From *Handbook of Language Analysis in Psychology*.
Edited by Morteza Dehghani and Ryan L. Boyd.

CHAPTER 13

Lies and Language

A Context-Contingent Approach to Verbal Cues of Deceit

David M. Markowitz and Jeffrey T. Hancock

INTRODUCTION

How can you catch a liar? Ask a passerby, and he or she will likely say that eye gaze aversion (Akehurst, Köhnken, Vrij, & Bull, 1996), fidgeting (Zuckerman, Koestner, & Driver, 1981), or stuttering reveal dishonesty (Global Deception Research Team, 2006). Ask a deception researcher, and he or she will tell you that most cues are systematically unreliable indicators of deception (DePaulo et al., 2003; Luke, 2019). Why are cue theories problematic for deception, and why do behavior patterns—from verbal to nonverbal—often fail to replicate?

Recent evidence suggests that most cues are not universal, because the context affects how lies and truths are communicated. To better understand the complexities of a deception and why people choose to lie versus tell the truth, a comprehensive view of the deception context is needed and may help to understand mixed empirical findings (Blair, Levine, & Shaw, 2010; Levine, 2018; Markowitz & Griffin, 2020; Markowitz & Hancock, 2019; Nahari et al., 2019). A recent meta-analysis that supports this claim revealed context-related moderators (e.g., the production mode, motivation) affect the relationship between deception and language (Hauch, Blandón-Gitlin, Masip, & Sporer, 2015). For example, liars often use fewer words than truth tellers, though this effect changes at different interaction levels. Liars often use more words than do truth tellers when communicating via computer-mediated channels, but they use fewer words when the interaction is an interview, a person-to-person interaction, or without an interaction.

If aspects of the context moderate how deception affects word patterns, they should be modeled in our conceptualization of how deception modifies verbal patterns. Recent deception theories have answered this call and outlined three contextual factors that relate to how deception affects communication behavior, specifically, language. We, Markowitz and Hancock (2019), proposed the Contextual Organization of Language and Deception (COLD) framework, which highlights how three factors about deception that are context-dependent—*psychological dynamics, pragmatic goals,* and *genre conventions*—each can influence how deception affects language.

Psychological dynamics refers to the emotional and cognitive aspects of a deception that may matter given that evidence suggests liars and truth tellers have different psycho-

logical experiences. For example, in some cases, liars may overuse emotion in their speech when talking about abortion or their friends relative to truth tellers (Markowitz & Griffin, 2020; Newman, Pennebaker, Berry, & Richards, 2003). In other cases, liars cannot approximate a genuine emotional experience and their language is less emotional than that of truth tellers (Burns & Moffitt, 2014). Similarly, lying can be a cognitively taxing task in some cases (Blandón-Gitlin, Fenn, Masip, & Yoo, 2014; Newman et al., 2003) but not others (Hancock, Curry, Goorha, & Woodworth, 2007; Schober & Glick, 2011). Emotional and cognitive responses are foundational to understand a deception (Ekman, 2001; Hauch et al., 2015). The COLD framework argues that applying a universal model of language across deceptions fails to account for the psychological nuances of each setting.

The second aspect of the COLD model argues that pragmatic goals shape how people communicate their lies. People lie for a reason and when the truth is problematic, not just because they can (Ariely, 2012; Levine, 2020; Levine, Kim, & Hamel, 2010; Markowitz & Levine, 2021). Early, yet still relevant research by Turner, Edgley, and Olmstead (1975) suggests that people lie to save face, to prevent embarrassment, among other reasons, and deception goals can span cultures (Levine, Ali, Dean, Abdulla, & Garcia-Ruano, 2016). Linguistic indicators of deception can also be affected by pragmatic goals as we report in an evaluation of lies by American presidents (Markowitz & Hancock, 2019). Presidents who lied about policy (e.g., the Bush Administration and weapons of mass destruction, Lyndon B. Johnson and the Gulf of Tonkin incident) had a different linguistic profile than presidents who lied to safe face and prevent embarrassment after a scandal (e.g., Bill Clinton and the Monica Lewinsky scandal, Richard Nixon and Watergate). Policy lies often contain a reduced rate of self-references, which follows cue-based theories that suggest liars often experience a psychological distancing effect compared to truth tellers (Markowitz & Griffin, 2020; Newman et al., 2003). However, lies to prevent embarrassment featured an amplified rate of self-references compared to truths, perhaps out of a need to obfuscate one's own behaviors and motives. These data suggest that a one-size-fits-all approach to lies and language may miss the fact that deceptive language is in part a function of what a liar is trying to accomplish.

