the findings from the validation study, the linguistic style matching analysis, and the distal language analysis. In addition, this chapter addresses limitations of the study, and topics for future research that were not addressed in previous chapters. In sum, the chapter responds to the research question and features the Gramulator as a versatile computational tool appropriate for deception analysis.

Relevance of the Dissertation. This dissertation is relevant to researchers and practitioners from a wide range of fields and industries concerned with deception research and with contemporary approaches to deception research. Fields ranging from applied natural language processing (ANLP) and second language acquisition through to criminal intelligence, and industries ranging from fiction writing and data analytics through to politics have contributed to the growing interest in deception research. Even the popular TED talks (www.ted.com), a web-based forum for sharing innovative and groundbreaking ideas has featured talks by deception analysts (e.g., Meyer, 2011). However, popularity is by no means a measure of validity. Instead, the growing interest in deception research has generated diminishing interest in internal validation procedures in deception research. It therefore follows that the deception analysis of this dissertation involves rigorous internal validation assessments and the use of computational algorithms that incorporate statistical procedures with linguistic theory. This deception analysis is concerned with redirecting researchers and practitioners toward more reliable empirical research and away from often unreliable popular techniques. As such, this deception analysis is relevant to fields and industries interested in advances in technology (e.g., Gramulator) used to detect the prominent features of deceptive discourse. Such technology facilitates understanding how subtle changes in word choice and the framing of those choices characterize the language of deceivers and truth tellers relative to one another and the language of liberal deceivers and conservative deceivers relative to one another.

Chapter 2: Literature Review

Deception research is dedicated to the analysis of deceptive acts, particularly lying. More specifically, deception research is concerned with the language of lies (e.g., structures, features, and properties), and the form of language used to communicate those lies (e.g., spoken, written, or gestural). The purpose of this review of deception research is two-fold: to summarize studies that address a diversity of approaches and methods used in deception research, and to provide an account of the recent advances that have shaped and directed those approaches and methods. To this end, I have organized this review into two main sections: (1) Background and (2) Summaries. In the Background section, I describe the pioneering approaches and methods used in deception research. In the Summaries section, I describe and discuss several of the key studies that help to define the field.

Background

Deception research has traditionally focused on the credibility of written statements (e.g., witness statements, victim statements, criminal statements) and the accuracy of human judges (including interrogation officers, and experimental participants). It follows that researchers interested in analyzing written statements for the benefit of human judges have introduced and tested several methods for such purposes. Generally, these methods build from the hypothesis attributed to Undeutsch (1967), which predicts that the characteristics of truthful accounts of actual experiences vary from those of deceptive accounts of fabricated experiences.

Among the widely used methods based on the *Undeutsch Hypothesis* (coined by Steller, 1989) are *criteria-based content analysis (CBCA)*, and *scientific content analysis*

(SCAN) (Gordon & Fleisher, 2006). Criteria-based content analysis (CBCA) is based on a number of criteria proposed by Steller and Köhnken (1989), which include general characteristics (e.g., *structured* and *unstructured content*), specific content (e.g., *contextual embedding, reproducing conversations, describing specifics of events*), peculiarities of content (e.g., *describing unusual* or *superfluous details*), motivation related content (e.g., *admitting failure to recall all details, self-doubt about testimony*), and offence specific elements (e.g., *details only available to witnesses*). Turning to the second method, Scientific content analysis (SCAN) differs from CBCA perhaps only in its focus, rather than its goals. SCAN was developed by Avinoam Sapir (Sapir, 1987) and is applied to written statements. SCAN is also based on a number of criteria that include a change in the language addressing important references (e.g., *family members, other people, modes of transportation*), the placement of emotions in the statement, inconsistent pronoun usage (e.g., switching from *he* to *they* when referring to the same subject), failure to remember details, and details that are out of sequence (Sapir, 1987). Although SCAN and CBCA are considered methods of detecting deception, SCAN does not claim to distinguish deceivers from truth tellers. Instead, SCAN is intended to guide interviewers as they challenge and question details of a statement (Shuy, 1998).

As CBCA and SCAN became widely used in law enforcement and academic research, alternative methods emerged including Reality Monitoring and Interpersonal Deception Theory. Reality Monitoring is the process by which accounts from memory are distinguished as being either externally derived experiences or *internally* derived experiences (Johnson & Raye, 1981). According to Reality Monitoring, externally derived memories are described using more sensory words and references (e.g., *visual* *stimulants, voices, sounds, odors, climate*), temporal and spatial references (e.g., *day of the week, part of town, the floor of a building*), and affect words (e.g., *fear, anger, sorrow, joy, disgust, frustration*). By contrast, internally derived memories are described using language that involves cognitive processes (e.g., *reasoning, thoughts, and imagining*). Interpersonal Deception Theory defines deception as a strategic communicative act between deceivers and respondents. Interpersonal Deception Theory is based on a set of 18 propositions that involve distinguishing between what deceivers and respondents think and do, the different strategic moves used by deceivers and respondents, and interaction between the role of suspicion and the behaviors of deceivers and respondents.

Summaries

The deception studies summarized in this literature review introduce or investigate approaches and methods of distinguishing truthful language from deceptive language (e.g., Buller & Burgoon, 1996; Stellar & Köhnken,1989). The approaches and methods considered in these deception studies are organized into three themes: *behaviorism*, *natural language processing (NLP)*, and *applied natural language processing (ANLP)*.

Deception detection studies in behaviorism analyze behavioral patterns of deceivers relative to truth tellers (e.g., Broadhurst & Cheng, 2005; Burgoon & Qin, 2006; DePaulo, Kashy, Kirkendol, Wyer, & Epstein, 1996). The types of behavioral patterns analyzed may be non-verbal patterns (e.g., emotions and gestures) and/or verbal patterns (e.g., pronoun usage and vocabulary diversity). Deception detection studies in behaviorism often contribute to the theoretical framework used by deception detection studies in natural language processing (NLP). Although often related in terms of *content*, deception studies in behaviorism and in NLP are less related in terms of their respective *goals*. More specifically, deception detection studies in behaviorism tend to have a more *altruistic* goal, whereas deception detection studies in NLP tend to have a more *directed* goal. That is, deception detection studies in behaviorism is more about deepening our understanding of the topic at hand, which is *deception*, for the general benefit of researchers and practitioners of the field (e.g., Broadhurst & Cheng, 2005; Burgoon & Qin, 2006; DePaulo, Kashy, Kirkendol, Wyer, & Epstein, 1996). By contrast, deception detection studies in NLP is more about deepening our understanding of the topic at hand for the benefit of programmers. More specifically, deception detection studies in NLP have the goal of facilitating programmers interested in *developing* deception detection software (e.g., Carlson, George, Burgoon, Adkins, & White, 2004; Hancock et al., 2004; Zhou, Burgoon, Nunamaker, & Twitchell, 2004). As such, the goal of deception detection studies in NLP is more related to those in ANLP than in behaviorism. The goal of deception detection studies in ANLP is to use *existing* computational tools to advance how we approach understanding the topic at hand (e.g., Duran et al., 2009; Newman, Pennebaker, Berry, & Richards, 2003; Toma & Hancock, 2010;). Thus, the combined deception detection studies across the themes of behaviorism, NLP, and ANLP provide a more holistic overview of the literature.

Although interconnections exist between the studies on deception detection in behaviorism, NLP and ANLP, the literature addressing each of these themes has largely evolved independently. For a greater appreciation of the breadth of the literature, I summarize key studies within each of these three themes as a means for advancing the level of *interdiscursivity* (Bhatia, 2007) in deception research. That is, I draw evidence from disparate disciplines to encourage a more comprehensive awareness of the complexities of deception research.

Deception Detection Studies in Behaviorism

DePaulo, Kashy, Kirkendol, Wyer, and Epstein (1996). Central to the literature on deception detection in behaviorism is the study on lies in daily communications conducted by DePaulo and her colleagues (1996). The authors analyze participants' self-reported ratings of emotion and stress levels when communicating lies and other social interactions. The purpose of their study is to provide evidence in support of the claim that the act of lying is a social norm that is prevalent in our daily communications. To this end, the authors test four hypotheses on the behavioral patterns of deceivers relative to truth tellers. The hypotheses are that the participants 1) report low levels of emotional stress when lying; 2) claim to tell fewer lies than their peers; 3) pretend to feel more positive than negative when lying; and 4) lie more for psychological reasons than for material reasons or personal gain. The results of both independent experiments support all of the authors' predictions.

Approach. The approach used by DePaulo and her colleagues (1996) draws primarily on social theory, particularly as it relates to identity, self-presentation, impression management, and politeness (e.g., Brown & Levinson, 1987; Goffman, 1959; Schlenker & Weigold, 1989). This approach is used to identify and measure non-verbal behavioral indicators of deceptive discourse. More specifically, this approach emphasizes the behaviors used to change or transform the "self" for the communicative purposes (e.g., truth telling, or lying) on a daily basis.

Experiment. DePaulo and her colleagues (1996) conduct two experiments (Experiment 1 and Experiment 2). More specifically, the authors compare and contrast the results of a single experiment taken by two independent groups of participants. In Experiment 1, participants comprise traditional college students. In the Experiment 2, participants comprise diverse members of the local community college.

The number of participant from Experiment 1 and Experiment 2 are 77 and 70, respectively. The participants of Experiment 1 include 30 male students and 47 female students ranging from 17 to 22 years of age. These participants received partial course credit for an introductory psychology course. The participants of Experiment 2 included 30 male and 40 female community college recruits who were paid $35 to participate in the experiment. The participants from both experiments were instructed to record all of their social interactions and all of the lies they told for seven days. The social interactions recorded were to last at least 10 minutes, but no time limits were imposed on the length of interactions involving lies. Interactions were categorized as having occurred in writing, by telephone, or in person.

The participants were given explicit instructions on how to qualify a social interaction as a lie. According to the authors, intentionality and deception were required to qualify a social interaction as a lie. They defined a lie as any instance in which participants intentionally mislead some other person(s), using verbal or non-verbal communication. Participants were given many examples of lies. A typical example that *did not* qualify as a lie, and that was not to be recorded as such, was responding ‘fine,’ when indeed you were not, to common greetings (i.e., *How are you?*). In the recorded entry, participants were allowed to exclude the contents of a lie that they did not want to share. This recorded entry allowed the authors to count the lie even though they were not able to analyze the contents of the lie. All recorded entries of social interactions and lies were anonymous and tracked using an identification number chosen by the participants.

The format used to record the social interactions and lies was based on the Rochester Interaction Record (RIR) developed by Wheeler and Nezlek (1977). A 9-point scale was used to record the level of intimacy per interaction, which ranged from (1) superficial to (9) meaningful. The general quality of the interaction was rated with the lowest level of (1) unpleasant to the highest level of (9) pleasant. The recorded entries for lies included additional 9-point rating scales that measured the degree of planning involved with each lie, the importance of not getting caught telling the lie, the level of comfort experienced in telling the lie, the gravity of the lie, and the reaction to the lie on the part of the recipient(s). Participants also rated how they (the participants) would have felt if the truth were shared instead of the lie, and how the recipient(s) would have felt if the truth were instead shared.

After recording their social interactions and lies for seven days, the participants indicated whether the lies recorded were discovered by the recipients of the lies (i.e., *no; not yet; don't know;* or *yes*) and whether they would retell the lie if they were able to relive the moment (i.e., *no,* or *yes*). On a 9-point scale, participants also rated how frequently they lied compared to their own expectations, and compared to their peers. The lies recorded by the participants of both experiments totaled 215 and were categorized by type (*outright; exaggeration; subtle*), by content (*feelings; achievements/failures; actions/plans; explanations for behavior; facts/events*), by the referent of the lie (*liar; target; other person; object/event*), and by the reason for the lie (*self-oriented to advantage or protect the liar; other-oriented*).

Results. The authors reported that participants from Experiment 1 lied in approximately one out of three recorded social interactions and the participants from Experiment 2 lied in approximately one out of five recorded social interactions. According to the authors, these results suggested that lying is an everyday event in life. The results from both experiments provide evidence in support of the first hypothesis (*Participants report low levels of emotional stress when lying.*), and the second hypothesis (*Participants claim to tell fewer lies than their peers.*). Overall, the mean rating for the participants of Experiment 1 and 2 were $t(76) = 2.57$, $p = .01$, and $t(70) = 2.22$, $p = .03$, respectively. In terms of the third hypothesis, the results provided evidence in support of the authors' prediction: *Participants pretended to feel more positive than negative when lying*. The mean score for the participants of Experiment 1 pretending to feel more positive was 25.23%, and the mean score for the participants of Experiment 2 was 23.79% . These mean scores were significantly greater than the mean scores for the participants pretending to feel more negative, which corresponded to 2.98% of the lies told by participants of Experiment 1, and 1.19% of the lies told by the participants of Experiment 2. Women were more likely than men to pretend to feel positive when lying (Women = 28.87%, Men = 19.33%).

The results from a mixed-model ANOVA provided evidence in support of the fourth hypothesis: *Participants lie more for psychological reasons than for material reasons or personal gain*. The mean scores for psychological lies communicated by the participants of Experiment 1 and Experiment 2 was $M = 26.78$ ($SD = 20.83$), and $M =$

36.72 (*SD* = 31.26), respectively. These results were lower than the mean scores of personal advantage lies, which were *M* =18.61 (*SD* = 22.24), and *M* =19.95 (*SD* = 22.69), respectively.

Summary. The results reported by DePaulo and her colleagues (1996) supported all four hypotheses. Thus, participants experienced low stress levels were when lying. Their self-reported lie frequencies were less than their peers. The participants were more likely to pretend to feel positive when lying. And, finally, participants told more lies to benefit others or to protect the feelings of others than for material reasons or personal gain. The study provides evidence for the claim that people lie more than they might think; that they are comfortable with lying; and that the act of lying is generally a normal part of daily life.

Burgoon and Qin (2006). Burgoon and Qin (2006) contribute to deception detection in behaviorism with their study on interpersonal deception. The authors focus on *verbal patterns* of deceptive discourse (e.g., number of words, personal pronoun usage). This focus on verbal patterns contrasts with the previous study by DePaulo and her colleagues (1996), which focused on *non-verbal patterns* of deceptive discourse (e.g., emotion and stress levels). Burgoon and Qin (2006) examine the effect of temporal changes and sequential order on truthful and deceptive interactions. The purpose of their study is to provide evidence in support of the claim that verbal patterns differ across time and sequence. To this end, the authors test two hypotheses in their study: (1) The rates of verbal behavioral indicators are more divergent at the onset of truthful and deceptive discourse, than at the end of those same discourses; and (2) behavioral indicators differ depending on the articulation order of truthful and deceptive discourse. The results of the study provide evidence in support of the authors' argument that time and sequence effect the rate of behavioral indicators (e.g., *sentence length, repetition, and third-person pronouns*).

Approach. The approach used by Burgoon and Qin (2006) builds from the research on powerful and powerless speech styles, and message effectiveness (e.g., Bradac, 1990; Bradac & Mulac, 1984). This approach is used to analyze verbal behavioral indicators of deceptive discourse. More specifically, this approach is used to analyze 23 verbal behavioral indicators, and compare verbal patterns *at the onset* of truthful and deceptive communications, with verbal patterns *at the end* of truthful and deceptive communications.

Experiment. Burgoon and Qin (2006) used the data collected from an experiment conducted by Burgoon, Buller, White, Afifi, and Buslig (1999). The participants comprised graduate students and a mix of community members based in a large southwestern city. The participants were 26 years old and older in age. They were paired in random order, forming 61 partners (32 male-female pairs; and 29 same gender pairs, of which 14 were male-male pairs, and 15 were female-female pairs). The participants answered 12 interview questions that included questions on educational background, occupation, career aspirations, political issues of concern, situational questions (i.e., *What would you do with a wallet containing $1,000 that you found with no identification?*), and moral questions (i.e., *What would you do if your best friend was cheating on his or her spouse*). Responses to the interview questions were solicited in one of the following alternating sequences: truth-deception-truth-deception (referred to as *truth-first*), or deception-truth-deception-truth (referred to as *deception-first*). Three responses corresponded to each *Block* of truthful and deceptive discourse. The Block number corresponds to the order in which the discourse was delivered (Block 1 represents the discourse articulated before Block 2, etc.).

Results. Burgoon and Qin (2006) report results for a series of verbal indicators within categories that include: (1) *Quantity* (2) *Complexity*, (2) *Diversity*, and (3) *Specificity*. Each category corresponds to a series of measurable indicators. For example, number of words, number of verbs, and number of sentences are measures of *quantity*.

In terms of *Quantity*, deceivers across sequence generally used fewer words than truth tellers across sequences. However, truth-first participants did not display the same behavioral patterns over time. In Block 1, truth-first sequence truth tellers used *fewer* words than deceivers ($p < .0001$). Later, in Block 2, those same truth-first participants *increased* the number of words used when telling the truth than when deceiving ($p = .009$). In other words, truth tellers at the onset used fewer words than deceivers, but increased the number of words they articulated by the end of the discourse. This increase over time for the truth-first truth tellers is evidence in support the first hypothesis of the study: *The rates of verbal behavioral indicators are more divergent at the onset of truthful and deceptive discourse, than at the end of those same discourses*.

In terms of *Complexity*, the results revealed that truth tellers in the deception-first sequence used significantly shorter sentences on average than truth tellers in general ($p = .001$). The authors presumed that participants were more comfortable as time passed and therefore used a more simplified sentence structure. In other words, if sentences were on average shorter, then participants were presumably more comfortable. By contrast, the authors report that deceivers were less likely to use longer sentences in a deception-first sequence than deceivers of a truth-first sequence. These results are evidence in support of the second hypothesis: *Behavioral indicators differ depending on the articulation order of truthful and deceptive discourse.* Based on these results, the authors claim that for truth-first participants, behavioral indicators that are greater in truthful discourse are more likely to be approximated when those same participants switch to deceptive discourse.

In terms of *Diversity*, the results revealed that truthful discourse was more representative of diverse language than deceptive discourse ($p < .05$). By the end of the discourse, the authors reported that truthful discourse did not have as diverse vocabulary as deceptive discourse because the vocabulary used over time by truth tellers was less diverse than the vocabulary used over time by deceivers. These results support the first hypothesis of the study: *The rates of verbal behavioral indicators are more divergent at the onset of truthful and deceptive discourse, than at the end of those same discourses*. Although not the focus of the hypothesis, it appears that deceivers, on the one hand, become more comfortable and use simplified language over time (less complexity) and, on the other, use more diverse vocabulary (more diversity) than truth tellers over time. This observation suggests that deceivers use simplified, more diverse language over time.

In terms of *Specificity*, the results revealed that truth tellers generally used more first-person pronouns than deceivers. Deceivers used more second-person pronouns than truth tellers in Block 1 ($p < .0001$), but showed no difference in Block 2. The latter part of this observation does not support the authors' first hypothesis: *The rates of verbal behavioral indicators are more divergent at the onset of truthful and deceptive discourse, than at the end of those same discourses*. By contrast, participants of the deception-first sequence used more third-person pronouns in their truthful discourse than in their deceptive discourse (p = .039), but there was no significant differences reported in the truth-first sequence. These results support the authors' second hypothesis: *Behavioral indicators differ depending on the articulation order of truthful and deceptive discourse*.

Summary. Taken as a whole, Burgoon and Qin (2006) provide evidence in support of the claim that behavioral patterns in truthful and deceptive discourse are complex and rarely constant. Generally, the results supported the authors' first hypothesis that behavioral patterns vary from the onset to the end of truthful and deceptive discourse. However, interpreting these varying patterns proved to be challenging, particularly in terms of the *complexity* and *diversity* of language in use. Furthermore, the results did not generally support the second hypothesis. Instead, they suggested that the deception-first sequence unwittingly served as a pre-conditioning environment for any discourse that followed. The authors recognized that their study did not provide a more accurate method for deception detection. Instead, their study underscored the importance of analyzing linguistic measures over time and sequence to gain a deeper understanding of behavioral patterns in truthful and deceptive discourse.

Broadhurst and Cheng (2005). Broadhurst and Cheng (2005) divert from the common practice of analyzing the communicative practices of English speakers, and mono-lingual English speakers in particular. Instead, the authors analyze the communicative practices of bilingual speakers. The purpose of their study is to provide evidence in support of the claim that first and second language abilities have important implications on the predictability of human judges to identify lies in a bilingual context. The authors tested two hypotheses: 1) Speakers produce different verbal and non-verbal deception cues depending on the language spoken, in this case Cantonese and English;

and 2) judges are better at identifying deceptive discourse of participants speaking in Cantonese, their first language, than in English, their second language. Generally, the results support the first hypothesis because behavioral cues displayed by speakers were different in Cantonese than in English. The results did not support the second hypothesis because judges were better detectors of deception expressed in English than in Cantonese.

Approach. Broadhurst and Cheng (2005) use Ekman's (1999) universalist theory on emotion to analyze deceptive communications expressed by bilingual speakers of English and Cantonese. Ekman claims that basic emotions (e.g., *contempt, guilt, relief*) are associated with particular facial expressions, and are not specific to culture. Ekman's claim suggests that emotions and their associated facial expressions may be used to detect deceptive discourse from truthful discourse. In response to this claim and what it suggests, Broadhurst and Cheng asses the reliability of verbal and non-verbal indicators to detect deception communicated in English and Cantonese by bilingual speakers.

Experiment. The experiment conducted by Broadhurst and Cheng (2006) is based on a modified Mehrabian false-opinion paradigm proposed by Frank and Ekman (1997). Broadhurst and Cheng conduct an empirical analysis of verbal and non-verbal behavioral cues that are measured with the purpose of detecting deceptive discourse.

The stimulus material for the experiment consisted of a total of 20 video tapes of undergraduate students (20 – 22 years). The students self-rated their fluency in English and Cantonese at 4 or better on a 7-point Lickert scale (1 = very poor, 7 = very good), and were randomly selected to share deceptive and truthful opinions on debatable issues including capital punishment and the rights of same sex partners. One Cantonese-English bilingual experimenter posed the same set of questions to each participant. The experiment also called for 27 post-graduate volunteers to assess each video segment as truthful or deceptive. Students and judges were asked to complete a questionnaire that directed them to rate the reliability of 9 verbal and non-verbal indicators used to detect deception.

Results. The results reveal that overall judgment accuracy was above the level of chance. However, the results did not support the authors' detection ability hypothesis as it relates to discourse in Cantonese. Judges were better at identifying deception among participants speaking in English. In other words, judges were less accurate at identifying deception among native Cantonese speakers. The data from the self-rating survey indicated that when deceiving, particularly in English, speakers perceived that they had less control over smiling and laughing; leg and foot movements; head and body movements; and facial expressions. The authors claim that judges were better at detecting deception in English because the speakers had less control over these behavioral cues when speaking English. They established a cause and effect relationship between the data from the self-rating survey and the detection ability of judges. The authors observed a correlation and misinterpreted this correlation as a cause and effect relationship. The actual data from the self-rating survey did not cause the judges to better detect deception in English.

The results also revealed that participants generally displayed more verbal and non-verbal behaviors when speaking English than when speaking Cantonese, whether truth telling or deceiving. Judges identified Cantonese speaking deceivers 66.9% of the time, and English speaking deceivers 73.1% of the time. Participants reported that, when telling the truth in English, they were aware of controlling the amount of eye contact, hesitations in speech, and changes in pitch. The authors used this finding to support conflicting arguments. On the one hand, they concluded that the increased use of these behavioral cues helped judges detect deceivers in English better than truth tellers in English. On the other hand, the authors claimed that the increased use of these behavioral cues may have also confused the judges when detecting deception of participants speaking in Cantonese.

The authors observed that when participants engaged in both deceptive and truthful discourse, they displayed more non-verbal indicators in English than in Cantonese. Hand and arm movements more than doubled for participants speaking in English than for those speaking Cantonese, in both deceptive and truthful discourse. Gaze aversion more than doubled for English speaking participants engaged in deceptive discourse. These results suggest these non-verbal indicators were more characteristic of participants deceiving in English than in Cantonese. However, truth tellers of Cantonese had a higher frequency of gaze aversion than truth tellers of English, which serves as evidence partially in support of the authors' first hypothesis that non-verbal behaviors differ in first language discourses from second language discourses.

The authors were also concerned with how the reported use of emotions and reasoning correlated with detecting deception. Participants reported having expressed higher levels of 'disgust' and 'surprise' when deceiving than when truth telling. Participants also reported having more difficulty coming up with reasons to support their opinions when deceiving than when truth telling in both languages. In examining the languages separately, participants believed that deceiving in English required more cognitive ability than truth telling in English. The authors predicted that an increase in cognitive ability when deceiving in a second language would be matched by an increase in non-verbal behavioral cues, such as gaze aversion and leg and foot movements.

An analysis of instances where code-switching was used by the participants revealed a linguistic measure potentially correlated to deception. The authors observed more code-switching when participants were deceiving than when they were truth telling. Code-switching was observed in two video interviews when participants were deceiving in Cantonese and in one video interview when the participant was deceiving in English. Based on these results, code-switching was not considered a reliable cue for detecting deception. Moreover, code-switching between English and Cantonese is a common practice of Chinese people from Hong Kong (Chan, 1993). Although no conclusions could be made based on the limited instances of code-switching in the study, the authors argued that code-switching in deceptive discourse is worthy of further research.

Summary. Broadhurst and Cheng (2005) provide evidence for the claim that deception detecting correlates with increases in certain non-verbal and verbal behavioral indicators. However, the authors merely alluded to such behavioral indicators, without explicitly providing a list of behavioral indicators that actually supported their claim. As such, it was somewhat unclear whether the gaze aversion was such an example. Closer inspection of the results revealed that the high frequency reported on gaze aversion among Cantonese speakers corresponded to truth tellers, and not deceivers, as claimed by the authors. At the same time, the authors reported that non-verbal *and* verbal behaviors increase in frequency when truth tellers *and* deceivers speak in a second language (and not the predicted first language), which is evidence against their second hypothesis. In the end, this study underscored the complexities involved with human detectors of deception and suggested that studies on deception have important implications for approaches to research on second language acquisition.

Summary of Deception Detection Studies in Behaviorism. The approaches used in the summarized deception detection studies in behaviorism varied in terms of their theoretical framework and/or their experimental application. In instances where approaches converged in terms of their theoretical framework, these approaches diverged in terms of their experimental application. In instances where approaches overlapped in terms of experimental application, these approaches contrasted in terms of their theoretical framework.

The approaches used by DePaulo and her colleagues (1996), and Burgoon and Qin (2006) were influence by social theory. However, there was some variation in the approaches used for each study. DePaulo and her colleagues analyzed non-verbal patterns (e.g., *emotion*, and *stress levels*) of deception reported over the course of the experiment. That is, DePaulo and her colleagues did not factor in the time of day that participants deceived, or whether the deceptive communication was expressed during a sequence initiated by the participant or by the other person in the conversation. By contrast, Burgoon and Qin analyzed verbal patterns of deception (e.g., *complexity*, and *diversity*) across time and sequence.

Unlike the first two deception detection studies in behaviorism, Broadhurst and Cheng (2005) applied Ekman's (1999) universalist approach to emotional behavior. In terms of the experimental application, Broadhurst and Cheng focused on *observed data* (i.e., Burgoon & Qin, 2006), and not on *self-reported data* (i.e., DePaulo et al., 1996).

However, Broadhurst and Cheng analyzed both *verbal* (i.e., Burgoon & Qin, 2006) and *non-verbal patterns* (DePaulo et al., 1996) of deception.

Taken as a whole, these deception detection studies in behaviorism reveal that the deceptive act of lying is a social norm characterized by complex verbal and non-verbal behavioral patterns that persist across languages. Overall, the results suggest that some advances have been made in the identification of behavioral indicators of deceptive discourse. However, the results also suggest that human judges are unreliable for detecting deception. Combined, these studies suggest that new approaches are needed that can be applied to diverse linguistic contexts, while minimizing human error.

Deception Detection Studies in NLP

Hancock, Thom-Santelli, and Ritchie (2004). Hancock and his colleagues (2004) advance deception detection in natural language processing (NLP) with their study on lies in daily communications. Their study focuses on lies communicated across communication mediums. More specifically, the authors analyze participants' daily reports of their communications during face to face interactions, telephone conversations, email correspondences, and instant messaging. The purpose of their study is to provide evidence in support of the claim that communication technologies have an effect on lying behavior. To this end, the authors test three contrasting hypotheses: (1) Richer mediums are ranked higher than leaner mediums to communicate lies: face to face, telephone, instant messaging and email (i.e., Media Richness Theory), (2) Leaner mediums are ranked higher than richer mediums to communicate lies: email, instant messaging, telephone, and face to face (i.e., Social Distancing Theory), and (3) The features of *synchronicity, recordability,* and *distribution* of communicative acts are better predictors of the preferred medium with which to communicate lies: telephone, face to face, instant messaging, and email. (i.e., feature-based model). The results revealed that participants communicated the most lies by telephone and the least lies by email. These results provide evidence in support of the first hypothesis based on Media Richness Theory, and the third hypothesis based on the feature-based model. The results also revealed that the rate of lies in face to face interactions did not differ significantly from the rate of lies in instant messages, evidence in contrast with the second hypothesis based on Social Distancing Theory.

Approach. Hancock and his colleagues (2004) use a feature-based approach to analyze deception across communication mediums. Their approach draws from *Media Richness Theory* (Daft & Lengel, 1984), and from *Social Distancing Theory* (DePaulo et al., 1996). Media Richness Theory facilitates ranking the level at which communicative acts may be reproduced across mediums. For example, Media Richness Theory predicts that instant messaging ranks below video face to face conversations in terms of richness, because with instant messages, communicators are limited to emoticons to reproduce the facial expressions and gestures that are otherwise easily reproduced in face to face conversations. Using the same example, Social Distancing Theory predicts that instant messaging facilitates communicating lies better than face to face conversations because communications are filtered down to short segments of words and gestures cannot contribute to the distinguishing lies from the truth. Hancock and his colleagues (2004) use these theories to support their recommendations to developers of automated deception detection techniques. More specifically, the authors encourage such developers to consider the effect of various mediums on deceptive communications in the design of deception detection software.

Experiment. Following the general guidelines presented in the seven day journal experiment conducted by DePaulo and her colleagues (1996), Hancock and his colleagues (2004) conduct a similar experiment on 30 undergraduate students (17 females, 13 males) from a northeastern American university. These students are enrolled in upper-level courses in the Communications department. The average age of the participants was 21 and course credit was granted for participation. The data from two of the participants were discarded because of failure to follow the instructions. As a result, the data analyzed for this study came from 28 participants.

Participants were given a 4-page instruction booklet with information on how to record their communications over a 7-day period. All communications were categorized as either a social interaction or a lie and recorded on *Social Interaction and Deception (SID)* forms. Social interactions and lies considered in the study by Hancock and his colleagues had to last a minimum of ten minutes. Participants identified the communication medium of each communicative event (face to face, telephone, email, or instant messaging). They also recorded the number of other people in the discourse and the gender of each person.

Specific instructions were also given to participants for recording email messages and for assessing each lie. Participants entered the length of time needed to compose the text of the email messages for social interaction and lies. When recording lies, participants shared the reason for the lie, the content of the lie, and the subjects of the lie. All lies were assessed by participants using a 9-point Lickert scale (1 = not at all;

9 = completely). For each lie, participants rated how planned the lie was, the degree to which they thought the recipient believed the lie, and the importance of the lie.

Results. Over the 7-day period, participants submitted a total of 1,508 journal entries (1,198 social interactions and 310 lies). The data revealed that participants lied an average of 1.6 times per day. This data supports the average of 1.9 times per day reported by DePaulo and her colleagues (1996). A repeated measures ANOVA indicated a significant overall difference in the measure of the rate of lies across the four settings of face to face, telephone, email, and instant messaging. Planned paired-sample *t*-tests comparing each medium with face to face conversations revealed that participants were more likely to lie during telephone conversations than during face to face conversations ($t(26) = 2.18$, $p < .05$). However, significantly fewer lies were communicated in email messages than in face to face interactions ($t(21) = 2.66$, $p < .05$). The rate of instant messages that were lies were not significantly different from face to face lies.

Planned paired-sample *t*-tests were also conducted on the data from the participants' self-reported perceptions of each lie (i.e., *How planned; How believed; How important*). These *t*-tests compared the perceptions of lies in each medium to the face to face results. The degree to which lies were planned differed significantly across the mediums, $F(3,21) = 3.96$, $p < .05$. Participants planned their lies communicated by email significantly more than their lies communicated during face to face interactions, $t(8) = -2.65$, $p < .05$. There was no significant difference observed between the planning of lies communicated during face to face interactions and the planning of lies communicated by telephone, or by instant messaging. There was also no significant difference observed across mediums in the degree to which participants perceived their lies to be believed. Generally, participants rated their lies as not very important ($M = 5.24$ on a 9-point scale).

The highest proportion of lies was communicated during telephone conversations (37%) followed by lies communicated during face to face interactions (27%). These results support the predictions of the feature-based model that lies were more likely to be communicated during telephone conversations than during face to face interactions. The Social Distance Hypothesis did not predict that lies would be communicated during telephone conversations above all of the other mediums considered, but it did predict that lies would be communicated during telephone conversations above face to face interactions.

Generally, the results of this study were inconsistent with the predictions of the Social Distance Hypothesis and the Media Richness Theory. Contrary to the predictions of the Social Distance Hypothesis, significantly fewer lies were communicated using email (14%) than during face to face interactions, and lies communicated by instant messaging (21%) were not observed to be significantly different than lies communicated in face to face interactions. The Media Richness Theory predicted that the preferred mediums to communicate lies places the richer mediums first and the leaner mediums last (face to face, telephone, instant messaging and email). However, the results observed in this study indicated that the proportion of lies communicated during telephone conversations was significantly more than the proportion of lies communicated during face to face interactions.

The authors maintain that the above mentioned results may generalize to assist natural language processing approaches to the automation of deception detection.

According to their recommendation, a system derived on the statistics that lying occurs in 14% of a user's email communications and in 21% of a user's instant messages, may be able to approximate the 14% of emails, and the 21% of instant messages that include lies. However, the authors do acknowledge that their results are limited to a student population.

Summary. The results provided by Hancock and his colleagues (2004) provided evidence in support of the predictions of the feature-based model. Specifically, the results suggested that lies are more likely to be communicated by telephone above all other mediums considered, and that lies are least likely to be communicated by email. The results did not support predications generated from the Media Richness Theory and the Social Distance Hypothesis. The authors recommend that the results be considered as a statistical baseline for the automated deception detection techniques and technologies.

Zhou, Burgoon, Nunamaker, and Twitchell (2004). Zhou and her colleagues (2004) contribute to deception detection in NLP with a focus on computer-mediated communication (CMC). More specifically, the authors use 27 viable linguistic-based cues to analyze deceptive email messages. The purpose of their study is to provide a model based on the linguistic-based cues that may be used to automate deception detection in computer-mediated communication (CMC), and email messages in particular. To this end, the authors test two hypotheses on participants engaged in *dyadic*, or two-person, email communications. The first hypothesis predicts that email messages sent by deceivers (as opposed to truth tellers) will have significantly *higher* levels of (a) *quantity*, (b) *expressivity*, (c) *positive affect*, (d) *informality*, (e) *uncertainty*, and (f) *nonimmediacy*. The first hypothesis goes on to predict that email messages sent by deceivers (as opposed to truth tellers) will have significantly *lower* instances of (g) *complexity*, (h) *diversity*, and (i) *specificity*. The second hypothesis forms the same predictions, but this hypothesis is based on a comparison of email messages sent by deceivers to those sent by naïve receivers. With the exception of the prediction concerning *specificity*, the results provide evidence in support of the viability of all of the linguistic-based cues used to detect deceptive email messages.

Approach. Zhou and her colleagues (2004) use a feature-based approach to analyze deception in email messages. Their feature-based approach contrasts with the feature-based approach used by Hancock, Thom-Santelli, and Ritchi (2004) in terms of the number of features analyzed. Specifically, Hancock and his colleagues analyzed three features: (1) *synchronicity,*(2) *recordability,* and (3) *distribution* of communicative acts; whereas Zhou and her colleagues analyze nine features, or constructs: (1) *affect*, (2) *complexity,* (3) *diversity*, (4) *expressivity*, (5) *informality*, (6) *non-immediacy*, (7) *quantity*, (8) *specificity*, and (9) *uncertainty*. These nine features comprise a combination of 27 verbal indicators from five proposed coding systems: Criteria-based Content Analysis (CBCA), Interpersonal Deception Theory, Reality Monitoring, Scientific Content Analysis (SCAN), and Verbal Immediacy (see Background section of this chapter).

Experiment. Zhou and her colleagues (2004) conduct a 2 x 2 repeated measures experiment. The conditions of the experiment differ in terms of the roles (*sender* or *receiver*) assigned to participants, and the message type (*truthful* or *deceptive*) produced by the participants. The experiment takes place over three consecutive days under the assigned conditions.

Participants comprised 60 college students, forming 30 *dyads*, or pairs. These participants received extra-credit in a Management Information Systems course at a large southwestern university for participating in the experiment. In total, the participants produced 16 deceptive messages and 14 truthful messages.

The experiment involved the use of a computer (public university owned computers, or private participant owned computers). Participants were instructed to use a designated web server and were randomly assigned the role of *sender* or *receiver*. The task assigned to participants was based on the Desert Survival Problem (Lafferty & Eady, 1974). The task has two main parts: an accident, and a scavenger hunt. The accident involved a jeep that crashes in the Kuwaiti desert. Given this scenario, participants were asked prioritize scavenged items remaining from the accident that would be most useful for survival. More specifically, paired senders and receivers had to agree on a final ranking of those prioritized items. However, unbeknownst to the receivers, each *sender* was assigned to the *truthful* or *deceptive* condition.

The instructions of the experiment varied over the three day period. On the first day, senders sent their truthful or deceptive ranking of each salvaged item, along with their truthful or deceptive reasons for the ranking, to naïve receivers. The naïve receivers re-ranked the items. These naïve receivers then sent the re-ranked items along with their honest reasons for their ranking. On the second and third days, one of the salvaged items was randomly removed from the experiment. The purpose for removing one item was to create a sense of urgency and to facilitate discussions between the dyads.

Results. As predicted in both hypotheses, deceptive messages had higher frequency of words, verbs, noun phrases, and sentences than messages from truthful or naïve recipients. Deceptive messages had more typographical errors, modal verbs, and modifiers, and more instances of the first-person plural pronoun. Deceptive messages revealed fewer instances of the first-person singular pronoun, and less punctuation than truthful messages. The higher frequency of words and sentences in deceptive email messages contradicts the prediction of Interpersonal Deception Theory that characterizes deceptive communications as having less *talk time*. The absence of evidence in support of *specificity* weakens its viability as purported by Criteria-based Content Analysis, Reality Monitoring , and Verbal Immediacy.

The results on the viability of nonimmediacy cues were inconsistent. On the one hand, deceptive messages had higher instances of first person plural pronouns than truthful messages, but passive voice, objectification and generalizing terms were not viable predictors of deception. The authors suggest that the viability of nonimmediacy cues may depend on the whether the text results from asynchronous computer-mediated communication (e.g., email), or from other means of communication (verbal testimony, face to face, or instant messaging).

Summary. Zhou and her colleagues (2004) demonstrated that that the majority of the linguistic-based cues are reliable predictors of deceptions. However, the results revealed that the *specificity* indicators were not a viable predictor of deception. The results also revealed that the *nonimmediacy* indicators were a weak predictor of deception. Combined, these results serve as a guide for researchers in NLP who might be interested in automated systems of deception detection.

Carlson, George, Burgoon, Adkins, and White (2004). Carlson and his colleagues (2004) contribute to deception detection in NLP with a focus on computer-mediated communication (CMC). The purpose of their study is the same as that of the previous study by Zhou and her colleagues (2004): To present a model for automating deception detection in computer-mediated communication (CMC), and email messages in particular. The model by Carlson and his colleagues is based on the presupposition that deception in CMC differs from deception communicated in face to face interactions. Central to the model is the identification of hypotheses predicting how to detect lies communicated in CMC environments.

Approach. The qualitative approach taken by Carlson and his colleagues (2004) is based on two theories: Interpersonal Deception Theory, and Channel Expansion Theory. Interpersonal Deception Theory is concerned with the communication dynamics of face to face interactions and explores the roles of deceivers and detectors engaged in deceptive discourse (Buller & Burgoon, 1996). Channel Expansion Theory is concerned with the experiences of user when using communication mediums (Carson & Zmud, 1999). The authors identify four hypotheses based on behavioral indicators of deceivers. These behavior indicators are measures of following constructs: (1) *motivation* of the deceiver; (2) *instrinsic ability* (e.g., communication skills) of the deceiver, (3) *complexity* of the task; and (4) *capability* of the communicative medium.

Experiment. Unlike previous studies summarized in this chapter, Carlson and his colleagues (2004) conduct *qualitative* analyses of their data. More specifically, the authors did not conduct an experiment, and they do not have quantitative results to present and describe. Instead, following established qualitative practice (e.g., Denzin &

Lincoln, 2000), the authors review and analyze relevant studies, and present a series of hypotheses based on those studies.

Results. The hypotheses of the proposed model are described in terms of the construct under analysis: (1) *motivation;* (2) *intrinsic ability*; (3) *complexity*; and (4) *capability*. In terms of *motivation*, the central hypothesis predicts that the correlation between the levels of motivation expressed by the deceiver, and the deceiver's ability to successfully deceive, is represented by a curvilinear relationship. More specifically, undetected deception is more likely to occur when the deceiver's motivation level is moderate to moderately high than when their level is extremely high and extremely low.

In terms of *intrinsic ability*, the communication skills and abilities of the deceiver are the focus of a central hypothesis. The hypothesis predicts that the more matched a deceiver's communication abilities are with the communication medium, the fewer deceptive messages are likely to be detected. In other words, this hypothesis predicts that a deceiver who is a gifted orator, but a weak writer is likely to be more successful at deceiving by phone than by email.

In terms of *complexity*, the central hypothesis predicts that deceptive messages are more likely to be detected when the communication demands of the task varies greatly. That is, deceivers have to keep track of their deception. Communications that involve long term and short term goals, or contradictory and complicated goals are more likely to be detected.

In terms of the *capability*, there are two central hypotheses that correlate deception detection with the technical capability of the communication medium. The first hypothesis predicts that deceptive messages were less likely to be detected when deceivers use a communication medium with high levels of *social presence*. In other words, deceivers were more successful at deceiving using Facebook than texting because Facebook has a higher *social presence* than texting. That is, Facebook allows users to express themselves using an array of files, formats, images, audio, and video, whereas texting is restricted to a limited file size of text only messages. The second hypothesis predicts that deceptive messages were less likely to be detected using a communication medium with the greatest proportion of manageable channels. This hypothesis is closely related to the first hypothesis. More specifically, in the first hypothesis, social presence is the predictor of deception detection and is understood as the combined output of the available channels (visual, audio, textual). By contrast, in the second hypothesis, the proportion of channels offered by the communication medium is the predictor of deception detection.

The authors claim that the type of communication medium used in deceptive messages is a factor in automating systems that detect deception. The authors assert that text-based CMC, particularly email, is by and large deceivers preferred communication medium for several reasons. However, the main reason for this preference is that email allows users in general, and deceivers in particular, to revise and refine messages before sending them. This feature facilitates manipulating deceptive messages, which is helpful to deceivers as they attempt to keep track of the details in their deceptive messages.

Summary. The proposed model by Carlson and his colleagues (2004) combines Interpersonal Deception Theory and Channel Expansion Theory. The former theory suggests that deception is better understood by analyzing the interactions between deceivers and receivers. The latter theory suggests that deception is better understood by analyzing the interactions between deceivers and their communication medium of choice. By focusing on both theories, the authors propose a predictions-based model that is directed to the attention of researchers interested in the cross examination of deception detection techniques, and the influence of communication mediums on deceptive discourse.

Summary of Deception Detection Studies in NLP. The deception detection studies in NLP summarized above used contrasting approaches to achieve the shared goal of influencing the design and development of deception detection software. Their approaches also contrast at the level of the experiment. Combined, these studies demonstrate that divergent theoretical approaches and experimental applications may converge in terms of their research goals.

In terms of the theoretical approaches, Hancock and his colleagues (2004), and Zhou and her colleagues (2004) use a *quantitative* approach to analyzing the features of deceptive discourse, whereas Carlson and his colleagues (2004) use a *qualitative* approach to present a proposed model for detecting deception. Moreover, the former studies analyze *experimental data*, which is data generated from empirical experiments. By contrast, the latter study analyzes disparate theories and studies of deception, and *not experimental data*. Despite these differences, all three of the studies have the shared goal of advancing the field for the benefit of programmers who may be interested in developing deception detection software.

Taken as a whole, these deception studies in NLP suggested that deception detection software needs to be *design-specific*, and *feature-specific*. While all software is essentially *design-specific*, these studies suggested that there are two specific areas that may influence the design of deception detection software. Those two areas are the communication medium of choice, and the interaction between senders and receivers (Carlson et al., 2004). The results revealed that deceivers presumably prefer to use the telephone over face to face interactions, email correspondences, and instant messaging (Hancock et al., 2004). Therefore, programmers interested in speech recognition software for the purposes of detecting deception might also consider the interactions between deceivers and the other participant of the conversations (i.e., *senders* or *receivers*). By contrast, all software may not be considered *feature-specific*. The results of these studies suggested that linguistic features of deception might influence the design of deception detection software. More specifically, the results suggested that the features of *affect*, *complexity, diversity*, *expressivity*, *informality*, *quantity*, and *uncertainty* are more reliable predictors of deceptive email messages than *specificity*, and *nonimmediacy* (Zhou et al., 2004). Overall, these NLP studies encouraged programmers to apply a design-specific, feature-based approach to developing deception detection software.

Deception Detection Studies in ANLP

Newman, Pennebaker, Berry, and Richards (2003). Newman and his colleagues (2003) offer new directions in deception detection in applied natural language processing (ANLP). The authors apply the computational textual analysis tool of *LIWC (Linguistic Inquiry Word Count)* to deceptive discourse. LIWC was designed to facilitate the identification of psychological phenomena by analyzing word choice across groups of texts (Pennebaker, Chung, Ireland, Gonzales, & Booth, 2007). Newman and his colleagues (2003) use LIWC to assess whether linguistic styles may be used to distinguish deceptive discourse from truthful discourse. The purpose of their study is to demonstrate that LIWC is appropriate for detecting deception. To this end, the authors test three hypotheses: (1) Deceptive discourse has a lower frequency of first person pronouns than truthful discourse; (2) linguistic markers of deception generalize from one experiment to another; and (3) LIWC is more accurate than untrained human judges in detecting deception. Generally, the deceptive communications across experiments did have fewer instances of *I*, *me*, and *my*. Across the experiments, deceptive communications were marked by fewer first-person singular pronouns, fewer third person pronouns, more negative emotion words, fewer exclusive words, and more motion verbs. Overall, LIWC is better than human judges in identifying deceptive discourse from truthful discourse.

Approach. Newman and his colleagues (2003) use a computerized word count approach to derive linguistic profiles of truthful and deceptive discourse. They assess the extent to which these profiles can be generalized to other independent samples. Finally, they contrast their empirical results with the results from human judges.

The approach used by Newman and his colleagues is influenced by related studies in social psychology (e.g., Pennebaker & Francis, 1996; Pennebaker, Mayne, & Francis, 1997). Such studies examine the linguistic style used by people when communicating emotional and personal experiences. These studies suggest that the language used to communicate personal topics is characterized by identifiable linguistic patterns. As such, deceptive discourse is predicted to reflect linguistic patterns that are markedly different from truthful discourse.

Experiment. Newman and his colleagues (2003) conduct five experiments on independent samples. In all of the experiments, the participants were instructed to communicate in a convincing manner their views on abortion (Experiments 1, 2, and 3), feelings about friends (Experiment 4), and a criminal defense of a mock crime (Experiment 5). Participants from each experiment delivered one truthful communication and one deceptive communication. In Experiments 1, 2, and 3, the participants identified themselves as "pro-life" or "pro-choice" and were informed that their communications would be judged by another group for truthfulness. All 400 communications from Experiments 1, 2, and 3 were rated as truthful or deceptive by untrained judges.

The details of the experimental design vary across experiments. Participants from Experiments 1 through 4 comprised undergraduate students from Southern Methodist University. Experiment 5 comprised undergraduate students from the University of Texas at Austin. The participants from Experiment 1, 4, and 5 were videotaped. The participants from Experiment 2 entered their communications on a computer, and those from Experiment 3 wrote their communications. All of the videotaped sessions and handwritten communications were transcribed to separate text files. A total of 568 files from all of the experiments were analyzed by LIWC using 29 variables.

Experiment 1 included 101 participants (*54 men; 47 women*) consisting of 26 "pro-life" participants (*13 men; 13 women*) and 76 "pro-choice" participants (*42 men; 34 women*). A female experimenter instructed the participants to talk for at least 2 minutes while being videotaped. Experiment 2 included 44 participants (*18 men; 26 women*) and one female experimenter who instructed participants to type their true and false views on abortion for 5 minutes. Of the participants in Experiment 2, 13 (*8 men; 5 women*) identified themselves as "pro-life" and 31 (*10 men; 21 women*) identified themselves as "pro-choice." Experiment 3 included 55 participants (*15 men; 40 women*) consisting of 14 “pro-life” participants (*3 men; 11 women*), and 41 “pro-choice” participants (*12 men; 29 women*). Over a two week period, participants wrote their true and false views on abortion. The instructions to participants were enclosed in a take-home packet and the paper on which to write each essay was also included.

Experiments 4 and 5 were procedurally different from the first three experiments. In Experiment 4, 27 participants (*8 men; 19 women*) were videotaped talking about two friends that they liked as if they were disliked, and two individuals that they disliked as if they were liked. This experiment follows the paradigm developed by DePaulo and Rosenthal (1979) and also involved monitoring heart rate and skin conductance levels for each of the 3 minute segments.

Experiment 5 included 60 participants (*23 men; 37 women*) who began the experiment sitting in a crowded room. Half were instructed to look around the room and the other half were instructed to steal a dollar located inside of a book. All of the participants were informed that they would be accused of stealing a dollar by an interrogator and all were to flatly deny any such wrong doing. Participants were told that if the interrogator accepted their plea of innocence, they would be given one dollar. They were taken to a separate room that had a functioning video camera pointed at them. Each participant was connected to a fake polygraph system for an interrogation lasting just under 2 minutes. The interrogator pretended to believe all of the pleas and paid every participant one dollar.

Results. Newman and his colleagues (2003) synthesize computationally derived linguistic profiles generated from four out of five experiments. They use these profiles to predict deception in the discourse generated from the remaining experiment. For example, the results from Experiments 2 through 5 indicated that liars used more negative emotion words (i.e., *hate, worthless, enemy*), more motion verbs (i.e., *walk, move, go*), fewer exclusive words (i.e., *but, except, without*) and fewer sensation words (i.e., *see, touch, listen*) than truth tellers. When this linguistic profile was applied to Experiment 1, 60% of the liars were correctly identified, and 58% of the truth tellers were correctly identified for an overall accuracy rate of 59%. This analysis was applied to each study and Experiments 4 and 5 did not have results that were better than chance ($z = .60$, *ns*; $z = .37$, *ns*). However, the results were better than chance for Experiment 1 ($z = 2.25$, $p < .05$), Experiment 2 ($z = 1.80$, $p < .05$), and Experiment 3 ($z = 3.40$, $p < .001$).

The general equation of the combined studies correctly identified 59% of liars and 62% of truth tellers ($z = 5.50$, $p < .001$). The overall results were consistent with the first and second hypotheses that predicted deceptive discourse would have fewer first-person pronouns than truthful discourse, and that linguistic markers of deception generalize across experiments. Of Experiments 1 through 3 that addressed abortion views, LIWC was correct 62% of the time in determining whether the communications were true or not, and judges were correct 52% of the time. The authors suggest that these results demonstrate that a significantly higher performance at detecting deception by LIWC than by human judges.

Although there were no predictions made on the use of third-person pronouns, the results revealed that liars generally used fewer third-person pronouns than truth tellers. This data is inconsistent with previous research on deception including the results reported by Knapp, Hart, and Dennis (1974). The results provide evidence to suggest that truth tellers were referring to others more than liars. According to the authors' claims, the topic of abortion may have required truth tellers to refer to the experiences of others to support their views. By extension, such a claim also suggests that the topic of abortion may have required liars to also refer to the experiences of others to support their false views on abortion. However, the authors did not extend their claim to the case of liars. Nevertheless, the authors did demonstrate that pronoun usage is a measurable indicator of select psychological phenomena, including deception.

Summary. Overall, the results reported by Newman and his colleagues (2003) reveal that deceptive discourse has fewer instances of self-referencing, and more negative language than truthful discourse. The results also revealed that the linguistic markers of deception generalize across the experiments. However, the authors concluded that further research is needed to determine which linguistic markers are indicative of deceptive discourse, and which linguistic markers are indicative of discourse on abortion. Finally, Newman and his colleagues concluded that LIWC detects deception better than untrained judges.

Toma and Hancock (2010). In another deception detection study using LIWC, Toma and Hancock (2010) identify linguistic features of deception in online dating profiles. The authors offered three hypotheses on the linguistic features of deceptive profiles. First, the authors hypothesized that fewer instances of first person pronouns, but more instances of negations and negative emotion words are more characteristic of highly deceptive texts than less deceptive texts. Second, they hypothesized that fewer instances of exclusive words and more instances of motion words, but a lower overall word count are more characteristic of highly deceptive texts than less deceptive texts. Third, they hypothesized that emotional indicators (first-person pronouns, negations, and negative emotion words) account for more variance in the deception index than cognitive indicators (word count, exclusive words, and motion words). The results of their study provide evidence largely in support for the first and third hypotheses, but against the second hypothesis. The overall word count was the only significant predictor of highly deceptive texts, and exclusive words and motion words were not related to deceptive profiles.

Approach. Toma and Hancock (2010) use a linguistic approach to analyze deception in online dating profiles. This approach builds from Newman and his colleagues (2003). Both approaches use LIWC, a computerized word counting tool designed to facilitate the identification of psychological phenomena by providing word by word comparisons across groups of texts. However, Toma and Hancock (2010) use LIWC to facilitate identifying linguistic information *specific* to online dating profiles, whereas Newman and his colleagues (2003) assess the *generalizability* of linguistic information identified by LIWC.

Experiment. Toma and Hancock (2010) recruited the 80 participants (*40 male; 40 female*) of their study from four mainstream online dating services: *American Singles, Match.com, Yahoo Personals,* and *Webdate*. These online dating services were selected because of their popularity and the detailed profiles required of users. The participants were selected from 251 respondents of an advertisement seeking participants for a self-presentation study. There was no mention of the word ‘deception’ in the advertisement, which was distributed online and in print throughout the city of New York. All participants were paid $30 for their participation.

The online profile of each participant was the focus of this study. The profiles were printed and filed prior to participants' arrivals at a designated lab at New York University where the study was conducted. Participants received a print-out of their profile and were asked to rate the accuracy of each section of their profile. After the ratings were submitted, the experimenter used a standard measuring tape to measure the height of each participant, and used a standard scale to weigh each participant. The experimenter verified the age of participants using their driver's license. The mean of the participants' ages was 30.55.

A *deception index* was calculated by subtracting the observed measurements of weight, height, and age from the measurements indicated on the participants' profiles. Participants rated the accuracy of the textual information of their respective profiles. They were asked to determine "the extent to which this information [on the printed profiles] reflects the truth about you now." The assessments were made using a 5-point Lickert scale, with 1 being 'completely inaccurate,' and 5 being 'completely accurate.'

Results. The data from the self-assessments indicated that participants rated their profiles as highly representative of the truth. The profiles were rated as 4.79 out of 5 (SD = 0,41, min = 4.00, max = 5.00). All of the textual information from the profiles was analyzed using LIWC. The LIWC report indicates the word frequency for categories including first person pronouns, negations, and negative emotion words. The word frequency is a measure of the percentage of the total words per file. The average length of the profiles were 156.16 words (SD = 118.54).

Regression models were built using the emotional indicators (first-person pronouns, negations, and negative emotion words) and the cognitive indicators (word count, exclusive words, and motion words) as predictor variables with the *deception index* as the dependent variable. For the emotional indicators, the model was a good fit (F(3,74) = 5.67, p < 0.001) and explained 15% of the variance in the deception index (r = 0.43, R^2 = 0.19, R^2_{adj} = 0.15). The standardized coefficients provided evidence in support of the first hypothesis that fewer instances of first person pronouns (*Std* β = -0.254), but more instances of negations (*Std* β = 0.281) are more characteristic of highly deceptive texts than less deceptive texts. The prediction that more instances of negative emotion words are more characteristic of highly deceptive texts than less deceptive texts was not supported by the data (*Std* β = -0.296).

In terms of the cognitive indicators, the model was not a good fit (F(3, 74) = 1,37, *ns*). Exclusive words (*Std* β = 0.005, p = 0.97) and motions words (*Std* β = 0.024, p = 0.84) were not found to be significant predictors, which does not support the first part of the second hypothesis: Fewer instances of exclusive words and more instances of motion words are more characteristic of highly deceptive texts than less deceptive texts. The model was revised and the non-significant predictors of exclusive words and motions words were removed. The remaining predictor variable of word count was used with the deception index as the dependent variable. The revised model was a good fit (F(1, 76) = 4.19, p = 0.04) and explained 4% of the variance in the deception index (r = 0.52, R^2 = 0.27, R^2_{adj} = 0.23). The latter part of the second hypothesis predicted that a lower overall word count is more characteristic of highly deceptive texts than less deceptive texts. This prediction was supported by the data on word count (*Std* β = -0,228, p = 0.04).

The emotional indicators account for 15% of the variance in the deception index, which is more than the 4% of the variance accounted by the cognitive indicators. These data provide evidence in support of the third hypothesis, but excludes exclusive words and motion words from the model.

Summary. Toma and Hancock (2010) were concerned with the textual narrative of the online dating profiles, as well as the verifiable data entered in those profiles (*height, weight,* and *age*). The results provided evidence in support of the majority of the predictor variables of the first and third hypotheses. That is, fewer instances of first-person pronouns and more instances of negative words were more characteristic of highly deceptive texts than less deceptive texts. Against all predictions, the data suggested that participants used less negative emotion words when describing themselves in their online dating profiles. Exclusive words and motions words were not significant predictors of deception, which apparently weakened the first half of the second prediction. The data suggested that word count was the only cognitive indicator predictive of deception, evidence in support of the latter part of the second hypothesis.

Duran, Hall, McCarthy, and McNamara (2010). Duran and his colleagues (2010) analyze deceptive instant messages using two existing computational textual analysis tools. More specifically, the authors compare the efficacy of LIWC and Coh-Metrix in distinguishing deceptive instant messages from truthful instant messages. More specifically still, the authors test both tools on the data from Hancock, Curry, Goorha, and Woodworth (2008), and compare the results generated by both tools. The purpose of the study by Duran and his colleagues (2010) is to reveal the overlapping functionality of Coh-Metrix and LIWC, whereby the functionality of Coh-Metrix largely entails the functionality of LIWC. The purpose is also to demonstrate that entailment does not equate with convergence. Using similar indices, the two computational tools produced dissimilar results, which are arguably reflective of the different algorithms of each tool.

Approach. Duran and his colleagues (2010) use a contrastive computational approach to analyzing deceptive discourse. This contrastive approach is two-fold. On the one hand, this contrastive approach is used to demonstrate the benefits of using quantitative tools such as LIWC, and Coh-Metrix. On the other hand, this contrastive approach is used to reveal the apparent redundancy associated with using both tools.

Experiment. The experiment that generated the data analyzed by Duran and his colleagues (2010) was conducted by Hancock and his colleagues (2008). This experiment involved 66 participants (*36 females, 30 males*) from a private northeastern university in the United States. The participants were all upper-level undergraduate students who were recruited to engage in a study on how strangers converse with each other across topics. The participants were not aware that deception would be of interest to researchers.

The experiment took place in a laboratory setting with computers. All of the participants were seated in separate rooms where they were directed to use the computers to send instant messages. They were randomly assigned to be the *senders* or the *receivers* of the messages. Senders initiated the deceptive conversations, and receivers responded unaware of the senders deceptive motives. The topics included discussing a significant person, recent mistakes, unpleasant jobs, and the concept of responsibility. The interactions were tracked using the *Netmeeting* software and saved anonymously.

The instant messages comprised 264 transcripts, which were separated by topic from the sender and the receiver. Each conversation, or dyad, resulted in 4 transcripts by the sender, and 4 by the receiver. A total of four transcripts were deceptive (2 from the sender, and 2 from the receiver), and four were truthful.

Results. The results from Hancock and his colleagues (2008) are generated using eight of the 72 LIWC indices: *word quantity, question frequency, pronouns, negative affect terms, distinction markers, causal terms, sense terms, and linguistic style matching (LSM)*. The results from Duran and his colleagues (2010) are generated using Coh-Metrix, which computes cohesion in written texts using over 700 linguistic indices. Duran and his colleagues (2010) operationalized the Coh-Metrix indices into 6 measureable constructs of *quantity, immediacy, specificity, accessibility, complexity, and redundancy*.

Two quantity analyses were conducted using Coh-Metrix and LIWC. In terms of *the first quantity analysis*, Coh-Metrix and LIWC produced highly similar results. In the Coh-Metrix analysis, a significant *F*-value was found for total word count, $F(1, 33) = 8.87$, $p = .005$. More words were used to communicate deceptive messages ($M = 159.38$, $SE = 9.97$) than truthful messages ($M = 122.76$, $SE = 9.23$). Senders increased their word use to communicate deceptive messages to 158.16 words ($SE = 12.01$) from 123.15 words ($SE = 10.21$) to communicate truthful messages. Receivers increased their word use when engaged in deceptive conversations to 160.59 words ($SE = 16.12$) from 122.37 words ($SE = 10.39$) when engaged in truthful conversations. In the LIWC analysis, more words were used to communicate deceptive messages ($M = 156.53$, $SE = 13.73$) than truthful messages ($M = 122.32$, $SE = 10.45$), $F(1,31) = 6.86$, $p < .05$. The increase in word use to communicate deceptive messages was the same for senders (liars)

($M = 138.40$, $SE = 12.13$) and receivers (partners) ($M = 140.45$, $SE = 9.67$), $F(1, 31) < 1$, *ns*, suggesting that when the deceptive message was initiated by the sender (liar), both senders (liars) and receivers (partners) used more words.

In terms of *the second quantity analysis,* Coh-Metrix and LIWC produced mixed results. In the Coh-Metrix analysis, a significant *F*-value was found for total words per conversational turn, $F(1, 33) = 3.50$, $p = .05$. For each conversational turn, fewer words were produced to communicate deceptive messages ($M = 7.73$, SE = 0.27) than to communicate truthful messages ($M = 8.37$, $SE = 0.36$). Deceptive senders used fewer words per turn ($M = 7.98$, $SE = 0.42$) than truthful senders ($M = 8.19$, $SE = 0.55$). Receivers engaged in deceptive conversations used fewer words per turn ($M = 7.48$, $SE = 0.35$), than receivers engaged in truthful conversations ($M = 8.55$, $SE = 0.55$). The combined quantity results using Coh-Metrix demonstrated that the deceptive messages of senders and receivers are characterized by more overall word use, but by fewer word use per turn. These Coh-Metrix results do not support the LIWC results reported by Hancock and his colleagues. Instead, Hancock and his colleagues' marginal results (two tailed, $p = .06$) indicate that receivers (not both senders and receivers) used fewer words per turn. The mixed results may be explained by the way each tool computes the unit of a word. That is, LIWC uses a space between any sequence of alphanumeric characters to compute a word. However, Coh-Metrix uses the Charniak syntactic parser and corresponding word tags to compute words. The difference between LIWC and Coh-Metrix can be demonstrated using the example of "*He is sooo... lucky*." LIWC interprets *sooo...* as one word. However, Coh-Metrix distinguishes *sooo* from the ellipsis.

In terms of *immediacy*, Coh-Metrix and LIWC measured pronoun use and produced comparable results, with marginal differences. Coh-Metrix and LIWC did not produce statistically significant effects for first and second pronoun use. However, using Coh-Metrix, a marginally significant *F*-value was found for third person pronouns for speaker type (sender vs. receiver) and the interaction, $F(1, 33) = 3.84$, $p = .06$, and $F(1, 33) = 3.20$, $p = .08$, respectively. The Coh-Metrix results demonstrated that senders used more third person pronouns to communicate deceptive messages ($M = 2.93$, $SE = 0.32$) than to communicate truthful messages ($M = 1.94$, $SE = 0.22$), $F(1, 33) = 5.73$, $p = .02$.