The final dimension of the COLD model suggests that genre conventions affect how people communicate. Drawing on linguistics research from Biber, Connor, and Upton (2007), the genre is defined as "message-related conventions reflected in communication behavior that contain an 'internal structure' of a discourse community" (Markowitz & Griffin, 2020, p. 292). Since people belong to multiple communities and communities have distinct conventions, it is reasonable to assume that certain language patterns may matter or be conventional for some communities and not for others. For example, self-references are often reliable indicators of deception in some settings (Hauch et al., 2015; Markowitz & Griffin, 2020; Newman et al., 2003), but they are uncommon in science writing. Since pronouns are unconventional in academic papers, they should not be diagnostic of false speech in cases of science fraud, because the base rates of pronouns are too low (Markowitz & Hancock, 2014, 2016).

The COLD model is a theory-based attempt both to reconcile and rationalize mixed findings from the deception and language literature. We focus the remainder of this chapter on four areas of deception research that have received substantial treatment from a language perspective, particularly studies that use automated text analysis to understand word patterns of deceit. We discuss this work in relation to the COLD model, positioning why mixed effects may occur across studies and how the same language markers (e.g., self-references, adjectives, other details) are affected by contextual constraints. These areas highlight how deception plays a crucial role in online and offline life but manifests in context-contingent ways.

It is important to note that most researchers care about deception to the degree that it can be detected with above-chance accuracy using machine learning and natural language processing techniques. As social scientists, we take a different but complemen-

tary approach. We primarily focus on the language features that reveal deception and how they might indicate an important social or psychological process about the communicator. Detection accuracy is important to the degree that such features provide a reliable signal about what people are thinking, feeling, and experiencing psychologically. The features and their social and psychological meaning are paramount in our work; accuracy modeled with opaque feature sets is less of an interest.

A PRIMER ON CONTENT AND STYLE

Before we begin our review, it is important to identify classes of language that may reflect lies and truths. Language patterns can be segmented into two broad classes: (1) content words, and (2) style words (Boyd & Pennebaker, 2015; Chung & Pennebaker, 2007). Content words include nouns, verbs, and other descriptors that indicate the subjects of a sentence. If people were asked to recall the essence of a sentence, they would recall nouns and verbs. For example, in the sentence "Why did the chicken cross the road?" people would likely recall "chicken," "road," and "cross," since these words describe *what* people are communicating. Content words are important for revealing thematic differences across texts (Blackburn, Yilmaz, & Boyd, 2018; Markowitz & Hancock, 2014) and can describe what information people tend to focus on in a disclosure (Guillory, Hancock, Woodruff, & Keilman, 2015).

Style words, on the other hand, describe *how* a person is communicating, which is also different from syntax that describes the placement or position of words in sentence. Words such as articles (e.g., "a," "the"), prepositions (e.g., "above," "below"), and pronouns (e.g., "I," "our") are also called *function words*, and they account for more than half of our communication output. English speech contains less than 500 function words (Baayen, Piepenbrock, & Bulickers, 1995; Chung & Pennebaker, 2007; Tausczik & Pennebaker, 2010), but they are pervasive and serve as the connective tissue of conversation and writing. Function words are often associated with social and psychological processes, including interpersonal attraction (e.g., the more that partners match on function words, the more they have romantic interest; Ireland et al., 2011), social status (e.g., low-status people tend to use more self-references than do high-status people; Kacewicz, Pennebaker, Davis, Jeon, & Graesser, 2014; Markowitz, 2018), personality (e.g., extraversion is associated with high rates of collective references; Ireland & Mehl, 2014), and intelligence (e.g., in college admissions essays, the rate of articles and prepositions relative to storytelling words correlates with higher academic performance; Pennebaker, Chung, Frazee, Lavergne, & Beaver, 2014).

Both content and style words matter for revealing deception, and meta-analytic work by Hauch and colleagues (2015) provides an important list of linguistic cues that betray deceit. For example, deceptive discourse often contains more negative affect (content), fewer descriptions of cognitive processes (content), and fewer first-person pronouns (style) than does truthful discourse. Recall, however, the effects were often moderated by context-dependent characteristics, such as the event type (e.g., a first-person experience, attitude paradigm), the emotional valence of the situation (e.g., negative, neutral), the interaction level of the discourse (e.g., no interaction, an interview, a computer-mediated conversation, a face-to-face conversation), sender motivation (e.g., no motivation to lie, low to medium motivation, or high motivation), and the production mode (e.g., handwritten, typed, or spoken text). Therefore, such contextual dependencies must be considered in evaluations of deception and language. The following cases describe the need for a context-dependent view of deception and language, particularly the impact of deception on style and content as influenced by psychological dynamics, pragmatic goals, and genre conventions (Markowitz & Hancock, 2019).