In terms of *general specificity,* the Coh-Metrix results were partially supported by the LIWC results. Coh-Metrix computed a proportion score of *wh*-words (e.g., *what, where, why*) and LIWC computed the percentage of sentences ending with question marks. Using Coh-Metrix, a significant *F*-value was found for the interaction between message type (deception vs. truthful) and speaker type (sender vs. receiver) for the number of *wh*-adverbs used, $F(1, 33) = 6.83$, $p = .01$. Senders used fewer *wh*-adverbs, an indicator that they asked fewer questions when initiating deceptive conversations, ($M = 6.53$, $SE = 0.98$) than during truthful conversations ($M = 9.04$, $SE = 1.09$), $F(1, 33) = 4.19$, $p = .05$. However, receivers used marginally more *wh*-adverbs when engaged in deceptive conversations ($M = 10.34$, $SE = 1.23$) than when engaged in truthful conversations ($M = 7.33$, $SE = 1.02$), $F(1, 33) = 3.30$, $p = .08$. Using LIWC, the results supported the Coh-Metrix results for receivers. That is, receivers (partners) asked more questions when engaged in deceptive conversation. However, the LIWC results failed to support the Coh-Metrix results for senders. That is, there was no significant difference between the number of questions posed by senders across deceptive and truthful conversations. The difference between the results generated by Coh-Metrix and by LIWC may be explained by the different computational approaches for identifying question use. Coh-Metrix computes the proportion of *wh*-adverbs, whereas LIWC computes the percentage of question marks. Although not mentioned in the study, rules of punctuation are not standard in the instant messaging environment. That is, in the instant messaging environment, it is not uncommon for users to pose questions using *wh*-adverbs while omitting the actual question mark.

In terms of *accessibility,* Coh-Metrix produced significant results for word meaningfulness (or words with more associations) in conversations and word concreteness (or words that are explicitly grounded in perceptual experiences) that were unmatched by LIWC. The LIWC results reported by Hancock and his colleagues did not have equivalent indices. Using Coh-Metrix, a significant F-value was found for word meaningfulness in conversations, $F(1, 33) = 7.88$, $p = .008$. Deceptive conversations contained more meaningful words, ($M = 418.47$, $SE = 1.23$) than truthful conversations ($M = 412.76$, $SE = 1.75$). The use of meaningful words from senders increased from a rating of 415.21 ($SE = 2.30$) in truthful conversations to a rating of 418.78 ($SE = 1.47$) in deceptive conversations. The use of meaningful words from receivers increased from a rating of 410.31 ($SE = 2.60$) when engaged in truthful conversations to a rating of 418.78 ($SE = 2.00$) when engaged in deceptive conversations. There was no interaction observed between message type and speaker type.

In terms of *word concreteness*, a significant F-value was found for word concreteness, $F(1, 33) = 5.42$, $p = .02$. Senders used more concrete words in deceptive conversations ($M = 340.63$, $SE = 3.31$) than in truthful conversations ($M = 332.99$, *SE* = 2.69). There was no difference between message types for receivers.

In terms of *complexity*, Coh-Metrix produced significant results that were unmatched by LIWC. The LIWC results reported by Hancock and his colleagues did not have an equivalent index. Complexity is measured in Coh-Metrix by first automatically parsing each sentence using the Charniak parser. From the syntactic representation produced, the main verb of the main clause is identified. The number of words before the main verb of the main clause is then counted. Using Coh-Metrix, a significant *F*-value was found for complexity, $F(1, 33) = 5.63$, $p = .02$. Deceptive conversations had more words before the main verb ($M = 7.14$, $SE = 0.46$) than truthful conversations ($M = 5.79$, $SE = 0.37$). Receivers used more words before the main verb ($M = 7.50$, $SE = 0.60$) when engaged in deceptive conversations than when engaged in truthful conversations ($M = 5.41$, $SE = 0.43$). There was no interaction observed between message type and speaker type.

In terms of *redundancy*, no statistically significant effects were observed using Coh-Metrix. The LIWC results reported by Hancock and his colleagues did not have an equivalent index. However, using the LSA given-new value, a statistically significant main effect was observed for message type, $F(1, 33) = 9.32$, $p = .004$. The LSA given-new value was higher for deceptive conversations ($M = 0.25$, $SE = 0.005$) than for truthful conversations (M = 0.23, SE = 0.007). The given-new value of senders was higher for deceptive conversations (M = 0.26, SE = 0.007) than for truthful conversations ($M = 0.24$, $SE = 0.01$). The given-new value of receivers was higher when engaged in deceptive conversations ($M = 0.25$, $SE = 0.008$) than when engaged in truthful conversations ($M = 0.22$, $SE = 0.01$).

Although considered a more subtle measure, the results of the LSA given-new value provide evidence to suggest that deceptive conversations have more given information relative to the preceding context. In other words, between-sentence conceptual redundancy is more characteristic of deceptive conversations than for truthful conversations. Redundancy may be a technique exploited by senders to limit the amount of new information that they may need to track.

Summary. Of the measures examined by Duran and his colleagues (2010), the results produced by Coh-Metrix and LIWC converge on the *quantity* and *immediacy* measures. The total word count was higher for deceptive conversations than for truthful conversations, and senders used more third person pronouns in deceptive conversations. Some of the results converged on the general specificity measure. Using Coh-Metrix, receivers asked more questions when engaged in deceptive conversation, and senders asked fewer questions during deception. The LIWC analysis only revealed that receivers asked more questions when engaged in deceptive conversation. Evidence of mixed results may be explained by the differing algorithmic approaches applied to each of the computational tools. The LIWC analysis did not have equivalent indices for accessibility, redundancy, and complexity; suggesting that Coh-Metrix may have an advantage over LIWC in terms of these latter measures.

Duran, Hall, and McCarthy (2009). Duran and his colleagues (2009) advance deception detection in ANLP by addressing the observed shortcomings characteristic of computational analyses generated by LIWC and Coh-Metrix. More specifically, researchers interested in the thematic content representing the statistical outputs are limited by LIWC and Coh-Metrix. To address this need, Duran and his colleagues (2009) introduce the computational textual analysis tool of the *Gramulator* (McCarthy, Watanabe, & Lamkin, 2012) to analyze corpora from two previously published studies (Hancock et al., 2008; Newman et al., 2003). The Gramulator combines the conventional quantitative approach of counting the frequency of features (e.g., amount of cohesion, pronoun usage, total words) with a qualitative n-gram approach of extracting collocations.

Approach. Duran and his colleagues (2009) also use a contrastive computational approach (see previous summary of this section by Duran et al., 2010). Their contrastive computational approach builds from Hancock and his colleagues (2008) and from Newman and his colleagues (2003). More specifically, the contrastive computational approach takes into account the linguistic approach of using computerized word counting tools, such as LIWC and Coh-Metrix. At the same time, their contrastive computational approach extends such a linguistic approach to include the use of the *Gramulator,* which generates both quantitative, and qualitative analyses.

The Gramulator is an assessment tool that computes *differential n-grams* (see Chapter 3: Review of the Gramulator).[1] Differential n-grams are a set number of lexical items that co-occur throughout the corpora. The Gramulator operationalizes across two *contrastive corpora*, which are two highly related text sets. Differentials make up the 50% most frequent n-grams that are common to one text set, but uncommon to the other contrasting text set. That is, the qualitative outputs of differentials are quantitatively

[1] Duran and his colleagues (2009) use the term *significantly improbable features (SIFs)* to refer to *differentials*. For consistency with the chapters of this dissertation, I use the term differentials in this summary.

derived. Duran, Hall, and McCarthy (2009) focus on bi-grams (sequence of two consecutive words) and tri-grams (sequence of three consecutive words).

Experiment. The experiment that generated the data analyzed by Duran and his colleagues (2009) was conducted by Hancock and his colleagues (2008) (see previous summary of this section by Duran et al., 2010). The experiment involved 33 participants using instant messaging software. Duran and his colleagues use the Gramulator to analyze the 130 texts (*65 true*, and *65 lies*) produced from the instant messages.

Results. The Gramulator generated 44 differential bi-grams from the truth corpus and 65 differential bi-grams from the lie corpus, and an additional 22 differential tri-grams from the truth corpus and 26 differential tri-grams from the lie corpus. Following common practice, bi-grams that co-occurred in tri-grams were reduced to just one tri-gram, thus eliminating the redundant bi-gram (e.g., the bi-gram *my senior* occurs within the tri-gram *my senior year*, so only the tri-gram is counted). After removing redundant n-grams, 64 n-grams from the truth corpus and 76 n-grams from the lie corpus remained.

The entire data set of 130 texts were then processed using the Evaluator module of the Gramulator.[42] This module counts and normalizes the frequency of occurrences of all of the n-grams. The Evaluator generates two composite variables for each corpus. That is, each corpus has a value for the n-grams from the truth corpus (henceforth T-grams), and a value for the n-grams from the lie corpus (henceforth L-grams). The normalization process was based on the total number of words in the texts and the total number of T- and L- grams. The output was raised for the ease of reading. For example,

---

[2] Duran and his colleagues (2009) used an early version of the Gramulator that featured the Counter module. The Counter module has been renamed the Evaluator module. I use the term Evaluator in this study to be consistent with the chapters of this dissertation.

if Text A features 253 words, and the T-grams = 3 (out of 64), then the value for Text A is 3/253/64 * 100000 = 18.5277.

Due to the design of the original experiment, there did not exist a 1 to 1 ratio between each participant and each text. That is, participants generated multiple texts. Consequently, the authors conducted a *Hierarchical Linear Model* (HLM) analysis to account for participant to text ratio. The result provided evidence to suggest that the indicative language used in the LIE corpus was used more broadly than the indicative language used in the TRUE corpus (L-grams: $M = 42.12$, $SE = 2.92$; T-grams: $M = 34.41$, $SE = 2.92$; $F(1, 192.51) = 1.91$, $p < .05$).

The authors conducted a discriminant analysis to predict the membership of the T- and L-grams in the TRUE and LIE text types. The discriminant analysis was conducted multiple times and a leave-one out classification was applied to the analysis. A significant model emerged ($\chi^2 = 44.52$, $p < .001$), with an accuracy of 79.2%. Of the files, 50 out of 65 for TRUE-values, and 53 out of 65 for LIE-values were correctly categorized. The results suggest that the Gramulator does what it is designed to do. That is, the analysis of the differential n-grams produced by the Gramulator generates n-grams that are common to one text, but uncommon to the other contrasting text.

Having established quantitative confidence, the authors conducted a quantitative analysis of the differential n-grams from the first study conducted by Hancock, Curry, Goorha, and Woodworth (2008) (henceforth referred to as the Hancock corpus). They discuss a few noteworthy results from the two text types. The differential bi-gram with the highest frequency in the LIE corpus was *high school*, which co-occurred in the differential tri-gram with the highest frequency, *in high school.* According to the authors, this result suggested that people are comfortable lying about high school experiences. The differential bi-grams with the fifth highest frequency in the TRUE corpus was a familial referent: *my mom, my mother,* and *my dad*. However, the differential bi-gram with the fourth highest frequency in the LIE corpus was not a familial referent: *my boyfriend.* The differential bi-gram of *my friend(s)* was twice as frequent in the LIE corpus (11 out of 20 instances) compared to the TRUE corpus (3 out of 10 instances). The results suggest that people are less likely, or less willing to lie about their parents than they are to lie about their *boyfriends* and *friends*. From a more cognitive perspective, the tri-gram *I think that* was more indicative of the TRUTH corpus, whereas *I feel that* was more indicative of the LIE corpus.

The authors then use the Gramulator to conduct comparable quantitative and qualitative analyses on the abortion pro-life and pro-choice short essay from the typed abortion attitudes study published by Newman, Pennebaker, Berry, and Richards (2003). The 44 participants of the study (henceforth referred to as the Newman corpus) were instructed to write two short essays each addressing both their truthful and deceptive views on abortion. The participants identified themselves as pro-choice (31 participants) or pro-life (13 participants) prior to preparing their short essays. Duran and his colleagues used the Gramulator to process the differential n-grams for the pro-choice TRUE and LIE texts, and for the pro-life TRUE and LIE texts.

The output of the *pro-choice* corpus was 121 true differential bigrams and 16 lie differential bi-grams, and an additional 101 true differential tri-grams and 6 lie differential tri-grams. As with the Hancock corpus, all bi-grams that co-occurred in tri-grams were removed. The final output was 190 true differential n-grams and 20 lie differential n-grams (hereafter referred to as PC_T-grams and PC_L-grams, respectively). Following the procedures outlined for the Hancock corpus, all of the texts were processed through the Evaluator module of the Gramulator for occurrences of PC_T-grams and PC_L-grams.

The Evaluator generated two sets of evaluations (TRUE and LIE), each with variables of composite normed T-grams and composite normed L-grams. The pro-choice TRUE and LIE texts were prepared by different participants. As such, a paired t-test was conducted on the data in lieu of a *Hierarchical Linear Model* (HLM) analysis. The results from a paired t-test suggested that the composite normed PC_T-grams were more broadly used in the pro-choice texts than the composite normed PC_L-grams (L-grams: M = 45.538, SE = 7.450; T-grams: M = 73.4288, SE = 3.634; $t(1, 43) = -2.883$, $p = .006$). A total of 3 out of the 44 texts had equal values for PC_T-grams and PC_L-grams; 1 LIE text produced a higher TRUTH value, and the remaining 40 texts produced values in the correct directions (90.910%). The results suggest that the indicative language used in the TRUE corpus was more broadly used than the indicative language used in the LIE corpus.

As with the pro-choice corpus, the *pro-life* corpus was processed using the Gramulator. The first analysis produced 19 true differential bi-grams and 89 lie differential bi-grams, and an additional 10 true differential tri-grams and 47 lie differential tri-grams. After removing redundant n-grams, 24 true differential n-grams and 122 lie differential n-grams remained (hereafter referred to as PL-T-grams and PL_L-grams, respectively).

The entire data set of pro-life texts were then processed through the Evaluator module of the Gramulator. A paired t-test was conducted on the two sets of evaluations (TRUE and LIE). The result was not significant. However, the PL_L-grams produced higher values than the PL_T-grams. The texts also produced values in the correct directions (86.667%), suggesting that the indicative language used in the LIE corpus was used more broadly.

The authors also conducted a discriminant analysis and applied a leave-one out classification to predict the membership of the T- and L-grams in the TRUE and LIE text types. A significant model emerged for the pro-choice text ($\chi^2 = 20.056$, $p < .001$), with an accuracy of 73.9%. Of the files, 33 out of 44 for TRUE-values, and 32 out of 44 for LIE-values were correctly categorized. The pro-life model was also significant ($\chi^2 = 5.657$, $p = .017$). However, the model had an accuracy of a mere 62.2%.

Among the most frequent differential bi-grams for the *pro-choice truth tellers* were *to choose, her body,* and *her decision*. One of the most frequent differential tri-grams was *a woman's right*. Some of the differential n-grams included the first person pronoun, a typical characteristic of persuasive essays (e.g., *I believe, I do, I would,* and *I feel that*). The use of the first person pronoun may suggest greater ownership of the content. By contrast, *pro-choice liars (or pro-life pretenders)* used more abstract language, suggesting a distal relationship with the content (e.g., *the United States, the law,* and *the freedom*). The person figure represented in the differential n-grams of pro-choice liars (pro-life pretenders) is that of a child (e.g., *child is, the child in, and for the child*).

The *pro-life truth teller*s also used differential n-grams that included first person pronouns (e.g., *I feel, I do,* and *I feel that*). In this study, the word *feel* has been observed in both the differential n-grams for both groups of truth tellers. This data contradicts the observations made by Hancock and his colleagues (2008) who found that the word *feel* was more characteristic of liars than truth tellers. The pro-life truth tellers also emphasized the notion of responsibility with their differential n-grams (e.g., *responsibility for* and to *take responsibility*). The person figure represented in the differential n-grams of pro-life truth tellers is that of a baby personified (e.g., *a baby, baby is,* and *a person*). By contrast, pro-choice liars (pro-life pretenders) used the person figure of a child and it was strictly defined as a fetus, not a person (e.g., *a child, the child is,* and *the fetus*).

Summary. Duran and his colleagues (2009) introduced the Gramulator to emphasize the importance of statistically-based qualitative analyses in addition to traditional quantitative analyses. The Gramulator extracts thematic content, allowing for a more comprehensive analysis contrastive corpora. Using the Gramulator, the authors found that L-grams were used more broadly than T-grams in the Hancock corpus. However, pro-choice truth tellers (from the Newman corpus) used T-grams more broadly than L-grams in their truth texts. The results of the Gramulator analyses also revealed that thematic content observed as characteristic of liars, may also be characteristic of truth tellers, depending on the context of the communication. That is, the context of persuasive essay writing imposes as set of typical lexical items available to both truth tellers and liars (e.g., *I feel that*). The Gramulator results demonstrated that one word quantitative analyses may be limiting. The Gramulator is designed to conduct quantitative and qualitative analyses to detect the nuances of language usage across contrasting corpora, thereby distinguishing the Gramulator from traditional calculator based quantitative tools.

Summary of Deception Detection Studies in ANLP

The deception detection studies in ANLP that are summarized in this section used analogous computational approaches to analyzing deception. More specifically, these studies tested the efficacy of one or more of three computational textual analysis tools: LIWC, Coh-Metrix, and the Gramulator. However, these studies varied in terms of their experimental data. Half of the studies analyzed data from experiments that the authors themselves conducted (Newman et al., 2003; Toma & Hancock, 2010), and the other half of the studies analyzed the Hancock (2008) data (Duran et al., 2009; Duran et al., 2010). The trend toward using existing data calls into question how such data are being validated. That is, Duran and his colleagues (2009) and Duran and his colleagues (2010) tested the validity of their approaches. However, these authors did not report if any validation assessments were conducted on the data used from Hancock and his colleagues (2008). Thus, there was no evidence to suggest that what the authors claimed to be assessing (i.e., truthful texts, and deceptive texts) was actually representative of the construct under assessment (see Chapter 5: Validation).

Taken as a whole, these deception detection studies in ANLP described the convergent, divergent, and unique capabilities of LIWC, Coh-Metrix, and the Gramulator. LIWC and Coh-Metrix performed comparable quantitative analyses. However, their differences lie in the number and types of measures and indices each tool used to analyze corpora. By contrast, the Gramulator performed the same basic quantitative analyses as LIWC and Coh-Metrix, but extended those simplified outputs to include more complex n-gram analyses. More importantly, the Gramulator generated qualitative results that were absent from LIWC and Coh-Metrix. Thus, the Gramulator offered an original contrastive computational approach to identifying linguistic features of deception that advances the field, and that emphasizes the importance of quantitative and qualitative analyses.

Chapter 3: Review of The Gramulator

The Gramulator

In this dissertation, I use the computational textual analysis tool called the *Gramulator* (McCarthy et al., 2012). This tool has contributed significantly to the advancement of *contrastive corpus analysis*, which is the corpus analysis form that I examine in this dissertation. This review introduces the Gramulator (i.e., purpose, performance, platform, and practice), and summarizes studies that use this tool to analyze contrasting data sets within a range of text types.

Purpose: What does the Gramulator do?

The Gramulator is a general purpose tool designed to automate *contrastive corpus analysis* (McCarthy et al., 2012; Min & McCarthy, 2010). That is, the Gramulator facilitates conducting textual analyses between two or more contrastive data sets. More formally, the Gramulator is a computational textual analysis tool designed to identify linguistic features of contrastive data sets.

The Gramulator functions by contrasting highly related data sets of *the same text type,* or the same discourse unit (e.g., abstracts, articles, expository, genre, narrative, register, or sections of text). For example, my dissertation is concerned with the text type of *discourse*, and *political discourse* in particular. The highly related data sets of my dissertation comprise *truthful views* and *deceptive views*. As such, I use the Gramulator to conduct a contrastive corpus analysis of truthful and deceptive views of political discourse.

Performance: How does the Gramulator work?

The Gramulator operationalizes the identification of prominent linguistic features within text types. This process can be explained from two perspectives: the perspective of the programmer, and the perspective of the user. The perspective of the programmer involves explaining *how the engine works*. The perspective of the user involves explaining *how the interface works*. In basic terms, the engine connects the programmer to the algorithms that operate the tool, and the interface connects the user to the components of the tool.

In terms of the engine, the Gramulator uses a series of *n-gram* analyses to identify prominent linguistic features within text types. An *n-gram* is a contiguous sequence of any number of units, which in this case are any number of words. A contiguous sequence of two words is a *bi-gram*, and three words is a *tri-gram*. For example, the phrase of "*The United States of America*" comprises 4 bigrams (i.e., *The United, United States, States of,* and *of America*), and 3 tri-grams (i.e., *The United States, United States of*, and *States of America*).

An n-gram analysis primarily involves an algorithm that identifies contiguous sequences of words that are potentially diagnostic of one of the two contrastive data sets being compared. The algorithm operationalizes a series of n-gram *probabilistic measures*, which specify the likelihood that a specific n-gram will occur in the corpus. The n-gram probabilistic measures identified by the Gramulator are called *typicals*, *reciprocals*, and *differentials*. *Typicals* are n-grams that occur with above 50% frequency, and are characteristic (or typical) of the each data set. *Reciprocals* are n-grams that also occur with above 50% frequency, and are shared (or reciprocal) across both data sets.

Differentials are n-grams that also occur with above 50% frequency, but are indicative of (or differentiate) one data set, relative to the other contrasting data set. Note that all reciprocals and differentials are always typicals. That is, n-grams that are highly featured in both data sets (i.e., reciprocals), are also highly featured in each data set (typicals). And by the same token, n-grams that are indicative of just one data set (i.e., differentials), are also typical of that data set. Essentially, the engine operationalizes the Gramulator n-gram analyses.

In terms of the interface, the Gramulator features ten modules that help navigate the user throughout the n-gram analyses and other quantitative and qualitative analyses. The modules have many varying functions (McCarthy, Booker, Guess, & Crabtree, 2012; McCarthy et al., 2012). While the platform of the interface varies depending on the version used (i.e., publicly available versions, or internal versions experimenting with developmental aspects), the performance of each module generally remains consistent. The modules can be divided into pre-processing (e.g., Converter, Cleanser, Sorter), processing (e.g., Main Module, Viewer, Concordancer, Evaluator, Parser, Balancer); and post-processing (e.g., Checker) components. Each module either facilitates preparing data for a Gramulator analysis (pre-processing), running the Gramulator analysis for the researcher's assessment (processing), or applying additional tests, measures, or filters to the outputs of the Gramulator analysis (post-processing).

Platform: How do the modules work?

The Gramulator modules facilitate conducting computational analyses, statistical tests, and lexical assessments. In my dissertation, I use six of the featured modules:

1) The Main Module, 2) Converter, 3) Cleanser, 4) Sorter, 5) Concordancer, and 6) Evaluator. These six modules are further discussed below.

The Main Module initiates the execution of the Gramulator analysis. As the name suggests, the Main Module is the opening screen of the Gramulator and the first point of contact for the user. The Main Module is the screen from which researchers can retrieve their files for analysis. More importantly, this is the module that initializes the operations for the n-gram analysis (i.e., typicals, reciprocals, and differentials). In general terms, the Main Module allows researchers to access all of the other modules and the n-gram probabilistic measures (e.g., typicals, reciprocal, and differentials) of the Gramulator.

The *Converter* module creates text files (.txt) from data stored in a spreadsheet (e.g., Microsoft Office Excel; Google Spreadsheets) or a word processing document (e.g., Microsoft Word, Google Docs). The Converter allows researchers to expedite the process of reformatting files; a process which otherwise involves manually cutting and pasting each file. The Converter displays three window of cut and paste instructions. Once the user has completed the three step process, the Converter automatically saves each reformatted text file in the "Converter" subfolder of the "Gramulator" folder.

The *Cleanser* module corrects, modifies, and standardizes the texts of any given corpus. More specifically, the Cleanser allows researchers to edit typographical errors (e.g., *ilegal* to *illegal*), to standardize representation of multiple form words (e.g., *didn't* vs. *did not*) and to reduce the inflected forms of a word to a single form (e.g., *walks, walked, walking* = *walk*). When texts are systematically cleaned, an improvement in the validity of the n-gram analysis is likely.

The *Sorter* module divides data into convenient sub-groups. For example, the researcher can make sub-groups of a training set data, test set data, and validation set data. When creating a training set and a test set, the percentage of files allocated to each data set needs to be defined. The Sorter has a "Percentage" window that allows researchers to define those percentages. The percentage default is set at 67% for the training set, and 33% for the test set, which is a standard ration in corpus analyses. The percent default is set at 0% for the validation set. The sorted data sets are saved in the Gramulator folder under a designated file name and sub-folder. The Sorter also allows researchers to further restrict the redistribution by utilizing key word searches. For example, researchers can opt to include only texts that contain words X and Y. Similarly, researchers may wish to exclude texts that contain X or Y.

The *Concordancer* module extracts examples (at the sentence level) from the texts of a corpus using word search parameters. The Concordancer allows researchers to capture key words in context. For example, using my LIE training set (see Chapter 4: Methods for details). I searched for "_us_; not_us_citizens" to help discriminate between the files in which "us" is used as a pronoun (e.g., the right to bear arms is provided to <<us>> in the bill of rights) and the files in which "us" is used as an abbreviated version of United States (e.g., current <<us>> citizens will not have their chances of getting a job). The Concordancer outputs the list of examples with relevant statistics. The Concordancer uses a quantitative approach to provide researchers with qualitative examples that may be analyzed or copied to support (or reject) an assumption, claim, or argument.

The *Evaluator* module runs numerous metrics on any selected data sets. The Evaluator allows researchers to conduct a series of statistical tests and apply statistical assessments from a single window of the Gramulator, thereby reducing the need for additional software (e.g., SPSS) or other suites of programs (e.g., Computerized Language Analysis). For example, in this dissertation, I use the Evaluator to conduct independent and paired t-tests (see Chapter 5). All test results can be saved in the "Evaluator Results" sub-folder of the "Stored" files folder. Taken as a whole, the Evaluator is a one-stop-shop for quantitative analysis.

Practice: What other studies use the Gramulator?

Recent studies demonstrate the efficacy of the Gramulator. That is, recent studies demonstrate that the Gramulator does what it was designed to do: identify linguistic features of contrastive data sets of the same text type. The studies use the Gramulator to identify linguistic features of contrastive data sets from a range of text types including *art critiques* (e.g., Hullender & McCarthy, 2011), *contested wills* (e.g., Briggs, 2012), *narratives* (e.g., Lamkin & McCarthy, 2012; Rufenacht, Lamkin, & McCarthy, 2012), *printed media* (e.g., Haertl & McCarthy, 2011; Terwilleger & McCarthy, 2011), *scientific journal articles* (e.g., McCarthy et al., in press; Min & McCarthy, 2010), and most importantly from the perspective of my dissertation, *political discourse* (e.g., Stanchevici, in press) and *deceptive discourse* (e.g., Duran et al., 2009; McCarthy et al., 2012). I briefly summarize one study from each of these text types to provide an overview of the analyses that the Gramulator is able to produce.

Hullender and McCarthy (2011) analyze art critiques to identify the linguistic features that distinguish photography critiques from modern art critiques. More specifically, the authors use the Gramulator to assess whether photography criticism can be considered a distinct genre from modern art criticism. The study reveals that the language of photography criticism contains more self-referential terms (e.g., *photography* and *photograph*) and more scientific-related structures (e.g., *test and retest*, and *the development of*) than the language of modern art criticism. By contrast, the language of modern art critiques contains more terms that refer to the process- or medium-related terms (e.g., *to paint*, and *the canvas*) and more hedges (e.g., *in a way, in terms of,* and *is a kind of*). The study concludes that the indicative features of photography critiques support defining photography criticism as a distinct genre from modern art criticism.

Briggs (2012) analyzes the language of contested wills. More specifically, Briggs uses the Gramulator to identify the linguistic features that distinguish contested holographic wills (i.e., handwritten wills) from contested professional wills. The study reveals that holographic wills are more narrative and lexically diverse than professional wills, and less formal and template driven than professional wills. Additionally, holographic wills contain more emotion, gender, pronominal, and relational terms than professional wills. The study concludes that the differences between the language of these types of wills may influence their interpretation, which may subsequently influence the probate court's ruling.

Rufenacht and her colleagues (2012) analyze narrative texts to assess their suitability as reading material for text-based language learning approaches. More specifically, the authors use the Gramulator to analyze the differences between fairy tales and tradition ESL texts to assess the suitability of fairy tales as reading materials for English language learners. Generally, the study reveals that fairy tales and ESL texts are highly correlated and the two text types share significant language structures. However, the *content* of fairy tales differs from the *content* of ESL texts (e.g., *the king* would never marry again vs. *the king* of rock and roll did sell a lot of records). Additionally, fairy tales have a significant amount of baseline structures (i.e., the structures common across an independent corpus of narrative, history, and science texts). The study concludes that fairy tales are similar to ESL texts and are reasonably suitable reading material for English language learners.

Haertl and McCarthy (2011) analyze printed media featuring immigration issues in the United States. More specifically, the authors use the Gramulator to conduct a contrastive corpus analysis of newspaper articles that were printed in the states that border Mexico and Canada, and that explicitly address the topic of *immigration*. The study reveals that the articles representing the southern border states are concerned with *security*, whereas the articles representing the northern border states are concerned with *citizenship* and *family*. The study concludes that newspaper articles are an important resource for capturing the language indicative of pervading regional perspectives.

Min and McCarthy (2010) analyze scientific journal articles published by American and Korean scientists. More specifically, the authors use the Gramulator to identify the systematic differences between the writing styles of American and Korean scientists. The study reveals that Korean scientists use *non-standard varietal forms* of English (i.e., grammatically correct forms that are less likely to be used by American scientists). For example Korean scientists are more likely than American scientists to use the structure of *issue + patient* when describing the subjects of their study, as in an *HIV patient*. By contrast, American scientists overwhelmingly use the standard varietal form which is patient + issue, as in a *patient with HIV*. The study concludes that the writing style of Korean scientists compared to American scientists is grammatically acceptable, but contains subtle lexical and structural differences that may signal a discourse style that is largely atypical of the field.

Stanchevici (in press) analyzes political discourse. More specifically, the author uses the Gramulator to conduct a discourse analysis of the *svodkas*, or reports from the intelligence departments of the Soviet Union. The svodkas under consideration were produced from 1918 through 1934 during Stalin's rule. The analysis of these svodkas was concerned with the degree to which the following social classifications of the peasantry were referenced: the *kulaks* (wealthier opposers of the Soviet regime), the *bednyaks* (underprivileged supporters of the Soviet regime), and the *serednyaks* (mix of middle class supporters and opposers of the regime). The study reveals that the term *kulaks* and its cognates appear in contexts marked by hostile activities against the regime. By contrast, the terms *bednyaks* and *serendnyaks* and their respective cognates appear in more neutral contexts marked by a combination of hostile and supportive actions. Taken as a whole, the term *kulaks* is prominent in the svodkas (i.e., higher frequency) and the construct of this term and its cognates involve the use of persuasive rhetoric (i.e., repetition of hostile actions). The author concludes that such features are indicative of the regime's ideological bias against the kulak peasantry.

Duran and his colleagues (2009) analyze deceptive discourse. More specifically, the authors analyze data published by Hancock and his colleagues (2008) and Newman and his colleagues (2003) to identify the indicative linguistic features of truthful and deceptive conversations and essays. Hancock and his colleagues (2008) collected truthful and deceptive conversations communicated in instant messages by naive participants. Newman and his colleagues (2003) collected truthful and deceptive short persuasive essays written on the subject of abortion by pro-life and pro-choice participants. In terms of the 2008 data, the Gramulator analysis reveals that when participants lie, they are less likely, or less willing to lie about their *parents* than they are to lie about their *boyfriends* and *friends*. In terms of the 2003 data, the Gramulator analysis reveals that *pro-choice truth tellers* were more likely to use specific terms including *her body*, and *her decision*. By contrast, *pro-choice liars* were more likely to use the general terms including *the law*, and *the freedom*. However, the word *feel* was observed in the language of both pro-life and pro-choice truth tellers. The study concludes that indicative linguistic features of truthful and deceptive texts vary across general conversations and persuasive arguments.

The studies included in the above summary use the Gramulator to identify prominent linguistic features within text types. In practice, the Gramulator appears to do what it was designed to do. In other words, studies such as these demonstrate the validity of the Gramulator (see Chapter 5 on "extrinsic validity"). Of importance to my dissertation, these studies provide evidence that the Gramulator is appropriate for use in a wide range of approaches to analyzing diversity within text types, presumably including the subject matter of my own dissertation: deception in political discourse.

Chapter 4: Methods

Mixed Methods Approach

In this dissertation, I use a *mixed methods approach* to collect and analyze qualitative and quantitative data (Croswell, 2003). More specifically, I conduct an online experiment to collect truthful and deceptive views of political discourse. These truthful and deceptive views form the corpus, on which I conduct a contrastive corpus analysis (see Chapter 3: Review of the Gramulator). In this chapter, I describe the experiment of my dissertation and its components. In addition, I introduce the tools used to administer the experiment, and to prepare and asses the data of the experiment.

The Experiment

The experiment of this dissertation builds from Newman, Pennebaker, Berry, and Richards (2003). The authors conduct three experiments on a major political issue: *abortion* (see Chapter 2: Literature Review). The general design of their experiments calls for self-identified supporters and opposers of abortion to share their truthful and deceptive views on the issue. My experiment follows this general format. However, I adjust their experimental design in several ways, particularly as it relates to the components of the experiment. This section briefly describes both experimental design of Newman and his colleagues, and my adaptation of that design.

Experimental Design. The experimental design is the plan of an information-gathering exercise that addresses the differing conditions under which participants are subjected (Montgomery, 2005). There are two basic components of the experimental design used by Newman and his colleagues (2003) that have been adjusted and applied to my experiment: *the factors*, and *the responses*. The factors are the inputs of the experiment, and in this case are the controlled independent variables of the experiment. The responses are the outputs of the experiment, and in this case are texts and supplementary information related to those texts. The responses generated from an experimental design may influence the factors of that design, and are therefore addressed within the descriptions of the factors.

The factors that differ across both experiments are the setting, the topic, number of topics, and the mode of communication. Newman and his colleagues conduct their experiments on one topic (abortion) in a laboratory setting. In general terms, Newman and his colleagues are concerned with the linguistic styles of deceptive discourse expressed in different environmental contexts. As such, they use three different laboratory settings to collect three different types of responses, or submissions: videotaped interviews, typed submissions, and handwritten submissions. With each setting, Newman and his colleagues depend on an experimenter to instruct and guide participants through the experiment.

In contrast to the experiment conducted by Newman and his colleagues, I conduct my experiment on two topics: *stricter gun control laws, and stricter immigration laws in the United States.* In general terms, my dissertation is concerned with the linguistic features of deceptive discourse expressed across political issues, with the environmental context being held constant. As such, I use an online laboratory setting to collect one type of response: computer-ready, typewritten texts. The online environment eliminates the need for an experimenter. Instead, I include a survey with clear and precise instructions for the participants.

Survey

Surveys are a basic research tool used to collect qualitative and quantitative data (Marsden & Wright, 2010). I use a survey to instruct and guide participants through an online experimental process, and to facilitate collecting text submissions and supplemental information related to those submissions. To design a survey appropriate for an online experimental process, I explore web-based survey building tools that feature a menu of question formats (i.e., text field, scale) and support downloading data into other programs and formats (i.e., Microsoft Excel, Word). After experimenting with four comparable online survey tools (docs.google.com; feedbackfarm.com; surveymonkey.com; and zoomerang.com)[1], I selected Google Docs, a web-based tool used to create original forms. While all of the online survey tools facilitate collecting quantitative and qualitative data, Google Docs provides a host of features that support new and experienced researchers.

Google Docs and its competitors offer researchers, without programming experience, the ability to create high quality surveys. They also offer researchers the ability to download data in Microsoft Excel, a software program that supports statistical analyses (i.e., t-tests; and averages) and data conversion (i.e., Microsoft Word, .txt files). The features that distinguish Google Docs from its competitors are its free and unlimited online service and its streamlining capabilities with gmail.com, my email service of choice. Google Docs has a restriction-free policy on the number of questions and the number of responses for each survey. In addition, researchers have unlimited access to

[1] Survey Monkey acquired Zoomerang on December 14, 2011 (Rusli, 2011). I conducted my assessment of four survey tools including Survey Monkey and Zoomerang in May 2010.

survey results using the Documents menu in gmail.com. Although any of the survey tools examined are likely to provide useful data, Google Docs appears to be a reasonable point of departure.

My survey entitled, *Truth and Lies in Politics*, has four parts that facilitate in the collection of qualitative and quantitative data (Duff, 2002). Contrastive corpus analyses are conducted on these data allowing for the qualitative and quantitative assessment of contrastive sets of corpora that differ minimally (McCarthy et al., 2012; Min & McCarthy, 2010). In addition to providing a copy of the survey in Appendix A, I provide a brief overview of the survey for convenience.

Following standard protocol in experimental research, prospective participants begin by submitting their consent to participate in the survey. In Part I, all of the willing participants are asked to rate their political position on the issues of *stricter gun control laws in the United States* and *stricter immigration control laws in the United States*. Ratings are submitted using a Likert scale from 1 (*For*) to 6 (*Against*) for each issue under consideration. A 6-point scale is used over a 5-point or 7-point scale because a 6-point scale requires that participants choose a side. A 5-point scale and a 7-point scale offer the option to select (3) and (4) respectively, which are perceived as neither *for* nor *against* an issue, but right in the middle. McCarthy and McNamara (2012) suggest that participants prefer to straddle the fence rather than make a clear *for* or *against* decision when the middle option is available. A 6-point scale eliminates the middle option and offers participants a range that represents being *for* an issue (1, 2, 3) or *against* an issue (4, 5, 6).

In Part II, the participants are asked to share their views concerning whether or not to have *stricter gun control laws in the United States* and whether or not to have *stricter immigration control laws in the United States*. Each participant submits their views on both issues with one submission randomly assigned to a truth condition and the other submission randomly assigned to a lie condition. For each submission, the participants provide one example to support of their views (i.e., truthful views and deceptive views). The supporting example refers to personal experiences of the participants or refers to experiences of some other person(s). For example, under a truth condition, one moderate conservative participant *against stricter gun control laws* (Likert rating of 5) refers to her own experiences in her *truthful views*. By contrast, under a lie condition, one moderate liberal participant *for stricter gun control laws* (Likert rating of 2) refers to the experiences of *a relative that the participant "did not wish to name"* in his *deceptive views*.

Parts III and IV are the final parts of the survey. Part III comprises of a series of follow-up questions to the supporting example under each condition. These follow-up questions are primarily concerned with the relational proximity (or distance) that exists between the participant and the identified *agent subjects* and the *locative subjects* of their example. These respective subjects refer to the person(s) and place(s) that figure prominently in the participant's supporting example (Berk, 1999). By contrast, Part IV comprises the demographics section of the survey for which participants identify their *age, sex, race, state of residency, political affiliation(liberal or conservative), voter registration status* and *preferred news outlet*.

Amazon Mechanical Turk (MTurk)

Surveys are an essential tool for collecting qualitative and quantitative data. However, the survey alone does not constitute an experiment. That is, a survey needs participants and a vehicle that connects participants to the survey in order to be considered an experiment, and more precisely, an online experiment. As such, I use *Amazon Mechanical Turk (MTurk)* to administer and disseminate my survey in the online environment. *Amazon Mechanical Turk (MTurk)* is a web-based vehicle that connects people willing to be experimental participants (called *Workers*) with researchers (called *Requesters*) who need them. More formally, MTurk provides a crowdsourcing platform that enables individuals to place an open call to independent workers to perform an array of tasks (e.g., data classification, editing, translating), complete surveys, and participate in online experiments. And in practical terms, MTurk is an online experimental laboratory (Strain & Booker, 2012).

I use MTurk to recruit participants, instead of using a traditional sample of college students, as in the case of the study by Newman and his colleagues (2003). Recent studies demonstrate that MTurk offers a reliable service, a speedy recruitment process, and a level of flexibility unmatched by the traditional laboratory experience (Strain & Booker, 2012). MTurk has the capacity to deliver a sample of applicants that is significantly more diverse than typical college samples (Buhrmester, Kwang, & Gosling, 2011; Ipeirotis, 2010b). Requesters can reasonably expect participants to complete posted projects in just minutes. Within 40 minutes, 30 participants responded to my *Truth and Lies in Politics* survey. Within a three-day period, I received over 350 responses to my survey. By using MTurk, researchers are no longer confined to students' schedules and laboratory hours to conduct experiments. Projects can be posted on-line at the convenience of the researcher and large batches of data can be collected within days.

MTurk offers researchers a user-friendly system that is compatible with other Amazon.com services. To become a Requester, individuals can create a new user account online at www.mturk.com. I became a Requester by using my already existing Amazon.com shopping account. The main requirement for posting a project is to create an account fund using a valid debit or credit card or by allowing access to a valid bank account with available funds to cover the cost to complete the project, including participants' earnings and Amazon's 10% service fee.

Projects can be posted on MTurk for a minimum of $.01 under any desired name (e.g., name of the requester or a company name). I posted my survey as an anonymous Amazon Requester and offered workers $1.00 for the successful completion of the entire survey. Recent studies on the compensation rate of MTurk Workers reveal that an increase in compensation highly correlates with a decrease in the completion time of experiments (Strain & Booker, 2012). For example, Buhrmester and his colleagues (2011) posted a 30-minute experiment on MTurk at a compensation rate of $.02. In the end, it took approximately 5 hours for 25 Workers to complete the experiment. Also in their study, the authors posted the same 30-minute experiment on MTurk at a compensation rate of $.50. The $.50 experiment took less than 2 hours for 25 Workers to participate. The authors demonstrated that MTurk Workers are willing to participate in experiments for minimal compensation. In addition, their study suggests that there are minimal differences in the quality of the two participant pools. However, the significant difference in the completion time across the experiments (i.e., 5 hours compared to 2 hours) is likely a deciding factor for researchers who are generally under strict time constraints in addition to strict budget constraints. In the case of my experiment, it took approximately 2 hours for 30 MTurk Workers to complete my survey at a compensation rate of $1.00. In total, it took approximately 24 hours, distributed over 4 days, for me to receive acceptable submissions from 353 participants.

MTurk reported the average hourly wage of my participants at $4.37. On the surface, it appears that my average hourly rate of $4.37 is high compared to the $1.40 average hourly rate that workers are typically willing to be paid (Horton & Chilton, 2010). However, my survey was equivalent to one project, or *HIT (Human Intelligence Task)*, that contained 17 questions. My survey is designed for one participant to complete the entire survey. It is common practice for Requesters to divide projects of comparable length (17 questions) into separate HITs. That is, each HIT is equivalent to one of my survey questions, which amounts to 17 total HITs. In this scenario, Workers would have the option to choose which HIT they wished to complete, a process that is not conducive to collecting survey data. A better interpretation of my rate of compensation is that Workers completing my survey are paid $1.00 to complete 17 HITs, which averages just under 6 cents per HIT. This rate of compensation is consistent with the typical rate of Workers' compensation. Ipeirotis (2010a) reported that 90% of the HITs examined in his study offered Workers less than 10 cents, and 70% of the HITs examined offered Workers 5 cents or less . In the end, Workers participating in my study did not have to spend unpaid time searching for several HITs that amounted to $1.00. With my study, Workers were able to earn $1.00 after successfully completing 17 questions. This rate of compensation is reasonable for such an online survey (Strain & Booker, 2012).

MTurk offers a menu of standard and customizable features to Requesters. An important standard feature is the use of a unique code to track Workers. This code is valuable to Requesters when using customizable features. For example, MTurk offers two payment options. Requesters may choose whether to schedule automatic payments or to review Workers' submissions prior to agreeing to pay. I chose to review the responses submitted to my survey prior to payment. My study is based on a one to one model such that any one code corresponds to one survey response. Using the unique code, I was able to separate quality work from substandard work and compensate participants accordingly without having to interact directly with those participants.

Another customizable feature is the ability to configure the parameters of the research pool. I restricted my pool to Workers located in the United States that have an approval rating of 50% or higher. My study is concerned with political issues particular to the United States. Therefore, it follows that my participants should be residents of the United States. MTurk records Workers' approval rates. In the Amazon Mechanical Turk Best Practices Guide (www.mturk.com), it is suggested that an approval rating of 95% or better is indicative of the best selection of Workers. I opened my approval rating to 50% or better because I was interested in a wide and diverse range of views and did not want to be limited to MTurk's approval rating. For additional measure, I selected the option to review survey submissions prior to payment and exercised my option to block any workers that submitted poor quality work (e.g., plagarism).

Another feature that I added to my survey is completion time of 40 minutes. In comparable studies conducted on the traditional college pool, an experimenter is present to explain the study and review the instructions (Hancock et al., 2008; Newman et al., 2003). In an on-line environment, participants rely exclusively on the written descriptions and instructions. That is, the experimenter is replaced by explicit instructions. Of the 40 minutes, 5 minutes are allotted in total for participants to read and follow each set of instructions. Using the time allotment proposed by Newman and his colleagues (2003), approximately ten minutes are dedicated to typing the qualitative data, 5 minutes for truthful views, and 5 minutes for deceptive views. Another five minutes are available for reviewing the text submissions, leaving a remainder of 20 minutes. Of the remaining 20 minutes, 10 minutes are available to complete the quantitative questions of the study and 10 minutes are available to return to the MTurk interface to review the required information needed for payment.

MTurk facilitates conducting online experiments including the recruitment and monitoring of participants, the dissemination of a survey, and the timely collection of survey responses. In short, MTurk is a web-based vehicle that may be used to collect original text submissions, which constitute the corpus.

The Corpus

A corpus is a sample of texts that meet a series of standard requirements (see Chapter 5: Validation). My corpus comprises truthful texts and deceptive texts of political discourse. This section addresses the corpus composition, the corpus assessment, pre-analysis preparation, and the method of analysis.

Corpus Composition. My corpus comprises 706 original typed submissions in English with an equal distribution of truthful texts (353) and deceptive texts (353). Each set of texts addresses either one of two political issues in the United States: *stricter gun* *control laws*, and *stricter immigration control laws*. The number of texts for each condition is presented in Table 3.

Table 3

The Corpora

Condition	**Truthful Texts**	**Deceptive Texts**
Stricter Gun Control Laws	179	174
Stricter Immigration Control Laws	174	179
Total Texts	353	353

Corpus Assessment. The corpus was assessed primarily on the text length of each submission. For the purpose of this experiment, I set out to analyze a corpus comprising 720 original typed submissions in English, 360 truthful texts, and 360 deceptive texts. After removing texts that did not satisfy the conditions of the corpus assessment, my actual corpus comprised 706 original submissions, 353 truthful texts, and 353 deceptive texts. This 706 total number of texts is within the range of the guidelines proposed for a deception analysis that includes rigorous validation assessments (McCarthy, Duran, & Booker, 2012).

Participants of the experiment were instructed to type two text submissions containing a minimum requirement of 8 full sentences (or 80 words) and to respond to questions related to the two text submissions (see Appendix A: Survey, Part I). Text submissions were rejected if both truthful and deceptive submissions per participant were less than the lower bound of 67 words, a calculation equal to one standard deviation (43) from the average of all of word count totals (109). There were two operating assumptions on which this lower bound parameter was based. The first assumption stipulated that acceptable participants following the instructions may not have needed 80 words under each condition to express their views. That is, submissions that are reasonably close to this number are likely accepted. The second assumption stipulated that rejected participants who did not meet the lower bound under either condition failed to follow the instructions properly. After removing the texts that did not satisfy the lower bound, 706 texts remained for analysis. In other words, less than 5% (14) of the 720 target were rejected, a commonly used rejection rate in experimental research (Brace, Kemp, & Snelgar, 2006).

A corpus size of 706 texts improves the chances that the 353 truthful texts and 353 deceptive texts are *representative* and *balanced* (see Chapter 5: Validation Results: The Corpus). As such, the corpus is more likely to contain the full range of language features characteristic of the population of text types (Biber, 1993; McNamara, Graesser, McCarthy & Cai, in press). As an additional measure, I apply the 20:1 rule, which advises researchers to have at least 20 texts per variable (McNamara et al., in press). With a minimum of 300 truthful texts and 300 deceptive texts, the 20:1 rule suggests that a researcher can comfortably conduct analyses across 15 variables. My corpus exceeds this minimum requirements of texts with 353 truthful texts and 353 deceptive texts. Moreover, I have a conservative number of 11 variables available for analyses: (1) for or against the issue; (2) relational proximity to the agent subject(s); (3) relational proximity to the locative subject; (4) temporal proximity to the event time; (5) age; (6) sex; (7) race;

(8) states of residency; (9) voter registration status; (10) political affiliation; and (11) preferred news outlet.

Pre-analysis Preparation. The process of preparing data for analysis is necessary, yet often tedious and time consuming. However, I use the Gramulator and its pre-processing modules (*Converter, Cleanser,* and *Sorter*) to facilitate the process. While an effective tool, the Gramulator is not a panacea for all difficulties associated with pre-analysis preparation. That is, even when using the Gramulator, some tasks may have to be completed manually (Briggs, 2012). In this section, I describe three phases used to prepare the corpus for analysis, and the modules of the Gramulator used to complete each phase.

The first phase of the pre-analysis preparation involves converting the format of the computer-ready text submissions into the format required for analyses produced by the Gramulator. The 706 texts submissions were automatically saved into an Excel spreadsheet, a standard feature of Google Docs. Each text submission (truthful and deceptive) was individually saved into a single cell in Excel. These texts were then converted to .txt files using the *Converter*, a module of the *Gramulator* that converts any set of strings in a single cell of an Excel spreadsheet into a single .txt file (see Chapter 3 for details). The text conversion process resulted in 706 .txt files, which is the format accepted by the Gramulator system.

The second phase of the pre-analysis preparation involves cleaning the 706 .txt files. I use a combination of automated processes and manual processes to clean the files. In term of the automated processes, I use the dictionary of the Gramulator to identify any unidentifiable strings. The dictionary identified a list of 660 unidentifiable strings that included acceptable terms (e.g., full names, common slangs, foreign words) and standard typographical errors (e.g., *ilegal* > *illegal*). A subsection of the spelling errors were changed manually because they were isolated examples that were context specific. For example, the dictionary identified two different sets of strings that represented the same entity (e.g., *al'qaida*; and *al-queda*). I manually changed all versions to one standard representation (e.g., *al'quaida* > *al-Qaeda*; *al-queda* > *al-Qaeda*). The remainder of the standard errors were changed using the *Cleanser* module of the Gramulator (see Chapter 3 for details). All cleansed files are saved in the Gramulator folder for sorting purposes.

The third phase of the pre-analysis preparation involves sorting the cleansed files. I use the *Sorter* module of the Gramulator to carry out this entire phase (see Chapter 3: for details). In the first step of this phase, I divide all of the files into their corresponding contrastive constructs. The most general contrastive construct of my data is TRUTH and LIE. That is, all truth tellers across both topics constitute the construct of TRUTH, and all liars across both topics constitute the construct of LIE. This contrastive construct is divided further into subgroups by political positioning (FOR or AGAINST an issue) and by political affiliation (LIBERAL or CONSERVATIVE). I also used the Sorter module to divide the files of each contrastive construct into training sets and test sets. I apply the default percentage of 67/33 for the training set and test set for each contrastive construct. All sorted files are saved and stored in the Gramulator folder for analysis.

The pre-analysis preparation of the corpus was completed using the Gramulator and its pre-processing modules of the Converter, Cleanser, and Sorter. These modules offer a systematic and efficient means of completing the majority of the tasks commonly associated with preparing data for analysis. In short, the Gramulator automates all phases of the pre-analysis process.

Method of Analysis

The method of analysis selected for my corpus builds from McCarthy, Duran, and Booker (2012), and from McCarthy, Watanabe, and Lamkin (2012). More specifically, my dissertation is a deception analysis that takes the form of a *contrastive corpus analysis (CCA)*, which is an analysis that involves contrasting data sets. CCA may be used to test hypotheses, whether those hypotheses are in the form of research questions or validation assessments (McCarthy, Watanabe, & Lamkin, 2012; Min & McCarthy, 2010; Rufenacht, McCarthy, & Lamkin, 2012).While I use CCA to test hypotheses in both forms (i.e., hypotheses, and research questions), I provide an example from the Internal Validation Process (IVP) introduced in Chapter 5 to describe CCA using the Gramulator. I have included this example so as to demonstrate the applicability of CCA to analyze corpora.

Contrastive Corpus Analysis. In textual analysis, including deception analysis, *contrastive corpus analysis (CCA)* is the method through which linguistic features are assessed using a process of relativity (McCarthy, Watanabe, & Lamkin, 2012). The primary component of CCA is the discourse unit (e.g., political discourse). According to the principle of CCA, to understand any discourse unit, a sample of the discourse unit is to be analyzed *relative* to another contrasting, yet independent sample of the same discourse unit. These relative samples constitute a contrastive data set of the corpus. For example, a contrastive data set of my corpus is truthful and deceptive discourse. That is, the texts that were submitted under the truth condition across both political issues of stricter gun control laws and stricter immigration laws comprise the data set of truthful discourse. By contrast, the texts that were submitted under the lie condition across the same political issues comprise the data set of deceptive discourse (see earlier section on The Corpus).

The Gramulator. The Gramulator is a computational textual analysis tool designed to automate CCA (see Chapter 3). The method for using the Gramulator in CCA involves conducting a series of n-gram analyses to identify prominent linguistic features of contrastive data sets (see Chapter 3). Four steps are used to process the n-gram analyses. I describe each of these steps, and identify the component of the Gramulator used for each step.

The first step is to locate the two folder corresponding to each training set of the contrastive data set under consideration. For the purposes of this description, I refer to the contrastive data set of TRUTH and LIE, and their corresponding training sets and test sets. All folders appear on the left side of the screen in the Gramulator's Main Module.

The second step is to process the n-grams. This process is initiated from the Main Module. For the purposes of this description, the n-gram analysis is concerned with bi-grams, which are any sequence of two consecutive words. There are two types of n-grams that can be processed from the Main Module: *differentials* and *typicals*. As described in Chapter 3, differentials are n-grams that are indicative of (or differentiate) one data set, relative to another contrastive data set. By contrast, typicals are n-grams that are characteristic (or typical) of each data set. The *Process* feature on the Main Module allows users to select the *Differentials* option. Once selected, the *Differentials* option displays a series of windows that guide users through the n-gram process. Ultimately, the

Gramulator saves two separate files. The first file contains the list of differential bi-grams from the TRUTH data set relative to the LIE data set. The second file contains the list of differential bi-grams from the LIE data set relative to the TRUTH data set. The file names appear with a D (for differentials), followed by the date, the name of the file, and then the format of the file (e.g., D_01-01-2012_LIE_training(TRUTH_training).txt).

The third step is to use the Evaluator module of the Gramulator to conduct the n-gram statistical analyses. The aim of this step is to use the differential bi-grams generated from the training set data to predict the test set data. In other words, the training set (containing differential bi-grams) form the index by which the test set data is measured. For example, I use the Evaluator to test whether the LIE training set predicts the LIE test set better than the TRUTH training set predicts the LIE test set (see Chapter 5: Validation). The results generated by the Evaluator are automatically saved in the Evaluator Results sub-folder of the Stored folder.

The fourth step is to use the *Statistics* feature of the Evaluator module to assess whether there is a significant difference between the Evaluator results. That is, I perform an independent t-test on the Evaluator results (i.e., n-gram analyses). The t-test results provide evidence to suggest that the LIE training set predicts the LIE test set better than the TRUTH training set predicts the LIE test set (in my example, the result is statistically significant: $p = .038$). As such, the result allow me to claim that the LIE training set index is a better predictor of the LIE test set than the TRUTH training set index is a predictor of the LIE test set.

In short, the Gramulator automates CCA, and I use CCA as a framework within which to analyze my data. More specifically, I use the Gramulator to conduct a series of n-gram analyses on contrastive data sets (e.g., TRUTH and LIE). More specifically still, I use the Gramulator components and features to process the n-grams, to generate the results from the n-gram analyses, and to test the significance of those results.

Summary of Mixed Methods Approach

In this chapter, I present the mixed methods approach used in this deception analysis. More specifically, I use an experimental design that combines standard research elements (e.g., survey, and participants) with new technologies (Google Docs, and MTurk) to collect online submissions of qualitative and quantitative data. As a result of this experimental design, I collect a corpus that satisfies newly proposed standards in deception analysis (McCarthy, Duran, & Booker, 2012). In addition, I analyze the corpus within the framework of CCA. That is, I use the Gramulator, which automates CCA, to conduct computational and statistical analyses of truthful and deceptive discourse.

Chapter 5: Validation

Almost by definition, studies that are conducted with computational textual analysis tools (e.g. LIWC, Coh-Metrix, the Gramulator) are studies that interpret *corpora* by the use of *assessment* approaches.[1] As such, it is reasonable that any computational textual analysis study begins by assessing the validity of the approach (i.e., the computational measures) and the validity of the data (i.e., the corpus). That is, it needs to be shown that the assessment approaches of the tool do actually assess what they are supposed to be assessing; and by the same token, it needs to be shown that the corpus being assessed is representative of the construct under assessment. In short, the issue at hand is validity: Specifically, we can ask *how much confidence can we have that we are assessing what we think we are assessing*?

In this chapter, I address the issue of validity. To this end, I discuss validity as it relates to the assessment approach used in this dissertation, and validity as it relates to the corpus used in this dissertation. To address these issues, the chapter is divided into two major sections: The first major section concerns the theoretical components of establishing validity. Several sub-sections are included, with these sub-sections featuring discussions on such issues as the measure, metric, index distinction, the difference between intrinsic and extrinsic validation, corpus validation, and the components of the *Internal Validation Process* (IVP) of the Gramulator. The second major section concerns the application of these theoretical concerns to the computational approaches and data

[1] Other communicative forms and body movements are analyzed in deception research (e.g., facial and body movements, Broadhurst & Cheng, 2005, Ekman & Friesen, 1974; motor patterns, Duran, Dale, & McNamara, 2008; and fMRI images, Phan et al., 2005; Spence et al., 2001).

that I have described in Chapters 3 and 4. Once again, several sub-sections are included, with these sub-sections featuring the results and the conclusions of the analysis.

Theoretical Consideration

Measures, Metrics, and Indices. The primary goal of a computational textual assessment tool is to assess text. In practical terms, this *assessment* means "turning words into numbers." Turning words into numbers has value: as Lord Kelvin famously put it "… when you can measure what you are speaking about, and express it in numbers, you know something about it; but when you cannot express it in numbers, your knowledge is of a meagre and unsatisfactory kind…" (Kelvin, 1883). But of course, not everyone would so readily agree. Indeed, Cameron (1963)[2] rejoinders with his equally famous "Not everything that counts can be counted and not everything that can be counted counts." Presumably, both Kelvin and Cameron are highlighting legitimate issues, but whatever the degree of importance we attribute to the numbers that we derive from text, it is probably wise that we at least have some common understanding of the terms that we use to describe the numericalization system that informs the assessment. In this regard then, three terms are critical: *measure*, *metric*, and *index*. The purpose of this section of the chapter is to establish working definitions of these assessment terms.

Despite the prevalence of the terms *measure*, *metric*, and *index*, the literature has surprisingly few references that attempts to distinguish them. On this note, Gervasi and Ambriola (2002) is a rare example of a paper that includes each of the following three search criteria: "*a measure is*," "*a metric is*," and "*an index is*." However, even that study

[2] Often mistakingly attributed to Albert Einstein: http://quoteinvestigator.com/2010/05/26/everything-counts-einstein/ (05/01/2012).

is of limited assistance as is evidenced by the conflation of all three terms in the following extract "… by computing an **index** based on at least a **metric** from each group, and by assigning weights to the different **measures** according to the needs that such an **index** is intended to satisfy" (pp. 224-225). Turning to the anthological Applied Natural Language Processing (McCarthy & Boonthum-Denecke, 2012), which offers chapters on three of the best established textual analysis tools (LIWC, Coh-Metrix, and the Gramulator), we see that none of the chapters actually distinguish the terms. More specifically, the LIWC chapter uses the terms *measure* and *metric* (seemingly inter-changeably) although the term *index* is not used; the Coh-Metrix chapter uses all three terms (also seemingly inter-changeably); and the Gramulator chapter uses only the term *index* (and as such, it doesn't explain how *index* may be different from *measure* or *metric*). In short then, it appears that there is little agreement in the literature as to these terms. Therefore, the definitions below will be used in this dissertation to serve as a guide for the various validation procedures and evaluations used throughout the dissertation. However, it should be understood that these definitions are heuristic in nature, and not meant as a straight-jacket into which future research should be placed.

With the above provisions in mind, this dissertation will use the term *measure* to mean a method of calculation, typically resulting in a numerical output. The measure may be a relatively simple formula (as in a ratio) or it may be a sophisticated set of procedures (as in latent semantic analysis; Landauer & Dumais, 1997); however, whenever the term *measure* is used, the focus will be the assessment technique rather than the resultant output. By contrast, the term *metric* will refer specifically to the numerical output, a result that follows from the measure. More particularly, *metric* can refer to a group of highly related (and generally highly correlated) outputs. Meanwhile, an *index* (plural, *indices*) will refer to any instance of one of these outputs. Thus, if a metric only has one instantiation then the index and the metric will be synonymous.

Measurement Validation. The validity of the assessment approach (i.e., the combination of measure, metric, and indices) may take two forms, which I will refer to here as *intrinsic* and *extrinsic* (see also McNamara et al., in press). I will use intrinsic validation to refer to the provision of evidence for the approach itself: a demonstration that the approach is actually doing what it is supposed to be doing. I will use extrinsic validation to refer to the provision of evidence in terms of widespread use and/or acceptance by the discourse community. Thus, intrinsic validity establishes that X is suitably representative of Y, regardless of whether anyone treats it as such; whereas extrinsic validity demonstrates that X is treated by the discourse community as suitably representative of Y, regardless of whether it actually is. Needless to say, a combination of both intrinsic and extrinsic validity is desirable to establish confidence in an evaluation approach.

Based on the above definitions, *intrinsic validity* refers to a demonstration that the credentials of a particular empirical process are well-founded. In this sense, intrinsic validity contrasts with *extrinsic validity*, which I will use to refer merely to claiming (rather than demonstrating) that the credentials of a particular empirical process are well-founded. In other words, I will use intrinsic to describe *how-it-works* whereas I will use extrinsic to describe *that-it works* (see also McCarthy & Jarvis, in press).

There is no single approach to establishing intrinsic validity. That is, each study has multiple aspects, and establishing intrinsic validity requires providing evidence that suits each aspect and the role it plays in the study (Gulliksen, 1950). Even highly related studies may vary considerably in how they establish validity. For example, both McCarthy and Jarvis (2007) and, three years later, McCarthy and Jarvis (2010) seek to establish intrinsic validity for a set of lexical diversity indices. However, while the studies are highly related in terms of their corpus, their measures, their goals, and even the authors themselves, the intrinsic validation process is not identical. Specifically, in McCarthy and Jarvis (2007), the authors establish the intrinsic validity of a lexical diversity by conducting an *internal* validity assessment. The internal validity assessment (described in more detail later in this chapter[3]) is used to demonstrate the degree to which the output of the index is affected by text length. By contrast, McCarthy and Jarvis (2010) establish intrinsic validity using a set of four types of validity: *convergent, divergent, internal*, and *incremental* (also discussed later in this chapter). The intrinsic validation process differs across the two studies because of the degree of establishment of the lexical diversity indices in question. In the first study, a well-established lexical diversity index, *vocd*, is the focus. Because *vocd* is already the industry standard (i.e., it has extrinsic validity), a single validation approach is sufficiently compelling to be of interest to the discourse community. However, in the second study, a very new lexical diversity measure is introduced, MTLD (Measure of Textual Lexical Diversity). Because the measure is new, the authors present a wide array of validation assessments.

Perhaps the most obvious difference between intrinsic validity and extrinsic validity is simply the amount of space in an article that is dedicated to its establishment.

---

[3] Note that *internal* validity assessment is one of the most common forms of *intrinsic* validity establishment. The similarity of their names can be confusing, but it is important to understand that *internal validity* is an example of *intrinsic validity*, not a synonym for it.

Intrinsic validation examples are typically the sole subject of the study. Consequently, intrinsic validity can be an extensive affair, often involving multiple experiments (e.g., the co-reference cohesion metric: McNamara, Louwerse, McCarthy, & Graesser, 2010) and sometimes even involving multiple experiments across multiple studies (e.g., the lexical diversity metrics: McCarthy & Jarvis, 2007, 2010, in press). By contrast, extrinsic validity can often be accomplished in a short section (e.g., Crossley, Greenfield, & McNamara, 2008) or even a single paragraph (e.g., McCarthy, Myers, Briner, Graesser, & McNamara, 2009). With extrinsic validity, the point being made is that the assessment has been scrutinized elsewhere; therefore, it does not require extensive scrutiny in the study of which it is a part (McNamara et al., in press).

In the remainder of this section, I provide a summary of four studies that demonstrate some of the key components of intrinsic validation procedures. These studies range from intrinsic validation at the level of the computational tool: LIWC (Linguistic Inquiry Word Count) and Coh-Metrix (Pennebaker & Francis, 1996; McNamara et al., 2010), to the level of the measures that comprise the functional components of the tools: MTLD (Measure of Textual Lexical Diversity; McCarthy & Jarvis, 2010) and LSA (Latent Semantic Analysis; Landauer & Dumais, 1997).

LIWC. Pennebaker and Francis (1996) conduct intrinsic validity by testing the degree to which LIWC (Linguistic Inquiry Word Count) variables measure psychological constructs. Specifically, the authors were interested in measures of *emotion* (positive and negative), *cognitive processes* (causation, discrepancy, insight, and uncertainty), *thematic content* (body, family, friends, school), and *composition* (first person usage, past tense verb usage). The authors' approach involved establishing internal validity. To this end, the authors analyzed the writing samples of 72 college students. The experimental group (n = 35) submitted writing samples that addressed the emotions they experienced on their journey to college. The participants of the control group (n = 37) were required to submit writing samples about an event or object that they believed were void of any emotional context. The writing samples of both groups were assessed by LIWC and by a group of four judges on various emotional, cognitive, content, and composition dimensions. The results of the LIWC ratings and the judges' ratings were highly correlated. Thus, the authors argue that they have demonstrated that LIWC *does what it is supposed to do*, namely, measure psychological constructs.

Coh-Metrix. McNamara and her colleagues (2010) conduct intrinsic validity by testing the degree to which the Coh-Metrix computational tool measures text cohesion and text difficulty at the word, sentence, paragraph, and discourse levels. Like the LIWC study, the authors' approach involved establishing internal validity. In the Coh-Metrix study, the authors conduct a corpus analysis of texts from12 previously published studies on text cohesion (e.g., revising problems in narrative passages; inserting causal connectives; adding argument overlaps; increasing connections between subtopics and main topics). The results of the corpus analysis using Coh-Metrix provided evidence to suggest that an increase in cohesive elements is more beneficial to low-knowledge readers than to high-knowledge readers. The authors argue that they have demonstrated that Coh-Metrix *does what it was designed to do*, namely, measure text cohesion and text difficulty.

MTLD. McCarthy and Jarvis (2010) conduct intrinsic validity by testing the degree to which the MTLD (Measure of Textual Lexical Diversity) indices measure lexical diversity. The authors' approach involved establishing *convergent, divergent, internal, and incremental validity*. That is, to assess the validity of MTLD, the authors test the degree to which the MTLD indices 1) approximate the standard indices of vocd-D, HD-D, K, and Mass (convergent validity); 2) differ from the competing index of Type-Token Ratio (divergent validity); 3) correlate with text length (internal validity); and 4) have a predictive ability over the standard indices (incremental validity). The results suggest that MTLD satisfies convergent validity to at least the same degree as the widely accepted indices of vocd-D, HD-D, K, and Mass. That is, MTLD approximates the four other indices. In terms of divergent validity, the results suggest that greater differences exist between MTLD and Type-Token Ratio than similarities. That is, MTLD diverges from the Type-Token Ratio index. In terms of internal validity, the results suggest that there is no correlation between the LD index and text length. That is, the LD index is independent of text length. In terms of incremental validity, the results suggest that three of the LD variables were able to capture unique lexical diversity information. More specifically, the results suggest that incremental validity was established for MTLD, vocd-D, and Maas. Taken as a whole, the results provide evidence to suggest that MTLD *does what it was designed to do*, namely, measure lexical diversity.

LSA. Landauer and Dumais (1997) conduct intrinsic validity by testing the general mechanism of LSA (Latent Semantic Analysis) to learn word similarities from texts. This general mechanism is a computational model that learns by exposure to a given text and has no access to word similarity information based on linguistic or perceptual knowledge. As in the previous summaries, the authors' approach involved establishing internal validity. To this end, the authors include a simulation study to assess the model's performance on a synonym test taken from portions of the Test of English as a Foreign Language (TOEFL), and on a competence test that assesses the model's ability to learn vocabulary knowledge compared to school age children. The results of the synonym test suggest that the model performs as well as the average of a large sample of TOEFL test takers, all of which were foreign students. The results of the competence test suggest that the model learned vocabulary knowledge approximating the rate of school age children. The model performed better than school aged participants of studies that attempted to teach vocabulary by context. Taken as a whole, the results provide evidence to suggest that the general mechanism of LSA *does what it was supposed to do*, namely, learn word similarities from texts.

In sum. The four summaries given above demonstrate the similarities and differences that can be a part of establishing intrinsic validity. Specifically, internal validity is a major approach in all four studies; however, the studies differ considerably in how they go about establishing this validity. Thus, the LIWC study compares automated scoring to human expert judges; the Coh-Metrix study compares automated scoring to previously documented results; the MTLD study compares automated scoring to various manipulations of texts; and the LSA study compares automated scoring to non-expert test takers. Although each of the studies varies in its approach, the studies still overlap inasmuch as they are all demonstrations that establish that something works as it is supposed to work.

Corpus Validation. Just as demonstrating the validity of assessment approaches (i.e., the measures, metrics, and indices) is a fundamental aspect of a textual analysis study, so too is validating the corpus (Granger & Petch-Tyson, 2003). But because all things fundamental are not necessarily universal, the question becomes *do researchers recognize the need for validating the corpus to a similar degree as they recognize the need for validating the assessment approach*?

As a preliminary step to addressing the question of corpus validity, I conducted a rudimentary Google Scholar search using the key term of *deception* (which is the focus of my dissertation), along with the names of arguably the three most prominent computational textual analysis tools: LIWC, Coh-Metrix, and the Gramulator. This search facilitated the identification of six prominent studies (see Table 4). I then searched this sample for the various forms of the word *validity*. The results of the search indicated that instances of the forms of *validity* appear in half of the studies, specifically as *converging validity*, *conceptual validity*, *algorithmic validity*, and *Statement Validity Analysis*. However, none of these four examples of validity refer to an assessment of the corpus in the study (as opposed to the assessment approach). In other words, while researchers in deception studies may have been careful to assess the validity of assessment approaches used to measure the data, and even the techniques used to derive the reliability of the assessments of the data, there is little evidence that researchers rigorously question the data itself, beyond describing a reasonable process of its initial collection.

Table 4

Internal Validation Investigation

Study	Texts	Average Words	*Validity* Instances	ANLP Tool(s)
Duran, Crossley, Hall, McCarthy, & McNamara (2009)	264*	122.32 (T) 156.53(L)	1	Coh-Metrix LIWC
Duran, Hall, McCarthy, & McNamara (2010)	264*	122.32 (T) 156.53 (L)	5	Coh-Metrix LIWC
Duran, Hall, & McCarthy (2009)	130* 88†	140 290	0 0	Coh-Metrix LIWC Gramulator
Hancock, Curry, Goorha, & Woodworth (2008)	264	122.32 (T) 156.53 (L)	1	LIWC
Newman, Pennebaker, Berry, & Richards (2003)	88	290	0	LIWC
Toma & Hancock (2010)	80	156.16	0	LIWC

Note: * Hancock et al. (2008) data; †Newman, Pennebaker, Berry & Richards (2003) data

It is certainly the case that each of the deception studies given in Table 4 have contributed greatly to our understanding of the linguistic features of deceptive discourse. Each study has also been effective in demonstrating the capabilities of the computational tools (LIWC, Coh-Metrix, and the Gramulator) used for analyzing deceptive discourse. Furthermore, none of the studies were conducted with the sole purpose of establishing the validity of a corpus. And perhaps most importantly of all, it is not completely clear from the designs of these studies how corpus validity, even if desired, could have been established. But even allowing for each of these points, the assumption of the validity of the corpus is cause for concern because without a validity assessment, we cannot establish high confidence that the features identified as illustrative of the construct are actually drawn from the corpus under investigation.

Criteria for Establishing the Validity of a Corpus. Central to establishing a high degree of corpus validity is having a large enough corpus to allow for validity assessment to take place (McCarthy & Jarvis, 2007; McCarthy et al., 2012). But how large is large enough? As a very general starting point for the issue of *large enough*, McNamara and her colleagues (in press) suggest researchers might begin by thinking in terms of 300 texts of about 300 word each. That having been said, McNamara and her colleagues emphasize that the primary concerns for corpus size are that the collection be 1) large enough to reflect the construct, and 2) large enough that the results of the analysis can drive the research forward. More specifically, the first point is that the size of the corpus needs to reflect the scope of the construct. So, a very broad construct such as "Learner English" requires a very large corpus (because there are many kinds of learners of English). By contrast, a much more narrowly defined construct might require substantially fewer items. For example, Briggs (2012) uses a corpus of 72 examples to examine the construct of contested hand-written wills and legal wills from the state of Tennessee. Turning to the second point, the size of the corpus helps to establish the validity of any claims made from its analysis. Thus, we can presume that the larger the corpus, the closer it is to the actual construct, meaning that the less the generalization that has to be made. As such, larger corpora generate more confidence, which helps to drive forward the research in that particular area.

Internal Validity Process (IVP) of the Gramulator. The remainder of this chapter focuses on the application of what I will refer to as the *Internal Validation Process (IVP)* of the Gramulator. To address the IVP, I begin with a description of internal validity.[4] This description is followed by an explanation of the six assessments that comprise the IVP. Having explained the assessments themselves, I then offer an explanation as to the interpretation of the results of such an analysis. Next, the data used in the assessments is described, and predictions of the results are presented. Finally, the results of the analysis are given along with a summary and discussion.

Internal Validity. The attempt to establish internal validity is perhaps the most commonly conducted test in applied natural language processing. As we see from the summaries given above, while other validation types may also be tested, the internal variety is seldom absent. Perhaps the ubiquity of the validation type stems from the combination that it is a relatively straightforward procedure (at least in theory) and yet its outcome is generally quite compelling. Although numerous resources discuss and debate internal validity, in this dissertation, I use McCarthy and Jarvis (2010) as the base for my application of the construct. McCarthy and Jarvis is an appropriate point of departure as McCarthy has published extensively on applied natural language processing (Boonthum-Denecke, McCarthy, & Lamkin, 2012; McCarthy & Boonthum-Denecke, 2012) and Jarvis has contributed equally on the subject of validation (Jarvis, 2002; Jarvis & Daller, in press). From this starting point then, I will use the term internal validity to refer to an evaluation of the sensitivity of an index. More specifically, the evaluation is the result of a series of manipulations conducted on one variable (i.e., the corpus) by another variable

[4] Recall that internal validity is an example of intrinsic validity, not a synonym for it.

(i.e., the index). The manipulations generate variations in assessment, and the sensitivity of the approach can be assessed by the degree to which its outcomes meet its predictions. Thus, the higher the degree to which the predictions are met, the higher is the claim of internal validation.

The Six Assessments of the IVP. The IVP involves six assessments. These assessments address the *homogeneity* and *markedness* of the corpus data and the derived Gramulator indices (see Chapter 4: Methods for details). The representation and purpose of each assessment is shown in Table 5. It is important to note that the IVP was specifically designed with Gramulator analysis in mind. As such, the process assumes the need to assess both the data sets as well as the indices that are derived from the data sets.

Table 5

The Six Assessments of the Internal Validation Process (IVP)

Levels	Assessments	t-test	Purpose
1	Data 1 → Index 1 Data 2 → Index 1	Independent	Homogeneity of Data
2	Data 1 → Index 2 Data 2 → Index 2	Independent	Homogeneity of Data
3	Data 1 → Index 1 Data 1 → Index 2	Paired	Homogeneity of Index
4	Data 2 → Index 1 Data 2 → Index 2	Paired	Homogeneity of Index
5	Data 1 → Index 1 Data 2 → Index 2	Independent	Markedness Test of Index
6	Data 2 → Index 1 Data 1 → Index 2	Independent	Default Test of Data

Homogeneity of Data. The first two levels of the IVP assess *the homogeneity of the data.* That is, is the data across the data set consistent, or is the data set composed of pockets of varying signal, or simple noise? The assumption is that the indices are valid, and with these valid indices we are evaluating the consistency of the data sets. Specifically then, the assessments are used to evaluate whether the test set data (which is independent of the training set data, from which the training indices were derived) yield higher values for their corresponding training index. For example, the LIE test set is predicted to yield numerically higher values than the TRUTH test set for the LIE training index.

Homogeneity of Indices. The third and fourth levels assess *the homogeneity of the indices.* That is, are the differentials across the index consistent, or is the index composed of pockets of varying signal, or simple noise? The assumption is that the data set is valid, and with these valid data sets we are evaluating the consistency of the indices. Specifically then, the assessment is used to evaluate whether the training indices (created from the training set data) are more predictive of their corresponding test set data than their contrastive indices. For example, the LIE training index is hypothesized to be more predictive than the TRUTH training index of the LIE test set data.

Markedness Test of Index. The fifth level assesses *the markedness of the indices*. That is, is there more marked language in marked data than there is default language in default data? Specifically, the markedness test is used to demonstrate that the more distinctive, or *marked index* (LIE) is a better predictor of the *marked data set*, than the default, or *unmarked index* (TRUTH) is a predictor of the *unmarked default data set*.

Thus, the LIE index is predicted to measure LIE better than the TRUTH index is predicted to measure TRUTH.

Default Test of Data. The sixth level assesses the distribution of *the marked and default language across the data sets.* That is, is there more default language in marked data than there is marked language in default data? Specifically, the default test is used to demonstrate that the marked data have more of the language of the unmarked index than the unmarked data have the language of the marked index. For example, the LIE test set is predicted to have more of the language of the TRUTH index, than the TRUTH test set is predicted to have of the language of the LIE index, because there is presumably more truthful language in lying texts, than there is lying language in truthful texts. In other words, there is more 'truth' in lies than there are 'lies' in truth.

Interpretation of the Results. The IVP involves a series of six t-tests (four independent and two paired). Individual *p*-values for these tests are informative (and will be reported when significant at the traditional level of $< .05$); however, collectively, they do not offer a useful gauge of overall significance of the assessment because, obviously, a series of six tests is less likely to yield consistent predicted results than is just one. Note that the IVP is effectively the opposite of conducting six tests on one set of data, which would call for a Bonferroni adjustment (here .05/6). Instead then, we need to assess the probability of six tests yielding a given result. In this case, we can say that the mean for each group in each of the t-tests is either a) in the predicted direction or b) not in the predicted direction. As such, the probability of either result is .5, or 50/50. Using the binomial distribution function, we can say that the probability of all six tests resulting in means in the predicted direction is .016. And for five of the six, the probability is .094.

Setting alpha at .05, we can say that if all 6 tests are in the predicted direction, we will deem the result "significant." And if five of the six results are in the predicted direction, we will use common terminology and refer to the result as "approaching significance." Fewer than five results in the predicted direction can be attributed to chance, and will be deemed not significant.

The IVP involves six assessments that are applied to two different contrastive constructs: TRUTH and LIE; and FOR_IT and AGAINST_IT (with "IT" always referring to the issue under consideration; for example, against 'stricter gun control laws in the United States'). I test the validity of these two contrastive constructs because these contrastive constructs represent the two conditions of the experimental design (see Chapter 4: Methods). That is, participants are instructed to either share the *truth* or a *lie* about two political issues, and to explicitly indicate whether they are *for* or *against* the issues under consideration. As such, the IVP allows for evaluating the consistency of the results across each contrastive construct. If the results are not in the predicted direction and if they cannot be explained as a whole, then it may be appropriate to re-examine the approach to collecting the corpus. That is, if internal validity cannot be established, then the issue may not be one of *validity*, but one of experimental design.

Application

McCarthy, Duran and Booker (2012) propose using a *Devil's Advocate approach*. That is, use an experimental method that facilitates collecting truthful and deceptive political views from each participant. For example, the authors recommend using a survey that on the one hand, instructs participants to share their political views about two key political issues (e.g., gun control, and immigration control), and on the other, instructs them to suppress those political views to communicate deception. McCarthy and his colleagues suggest that this suppression of behavior offers an opportunity to better understand deceptive linguistic behavior. Although better techniques may have been envisaged and are encouraged, the suppression approach (or Devil's Advocate approach) facilitates dividing the corpus into the contrastive constructs needed for the IVP. Both the approach and the IVP is a reasonable point of departure toward establishing a standard for internal validity in deception analysis.

The IVP is relevant to this study because it offers an opportunity to do what previous deception detection studies in ANLP seem to avoid or fail to do: to establish internal validity using rigorous assessments. The deception detection studies considered in this chapter have produced insightful results, but it is unclear whether those results are actually reporting on the phenomena of interest. With no standards for establishing internal validity and no standards for reporting those results in deception detection studies, we cannot be sure that what is being observed is what we want to observe. In this chapter, I attempt to remedy this issue. I include the IVP in my deception analysis and demonstrate the value of adhering to the standards for using the IVP and the methods for evaluating the results of the IVP assessments.

Data. My data is a corpus of deceptive discourse (see Chapter 4: Methods for full details). Following McCarthy, Duran and Booker (2012), I posted my survey entitled *Truth and Lies in Politics* on Amazon Mechanical Turk (MTurk), the online experimental laboratory that connects willing participants to research projects (Strain & Booker, 2012). The survey instructed participants to express their truthful and deceptive views on *stricter gun control laws in the United States*, and *stricter immigration laws in the United States*.

I collected a total corpus size of 706 texts, 353 TRUTH texts, and 353 LIE texts. Both the number and size of the corpora meet the recommended standards for conducting the IVP assessments (see McCarthy et al., 2012; McNamara et al., in press).

I divided my 706 texts into two contrastive constructs: (a) TRUTH and LIE; and (b) FOR_IT and AGAINST_IT (with "IT" always referring to the issue under consideration; for example, against 'stricter gun control laws in the United States'). These contrastive constructs have been created based on the responses provided by the survey participants (see more on Survey in Chapter 4: Methods).

In preparation for establishing internal validity, I used the Gramulator[5] to randomly divide each set of corpora per construct into a training set and a test set (e.g., TRUTH stricter gun control laws training set, and TRUTH stricter gun control laws test set). Using the pre-processing *Sorter* module of the Gramulator, 67% of the corpora for each assessment was randomly selected to create a training set, and the remaining 33% formed the test set. I then used the *Evaluator* module of the Gramulator and its t-test feature to access the mean values needed to validate the training indices and the corresponding test sets of each of the contrastive constructs (see more about the modules in Chapter 3: Review of the Gramulator).

Predictions. The IVP of the Gramulator consists of two sets of six tests each. In this sub-section, I offer predictions for each of these sets of tests. In all cases, the distinction is between the marked case and the default case. That is, the marked case is understood to be the non-prototypical form, and therefore the most observable in terms of

[5] See Chapter 3 for more details about the Gramulator.

features. In turn, this aspect allows the data greater focus and therefore, presumably, a greater degree of power.

In the first set of analyses, the contrastive construct of TRUTH_LIE is assessed. For the homogeneity assessments, means are predicted to be in the direction of the corpus related to the dominant derived index. Thus, testing data is predicted to be in the direction of TRUTH if assessed by a TRUTH index, and in the direction of LIE if assessed by a LIE index. When TRUTH and LIE are compared, means are predicted to be in the direction of LIE because LIE is presumed to be the marked index.

In the second set of analyses, the contrastive construct of FOR_IT and AGAINST_IT is assessed. Like TRUTH and LIE, the homogeneity assessments of FOR_IT and AGAINST_IT are predicted to have means in the direction of the corpus related to the dominant derived index. Thus, testing data is predicted to be in the direction of FOR_IT if assessed by a FOR_IT index, and in the direction of AGAINST_IT if assessed by an AGAINST_IT index. When FOR_IT and AGAINST_IT are compared, means are predicted to be in the direction of FOR_IT because FOR_IT is presumed to be the marked index. FOR_IT is presumed to be the marked index based on the theory of issue ownership (Petrocik, Benoit & Hansen, 2003). That is, supporters (i.e., FOR_IT) *own* their issue and will deploy the specific language of their issue. By contrast, opposers will not have an established set of language, making that deployment of lexical features less consistent, and therefore weaker.

Results

TRUTH and LIE Results. The results of the IVP assessments for TRUTH and LIE were all in the predicted direction (see Table 6). Taken as a whole, the results broadly support the hypothesis that the TRUTH and LIE corpora are homogenous, and that the LIE corpus is marked. The results provide evidence to suggest that the language of lies is significantly different from the language of the truth. Although we cannot know for sure that people were actually "lying" when they produced their material, we do know that the two groups (liars and truth-tellers) differed only in that the instruction for one group was *to lie* and for the other group was *to tell the truth*. As such, the fact that the groups' results are different, and that only one instruction differed (i.e., truth vs. lie), I am able to claim there is reasonable evidence that the difference between the groups can be attributed to the language of deception. In sum, the results suggest that internal validity has been established for the data used in this study that forms the contrastive construct of TRUTH and LIE.

Table 6

TRUTH and LIE Results

Levels	Representation	Test	Predicted
1a	**$LIE_{test} \rightarrow LIE_{train}(TRUTH_{train})$** $TRUTH_{test} \rightarrow LIE_{train}(TRUTH_{train})$	Homogeneity - Data	**Yes (p = .038)**
2a	$LIE_{test} \rightarrow TRUTH_{train}(LIE_{train})$ **$TRUTH_{test} \rightarrow TRUTH_{train}(LIE_{train})$**	Homogeneity - Data	**Yes**
3a	**$LIE_{test} \rightarrow LIE_{train}(TRUTH_{train})$** $LIE_{test} \rightarrow TRUTH_{train}(LIE_{train})$	Homogeneity - Index	**Yes (p = .063)**
4a	$TRUTH_{test} \rightarrow LIE_{train}(TRUTH_{train})$ **$TRUTH_{test} \rightarrow TRUTH_{train}(LIE_{train})$**	Homogeneity - Index	**Yes**
5a	**$LIE_{test} \rightarrow LIE_{train}(TRUTH_{train})$** $TRUTH_{test} \rightarrow TRUTH_{train}(LIE_{train})$	Markedness - Index	**Yes**
6a	$TRUTH_{test} \rightarrow LIE_{train}(TRUTH_{train})$ **$LIE_{test} \rightarrow TRUTH_{train}(LIE_{train})$**	Default - Data	**Yes**

Note: The assessments in bold indicate the predicted direction.

FOR_IT and AGAINST_IT Results. The results of the IVP assessments for FOR_IT and AGAINST_IT were all in the predicted direction except for the 4b test (see Table 7). In the 4b test, the direction was counter to the prediction with the mean values favoring AGAINST_IT$_{test}$ → FOR_IT$_{train}$(AGAINST_IT$_{train}$): over AGAINST_IT$_{test}$ → AGAINST_IT$_{train}$(FOR_IT$_{train}$): M = 0.302 and M = 0.270 respectively.

Table 7

FOR_IT and AGAINST_IT Results

Levels	Representation	Test	Predicted
1b	**FOR_IT$_{test}$ → FOR_IT$_{train}$(AGAINST_IT$_{train}$)** AGAINST_IT$_{test}$ → FOR_IT$_{train}$(AGAINST_IT$_{train}$)	Homogeneity - Data	**Yes**
2b	FOR_IT$_{test}$ → AGAINST_IT$_{train}$(FOR_IT$_{train}$) **AGAINST_IT$_{test}$ → AGAINST_IT$_{train}$(FOR_IT$_{train}$)**	Homogeneity - Data	**Yes**
3b	**FOR_IT$_{test}$ → FOR_IT$_{train}$(AGAINST_IT$_{train}$)** FOR_IT$_{test}$ → AGAINST_IT$_{train}$(FOR_IT$_{train}$)	Homogeneity - Index	**Yes** ***p* = .01**
4b	AGAINST_IT$_{test}$ → FOR_IT$_{train}$(AGAINST_IT$_{train}$) **AGAINST_IT$_{test}$ → AGAINST_IT$_{train}$(FOR_IT$_{train}$)**	Homogeneity - Index	No*
5b	**FOR_IT$_{test}$ → FOR_IT$_{train}$(AGAINST_IT$_{train}$)** AGAINST_IT$_{test}$ → AGAINST_IT$_{train}$(FOR_IT$_{train}$)	Markedness - Index	**Yes**
6b	**AGAINST_IT$_{test}$ → FOR_IT$_{train}$(AGAINST_IT$_{train}$)** FOR_IT$_{test}$ → AGAINST_IT$_{train}$(FOR_IT$_{train}$)	Default - Data	**Yes** ***p* = .049**

Note: * Error analysis for the Test 4b suggests that the predicted direction may be accurate.

Closer inspection of the 4b result demonstrated no recorded values for 33 of the 84 test items for the predicted AGAINST_IT$_{test}$ → AGAINST_IT$_{train}$(FOR_IT$_{train}$) data set. Meanwhile, the non-predicted AGAINST_IT$_{test}$ → FOR_IT$_{train}$(AGAINST_IT$_{train}$) had just 2 such examples. A Fisher's Exact Test suggested that this difference in number was significant ($p < .001$). Moreover, of these 33 examples, 24 of them were from the LIE condition, a significantly high sample ($p < .001$). Examining texts that had any values above 0 for the condition of $AGAINST\_IT_{test} \rightarrow AGAINST\_IT_{train}(FOR\_IT_{train})$, I found that 26 were TRUTH and 25 were LIE, a non significant difference. If the 0 value items are ignored in the analysis, then the result switches to the predicted direction: M: $AGAINST\_IT_{test} \rightarrow AGAINST\_IT_{train}(FOR\_IT_{train}) = 0.445$; $AGAINST\_IT_{test} \rightarrow FOR\_IT_{train}(AGAINST\_IT_{train})$: 0.309. Obviously, the conditions for this result have been considerably modified, and so any interpretation must be treated with caution. However, the widespread 0 values for the combination of LIE and AGAINST conditions suggests that people who LIE do not form a powerfully broad AGAINST_IT set of linguistic features. In yet other words, the result suggests that although the language features of AGAINST_IT are not widespread, they are highly indicative when they are present. To draw an analogy, the language features of AGAINST_IT might be compared to the issue of snake bites: That is, it is not ubiquity that is the issue, it is the immediacy.

With the above error analysis is taken into consideration, the results can be said to broadly support the hypothesis that the FOR_IT and AGAINST_IT corpora are homogenous, and that the FOR_IT corpus is marked. The results provide evidence to suggest that the language of FOR_IT is significantly different from the language of AGAINST_IT. Although we cannot know for sure that people were actually "FOR_IT" when they produced their material, we do know that the two groups (FOR_IT and AGAINST_IT) differed only in as much as the participants themselves indicated whether they were "for it" or "against it." As such, the fact that the groups' results are different, and that only one variable differed (i.e., "for it" and "against it"), I am able to claim that there is reasonable evidence that the difference between the groups can be attributed to the language of "position" (i.e., *supporter* or *opposer*). In sum, the results suggest that internal validity has been established for the data used that form the contrastive construct of FOR_IT and AGAINST_IT.

Discussion

The Internal Validity Process (IVP) of the Gramulator was used to assess the corpus and its derived indices. The corpus comprised 706 texts (353 TRUTH texts and 353 LIE texts). The texts were all gathered from participants using Amazon Mechanical Turk (MTurk). All participants were asked to express their truthful and deceptive views on the subjects of either *stricter gun control laws in the United States* or *stricter immigration laws in the United States*. The corpus was randomly divided into *training* and *testing* data sets before being assessed using the computational textual analysis tool, the Gramulator. I used the Gramulator to generate indices of TRUTH, LIE, FOR_IT, and AGAINST_IT from the training set data. These indices were then used to evaluate the testing set data. The primary purpose of the six assessments of the IVP was to establish the validity of the data set. Taken as a whole, the results of the IVP provide compelling evidence for the validity of the corpus and the indices that are derived from the data.

At a finer grain level, the analyses also provided assessments of the range of strengths and weaknesses of the four major aspects of the data set (i.e., TRUTH, LIE, FOR_IT, and AGAINST_IT). Specifically, the results suggest that the language of liars is more distinctive than the language of truth tellers. And the language of supporters (i.e., FOR_IT) is more distinctive than the language of opposers (i.e., AGAINST_IT). As a consequence of these two results, the analysis also suggests that the most distinctive and far reaching linguistic features are generated from *liars who are supporting an issue*. The findings reported here support previous observations that liars use language in qualitatively different ways than truth tellers (Duran et al., 2010; Newman & Pennebaker, 2003). However, the results go further than previous studies inasmuch as they provide evidence to suggest that independent of whether the truth or lies are communicated, the language of supporters (i.e., FOR_IT) is more distinctive than language of issue opposers (i.e., AGAINST_IT).

In conclusion, the results suggest that TRUTH and LIE, and FOR_IT and AGAINST_IT are consistent contrastive constructs. However, within these constructs, *lying about an issue* and *supporting an issue* appears to be the most consistent aspects. As such, the analysis suggests that while liars appear to be good at lying, liars who are supporting an issue may be even better at lying. Some may argue that this is exactly what is at stake in politics: *lying consistently about your issue*.

Chapter 6: Linguistic Style Matching

Linguistic style matching refers to the degree to which discourse participants match their word use (Neiderhoffer & Pennebaker, 2002). That is, given that a coherent linguistic exchange is occurring, *linguistic style matching* predicts that the language of the discourse participants will necessarily share linguistic features. Furthermore, we can presume that there is a relationship between the degree of shared linguistic features and the degree of common ground (Clark, 1996; Clark & Brennan, 1991) that the discourse participants have established. Thus, the more that any two people share in their conversational linguistic features, the more they are likely to be alike: or, at least, the more they are likely to believe they are alike. On this note, research suggests that people may be able to adapt to the linguistic styles of others to meet specific communicative goals (Neiderhoffer & Pennebaker, 2002). Accordingly, linguistic style matching may be a deception strategy used to influence others into a false belief.

This chapter analyzes the degree of linguistic style matching between deceivers and truth tellers from divergent political groups. More specifically, I use measures of linguistic style matching to assess whether liberal liars and conservative liars are actually performing the experimental task of lying about their political views (Analysis 1), and whether conservatives are better than liberal at lying about such views. (Analysis 2). To this end, I have organized this chapter as follows. First, I summarize the data collection methodology and data analysis methodology (see also Chapters 3 and 4). Second, I conduct two linguistic style matching analyses on the issues of *stricter gun control laws in the United States*, and *stricter immigration control laws in the United States*. Analysis

1 tests four hypotheses, and Analysis 2 tests two hypotheses Third, I detail the results of my analyses on the use of linguistic style matching in deceptive discourse. And finally, I offer a discussion as to some of the interpretations and implications of the analyses.

Analysis

As with the previous chapter, all computational analyses described in this chapter are conducted using the Gramulator (see Chapter 3). The corpus data used in this chapter and the survey through which the corpus was collected are described in detail in Chapter 4. For convenience to the reader, I summarize below the critical details pertaining to the collection of the corpus using a survey method, and the method of analysis that are discussed in this chapter.

Corpus. The corpus used for the linguistic style matching analysis comprises truthful texts and deceptive texts collected from survey participants. Participants were instructed to rate whether they were supporters or opposers of two political issues using a 6-point Lickert scale (i.e., 1 = FOR to 6 = AGAINST). The political issues were *stricter gun control laws in the United States,* and *stricter immigration control laws in the United States.* Participants were randomly selected to share their views on each issue. Specifically, participants made one truthful contribution and one deceptive contribution. After submitting their contributions, participants were asked a series of demographic questions, one of which asks them to rate their political views using a 6-point Lickert scale (i.e., 1 = Extreme Liberal, 2 = Moderate Liberal, 3 = Liberal Leaning, 4 = Conservative Leaning, 5 = Moderate Conservative, 6 = Extreme Conservative). Ratings of 1, 2, and 3 were grouped under *liberals*, and rating of 4, 5, and 6 were grouped under *conservatives*. That is, in the context of this analysis, *liberals* comprise extremists, moderates, and leaners who ascribe their views as being representative of the worldview of liberals, and *conservatives* comprise extremists, moderates, and leaners who ascribe their views as being representative of the worldview of conservatives. The contributions, in combination with the ascribed political identities, form the corpus of *political discourse* that is used in this dissertation.

Linguistic Style Matching Nomenclature. Linguistic style matching necessarily involves a plurality of agents and actions. That is, for a linguistic style matching analysis to occur there must be at least two discourse participants (e.g., sender and a receiver) engaged in an exchange of communication. In the linguistic style matching analysis of this chapter, the discourse participants comprise affiliates of divergent political groups (e.g., liberals and conservatives) who are communicating under a truth condition, and a lie condition. The first challenge in such an analysis is simply how to represent it. To address that issue, this section describes the nomenclature used throughout the chapter for all analyses of linguistic style matching.

Consider the condition represented as $(L^{-} \# C^{+}) > (L^{+} \# C^{+})$. In this condition, the upper case letters (L) and (C) represent *liberals* and *conservatives*, respectively. The negation symbol ($^{-}$) and addition symbol ($^{+}$) represent *liars* and *truth tellers*, respectively. The pound sign (#) denotes a case of *linguistic style matching*. When the pound sign is between any of the two upper case letters (e.g., L and C), the pound sign can be read as "*match the linguistic style of.*" Thus, $(L^{-} \# C^{+})$ reads *liberal liars match the linguistic style of conservative truth tellers*. Similarly, $(L^{+} \# C^{+})$ reads *liberal truth tellers match the linguistic style of conservative truth tellers.*

The *greater-than* sign (>) is used to hypothesize the greater degree of a given phenomenon under compared conditions. Thus, the condition (or parenthetical equation) to the left of the greater-than sign is hypothesized to have a greater degree of a given phenomenon than the parenthetic equation to the right of the greater-than sign. For example, the representation (L^- # C^+) > (L^+ # C^+) reads "liberal liars match the linguistic style of conservative truth tellers *to a greater extent than* liberal truth tellers match the linguistic style of conservative truth tellers."

Unless otherwise stated, the left symbol within the parenthetical equation is an index formed from the training set data (see Chapter 4: Methods for more details on training and testing data sets). The right side of that same parenthetical equation refers to the measured items, which are the texts taken from the test set data. So, within the hypothesis (L^- # C^+) > (L^+ # C^+) the left parenthetical equation of (L^- # C^+) can be understood as *the training index of liberal liars* (L^-) *is the index by which the test set of conservative truth tellers* (C^+) *is measured*.

Typicals. As the example above demonstrates, the linguistic style matching analysis of this chapter uses equations to express the hypothetical relationship between deceivers and truth tellers ascribed to the political groups of liberals and conservatives. Consider the same example of (L^- # C^+) > (L^+ # C^+). The first element of this equation is an index of deceivers (e.g., L^-) assessing the test set data of truth tellers (e.g., C^+), whereas the second element is the corresponding index of truth tellers (e.g., L^+) assessing the same test set data of truth tellers (e.g., C^+). It is important to recognize from this example that the test set data of truth tellers (C^+) is assessed twice: first by (L^-) and then by (L^+). For both assessment indices, *typical bigrams* are used. Typicals are above average occurring collocations that are *characteristic* or *typical* of any one corpus (see Chapter 3). Typicals (as opposed to *differentials*) are used because this linguistic style matching analysis is broadly concerned with the linguistic features that are *characteristic* (*or typical*) of discourse participants (and not the linguistic features that are *indicative* of, or that *differentiate* discourse participants). As such, any evidence of linguistic style matching in this chapter (as when the results comprise relatively high values for equations) implies that instances of characteristic (or typical) language co-occur in the discourse of deceivers from one political group (e.g., liberals) and in the discourse of truth tellers from another divergent political group (e.g., conservatives).

Analysis 1: Are Liars Lying?

In Analysis 1, I conduct paired t-tests using the typical form of the relevant bigrams to assess whether liars are lying. To this end, I test four hypotheses (1a, 1b, 2a, and 2b) on linguistic style matching between deceivers and truth tellers from divergent political groups. The first two hypotheses (1a, and 1b) are tested on the liberal issue of *stricter gun control laws in the United States*. The remaining two hypotheses (2a, and 2b) are tested on the conservative issue of *stricter immigration control laws in the United States*.

Hypothesis 1a: $(L^- \# C^+) > (L^+ \# C^+)$. In full, Hypothesis 1a predicts "On the liberal issue of stricter gun control laws, the index of liberal liars (formed from the training set data) will match the linguistic style of conservative truth tellers (based on the texts from the test set data) *to a greater extent than* the index of liberal truth-tellers (formed from the training set data) will match the linguistic style of conservative truth tellers (based on the texts from the test set data)."

Hypothesis 1b: (C^- # L^+) > (C^+ # L^+). In full, Hypothesis 1b predicts "On the liberal issue of stricter gun control laws, the index of conservative liars (formed from the training set data) will match the linguistic style of liberal truth tellers (based on the texts from the test set data) *to a greater extent than* the index of conservative truth tellers (formed from the training set data) will match the linguistic style of liberal truth tellers (based on the texts from the test set data)."

Hypothesis 2a: (C^- # L^+) > (C^+ # L^+). In full, Hypothesis 2a predicts "On the conservative issue of stricter immigration control laws, the index of conservative liars (formed from the training set data) will match the linguistic style of liberal truth tellers (based on the texts from the test set data) *to a greater extent than* the index of conservative truth tellers (formed from the training set data) will match the linguistic style of liberal truth tellers (based on the texts from the test set data)."

Hypothesis 2b: (L^- # C^+) > (L^+ # C^+). In full, Hypothesis 2b predicts "On the conservative issue of stricter immigration control laws, the index of liberal liars (formed from the training set data) will match the linguistic style of conservative truth tellers (based on the texts from the test set data) *to a greater extent than* the index of liberal truth tellers (formed from the training set data) will match the linguistic style of conservative truth tellers (based on the texts from the test set data)."

Predictions

The prediction of Analysis 1 is that there will be a greater degree of linguistic style matching between liars and truth tellers of divergent political groups, than between truth tellers of one political group and truth tellers from another divergent political group. More specifically, liberal liars will match the linguistic style of conservative truth tellers to a greater extent than liberal truth tellers will match the linguistic style of conservative truth tellers (1a and 2b); and conservative liars will match the linguistic style of liberal truth tellers to a greater extent than conservative truth tellers will match the linguistic style of liberal truth tellers (1b and 2a). In other words, liars (liberals and conservatives) are lying regardless of whether or not they *own* the issue under consideration (Petrocik, Benoit, & Hansen, 2003), and regardless of whether or not they tend to *support* or *oppose* such an issue (Harris Interactive, 2010; The Pew Research Center, 2012).

Such predicted outcomes stem from linguistic style matching research and deception research. More specifically, linguistic style matching research suggests that discourse participants coordinate linguistic features to meet specific communicative goals (Niederhoffer & Pennebaker, 2002). Correspondingly, deception research suggests that liars use distancing strategies to communicate deceptive discourse (Duran, Hall, McCarthy, & McNamara, 2010; Hancock, Curry, Goorha, & Woodworth, 2008; Newman, Pennebaker, Berry, & Richards, 2003). Analysis 1 assesses whether linguistic style matching may be used by deceivers to communicate their deceptive views. That is, when people lie about their political views, they presumably distance themselves from their truthful views, and perhaps appropriate the views of others that do not represent their truthful views. For example, liberal liars may distance themselves from their truthful liberal views, and appropriate the language of conservative truth tellers to communicate deceptive liberal views. Likewise, conservative liars may distance themselves from their truthful conservative views, and appropriate the language of liberal truth tellers to communicate deceptive conservative views. Thus, a greater degree of linguistic style matching between liars and truth tellers of divergent political groups than between truth tellers from one political group and truth tellers from the other divergent political group implies that liars are lying. In sum, the results for all four hypotheses of Analysis 1are predicted to be in the direction of the first parenthetical equation, which in all cases are for liars (i.e., whether liberals or conservatives).

Results

The results of the paired t-tests (see Table 8) of the typical form of bigrams were all in the predicted direction of liars (i.e., the first parenthetical equation). Moreover, the results were significant for liberal liars on the issue of *stricter gun control laws* (a liberal issue), and for conservative liars on both issues of *stricter gun control laws* (a liberal issue), and *stricter immigration control laws* (a conservative issue). However, the results were not significant for liberal liars on *stricter immigration control laws* (conservative issue).

Table 8

Analysis 1: Results from the Evaluations of Hypotheses 1a, 1b, 2a, and 2b

H	Representation	*M*	*SD*	*t*	*d*
1a	$(\mathbf{L^- \# C^+}) > (L^+ \# C^+)$	**1.488;** 1.227	0.782; 0.615	2.698*	0.779
1b	$(\mathbf{C^- \# L^+}) > (C^+ \# L^+)$	**2.089;** 1.875	0.984; 0.902	2.849*	0.520
2a	$(\mathbf{C^- \# L^+}) > (C^+ \# L^+)$	**1.700;** 1.116	0.616; 0.362	5.075*	1.269
2b	$(\mathbf{L- \# C^+}) > (L^+ \# C^+)$	**1.686;** 1.603	0.788; 0.617	0.682	0.149

*Significant result ($p < .05$)

H1a and H1b: Stricter Gun Control Laws (A Liberal Issue)

The first issue under consideration is that of *stricter gun control laws*, a liberal issue. The results of the H1a analysis, (L^- # C^+) > (L^+ # C^+) were in the predicted direction and reached a level of significance: $t(11) = 2.698$, $p = .021$, $d = 0.779$. The effect size of 0.779 can be described as large but note that the relatively small degrees of freedom means that the result should be interpreted with some caution. With this proviso in mind, the result provides evidence to suggest that on the liberal issue of *stricter gun control laws*, liberal liars (L^-) match the linguist style of conservative truth tellers (C^+) better than liberal truth tellers (L^+) match the linguistic style of conservative truth tellers (C^+). In other words, these results suggest that liberal liars and liberal truth tellers are using significantly different linguistic features. This difference suggests that liberals are performing well on the task to which they were assigned. In this case, that task is lying about their views on the issue of *stricter gun control laws*, an issue that liberals tend to support.

Similar to the results of the H1a analysis, the results of the H1b analysis, (C^- # L^+) > (C^+ # L^+), were also in the predicted direction and reached a level of significance: $t(29) = 2.849$, $p = .008$, $d = 0.520$. The effect size of 0.520 can be described as medium. The results provide evidence to suggest that on the liberal issue of *stricter gun control laws*, conservative liars (C^-) match the linguistic style of liberal truth tellers (L^+) better than conservative truth tellers (C^+) match the linguistic style of liberal truth tellers (L^+). In other words, the results suggest that conservative liars and conservative truth tellers are using significantly different linguistic features. This difference suggests that conservatives are performing well at the task to which they were assigned. In this case, that task is lying about their views on stricter gun control laws, an issue that conservatives tend to oppose.

H2a and H2b: Stricter Immigration Control Laws (A Conservative Issue)

The second issue under consideration is that of *stricter immigration control laws*, a conservative issue. The results from the H2a analysis, (C^- # L^+) > (C^+ # L^+) were in the predicted direction and reached a level of significance: t (15) = 5.075, $p < .001$, d = 1.269. The effect size of 1.269 can be described as very large although the relatively small degrees of freedom means that the result should be interpreted with some caution. With this proviso in mind, the results provide evidence to suggest that on the conservative issue of *stricter immigration control laws*, conservative liars (C^-) match the linguist style of liberal truth tellers (L^+) better than conservative truth tellers (C^+) match the linguistic style of liberal truth tellers (L^+). In other words, these results suggest that conservative liars and conservative truth tellers are using significantly different linguistic features. This difference suggests that conservatives are performing well on the task to which they were assigned. In this case, that task is lying about their views on the issue of *stricter immigration control laws*, an issue that conservatives tend to support.

The results from the H2b analysis, (L^- # C^+) > (L^+ # C^+) were in the predicted direction, but the results did not reach a level of significance: t(20) = 0.682, p = .503, d = 0.149. These results do not provide compelling evidence for the greater presence of linguistic style matching between liberal liars and conservative truth tellers. If anything, these results suggest that liberals are weak at lying about *stricter immigration control laws*, an issue that liberals tend to oppose.

Analysis 1: Discussion

Taken as a whole, the results from Analysis 1 broadly support the prediction that there will be a greater degree of linguistic style matching between liars and truth tellers of divergent political groups, than between truth tellers of one political group and truth tellers from the corresponding political group. That is, the results provide evidence to suggest that *liars are lying*.

Analysis 2: Are Conservative Liars More Convincing than Liberal Liars?

At a finer grain level, the results of Analysis 1 suggest that conservative liars may be better at matching the linguistic style of liberal truth tellers than liberal liars are at matching the linguistic style of conservative truth tellers. In other words, liars are lying, but conservative liars may be more convincing than liberal liars. This implication stems from two pieces of evidence. First, conservative liars outnumbered liberal liars (2 to 1) at matching the linguistic style of their corresponding truth tellers. Second, conservative liars matched the linguistic style of liberal truth tellers on both the liberal issue (i.e., *stricter gun control laws*) and the conservative issue (i.e., *stricter immigration control laws*). More specifically, the *p*-values were significant for conservative liars on both issues. By contrast, liberal liars matched the linguistic style of conservative truth tellers on the liberal issue only (i.e., stricter gun control laws). More specifically, the p-values were significant for liberal liars on the liberal issue. To assess whether conservative liars are more convincing than liberal liars, I conduct Analysis 2. More specifically, I test two hypotheses to assess whether conservative liars match the linguistic style of liberal truth tellers to a greater extent than liberal liars match the linguistic style of conservative truth tellers.

In Analysis 2, I conduct independent t-tests using the typical form of the relevant bigrams to assess whether conservative liars are more convincing than liberal liars. To this end, I test two hypotheses (3a, and 3b) on linguistic style matching between conservative deceivers and liberal truth tellers. Hypothesis 3a and 3b have identical representations, **$(C^- \# L^+) > (L^- \# C^+)$**. However, Hypothesis 3a is tested on the liberal issue of *stricter gun control laws* (abbreviated to *GUNS*), and Hypothesis 3b is tested on the conservative issue of *stricter immigration control laws* (abbreviated to *IMMIGRATION*).

Hypothesis 3a: $(C^- \# L^+) > (L^- \# C^+)$. In full, Hypothesis 3a predicts "On the liberal issue of stricter gun control laws, the index of conservative liars (formed from the training set data) will match the linguistic style of liberal truth tellers (based on the texts from the test set data) *to a greater extent than* the index of liberal liars (formed from the training set data) will match the linguistic style of conservative truth tellers (based on the texts from the test set data)."

Hypothesis 3b: $(C^- \# L^+) > (L^- \# C^+)$. In full, Hypothesis 3b predicts "On the conservative issue of stricter immigration control laws, the index of conservative liars (formed from the training set data) will match the linguistic style of liberal truth tellers (based on the texts from the test set data) *to a greater extent than* the index of liberal liars (formed from the training set data) will match the linguistic style of conservative truth tellers (based on the texts from the test set data)."

Prediction

The prediction of Analysis 2 is that conservative liars are more convincing than liberal liars. More specifically, there will be a greater degree of linguistic style matching between conservative liars and liberal truth tellers than between liberal liars and conservative truth tellers, regardless of the discourse issue. Such a predicted outcome stems from Analysis 1, but also from research on the discourse of liberals and conservatives, which suggests that conservatives are better *framers* than liberals (Feldman, 2007; Lakoff, 2002, 2004). That is, conservatives are better at communicating their political views so as to assert their worldview over any other worldviews (Feldman, 2007). Accordingly, if conservatives are better framers than liberals, then conservative liars are also likely to be better than liberal liars at framing deceptive political views. Thus, conservative liars are predicted to be more convincing than liberal liars, regardless of the issue under consideration. In sum, the results are predicted to be in the direction of conservative liars for the two hypotheses of Analysis 2 (i.e., the first parenthetical equation).

Results

The results of the independent t-tests (see Table 9) of the typical form of bigrams were all in the predicted direction of conservative liars (i.e., the first parenthetical equation). However, the results merely approached significance in H3a, and did not reach a level of significance in H3b.

Table 9

Analysis 2: Results from the Evaluations of Hypotheses 3a and 3b

H	**$(C^- \# L^+) > (L^- \# C^+)$**	***M***	***SD***	***t***	***d***
3a	GUNS	**2.089;** 1.488	0.984; 0.782	1.884	0.644
3b	IMMIGRATION	**1.700;** 1.686	0.616; 0.788	0.056	0.018

H3a: Stricter Gun Control Laws (A Liberal Issue)

The first issue under consideration is that of *stricter gun control laws*, a liberal issue. The results of the H3a analysis were in the predicted direction and was approaching a level of significance: $t(40) = 1.884$, $p = .067$, $d = 0.644$. The effect size of 0.644 can be described as medium. The results provide some evidence to suggest that on the liberal issue of *stricter gun control laws*, conservative liars (C^-) match the linguistic style of liberal truth tellers (L^+) to greater extent than liberal truth tellers (L^-) match the linguistic style of conservative truth tellers (C^+). This difference suggests that conservatives might be performing better than liberals at the task to which they were assigned. In this case, that task calls for conservatives and liberals to lie about their views on stricter gun control laws, which is an issue that conservatives tend to oppose.

H3b: Stricter Immigration Control Laws (A Conservative Issue)

The second issue under consideration is that of *stricter immigration control laws*, a conservative issue. The results from the H3b analysis were in the predicted direction, but the results did not reach a level of significance: $t(35) = 0.056$, $p = 0.956$, $d = 0.018$. The effect size of 0.018 can be described as negligible. The results did not provide compelling evidence for a greater degree of linguistic style matching between conservative liars and liberal truth tellers on stricter immigration control laws, which is an issue that conservatives tend to support.

Analysis 2: Discussion

Taken as a whole, the results from Analysis 2 provide intriguing evidence as to potential support for the prediction that conservative liars are more convincing than liberal liars. That is, conservative liars might be matching the linguistic style of liberal truth tellers slightly better than liberal liars are matching the linguistic style of conservative truth tellers. Thus, conservative liars may be better able to frame deceptive political views than liberal liars.

General Discussion

The results from Analyses 1 and 2 from this chapter provide evidence to suggest that linguistic style matching can be used as a theory to analyze characteristic linguistic features of deceivers and truth tellers from divergent political groups. In terms of the results from H1a and H1b of Analysis 1, liberal liars *and* conservative liars match the linguistic style of their corresponding truth tellers. In terms of the results from H2a and H2b of Analysis 1, conservative liars are better than liberal liars at matching the linguistic style of their corresponding truth tellers. In terms of the results from H3a and H3b, the results were intriguing and demonstrate the importance of further analyses to assess whether conservatives liars are potentially more convincing than liberal liars, and potentially better at framing deceptive discourse than liberals.

The combined results from Analysis 1 and 2 provide evidence in support of the prediction that liars are lying. More specifically, liberal liars and (perhaps to a greater extent) conservative liars are using characteristic features that are different from those of liberal truth tellers and conservative truth tellers, respectively. In general, these results support and expand the studies of Duran, Hall, McCarthy, and McNamara (2010), Hancock, Curry, Goorha, and Woodworth (2008), and Newman, Pennebaker, Berry, and Richards (2003). To my knowledge, no other study validates that liars are actually lying. Thus, this study may be the first to demonstrate that linguistic style matching can be used to validate that a lying condition has been observed by discourse participants.

This chapter demonstrates that linguistic style matching may be a deception strategy used to influence others into a false belief. Although the findings of the linguistic style matching analyses were encouraging, there were limitations. Specifically, the test set data for two (i.e., H1a, and H2a) out of the four linguistic style matching assessments of Analysis 1 contained less than 20 text files each, which is considered relatively low. Such a result may raise concerns as to the representativeness of the data; however, the relatively small degrees of freedom is a result of the randomization of the analysis sets. Future analysis will need to consider approaches to establishing greater confidence in this result. Future analysis will also need to determine whether conservatives are generally better than liberals at framing deceptive discourse, as the results of this chapter suggest, and whether that quality varies as a result of the political issues under consideration. However, even while the results here cannot be called conclusive, this chapter demonstrates an impressive development in deception research. That is, the Gramulator analysis of linguistic style matching facilitates detecting deception between deceivers and truth tellers from divergent political groups. Consequently, the Gramulator analysis might provide for greater comprehension of deception in political discourse.

Chapter 7: Distal Language

Research suggests that *distal language* is a linguistic feature of deception (Duran, Hall, McCarthy, & McNamara, 2010; Hancock, Curry, Goorha, & Woodworth, 2008; Newman & Pennebaker, 2003). More specifically, research suggests that distal language is more characteristic of deceivers because deceivers prefer to use distancing strategies in their discourse (e.g., references to someone other than the self). The purpose of this chapter is to assess the type of, and the degree to which distal language is used by deceivers relative to truth tellers.

In order to conduct an assessment of distal language, this chapter is arranged as follows. First, I briefly review the construct of *distal language* (see Chapter 1: Introduction). Second, I summarize the data collection methodology and data analysis methodology (see Chapters 3 and 4 for full details). Third, I detail the results of my analyses conducted on distal language use in deceptive discourse. And finally, I offer a discussion as to some of the interpretations and implications of the analyses.

Distal Language Nomenclature

Distal language is generally attributed to lexical items used by the *agent* to signal distance between that *agent* and *other figures* of the discourse (Ungerer & Schmid, 2006). The agent may be any *primary figure* of the discourse including the speaker or the writer. In the case of structured interactions (such as those discussed in Chapter 2: Literature Review), the agent may also be the sender, the deceiver, or the interlocutor. The *other figures* of the discourse may be the listener, the audience, the receiver, or subject(s) of the discourse.

While no agreed definition may be attributed to “distal language” per se, research in semantics and pragmatics uses the term *deixis* to refer to instances of distal language. Deixis refers to the interrelation between lexical items and the context in which those items are expressed (Levinson, 2006). For example, *distal deictic demonstratives* include terms such as *that, those, yonder, that one,* and *those ones*; and *distal deictic pronouns* include terms such as *their, theirs, its,* and *one's*. However, distal language is just one side of the deixis coin. That is, deixis is based on a notion of relativity. As such, distal language is understood in relation to *proximal language*. For example, proximal deictic demonstratives (e.g., *this* and *these*) are considered relational opposites of distal deictic demonstratives (e.g., *that* and *those*). Thus, analyses of distal language inherently also involves proximal language.

Based on the above notion of deixis, my dissertation assesses distal language relative to proximal language. In doing so, I follow a number of previous studies that have assessed relative distal language using deictic lexical features. These studies (see Chapter 2 for full details) include Duran, Hall, McCarthy, and McNamara (2010), Hancock, Curry, Goorha, and Woodworth (2008), and Newman, Pennebaker, Berry, and Richards (2003). Studies such as these suggest that deceivers (more so than truth tellers) tend to use second person, and third person pronouns (e.g., *you*, *his*, *hers*); and they tend *not* to use self-referencing first person pronouns (e.g., *me*, *ours*). Thus, the language of deceivers presumably entails distal deictic expressions in the form of pronouns.

Although such deception detection studies have advanced our understanding of distal language used by deceivers, these same studies have not moved beyond relatively simple regular pronoun counts (see Chapter 2: Literature Review). More specifically, available research reports results that are statistically significant but somewhat limited. That is, they do not extend pronoun usage into *pronominal* usage. Pronominals serve the same function as pronouns; however, they do not take the form of regular pronouns. For instance, in the sentence "*these people have the right to own guns*," the word *people* is a referent to an antecedent. As such, *people* is functioning as a pronoun, and is therefore pronominal. In this dissertation, I extend deictic pronoun assessment to include pronominals.

The lexical features of pronouns (and presumably also pronominals) can be used to assess distal language. However, whether a text is distal or proximal may also be assessed simply by the author's explicit indication. For example, the participants of my experiment (see Chapter 4) were instructed to respond as to whether the subjects of their views (e.g., deceptive or truthful) were about themselves (i.e., proximal), or about some other persons (i.e. distal). Specifically, participants either checked "Me" or "Some other person(s)," respectively. Thus, in my analysis, the use of "Me" is considered a measure of proximal language, whereas the use of "Some other person(s)" is considered a measure of distal language.

In sum, research suggests that deceptive language has characteristics that can be assessed. These characteristics of deception can be understood by an assessment of deixis. In deception studies, deixis represents the poles of distal and proximal language. One obvious way to determine whether text is characteristic of distal or proximal language is to consider the position of the primary figure of the text, the agent. If the primary figure of the text is someone other than the writer/speaker then the feature is distal rather than proximal. Turning to the explicit linguistic features of the text (as opposed to the agent), distal and proximal language in deception studies have traditionally considered only pronouns (e.g., *we* vs. *they*); however, pronominal features (e.g., *people*) may also be valuable indicators. In this dissertation, I conduct two analyses of distal language in deceptive discourse. The first analysis uses the self-indicated distance of the agent (i.e., *me* vs. *other*). And the second analysis uses the linguistic features of the text (i.e., *pronouns and pronominals*).

Analysis

All computational analyses described in this chapter are conducted using the Gramulator (see Chapter 3). The corpus data used in this chapter and the survey through which the corpus was collected are described in detail in Chapter 4. For convenience to the reader, I summarize below the critical details pertaining to the survey and the analyses that are discussed in this chapter.

Participants were asked to share their views on the issues of (1) *stricter gun control laws in the United States* and (2) *stricter immigration control laws in the United States*. Participants made one truthful contribution and one deceptive contribution. The contributions, in combination, form the corpus of *political discourse* that is used in this dissertation.

As part of the data collection process, participants were required to supplement their written views with a statement as to whether their contribution was mainly about "Me" (the participant) or "Some other person(s)." These responses combine to form the parameter allowing me to assess whether participants distanced themselves at the level of the agent. That is, if participants significantly tended to select "Some other person(s)" for the LIE condition, then there is evidence to suggest *distancing*.

In writing their views, participants used a wide variety of pronoun and pronominal terms. These terms can be considered distal or proximal, depending on context. The combination of pronouns and pronominals allow me to form a second parameter to assess whether participants distanced themselves at the level of the text. That is, if participants significantly tended to select distal pronouns and pronominals for the LIE condition then there is further evidence to suggest *distancing*.

In the first set of analyses, for self-indicated assessment of distancing at the level of the agent, I computed the number of times that participants indicated themselves as the primary figure of their views compared to the number of times that participants chose some other person(s) as the primary figure(s). I then used the statistical procedure of Fisher's Exact Test to assess whether there were significantly more instances of text where the frequencies of "Me" for liars was larger than the number for truth tellers.

In the second set of analyses, for an assessment of distancing at the level of the text, I used the pronouns and pronominals extracted from the respective TRUTH and LIE training set data. Ensuring context was appropriate for allocations, words such as *I*, *us*, and *Americans* were considered proximal, whereas words such as *they*, *them*, and *immigrants* are considered distal. I used the Gramulator procedures to compute the weighted frequencies of distal and proximal language across the corpus. I used t-tests to assess whether there was a significant trend in the direction of liars for the use of distal expressions.

Analysis 1: Explicit Test of Distal Language Usage in a Lie Condition

Analysis 1 features four conditions: GUNS_LIE, GUNS_TRUTH, IMM_LIE, IMM_TRUTH. The condition of GUNS_LIE includes participants who were randomly selected to discuss their deceptive views on gun control (n = 174). The condition of GUNS_TRUTH includes participants who were randomly selected to discuss their truthful views on gun control (n = 179). The condition of IMM_LIE includes participants who were randomly selected to discuss their deceptive views on immigration (n = 179). The condition of IMM_TRUTH includes participants who were randomly selected to discuss their truthful views on immigration (n = 174).

In the first of two assessments, the condition of GUNS_LIE is compared to the condition of GUNS_TRUTH. Additionally, the condition of IMM_LIE is compared to the condition of IMM_TRUTH. Note that in this assessment, the participants are independent for each contrasted condition. Thus, *no-one* in the GUNS_LIE condition appears in the GUNS_TRUTH condition.

In the next two assessments, the condition of GUNS_LIE is compared to the condition of IMM_TRUTH. Additionally, the condition of IMM_LIE is compared to the condition of GUNS_TRUTH. Note that in this assessment, the participants are paired for each contrasted condition. Thus, *everyone* in the GUNS_LIE condition also appears in the IMM_TRUTH condition.

Taken together, the two pairs of Fisher's Exact Test allow me to form an assessment as to whether liars tend to use distal language at the level of the text. The first assessment is within the same political issue considered, but between participants; the second analysis is between the different political issues considered, but within participants. If the LIE condition tends to produce higher frequencies for the distal language option (i.e., the primary figure of the discourse is *not* the writer of the discourse) then evidence will have been provided to support the distancing hypothesis.

Results. For the first analysis, the Fisher's Exact Test results were significant and in the predicted direction for both the independent and paired tests. On the issues of *stricter gun control laws* (GUNS LIE and GUNS TRUTH) and *stricter immigration control laws* (IMM LIE and IMM TRUTH), truth tellers chose themselves as the primary figure of their views significantly more often than liars (see Tables 10 and 11). That is, truth tellers shared more about themselves to support their truthful views than did liars to support their deceptive views. These results support the distal language hypothesis.

Table 10

Fisher's Exact Test: Independent Test Results

Issue	Frequency		Fisher's Exact Test
	Texts	ME	ME
GUNS LIE	174	56	*.038*
GUNS TRUTH	179	77	
IMM TRUTH	174	78	*.003*

Table 11

Fisher's Exact Test: Paired Test Results

Issue	Frequency		Fisher's Exact Test
	Texts	ME	ME
GUNS LIE	174	56	*.021*
IMM TRUTH	174	78	
GUNS TRUTH	179	77	*.008*

Analysis 2: Implicit Test of Distal Language Usage in a Lie Condition

Analysis 2 featured two conditions: LIE and TRUTH. The assessed data for the condition of LIE includes the test set examples (See Chapter 4: Methods for details of training and testing) of all participants' texts that were presented as deceptive views (n = 116). The assessed data for the condition of TRUTH includes the test set examples of all participants' texts that were presented as truthful views (n = 116).

In total, eight t-test assessments were conducted. In each of the eight tests, I compared the test set data from the LIE and TRUTH conditions. For each comparison, the assessment index was formed as a list of pronouns or pronominals (see Table 12). Pronouns were compiled as a straight-forward standard list (e.g., first person pronouns are *I, me, myself, mine, my, we, us, ourselves, ourself, ours, our*). [1] However, note that the non-standard 1st person plural pronoun *ourself* was also included. By contrast, pronominals were extracted from derived lists of typicals (see Chapter 4: Methods). Specifically, any pronoun (e.g. *you*, *us*, *themselves*), and any word functioning as a

[1] The data was modified for the purposes of clear identification: *the us* > *the_us* and *a us* > *a_us*. The underscore prior to the string *us* ensured that *us* as the country was not confused with the first person plural pronoun.

pronoun (e.g., *people*, *citizen*, *families*) were considered potential pronominals. TRUTH pronominals were derived from the TRUTH typicals, and LIE pronominals were derived from the LIE typicals. The Concordancer module of the Gramulator was used to ensure that all the pronominals were pronominals in context (see Chapter 3 for details about the modules). Note that typicals (as opposed to differentials) were used as the source material. This option provided the greatest diversity of potential candidate terms.

Table 12

Pronoun and Pronominal Terms

Test	Terms
1	*I me myself mine my we us ourselves ourself ours our*
	you yourself yours your thou thee thyself thine they yourselves ye he him
	himself his she her herself hers it itself its one oneself one's they them themself
2	*themselves theirself theirselves theirs their*
4	*themselves we they our us people their I you citizens*
5	*citizens their citizen they immigrants person our we families people*
7	*we our us I*
8	*themselves they people their you citizens immigrants person citizen families*

Note: Test 3 was a combination of Tests 1 and 2; Test 6 was a combination of Tests 4 and 5.

Taken together, the eight t-tests allow me to form an assessment as to whether liars use distal language at the level of lexical density (i.e., the proportion of the text that is constituted by the target language). If the LIE condition tends to produce higher values for the distal language option then evidence will have been provided to support the distancing hypothesis.

Pronoun Results. Test 1 assessed the conditions of LIE and TRUTH for the evaluation of 1st person pronouns (*I, me, myself, mine, my, we, us, ourselves, ourself, ours, our*). The means were in the direction of LIE; however, the difference between the groups was not significant. Test 2 also assessed the conditions of LIE and TRUTH. In this test, the evaluation was for a combination of 2nd and 3rd person pronouns (*you yourself yours your thou thee thyself thine they yourselves ye he him himself his she her herself hers it itself its one oneself one's they them themself themselves theirself theirselves theirs their*). The means were in the direction of LIE and the difference between the groups was significant (p = .030). Test 3 assessed the conditions of LIE and TRUTH for the evaluation of a combination of 1st, 2nd and 3rd person pronouns. In other words, Test 3 was a combination of Test 1 and Test 2. The means were in the direction of LIE and the difference between the groups was significant (p = .015).

Table 13

Pronoun Assessment Results

#	ASSESSMENT	*M*	*SD*	*t, p, d*
	LIE_{test} → 1st_PERSON_PRONOUNS	3.451	2.745	t (1,230) = 1.185,
1	$TRUTH_{test}$ → 1st_PERSON_PRONOUNS	3.035	2.599	p = .237, d = 0.156
	LIE_{test} → 2nd_and_3rd_PERSON_PRONOUNS	1.664	1.191	t (1,230) = 2.188,
2	$TRUTH_{test}$ → 2nd_and_3rd_PERSON_PRONOUNS	1.366	0.855	**p = .030,** d = 0.287
	LIE_{test} → 1st_2nd_and_3rd_PERSON_PRONOUNS	2.111	1.114	t (1,230) = 2.46,
3	$TRUTH_{test}$ → 1st_2nd_and_3rd_PERSON_PRONOUNS	1.783	0.901	p **= .015,** d = 0.323

Taken together, the results suggest that LIE texts have a greater consistency in terms of pronoun density. That is, regardless of the pronoun type (1st, 2nd, or 3rd), there may be a greater tendency for deceptive writers to use pronouns, as a proportion of overall lexical terms. Because results for 2nd and 3rd person pronouns are significant, and 2nd and 3rd person pronouns are likely to be more distal (i.e., less reflective of the self), the argument can be made that the pronoun results offer evidence in support of the distal language theory.

Pronominal Results. Tests 4, 5 and 6 evaluate the test data sets for the conditions of LIE and TRUTH (see Table 14). The tests are evaluated using indices corresponding to LIE and TRUTH pronominals, which were derived from the training set data. These three tests serve to validate the approach of proximal and distal pronominal assessments that are conducted in Tests 7 and 8 (see Table 15). Specifically, if Tests 4, 5, and 6 produce values consistent with previous analyses (i.e., there is higher consistency for the LIE condition) then the pronominal approach will have evidence of validation.

In Test 4, the conditions of LIE and TRUTH were evaluated by the index of LIE pronominals (*themselves, we, they, our, us, people, their, I, you, citizens*). The means were in the direction of LIE and the difference between the groups was significant (p = .031). In Test 5 the conditions of LIE and TRUTH were evaluated by the index of TRUTH pronominals (*citizens their citizen they immigrants person our we families people*). Again, the means were in the direction of LIE and the difference between the groups was significant (p = .021). In Test 6 the conditions of LIE and TRUTH were evaluated by the combined indices of LIE and TRUTH. The means were in the direction of LIE but the difference between the groups was merely approaching significance (p = .060).

Table 14

Pronominals: LIE and TRUTH Assessment Results

#	ASSESSMENT	M	SD	t, p, d
4	LIE_{test} → LIE_{train}_PRONOMINALS	6.554	3.563	*t* (1,230) = 2.173,
	$TRUTH_{test}$ → LIE_{train}_PRONOMINALS	5.612	3.019	***p* = .031**, *d* = 0.285
5	LIE_{test} → $TRUTH_{train}$_PRONOMINALS	6.025	3.796	*t* (1,230) = 2.325,
	$TRUTH_{test}$ → $TRUTH_{train}$_PRONOMINALS	4.991	2.924	***p* = .021,** *d* = 0.305
6	LIE_{test} → LIE_{train}_and_$TRUTH_{train}$_PRONOMINALS	5.333	2.783	*t* (1,230) = 1.888,
	$TRUTH_{test}$ → LIE_{train}_and_$TRUTH_{train}$_PRONOMINALS	4.699	2.315	*p* = .060, *d* = 0.248

Two results from these analyses require further investigation and explanation. First, the result from Test 5 suggests that TRUTH pronominals are more common in LIE texts than they are in TRUTH texts. The result suggests that TRUTH pronominals are more broadly dispersed (and therefore weak as indicators) whereas LIE pronominals are more concentrated (and therefore strong as indicators). The second enquiry concerns why Test 6 did not produce a significant result. Specifically, if the LIE index significantly predicts the LIE data (as in Test 4) and the TRUTH index also significantly predicts the LIE data (as in Test 5), then why does the combination of LIE and TRUTH not significantly predict the LIE data in Test 6? Further analyses revealed that the apparent weakening of the result is caused by the relative strength of the items that were added to the index. That is, the combination LIE_TRUTH index received 4 items from the LIE index that were not shared in the TRUTH index (*immigrants*, *person*, *citizen*, and *families*). Similarly, the combination LIE_TRUTH index also received 4 items from the TRUTH index that were not shared in the LIE index (*themselves*, *us*, *I*, and *you*). Although Test 6 shared all the items, the additional LIE items were significantly weaker at detecting TRUTH than the TRUTH items were at detecting LIE. The explaining results are as follows: $\text{TRUTH}_{\text{test}} \rightarrow$ NEW_LIE Pronominals (*immigrants, person, citizen, families*), $M = 2.415$; $\text{LIE}_{\text{test}} \rightarrow$ NEW_TRUTH Pronominals (*themselves, us, I, you*), $M = 3.604$; $p = .022$. These results also lend evidence to the claim that TRUTH pronominals are more broadly dispersed (i.e., weak), whereas LIE pronominals are more concentrated (i.e., strong).

Taken together, the results suggest that LIE texts have a greater consistency in terms of pronominal density. That is, regardless of whether the pronominals stem from LIE or TRUTH training data, there appears to be a greater tendency for liars to use pronominals as a proportion of overall lexical terms. The result supports previous findings reported in this dissertation as to the markedness of the LIE condition (see Chapter 5: Validation).

As with the six previous tests, Tests 7 and 8 evaluate the two conditions of LIE and TRUTH. However, in these final two tests, the indices correspond to PROXIMAL and DISTAL pronominals (see Table 15). In Test 7 the conditions of LIE and TRUTH were evaluated by the index of PROXIMAL pronominals (*we, our, us, I*). The means were in the direction of LIE but the difference between the groups was not significant ($p = .373$). In Test 8, the conditions of LIE and TRUTH were evaluated by the index of DISTAL pronominals (*themselves, they, people, their, you, citizens, immigrants, person, citizen, families*). The means were in the direction of LIE but the difference between the groups was merely approaching significance ($p = .086$).

Table 15

Pronominals: Proximal and Distal Assessment Results

#	ASSESSMENT	M	SD	t, p, d
	LIE_{test} → PROXIMAL_PRONOMINALS	7.988	6.534	t (1,230) = 0.893,
7	$TRUTH_{test}$ → PROXIMAL_PRONOMINALS	7.229	6.407	p = .373, d = 0.117
	LIE_{test} → DISTAL_PRONOMINALS	4.271	2.855	t (1,230) = 1.722,
8	$TRUTH_{test}$ → DISTAL_PRONOMINALS	3.686	2.285	p = .086, d = 0.226

Taken together, the results again suggest that LIE texts have a greater consistency in terms of pronominal density. That is, regardless of whether the pronominals stem from proximal or distal training data, there appears to be a possible tendency for liars to use pronominals as a proportion of overall lexical terms. Of course, there were only two results in this set of tests, and one of those results was merely directional and the other was not fully significant. As such, any conclusions must be interpreted with a high degree of caution. However, even with this provision, the results are tantalizing inasmuch as they appear to offer some support for previous analyses. Specifically, they suggest that LIE texts may be characterized by the use of distal pronominals: a result that would again offer some support to the distal language theory.

General Discussion

The primary purpose of this chapter was to assess the degree to which deceivers use distal language relative to truth tellers. Taken together, the results of the explicit and implicit analyses broadly support the distal language hypothesis. Specifically, the results suggest that people who are lying use more distal language. This distal language is made manifest in the participants' explicit choice to talk about others (as opposed to the self) and also in the participants' implicit choice to use a higher proportion of diverse deictic expressions (i.e., pronouns and pronominals).

Taken as a whole, the results provide evidence to suggest that deceivers employ *distancing strategies* (Duran et al., 2010; Hancock et al., 2008; Newman & Pennebaker, 2003). In terms of deceivers' explicit choice of primary figures in their deceptive discourse, deceivers appear willing to refer to the experiences of others rather than their own experiences to support their deceptive views, whereas truth tellers appear to be comfortable sharing their own experiences to support their truthful views. Such a conclusion supports the current theoretical framework.

An alternative explanation for why deceivers are more likely to refer to the experiences of others in their deceptive discourse is that deceivers may not be exploiting "distancing strategies" so much as exploiting "managing strategies" that facilitate creating and writing a lie. In other words, people may generally prefer to talk about themselves, and particularly when the task is relatively simple and involves a low level of cognitive load (i.e., tell the truth). However, when the task is more complex and involves a relatively higher level of cognitive load (i.e., tell a lie), people may need to go to a larger repository (i.e., others). Future research needs to consider how cognitive load relates to the performance of telling a lie compared to telling the truth, and what managing strategies are exploited to facilitate that performance.

In terms of deceivers' implicit choice to use of a higher proportion of diverse pronouns and pronominals, deceivers may be concerned with other referencing that begins with introducing the primary figure (i.e., explicit choice) and continues throughout the discourse with pronoun and pronominal usage (i.e., implicit choice). That is, deceivers may be concerned with exploiting distancing strategies that involve implicit, as well as explicit personal distance.

An alternative explanation for deceivers' implicit choice to use a higher proportion of diverse pronouns and pronominals is that deceivers may be exploiting distancing strategies that may not be restricted to *personal distance*. That is, deceivers' may be exploiting distancing strategies that also involve *ambiguous distance*. In other words, people may use a diversity of pronouns and pronominals to avoid explicitly specifying the referent (e.g., *they, them, people*), and thereby creating a level of ambiguity that perhaps hinders clear interpretations, and that perhaps redirects blame or ownership. That is, when deceivers use the pronominal *people*, they may do so without distinguishing *immigrants* from *Americans*, or *my (American) people* from *other (non-American) people*. Future research needs to consider how pronouns and pronominal usage may allow for a level of ambiguity that is characteristic or indicative of deceptive discourse relative to truthful discourse.

The explanations of the results broadly support the current theoretical framework. However, these explanations are more complex than the traditional pronoun usage explanations (Hancock, Curry, Goorha, & Woodworth, 2008; Newman & Pennebaker, 2003), which suggest that truth tellers are proximal pronoun users (e.g., *I , me, my*), and deceivers are distal pronouns users (e.g., *you, yours, they, them, theirs*). Based on the findings in this chapter, deceivers may use a diversity of pronouns and pronominals that includes proximal pronouns to distance themselves from their deceptive discourse. Such a diversity of pronoun and pronominal usage is also likely to be within the context of referring to others (not themselves).

This chapter challenged traditional word counting approaches that measure abstract lexical items (e.g., pronouns) and introduced the Gramulator's n-gram analyses that measure contextualized linguistic features (e.g., differentials and typicals). Although the findings of these analyses were encouraging, there were limitations. Specifically, the distal pronominal index was only a predictor of the LIE corpus in a 1-tailed test. Such a result raises concern as to how the distal and proximal pronominals were derived. In other words, the n-gram analyses used to derive the pronominals remained strong, but the process by which those derived pronominals were assigned to the distal and proximal categories was based on qualitative assessments that may not be easily replicable. Thus, the distal-proximal distinction needs further consideration to allow for standardizing the process.

Future research will focus on better assessing the parameters of the distal and proximal indexical measures discussed in this chapter. That is, it is important to assign confidence values to the distal and proximal indices so as to better assess the linguistic features of deceptive discourse relative to truthful discourse. Future analysis must also consider to what degree plural nouns are pronominally used by deceivers relative to truth tellers, and whether any significant differences may be attributed to distancing strategies (e.g., personal distance, and ambiguous distance) exploited by deceivers.

This chapter demonstrates the efficacy of the Gramulator to facilitate assessing distal language used by deceivers. The major results of the analyses certainly provide sufficient initial evidence that distal language better predicts deceptive discourse than truthful discourse. Furthermore, there have been no previous investigations of the degree to which pronominal references contribute to the language of deceptive discourse. The distal language assessments involving pronoun and pronominal usage provided an opportunity to better assess the distal strategies that are more characteristic of deceivers relative to truth tellers, and to better demonstrate the Gramulator design of identifying prominent linguistic features of correlative text types.

Chapter 8: Discussion

Overview

My dissertation set out to address the following research question: "What does *linguistic style matching* and *distal language* reveal about deception in political discourse?" More specifically, I investigated *which features of language uniquely characterize the views of deceivers and truth tellers, relative to each other; and of self-described liberals and self-described conservatives, relative to each other?*" To this end, I conducted a deception based contrastive corpus analysis to identify differences between deceivers and truth tellers, and differences between liberal deceivers and conservative deceivers. The deception analysis comprised two parts: a linguistic style matching analysis, and a distal language analysis. The results of the linguistic style matching analysis and the distal language analysis point toward a variety of distinctions that are characteristic of deceivers relative to truth tellers, and of conservative deceivers relative to liberal deceivers. Under a lie condition, the results produced evidence supporting the following:

1) There is a greater degree of linguistic style matching between liberals and conservatives who share the same rhetorical goal, but not the same experimental condition (e.g., liberals: lie condition; conservatives: truth condition) than between liberals and conservatives who share the same truth condition, but not the same rhetorical goal. Formally, we can represent this result as RG > TC, where RG = rhetorical goal, and TC = truth condition.

2) There may be a greater degree of linguistic style matching between conservative liars and liberal truth tellers who share the same rhetorical goal than

between liberal liars and conservative truth tellers who share the same rhetorical goal. Formally, we can represent this result as $RG(C^- \# L^+) > RG(L^- \# C^+)$, where RG = rhetorical goal, C^- = conservative liars, # = linguistic style matching, L^+ = liberal truth tellers, L^- = liberal liars, and C^+ = conservative truth tellers.

3) There is more distal language usage on the part of liberal and conservative liars than on the part of liberal and conservative truth tellers. Formally, we can represent this result as $d^- > d^+$, where d^- = distal language usage of liars, and d^+ = distal language usage of truth tellers.

In combination, we can formally represent the three results as $RG\ (C^-_d \# L^+) > RG\ (L^-_d \# C^+)$. This representation reads as follows: For a given rhetorical goal (RG), conservative liars (C^-) match the linguistic style of (#) liberal truth tellers (L^+) to a greater extent (>) than liberal liars (L^-) match the linguistic style of conservative truth tellers (C^+), with distal language (when present) being more prominent in the lie conditions of conservatives (C^-_d) and liberals (L^-_d).

Overall, the results support the claim that linguistic style matching and distal language usage may be deception strategies used by deceivers to influence others into a false belief. The results are important because they demonstrate, perhaps for the first time, a reliable relationship between linguistic style matching and shared rhetorical goals, and the distal language usage of liberals and conservatives, relative to each other.

Motivation for the Dissertation

Political discourse reveals much about the politicians, political groups, and the public. That is, while political discourse reveals information about the views, values, and voting trends of such political agents, political discourse also reveals information about the relationship between deceivers and truth tellers. In other words, divergent political groups such as liberals and conservatives have the same agenda. First, liberals and conservatives want to communicate their respective worldviews so as to gain supporters of such worldviews. Second, liberals and conservatives seek to demonstrate why their respective worldviews represent the truth and why opposing worldviews represent that which is not the truth (e.g., President Obama: "It's not true"), or essentially that which is a lie (e.g., Representative Wilson: "You lie!"). Thus, notions of the truth and of lies are fundamental to political discourse.

Political discourse is an abundant resource for both deception analysis and political discourse analysis. However, deception analysis and political discourse analysis differ in their considerations of political discourse. That is, although deception analysis (Duran, Hall, & McCarthy, 2009; Newman, Pennebaker, Berry, & Richards, 2003) may consider differences at the group level in terms of political positioning (e.g., pro-life and pro-choice), deception analysis may benefit by expanding such a consideration to include differences between divergent political groups (e.g., liberals and conservatives). And although political discourse analysis (Feldman, 2007; Lakoff, 2002; 2004) may consider differences between the views of divergent political groups, political discourse analysis may benefit by also considering differences between such views expressed under a truth condition and a lie condition. Thus, deception analysis needs to further assess political discourse across a diversity of political categories, and political discourse analysis needs to further assess political discourse across a diversity of experimental conditions.

The goal of this dissertation was to bridge such gaps in deception analysis and political discourse analysis. More specifically, this dissertation was a deception analysis of political discourse expressed by liberals and conservatives. The purpose of this deception analysis was to reveal differences between deceivers and truth tellers, and between liberals and conservatives. Such a purpose facilitated assessing liberal and conservative worldviews, and assessing each political group's ability to frame deceptive discourse.

Experimental Approach

This dissertation adopted a mixed methods approach that combined traditional experimental elements with new technologies to facilitate qualitative and quantitative analyses. More specifically, I conducted an online experiment using a web-based survey entitled, *Truth and Lies in Politics* (see Appendix A). This survey was designed with the document creation tools available through Google Docs and disseminated to willing participants using the crowdsourcing platform of Amazon Mechanical Turk (MTurk). Such an approach facilitated collecting a corpus of political discourse comprising truthful texts and deceptive texts submitted by self-described liberals and self-described conservatives.

The Corpus

Following the corpus standards for deception analysis proposed by McCarthy, Duran, and Booker (2012), I collected a corpus of 706 original submissions with an equal distribution of truthful texts (353) and deceptive texts (353). The truthful texts comprised typed submissions of truthful views on the issues of *stricter gun control laws in the United States* and *stricter immigration control laws in the United States.* Likewise, the deceptive texts comprised typed submissions of deceptive views on the same issues. A corpus of such a size and distribution allowed for conducting a contrastive corpus analysis (CCA) on several contrastive corpora in addition to the truth and lie corpora. Such contrastive corpora included the truth and lie corpora of liberals, the truth and lie corpora of conservatives, the truth and lie corpora of liberal supporters for stricter gun control laws, the truth and lie corpora of liberal opposers of stricter immigration control laws, the truth and lie corpora of conservative supporters of stricter immigration control laws, and the truth and lie corpora of conservative opposers of stricter gun control laws. All contrastive corpora were divided into training set and test set data for analysis.

Analysis Tool

The computational textual analysis tool used to analyze the contrastive corpora of this study was the Gramulator. This tool features modules that facilitated all phases of analysis including pre-assessment procedures (e.g., converting, cleaning, and sorting data) and the assessment procedures (e.g., conducting statistical tests, n-gram analyses, and qualitative analyses). More importantly, I used the Gramulator to process and analyze the typical form and the differential form of above average bi-grams, or collocations. That is, I used the Gramulator to assess prominent linguistic features in context that are characteristic, or typical of each corpus, and that are indicative of, or differentiate one corpus relative to the other contrastive corpus.

Analyses

This dissertation included three types of analyses of the data: 1) a validation of the corpus; 2) a linguistic style matching analysis; and 3) a distal language analysis.

The purpose of the validation was to demonstrate confidence in the analysis of the contrastive constructs under consideration (e.g., TRUTH and LIE; FOR_IT and AGAINST_IT). The purpose of the linguistic style matching analysis was to demonstrate that an experimental lie condition has been satisfied, and to assess the ability of liberals and conservatives to frame deceptive discourse. And the purpose of the distal language analysis was to demonstrate that explicit and implicit lexical choices by deceivers differ from such choices by truth tellers.

Validation Approach. I used the Internal Validity Process (IVP) to assess the corpus and its derived indices (McCarthy et al., 2012). In addition, I used the Gramulator to evaluate the data according to the IVP. The IVP facilitated in establishing the validity of the contrastive constructs, which represented the conditions of the experimental survey. That is, the participants were randomly assigned to a truth condition and a lie condition, under which they submitted their views on two political issues (i.e., *stricter gun control laws* and *stricter immigration control laws*). The submissions under the truth condition and the lie condition allowed for dividing the data into the contrastive construct of TRUTH and LIE. The participants were also instructed to rate their level of support or opposition for each issue. The results from those responses allowed for dividing the data into the contrastive construct of FOR_IT and AGAINST_IT. Thus, the primary purpose of the validity study was to assess the contrastive constructs of TRUTH and LIE, and FOR_IT and AGAINST_IT.

Validation Results. The validation results provided compelling evidence to suggest that the LIE data and the FOR_IT data comprised homogenous and marked language. More specifically, the results provided evidence to suggest the following characteristics about the LIE data and the FOR_IT data. First, each data set is consistent. Second, the differentials across the corresponding index are consistent. Third, such indices are the more distinctive, or marked index relative to their contrasting indices (i.e., TRUTH index, and AGAINST_IT index). Thus, I established intrinsic validity for the language used by deceivers, and for the language used by issue supporters, respectively.

With this validation study, the goal was to address a gap in deception analysis. More specifically, deception analysis typically involves the extrinsic validity of the data. Therefore, it is common for deception analysis to claim (rather than demonstrate) that their LIE data is in fact a corpus of texts representing the LIE construct, and that their TRUTH data is in fact a corpus of texts representing the TRUTH construct. Such a claim is often based on previous claims expressed by researchers who have used the same LIE and TRUTH data in earlier deception studies (e.g., data from Hancock, Curry, Goorha, & Woodworth, 2008). This dissertation worked beyond the claims, and making no a-priori assumptions, established validity for the LIE data, relative to the TRUTH data.

In contrast to traditional deception analysis (Duran et al., 2009; Hancock et al., 2008; Newman et al., 2003), this validation study also established validity for the language of supporters (i.e., FOR_IT), relative to the language of opposers (i.e., AGAINST_IT). Traditional deception analysis tends to limit its focus on *extrinsic validation* to the language of deceivers relative to truth tellers. Expanding such a focus to include the *intrinsic validation* of the language of deceivers relative to truth tellers and of the language of supporters relative to opposers is important for providing confidence that measured differences between deceivers and truth tellers from divergent political groups are indeed attributed to the validity (i.e., extrinsic and intrinsic) of the contrastive data sets and their corresponding indexical measures, and not to some other phenomenon.

Linguistic Style Matching Approach. I used the typical form of bi-grams produced by the Gramulator as a measure of linguistic style matching. This approach is based on the understanding that linguistic style matching is evidence of “harmonized” worldviews (Niederhoffer & Pennebaker, 2002). As such, linguistic style matching between deceivers and truth tellers from divergent political groups implies that such deceivers and truth tellers are using language characteristic of a shared worldview, albeit an appropriated deceptive worldview on the part of deceivers as compared to an authentic truthful worldview on the part of truth tellers.

To test for linguistic style matching between deceivers and truth tellers from divergent political groups, I rely on a concept of relativity. In other words, linguistic style matching involves measuring the typicals characteristic of two groups, A and B (e.g., conservative liars and liberal truth tellers) and the typicals characteristic of another set of two groups, C and D (e.g., conservative truth tellers and liberal truth tellers), and then conducting statistical tests to determine whether the differences between those measures are significant. Significant results provide evidence to suggest that there is a greater degree of linguistic style matching between A and B relative to C and D.

To measure the typicals characteristic of any two groups, I used the training set index of one group to assess the data set of the corresponding group. For example, in the representation of (C^{-} # L^{+}), I used the training set index of conservative liars (relative to conservative truth tellers) to measure the test set of liberal truth tellers. Such an approach relies on measuring language in context (e.g., typical bi-grams) and deviates from measuring words in isolation (Hancock et al., 2008; Niederhoffer & Pennebaker, 2002).

Linguistic Style Matching Results. The results from the linguistic style matching analysis revealed a greater degree of linguistic style matching between issue supporters who lie and issue opposers who tell the truth than between issue supporters who tell the truth and issue opposers who tell the truth. For example, on the conservative issue of stricter immigration control laws, conservative liars who supported stricter immigration control laws matched the linguistic style of liberal truth tellers who opposed the same issue better than conservative truth tellers who supported the issue of stricter immigration control laws matched the linguistic style of liberal truth tellers who opposed the same issue because conservative liars and liberal truth tellers share the same rhetorical goal, and conservative truth tellers and liberal truth tellers do not share the same rhetorical goal. In other words, conservative liars and liberal truth tellers expressed views against stricter immigration control laws, whereas as conservative truth tellers expressed views in favor of stricter immigration control laws and liberal truth tellers expressed views against those same laws. Thus, there is evidence of *convergence in the language* of contrasting political groups who share the same rhetorical goal, but one political group is under a lie condition and the corresponding political group is under a truth condition, and evidence *of divergence in the language* of contrasting political groups who share the same truth condition, but do not share the same rhetorical goal.

The results from the linguistic style matching analysis also revealed that conservative liars are potentially better than liberal liars at matching the linguistic style of their corresponding truth tellers. This result was primarily based on the evidence that conservative liars performed twice as well as liberal liars at matching the linguistic style of their corresponding truth tellers, and such a performance occurred on both the liberal issue under consideration, and the conservative issue under consideration. Thus, conservatives may potentially be better than liberals at framing deceptive discourse.

Taken as a whole, the linguistic style matching results revealed that evidence of linguistic style matching is independent of issue ownership. That is, liberal liars were better than liberal truth tellers at matching the linguistic style of conservative truth tellers on the liberal issue of stricter gun control laws, and conservative liars were better than conservative truth tellers at matching the linguistic style of liberal truth tellers on the liberal issue of stricter gun control laws, and on the conservative issue of stricter immigration control laws. Thus, a shared rhetorical goal is a better predictor of linguistic style matching between deceivers and truth tellers from divergent political groups than any perceived issue handling advantages on the part of one group, relative to the other contrasting group.

Distal Language Approach. I used the Gramulator to conduct the distal language analysis, which assessed the agents' explicit and implicit lexical choices. In terms of the explicit lexical choices, I used the Fisher's exact test to measure the degree to which deceivers chose some other person(s) as the primary figure(s) of their deceptive discourse relative to the degree to which truth tellers chose some other person(s) as the primary figure(s) of their truthful discourse. In terms of the implicit lexical choices, I created an index of pronouns and pronominals from the Gramulator's user-defined option. Such an index was used to measure distal language usage of deceivers relative to truth tellers. The pronoun assessment considered words in isolation as a means of comparing the Gramulator results to the results of previous studies that use word counting tools. The pronominal assessment considered words in context as a means of challenging and expanding previous studies that use word counting tools. Such an approach provided an opportunity to better understand how explicit and implicit lexical choices contribute to the language usage of deceivers relative to truth tellers.

Distal Language Results. The results from the distal language analysis revealed that deceivers were more likely than truth tellers to make the explicit choice to talk about others (as opposed to the self) in their political discourse. That is, deceivers indicated that the primary figure(s) of their deceptive discourse was some person(s) other than themselves. Likewise, truth tellers were more likely to make the explicit choice to talk about themselves (as opposed to others) in their truthful discourse. That is, truth tellers indicated that they chose themselves as primary figure(s) of their truthful discourse. Thus, deceivers tend to distance themselves from the primary figure of their deceptive discourse, whereas truth tellers tend to avoid such distancing strategies.

The results from the distal language analysis also revealed that deceivers were more likely than truth tellers to make the implicit choice to use a higher proportion of a diversity of pronouns and pronominals. That is, deceivers used proximal pronouns (e.g., first person pronouns) in addition to distal pronouns (e.g., second and third person pronouns) and distal pronominals in context (e.g., *people*, *citizens*). Thus, deceivers appear to appropriate language that is characteristic of truth tellers (e.g., first person pronouns) in combination with other distal pronouns and pronominals so as to communicate deceptive discourse.

Taken as a whole, the distal language results revealed that distal language usage may be a distancing strategy used by deceivers. However, the way in which deceivers distance themselves from their deceptive discourse is more complex than the traditional assessments that tend to focus on the polarities of pronoun usage (e.g., first person pronouns versus second and third person pronouns). That is, deceivers use a combination of techniques as part of their distancing strategy that involves their explicit and implicit choices of distal language.

Interpretations

Taken as a whole, the results may be interpreted as evidence in general support of the claim that linguistic style matching and distal language usage may be deception strategies characteristic of deceivers (Duran et al., 2009; Neiderhoffer & Pennebaker, 2002; Newman et al., 2003). This interpretation is based on evidence of a greater degree of linguistic style matching between deceivers and truth tellers across the political issues under consideration, and on evidence of a greater degree of distal language usage by deceivers. The results may also be interpreted as evidence in general support of the claim that deceivers use fewer self-references than truth tellers (Hancock et al., 2008; Newman et al., 2003). However, closer examination suggests that such an interpretation does not account for the full results. A more appropriate interpretation of the results is that deceivers use fewer self-references than truth tellers, but a greater diversity of self-references and other-references. This interpretation is based on evidence of a higher proportion of pronouns (e.g., first, second, and third person) and pronominals used by deceivers relative to truth tellers. Thus, the combined results provide compelling evidence in support of the claim that linguistic style matching and distal language may be used by deceivers to influence others into a false belief.

Implications

The implications of the findings of this dissertation are that linguistic style matching and distal language usage may be used to assess contrasting political group's ability to frame deceptive discourse. More specifically, the linguistic style analysis of this dissertation provided evidence to suggest that both liberals and conservatives satisfied the experimental lie condition, and potential evidence to suggest that conservative liars are more convincing than liberal liars under such a lie condition. However, such evidence does not imply that conservatives and liberals are categorical liars. Such a conclusion simply implies that when an experimental lie condition has been satisfied, conservatives and liberals perform well at framing deceptive discourse, and conservatives are potentially more convincing than liberals at framing deceptive discourse. Likewise, the distal language analysis of this dissertation provided evidence to suggest that distal language is more characteristic of deceivers. However, such evidence does not imply that deceivers completely avoid proximal language. As the results of this dissertation have revealed, deceivers use a combination of distal language and proximal language to frame deceptive discourse. Taken as a whole, linguistic style matching and distal language usage contribute to our understanding of the prominent linguistic features of deceivers relative to truth tellers, and of conservative deceivers relative to liberal deceivers.

Relevance of the Dissertation

This dissertation is relevant to a diversity of fields and industries that use computational and corpus-based approaches, and that are concerned with communicating to groups from contrasting worldviews. Such fields range from ANLP (applied natural language processing) to second language acquisition through to criminal intelligence.

And such industries range from fiction writing to data analytics through to politics. These fields and industries may benefit from understanding prominent linguistic features of deceivers relative to truth tellers, and of issue supporters relative to issue opposers. Moreover, such fields and industries may benefit from exposure to new technologies that facilitate such an understanding. Thus, this dissertation is relevant to a diversity of fields and industries largely because of the research approaches, methods, and techniques applied to this deception analysis of political discourse.

In terms of the emerging field of ANLP, researchers of this field are concerned with computational approaches to analyzing language and language-related issues and with validation procedures that provide confidence in such analyses. This study used the computational textual analysis tool of the Gramulator in combination with web-based technologies (e.g., Google Docs) and crowdsourcing techniques (e.g., Amazon Mechanical Turk) to collect and assess political discourse. This study also demonstrated the importance of establishing intrinsic validity in addition to extrinsic validity, and how the internal validation process (IVP) can be used to establish intrinsic validity. The IVP allowed for assessing the contrasting constructs under consideration (e.g., TRUTH and LIE; FOR_IT and AGAINST_IT), the homogeneity of the data and the derived indices, and the marked and default language constituting the data sets and the corresponding indices. Thus, new applications of such computational and validity assessment approaches contribute to the future of ANLP research.

In terms of the field of second language acquisition, researchers and practitioners of this field are concerned with the process of learning a second language, including understanding the social environment (e.g., context, setting, and structure) of the learner.

This dissertation suggested that distal language usage by deceivers and truth tellers relative to one another, and by liberals and conservatives relative to one another may be a predictor of the social distance between the contrasting groups. In other words, this study suggested that social distance between such contrasting groups is made manifest by their explicit and implicit choices of distal language. Thus, second language acquisition may benefit from assessing distal language usage as a measure of social distance.

In terms of the field of criminal intelligence, researchers and practitioners of this field are concerned with the language characteristic of deceivers in order to conduct criminal investigations. This dissertation revealed that linguistic style matching and distal language usage may be deception strategies used by deceivers to convince others into a false belief. Thus, criminal intelligence investigations may consider conducting appropriate linguistic style matching and distal language analyses to facilitate detecting deception in recorded conversations made by criminal suspects.

Turning to industry, this dissertation has implications for fiction writers interested in creating deceptive characters. Such a process involves developing monologues and dialogues that are typical of deceivers. This dissertation revealed that the language of deceivers is rather complex. That is, deceivers tend to refer to others in their deceptive discourse, but use a diversity of pronouns and pronominals in that discourse. Such an understanding of the language of deceivers facilitates creating authentic characters that appeal to audiences.

This dissertation is particularly relevant to the growing industry of data analytics. Researchers of this industry are concerned with applying new technologies to facilitate assessing data for predictive purposes. To this end, such researchers analyze data across a diversity of positive and negative psychological phenomenon to better understand the relationship between customers and products or services. In terms of negative psychological phenomenon, such researchers are beginning to explore how deception is used by customers, and how understanding that relationship can assist companies with improving customer relations and appropriately addressing customer claims.

This dissertation is also relevant to the industry of politics, including pollsters, speech writers, and campaign advertisers. Pollsters conduct surveys of public opinion. This dissertation demonstrated the importance of including a lie condition in the experimental design, and the use of linguistic style matching to test whether a lie condition has been satisfied. Pollsters may benefit from such an experimental design and assessment approach to better communicate the views of the public. This dissertation revealed that notions of the truth and of lies are communicated differently depending on the worldview of the communicators. Speech writers are concerned with communicating one worldview over any other worldviews. Understanding how the opposition communicates deceptive views may be beneficial to gaining the support of "swing voters" and "leaners" who tend to share some views that are aligned with the opposing worldview and who may communicate their political views according that worldview. This dissertation also revealed that issue supporters and issue opposers coordinate their language to meet specific communicative goals. Understanding the linguistic features characteristic of issue supporters and issue supporters may benefit campaign advertisers who are required to use short sound bites to communicate the advantages of their candidates and the disadvantages of the opposing candidates.

In sum, although deception is common in daily discourse, a diversity of fields and industries remain interested in contemporary approaches to deception analysis. This dissertation revealed how linguistic style matching and distal language analyses facilitate in assessing deceptive discourse between contrasting groups. Moreover, this study demonstrated that new technologies (e.g., Amazon Mechanical Turk, Google Docs, and the Gramulator) contribute to the advancement of field-specific and industry-specific deception research.

Limitations

Although this dissertation produced some significant and potentially important findings, there were limitations. Of particular concern are the limitations that pertain to the internal validation process (IVP) and the number of political issues under consideration. I present in this section the limitations to be addressed in future analyses.

The limitation of the internal validation process (IVP) is that although such a process contributes to providing confidence in the constructs under consideration, the IVP needs to be correlated with a number of other sources of evidence of validity to provide greater confidence in those constructs (Miller, McIntire, & Lovler, 2011). In other words, the IVP tests for the homogeneity and the markedness of the data, which signal any inconsistencies in the data representing the constructs under consideration (e.g., TRUTH and LIE; FOR_IT and AGAINST_IT). As such, the IVP constitutes one set of measures of a construct and may be considered as one source of evidence of validity. While the IVP represents a reasonable point of departure, future validation assessments need to include the IVP along with other measures that belong to a *nomological network* , or a system of interrelated tests (Cronbach & Meehl, 1955). For example, a confirmatory factor analysis (Jöreskog, 1969) can be used to examine construct validation. Thus, future analyses need to include statistical tests that facilitate assessing the degree to which measures of the nomological network are theoretically related.

The limitation of the number of political issues under consideration is that there were just two of them: the liberal issue of stricter gun control laws in the United States and the conservative issue of stricter immigration control laws in the United States. Consequently, comparative assessments between liberal deceivers and conservative deceivers were based on a 2:1 ratio. A 4:1 ratio would provide greater confidence in any conclusions drawn from the results. To allow for such a ratio, two other issues need to be considered, one being a liberal issue and the other being a conservative issue. For example, the liberal issue of legalizing same sex marriage in the United States and the conservative issue of imposing the death penalty in capital murder cases in the United States are political issues for future consideration. Thus, future analyses need to better assess the language of deceivers and of truth tellers across four political issues.

Future Research

In addition to the analyses addressed above, future research needs to also consider the supplemental data collected from survey participants that was not analyzed in this study. This supplemental data is concerned with three social factors: *social distance, social demographics,* and *media influences*. Thus, the foremost purpose of future research is to extend the deception analysis to include such social factors that may offer greater insight into the language and worldviews of deceivers and truth tellers, and of liberals and conservatives.

The first social factor is concerned with the social distance between the agents of the discourse and the primary figures of that discourse. More specifically, although the distal language analysis of this study was concerned with agents' explicit choice of primary figures in their political discourse, such analysis did not consider the supplemental data suggestive of the social distance between those agents and the primary figures of their political discourse (see Appendix A). This supplemental data comprised participants' responses to a series of questions using a 6-point Likert scale. Such responses included self-ratings of the relationship between the agents and the primary figures of their political discourse (e.g., 1 = Very close relationship and 6 = Very distant relationship), the time in which described events took place (e.g., 1 = Very recent past and 6 = Very distant past), and the location of such events (e.g., 1 = In my hometown and 6 = Extremely far from my hometown). Thus, future research on the explicit choice of primary figures needs to include the quantitative data representing relational, temporal, and spatial distance.

The second social factor is concerned with the social demographics of participants (see Appendix A). More specifically, although this study was concerned with the difference between the language and worldviews of liberals and conservatives, this study did not address the differences in the age, sex, race, geographic location, and voter registration status of such political groups. However, participants of this study did submit this demographic information. Thus, future research needs to analyze the prominent linguistic features of deceivers and truth tellers and of liberals and conservatives across the given social demographic categories.

The third social factor is concerned with media influences. More specifically, participants identified their preferred television news outlet from a list of options, which included MSNBC, CNN, and FOX (see Appendix A). Participants also specified other preferred television news outlets that were not on this list (e.g., PBS; ABC News). Future research needs to assess whether there is a correlation between liberals and specific television news outlets, and between conservatives and specific television news outlets.

Conclusion

The linguistic style matching analysis and distal language analysis conducted in this study of deception in political discourse revealed three main findings. First, deceivers and truth tellers from contrasting political groups who share rhetorical goals are more likely to coordinate their language. The corollary of this first finding is that linguistic style matching analysis may be used to test whether a lie condition has been satisfied. Second, conservative liars are potentially better than liberal liars at matching the linguistic styles of their corresponding truth tellers. The corollary of this second finding is that linguistic style matching between such political groups may be a phenomenon independent of issue ownership. Third, deceivers are more likely than truth tellers to choose primary figures other than themselves as the subject of their political discourse and to use a greater proportion of a diversity of pronouns and pronominals. The corollary of this third finding is that distal language analysis may be used to distinguish between deceivers' and truth tellers' explicit and implicit lexical choices. Taken as a whole, this study revealed prominent features of deception in political discourse while challenging traditional computational approaches to deception analysis. That is, traditional computational approaches tend to use standard word counting calculators (e.g., LIWC, Coh-Metrix) that measure abstract lexical items (e.g., pronouns). Such measures may lack the sophistication necessary to understand contextualized language. Instead, this study used the Gramulator to consider more carefully the contextualized language of deceptive discourse. Thus, this study demonstrated that the Gramulator is an effective computational textual analysis tool for analyzing deception in political discourse while emphasizing the importance of a combination of methods and new technologies to facilitate deception analysis.

References

Allen, C. S., & Bryant, B. R. (1996). Learning a user's linguistic style: using an adaptive parser to automatically customize a unification-based natural language grammar. *Proceedings of UM-96, The Fifth International Conference on User Modeling,* Kailua-Kona, Hawaii, 35-42.

Arrivé, M. (1992). *Linguistics and Psychoanalysis*. Amsterdam, The Netherlands: John Benjamins Publishing Co.

Battistella, E. L. (1990). *Markedness: The evaluative superstructure of language*. Albany, NY: State University of New York Press.

Berk, L. M. (1999). *English Syntax: From Word to Discourse.* New York, NY: Oxford University Press.

Bernieri, F., Davis, J., Rosenthal, R., & Knee, C. R. (1994). Interactional synchrony and rapport: Measuring synchrony in displays devoid of sound and facial affect. *Personality and Social Psychology Bulletin*, *20,* 303-311.

Bernieri, F., & Rosenthal, R. (1991). Interpersonal coordination: Behavior matching and interactional synchrony. In R. Feldman & B. Rime (Eds.), Fundamentals of nonverbal behavior. *Studies in emotion and social interaction,* (pp. 401-432). New York, NY: Cambridge University Press.

Bhatia, V. K. (2007). Interdiscursivity in critical genre analysis. *Proceedings from the 4th International Symposium on Genre Studies (SIGET)*. Tubarão, Santa Catarina, Brazil.

Biber, D. (1993). Representativeness in Corpus Design. *Literary and Linguistic Computing*, *8*(4), 243-257.

Boonthum-Denecke, C., McCarthy, P. M., & Lamkin,T. A. (Eds.) (2012). *Cross-disciplinary advances in Applied Natural Language Processing: Issues and Approaches*. Hershey, PA: IGI Global.

Brace, N., Kemp, R., & Snelgar, R. (2006). *SPSS for Psychologists: A guide to data analysis using SPSS for Windows*. Mahwah, NJ: Lawrence Erlbaum Associates, Publishers.

Bradac, J. (1990). Language attitudes and impression formation. In H. Giles & W. Robinson (Eds.), *Handbook of language and social psychology* (pp. 387-412). Chichester, UK: Wiley.

Bradac, J., & Mulac, A. (1984). A molecular view of powerful and powerless speech styles. *Communication Monographs, 51,* 307-319.

Briggs, B. (2012). *Determining intent: A quantitative and qualitative linguistic analysis of holographic and professional wills*. (Doctoral dissertation, The University of Memphis). Retrieved from https://umwa.memphis.edu/etd/.

Broadhurst, R. G., & Cheng, K. H. W. (2005). The detection of deception: The effects of first and second language on deception detection ability. *Journal of Psychiatry, Psychology and Law*, *12,* 107-118.

Brown, P., & Levinson, S. C. (1987). *Politeness: Some universals in language usage*. Cambridge, England: Cambridge University Press.

Buhrmester, M. D., Kwang, T., & Gosling, S.D. (2011). Amazon's Mechanical Turk: A new source of inexpensive, yet high-quality, data?. *Perspectives in Psychological Science*, *6,* 3-5.

Buller, D. B., & Burgoon, J. K. (1996). Interpersonal Deception Theory. *Communication Theory*, *6*(3), 2003-242.

Burgoon, J. K., Buller, D. B., White, C. H., Afifi, W. A., & Buslig, A. L. S. (1999). The role of conversational involvement in deceptive interpersonal communication. *Personality and Social Psychology Bulletin*, *25,* 669-685.

Burgoon, J. K., & Qin, T. (2006). The dynamic nature of deceptive verbal communication. *Journal of Language and Social Psychology*, 25(1), 76-96.

Cameron, W. B. (1963). *Informal Sociology: A Casual Introduction to Sociological Thinking,* (p. 13). New York, NY: Random House.

Carlson, J., George, J., Burgoon, J., Adkins, M., & White, C. (2004). Deception in Computer-Mediated Communication. *Group Decision and Negotiation*, *13,* 5-28.

Chilton, P. (2004). *Analysing Political Discourse: Theory and Practice*. New York, NY: Routledge.

Clark, H. H. (1996). *Using Language*. Cambridge, UK: Cambridge University Press.

Clark, H. H., & Brennan, S. A. (1991). Grounding in communication. In L. B. Resnick, J. M. Levine, & S. D. Teasley (Eds.). *Perspectives on socially shared cognition*. Washington: APA Books.

Cronbach, L. J, & Meehl, P. E. (1955). Construct validity in psychological tests. *Psychological Bulletin*, *52*(4), 281-302.

Crossley, S. A., Greenfield, J., & McNamara, D. S. (2008). Assessing text readability using cognitive based indices. *TESOL Quarterly*, *42*(3), 475-493.

Croswell, J. W. (2003*). Research Design: Qualitative, Quantitative, and Mixed Method Approaches (Second Edition)*. Thousand Oaks, CA: Sage.

Danescu-Niculescu-Mizil, C., Gamon, M., & Dumais, S. (2011). Mark my words! Linguistic style accommodation in social media. In *Proceedings of the International World Wide Web Conference, Hyderabad, India* (pp. 745-754).

Denzin, N. K., & Lincoln, Y. S. (Eds.) (2000). *Handbook of qualitative research: Second edition*. Thousand Oaks, CA: Sage Publications, Inc.

DePaulo, B. M., Kashy, D. A., Kirkendol, S. E., Wyer, M. M., & Epstein, J. A. (1996). Lying in everyday life. *Journal of Personality and Social Psychology*, *70*(5), 979-995.

DePaulo, B. M., Lindsay, J. J., Malone, B. E., Muhlenbruck, L., Charlton, K., & Cooper, H. (2003). Cues to deception. *Psychological Bulletin*, 129, 74-112.

Duff, P. A. (2002). Research Approaches in Applied Linguistics. In R. B. Kaplan (Ed.), *The Oxford Handbook of Applied Linguistics.* New York, NY: Oxford University Press.

Duran, N. D., Crossley, S. A., Hall, C., McCarthy, P. M., & McNamara, D. S. (2009). Expanding a catalogue of deceptive linguistic features with NLP technologies. *Proceedings of the 22nd International Florida Artificial Intelligence Research Society (FLAIRS) Conference,* Sanibel Island, FL. AAAI Press.

Duran, N. D., Dale, R., & McNamara, D. S. (2008). Does the arm know when you're lying? [Abstract]. *In the Proceedings of the 49th Annual Meeting of the Psychonomic Society* (p. 22). Austin, TX: Psychonomic Society Publications.

Duran, N. D., Hall, C., & McCarthy, P. M. (2009). Using statistically improbable n-gram features to reveal the thematic content of deception. *Paper presented at the 39th Annual Meeting of the Society for Computers in Psychology*, Boston, MA.

Duran, N. D., Hall, C., McCarthy, P. M., & McNamara, D. (2010). The linguistic correlates of conversational deception: Comparing natural language processing technologies. *Applied Psycholinguistics*, *31,* 439-462.

Ekman, P., & Friesen, W. V. (1974). Detecting deception from the body or face. *Journal of Personality and Social Psychology*, *29*(3), 288-298.

Feldman, J. (2007). *Framing the debate: Famous presidential speeches and how Progressives can use them to change the conversation (and win elections).* Brooklyn, NY: IG Publishing.

Firth, J. R. (1957). A synopsis of linguistic theory, 1930-55. Studies in linguistic analysis, pages 1-32. The Philological Society, Oxford. [Reprinted in Palmer, F. R. (Ed.) (1968) *Selected Papers of J. R. Firth 1952-59*, 168-205. Longmans, London.]

Frank, M. G., & Ekman, P. (1997). The ability to detect deceit generalizes across different types of high-stake lies. *Journal of Personality and Social Psychology*, *72,* 1429-1439.

Fowler, C. A., Richardson, M., Marsh, K., & Shockley, K. (2008). Language use, coordination and the emergence of cooperative action. In A. Fuchs & V. Jirsa (Eds.), *Coordination: Neural, Behavioral and Social Dynamics* (pp. 261-280). New York, NY: Springer.

Gervasi, V., & Ambriola, V. (2002). Quantitative assessment of textual complexity. In L. Merlini Barbaresi (Ed.), *Complexity in Language and Text* (pp.197-228). PLUS-University of Pisa, Pisa, Italy, 197-228.

Giles, H., Coupland, J., & Coupland, N. (Eds.). (1991). *The contexts of accommodation*. New York, NY: Cambridge University Press.

Giles, H., & Smith, P. (1979). Accommodation Theory: Optimal Levels of Convergence. In H. Giles & H. St. Clair (Eds.), Language and social psychology (pp. 45-65). Oxford, UK: Blackwell.

Givon, T. (1991). Markedness in grammar: distributional, communicative and cognitive correlates of syntactic structure. *Studies in Language*, 15, 2, 335-370.

Goffman, E. (1959). The presentation of self in everyday life. Garden City, NY: Doubleday, Anchor Books.

Gonzales, A. L., Hancock, J. T., & Pennebaker, J. W. (2010). Linguistic style matching as a predictor of social dynamics in small groups. *Communication Research*, 37(1), 3-19.

Gordon, N. J., & Fleisher, W. L. (2006). *Effective interviewing and interrogation techniques*. Burlington, MA: Academic Press.

Graesser, A. C., McNamara, D.S., Louwerse, M. M., & Cai, Z. (2004). Coh-Metrix: Analysis of text on cohesion and language. *Behavior Research Methods, Instruments, and Computers*, *36,* 193-202.

Granger, S., & Petch-Tyson, S. (Eds). (2003). *Extending the scope of corpus-based Research: New applications, new challenges*. Amsterdam: Rodopi.

Gulliksen, H. (1950). Instrinsic Validity. Adapted from *History of, and present trends in testing* (pp. 1-20). *The Sixth Yearbook of the National Council on Measurements Used in Education, 1948-1949.* National Council on Measurements Used in Education.

Haertl, B., & McCarthy, P.M. (2011). Differential Linguistic Features in U.S. Immigration Newspaper Articles: A Contrastive Corpus Analysis Using the

Gramulator. In C. Murray & P. M. McCarthy (Eds.). *Proceedings of the 24rd International Florida Artificial Intelligence Research Society Conference*. (pp. 349-350). Menlo Park, CA: The AAAI Press.

Hagège, C. (2005). On the part played by human conscious choice in language structure and language evolution. In Z. Frajzyngier, A. Hodges, & D. S. Rood (Eds.), *Linguistic Diversity and Language Theories*, (pp. 105-118). Amsterdam, The Netherlands: John Benjamins Publishing Co.

Hancock, J. T., Curry, L. E., Goorha, S., & Woodworth, M.T. (2008). On lying and being lied to: A linguistic analysis of deception. *Discourse Processes*, *45,* 1-23.

Hancock, J. T., Thom-Santelli, J., & Ritchie, T. (2004). Deception and design: The impact of communication technologies on lying behavior. *Proceedings of the Conference on Computer Human Interaction*, *6,* 130-136. New York, NY: ACM.

Harris Interactive. (2010). *The Harris Poll #17*. http://www.harrisinteractive.com.

Hayworth, J. D. (2005). *Enforcement First Immigration Reform Act: H.R. 3938*. 109th Congress. http://www.govtrack.us/congress/bills/109/hr3938

Hogg, M. A., & Vaughan, G. M. (2002). Social Psychology: Third Edition. London: Prentice-Hall.

Hullender A., & McCarthy, P.M. (2011). A Contrastive Corpus Analysis of Modern Art Criticism and Photography Criticism. In C. Murray & P.M. McCarthy (Eds.), *Proceedings of the 24rd International Florida Artificial Intelligence Research Society Conference*. (pp. 350-351). Menlo Park, CA: The AAAI Press.

Jarvis, S. (2002). Topic continuity in L2 English article use. *Studies in Second Language Acquisition,* 24, 387-418.

Jarvis, S., & Daller, M. (Eds.) (in press). *Vocabulary knowledge: Human ratings and automated measures*. Amsterdam: John Benjamins.

Johnson, M. K., & Raye, C. L. (1981). Reality monitoring. *Psychological Review*, *88*(1), 67-85.

Johnstone, B. (2002). *Discourse analysis*. Malden, MA: Blackwell Publishers Inc.

Jöreskog, K. G. (1960). A general approach to confirmatory maximum likelihood factor analysis. *Psychometrika*, *34*(2), 183-202

Joseph, J. E. (2006). *Language and Politics*. Edinburgh University Press, Ltd.

Kelvin, W. T. (1833). Electrical Units of Measurement. *Popular Lectures*, 1, pg. 73; quoted in *Encyclopaedia of Occupational Health and Safety* (1998).

Lafferty, J. C., & Eady, P. M. (1974). *The Desert Survival Problem*. Plymouth, Michigan: Experimental Learning Methods.

Lakoff, G. (2002). *Moral Politics: How liberals and conservatives think (2nd ed.)*. Chicago, IL: University of Chicago Press.

Lakoff, G. (2004). *Don't think of an elephant: Know your values and frame the debate-The essential guide for progressives.* White River Junction, VT: Chelsea Green Publishing Company.

Lamkin, T.A., & McCarthy, P.M. (2012).The Hierarchy of Detective Fiction. In C. Murray & P. M. McCarthy (Eds.), *Proceedings of the 24rd International Florida Artificial Intelligence Research Society Conference* (pp. 257-262). Menlo Park, CA: The AAAI Press.

Landauer, T. K., & Dumais, S. T. (1997). A solution to Plato's problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review, 2,* 211-240.

Langacker, R. W. (2000). *Grammar and Conceptualization*. Berlin: Mouton de Gruyter.

Langacker, R. W. (2001). Discourse in cognitive grammar. *Cognitive Linguistics*, *12,* 143-188.

Leal, S., Vrij, A., Mann, S., & Fisher, R. P. (2010). Detecting truth and false opinions: The Devil's Advocate approach as a lie detection aid. *Acta Psychologica*, *134*(3), 323-329.

Levinson, S. C. (2006). Deixis. In L.R. Horn & G. Ward (Eds.), *The Handbook of Pragmatics*, (pp. 97-121). Malden, MA: Blackwell Publishing Ltd.

McCarthy, P. M., Booker, L. M., Guess, R., & Crabtree, L. (2012). The Manualator. A Manual for the Gramulator. http://tinyurl.com/84575mf. Retrieved 03/30/2012.

McCarthy, P. M., & Boonthum-Denecke, C. (Eds.).(2012). Introduction. *Applied Natural Language Processing: Identification, Investigation, Resolution.* Hershey, PA: IGI Global.

McCarthy, P. M., Duran, N. D., & Booker, L. M. (2012). The Devil is in the Details: New Directions in Deception Analysis. In P. M. McCarthy & M. Youngblood (Eds.), *Proceedings of the 25th International Florida Artificial Intelligence Research Society Conference.* Menlo Park, CA: The AAAI Press.

McCarthy, P. M., Hall, C., Duran N. D., Doiuchi, M., Duncan, B., Fujiwara, Y., & McNamara, D. S., (in press). A computational analysis of journal abstracts written by Japanese, American, and British scientists. *The ESPecialist.*

McCarthy, P. M., & Jarvis, S. (2007). A theoretical and empirical evaluation of *vocd*. *Language Testing*, *24,* 459-488.

McCarthy, P. M., & Jarvis, S. (2010). MTLD, vocd-D, and HD-D: A validation study of sophisticated approaches to Lexical Diversity Assessment. *Behavior Research Methods, 42,* 381-392.

McCarthy, P. M., & Jarvis, S. (in press). From intrinsic to extrinsic issues of lexical diversity assessment: An ecological validation study. In S. Jarvis & M. Daller (Eds.), *Vocabulary knowledge: Human ratings and automated measures*. Amsterdam: Benjamins.

McCarthy, P. M., Myers, J. C., Briner, S. W., Graesser, A. C., & McNamara, D. S. (2009). A psychological and computational study of sub-sentential genre recognition. *Journal for Language Technology and Computational Linguistics*, *24,* 23-55.

McCarthy, P. M., Watanabe, S. & Lamkin, T. (2012). The Gramulator. In P.M. McCarthy and C. Boonthum-Denecke (Eds.), *Applied Natural Language Processing: Identification, Investigation, Resolution.* Hershey, PA: IGI Global.

McDowall, J. J. (1978). Interactional synchrony: A reappraisal. *Journal of Personality and Social Psychology*, *36,* 963-975.

McNamara, D. S., Graesser, A. C., McCarthy, P. M., & Cai, Z. (in press). *Coh-Metrix: Automated Evaluation of Text and Discourse*. Cambridge University Press.

McNamara, D. S., Louwerse, M. M., McCarthy, P. M., & Graesser, A. C. (2010). Coh-Metrix: Capturing linguistic features of cohesion. *Discourse Processes*, *47,* 292-330.

Meyer, P. (2011). How to spot a liar: Pamela Meyer on TED.com.

Miller, L. A., McIntire, S. A., & Lovler, R. L. (2011). *Foundations of psychological testing: A practical approach (3rd Edition)*. Thousand Oaks, CA: Sage Publications, Ltd.

Min, H. C., & McCarthy, P. M. (2010). Identifying varietals in the discourse of American and Korean scientists: A contrastive corpus analysis using the Gramulator. In H. W. Guesgen & C. Murray (Eds.), *Proceedings of the 23rd International Florida Artificial Intelligence Research Society Conference* (pp. 247-252). Menlo Park, CA: The AAAI Press.

Newman, M. L., Pennebaker, J.W., Berry, D. S., & Richards, J. M. (2003). Lying words: Predicting deception from linguistic styles. *Personality and Social Psychology Bulletin*, *29,* 665-667.

Niederhoffer, K. G., & Pennebaker, J.W. (2002). Linguistic style matching in social interaction. *Journal of Language and Social Psychology*, *21,* 337-360.

Nyhan, B., & Reifler, J. (forthcoming). The persistence of political misperceptions. *Political Behavior*.

Obama, B. H. (2009). *Remarks by the President to a joint session of Congress on healthcare.* U.S. Capitol, Washington, D.C. http://www.whitehouse.gov.

Pennebaker, J. W., Chung, C. K., Ireland, M., Gonzales, A., & Booth, R. J. (2007). *The development of the psychometric properties of LIWC2007*. Austin: www.liwc.net

Pennebaker, J. W., & Francis, M. E. (1996). Cognitive, emotional, and language processes in disclosure. *Cognition and Emotion*, *10,* 601-626.

Pennebaker, J. W., Mayne, T. J., & Francis, M. E. (1997). Linguistic predictors of adaptive bereavement. *Journal of Personality and Social Psychology,* 72, 863-871.

Petrocik, J. R. (1996). Issue ownership in presidential elections, with a 1980 case study. *American Journal of Political Science*, *40*(3), 825-850.

Petrocik, J. R., Benoit, W. L., & Hansen, G. T. (2003). Issue ownership and presidential campaigning, 1952 – 2000. *Political Science Quarterly*, *118*(4), 599-626.

Phan, K. L., Magalhaes, A., Ziemlewicz, T. J., Fitzgerald, D. A., Green, C., & Smith, W. (2005). Neural correlates of telling lies: A functional magnetic resonance imaging study at 4 Tesla. *Academic Radiology*, *12,* 164-172.

Rauh, G. (Ed). (1983). Aspects of deixis. *Essays on deixis*. Tübingen: Narr.

Rufenacht, R. M., Lamkin, T. A., & McCarthy, P. M. (2012). An Analysis of Linguistic Features of Fairy Tales and ESL Texts. In C. Boonthum-Denecke, P. M. McCarthy, & T. A. Lamkin, (Eds.), *Cross-Disciplinary Advances in Applied Natural Language Processing: Issues and Approaches*, Hershey, PA: IGI Global.

Sapir, A. (1987). Scientific Content Analysis (SCAN). Pheonix, AZ: Laboratory of Scientific Interrogation.

Schlenker, B. R., & Weigold, M. F. (1989). Goals and the self-identification process: Constructing desired identities. In L. A. Pervin (Ed.), *Goal concepts in personality and social psychology* (pp. 243-290), Hillsdale, NJ: Erlbaum.

Shuy, R. W. (1998). *The language of confession, interrogation, and deception.* Thousand Oaks, CA: Sage Productions, Inc.

Spence, S. A., Farrow, T. F. D., Herford, A. E., Wildinson, I. D., Zheng, Y., & Woodruff, P. W. R. (2001). Behavioural and functional anatomical correlates of deception in humans. *Neuroreport*, *12*(13), 2849-53.

Stanchevici, D. (in press). The rhetorical construction of social classes in the reports of Stalin's secret police. *Journal of Technical Writing and Communication*.

Stellar, M. (1989). Recent developments in statement analysis. *Credibility assessment*. Deventer, the Netherlands: Kluwer.

Stellar, M., & Köhnken, G. (1989). Criteria-based content analysis. In D. C. Raskin (Ed.), *Psychological methods in criminal investigation and evidence*. New York, NY: Springer-Verlag.

Strain, A. C., & Booker, L. M. (2012). Amazon Mechanical Turk: A web-based tool for facilitating experimental research in ANLP. In C. Boonthum-Denecke, P. M. McCarthy, & T. A. Lamkin (Eds.), *Cross-Disciplinary Advances in Applied Natural Language Processing: Issues and Approaches*. Hershey, PA: Information Science Reference (an imprint of IGI Global).

Sweller, J. (1994). Cognitive load theory, learning difficulty and instructional design. *Learning and Instruction, 4*(4), 295-312.

Tajfel, H., & Turner, J. (1986). The social identity theory of intergroup behavior. In S. Worchel and W.G. Austin (Eds.), *Psychology of Intergroup Relations*, (pp. 7-24). Chicago, IL: Nelson-Hall.

Terwilleger, B., & McCarthy, P. M. (2011). Bias in Hard News Articles from Fox News and MSNBC: An Empirical Assessment Using the Gramulator. In C. Murray & P. M. McCarthy (Eds.). *Proceedings of the 24rd International Florida Artificial*

Intelligence Research Society Conference (pp. 361-362). Menlo Park, CA: The AAAI Press.

The Pew Research Center. (2012). *Trends in American values: 1987-2012: Partisan polarization surges in Bush, Obama years*. http://www.people-press.org.

Toma, C., & Hancock, J. T. (2010). Reading between the lines: Linguistic cues to deception in online dating profiles. Paper presented at the Computer Supported Cooperative Work Conference, Savannah, GA.

Trosborg, A. (1997). Text typology: Register, genre and text type. Text typology and translation, (pp. 3-23). Amsterdam: John Benjamins.

Undeutsch, U. (1967). Beurteilung der Glaubhaftigkeit von Aussagen. *Handbuck der Psychologie*, *11,* 26-181.

Ungerer, F., & Schmid, H. (2006). *An introduction to cognitive linguistics: Second edition*. London – New York, NY: Pearson Education.

van Dijk, T. A. (1998). What is political discourse analysis? In J. Blommaert & C. Bulcaen (Eds.), *Political linguistics* (pp. 11-52). Amsterdam: John Benjamins.

Vrij, A., & Winkel, F. W. (1993). Objective and subjective indicators of deception. In N. K. Clark & G. M. Stephenson (Eds.), *Children, Evidence, and Procedure*. Leicester, UK: The British Psychological Society.

Zhou, L., Burgoon, J. K., Nunamaker, J. F., & Twitchell, D. (2004). Automating linguistic-based cues for detecting deception in text-based asynchronous computer-mediated communication. *Group Decision and Negotiation, 13,* 81-106.

Appendix A: Survey

Truth and Lies in Politics

Informed Consent

You are being asked to participate in a survey to explore how truth and lies are communicated. The results of this survey will be used for educational research purposes that include developing lie detection software. The results of this survey may also be published.

By agreeing to participate in this survey, you also agree to complete the demographics section. As a Mechanical Turk worker, your identity remains anonymous. Your participation in this study is voluntary and risk-free.

[ ] I agree to participate in this survey

[ ] I do not agree to participate in this survey

Part I: On a scale of 1 (*For*) to 6 (*Against*) with the distance between each number being equal, rate your views concerning the following issues:

Stricter gun control laws in the United States

For					**Against**
1	2	3	4	5	6

Stricter immigration control laws in the United States

For					**Against**
1	2	3	4	5	6

Part II: For the following questions, you have been randomly selected to type the TRUTH or a LIE. When you see **[TRUTH]**, you are to type information that represents what you believe to be 'the truth, the whole truth, and nothing but the truth'. When you see **[LIE]**, you are to type information that does NOT represent what you believe to be 'the truth, the whole truth, and nothing but the truth'. Try to be as convincing as possible.

1. Share your views concerning whether or not to have **stricter gun control laws in the United States**. Provide one example that involves you or another person to support your views. (Minimum requirement of 8 full sentences or 80 words) **[TRUTH]**

2. Share your views concerning whether or not to have **stricter immigration laws in the United States**. Provide one example that involves you or another person to support your views.
(Minimum requirement of 8 full sentences or 80 words)
[LIE]

Part III:
Please type the **[TRUTH]** for all of the following questions and answer all questions that apply. If you are asked to give a response based on a scale from 1 to 6, please note that the distance between each number is equal.

For questions 3 – 6 below, consider the example that you shared to support your views concerning whether or not to have **stricter gun control laws in the United States**…

3. If asked, who would you say your example is mainly about?

   [ ] Me [ ] Some other person(s)

4. If you selected "Some other person(s)" for question 3 above, please answer the following question. If not, please continue to question 5.
   Would you say that you have a close relationship with the other person(s) of your example?

Very Close Relationship					**Very Distant Relationship**
1	2	3	4	5	6

5. When would you say your example took place?

Very Recent Past					**Very Distant Past**
1	2	3	4	5	6

6. Where would you say your example took place?

In My Hometown					**Extremely Far From My Hometown**
1	2	3	4	5	6

For questions 7 – 10 below, consider the example that you shared to support your views concerning whether or not to have **stricter immigration laws in the United States**…

7. If asked, who would you say your example is mainly about?

   [ ] Me [ ] Some Other Person(s)

8. If you selected "Some Other Person(s)" for question 7 above, please answer the following question. If not, please continue to question 9.
   Would you say that you have a close relationship with the other person(s) of your example?

Very Close Relationship					**Very Distant Relationship**
1	2	3	4	5	6

9. When would you say your example took place?

Very Recent Past					**Very Distant Past**
1	2	3	4	5	6

10. Where would you say your example took place?

In My Hometown					**Extremely Far From My Hometown**
1	2	3	4	5	6

Part IV:
Please type the **[TRUTH]** for all of the following questions about you.

11. What is your age?

12. What is your sex? [ ] Female [ ] Male

13. What is your race?

14. What state do you live in?

15. Are you a registered voter in the United States?

    [ ] Yes [ ] No

16. How would you rate your political views?

Extreme Liberal	Moderate Liberal	Liberal Leaning	Conservative Leaning	Moderate Conservative	Extreme Conservative
1	2	3	4	5	6

17. Which of the following television news outlets do you watch most often?

[ ] MSNBC
[ ] CNN
[ ] FOX
[ ] Other (please specify) ____________________________________