ONLINE REVIEWS

Nearly 90% of consumers use online reviews before purchasing a product or visiting a business (Salen, 2015). Online user-generated content is therefore a persuasive source of information that people use to gain opin-

ions about products or services before making a final consumer decision (Margolin & Markowitz, 2018; Walther & Jang, 2012). A nontrivial concern for online consumers is the number of reviews that may be fake or paid for (Dwoskin & Timberg, 2018), which reduces the trust people may have in online reviews and questions the authenticity of review systems in general (e.g., whether those who are reviewing actually bought, used, or experienced the product).

Early research by Ott, Choi, Cardie, and Hancock (2011) explored how fake reviews are communicated linguistically. Ott and colleagues recruited participants to write fake hotel reviews and compared them to genuine reviews curated from Tripadvisor. Several patterns that emerged were consistent with traditional deception research (Johnson & Raye, 1981; Vrij et al., 2009), including the result that writers of fake reviews failed to approximate the level of details in their text compared to truthful reviews. Fake reviews contained fewer spatial terms than truthful reviews, which is reasonable, because writers of fake reviews never stayed at the hotel they were describing. Fake reviewers, therefore, would have few spatial reference points to use in their text (e.g., the distance between the lobby and the elevators or how wide the hallways were), but truthful reviewers would have this richness in detail because they stayed at the hotel.

Not all language cues were consistent with cue-based models from prior work, however. The data from Ott and colleagues (2011) suggest that self-references were indicative of false relative to truthful speech, which conflicted with the psychological distancing hypothesis by Newman et al. (2003) and others (ten Brinke & Porter, 2012). Considering the pragmatic goals of the liar may help to reconcile this finding. Fake reviewers were told to "to write a fake review (as if they were a customer)" and it must sound "realistic" while portraying the "hotel in a positive light" (Ott et al., 2011, p. 311). Prior research suggests that people can amplify their credibility by asserting a personal tone and taking an extreme stance (Margolin & Markowitz, 2018). Therefore, in cases where people are trying to appear credible and lying to accomplish this goal, they may overuse credibility markers such as self-references, especially given the instructions to write as if they had actually been there. This potentially inconsistent finding compared with prior work (Markowitz & Griffin, 2020; Newman et al., 2003; ten Brinke & Porter, 2012), when viewed from a context-contingent perspective via the COLD model makes sense after accounting for the pragmatic goals of the speaker.

SCIENCE FRAUD

Science, like most institutions, is based on trust (LaFollette, 1992). After publication, people largely believe that the information they read in journals is genuine, not fabricated, and the data are reliable (e.g., Funk & Kennedy, 2019). There are cases, however, when researchers break this trust, engage in fraud, and publish science that contains misconduct of some kind. The number of papers retracted due to fraud has risen swiftly (Fang, Steen, & Casadevall, 2012; Lu, Jin, Uzzi, & Jones, 2013), calling for additional work into the behavioral traces that may reveal scientific deception and identify fraudulent from genuine research.

A case study of Diederik Stapel, a former social psychologist in the Netherlands, was the first investigation to understand the linguistic traces of fraud in academic papers (Markowitz & Hancock, 2014). Stapel was a prolific researcher who committed over 50 acts of data fraud and whose publications were thoroughly vetted for veracity by independent researchers (Levelt, Drenth, & Noort, 2012). A linguistic evaluation of Stapel's papers compared his fraudulent and genuine first-authored publications and revealed deception-related and science-related markers of fraud. Stapel's fraudulent writing contained a reduced rate of adjectives and descriptive details compared to his genuine writing, a pattern consistent with most deception and language theory and empirical evidence (Johnson & Raye, 1981; Ott et al., 2011). However, Stapel's language patterns were also modified by goals to make his fraudulent science appear credible and rigorous. In his fraudulent papers, compared to his genuine papers, Stapel also overused science-related terms (e.g., words related to methods, investigation, and certainty).

This evidence suggests that when liars try to appear reliable, honest, and thorough, they may fail to approximate the number of genre-specific terms in the speech relative to truth tellers as a reflection of their deception goals.

Investigations of language patterns in science fraud (Markowitz & Hancock, 2014, 2016) also highlight the importance of genre conventions when evaluating whether certain markers are important for deception or not. Prior work suggests that self-references and personal pronouns, in general, are often viewed as unconventional in science (Hyland, 2001, 2003) because science papers should be objectively written and unbiased. Pronouns are therefore unlikely markers of deception in science fraud, because they are used infrequently and are non-normative.

ONLINE PERSONALS: RÉSUMÉS AND DATING PROFILES

People create online profiles to date (Ellison, Hancock, & Toma, 2012; Markowitz, Hancock, & Tong, 2018), receive an online loan (Larrimore, Jiang, Larrimore, Markowitz, & Gorski, 2011; Markowitz & Shulman, 2021), and connect with others in social networks (Back et al., 2010). Deception can play an important role in the creation of online profiles, because people have control over the information that is publicly presented and may want to develop enhanced perceptions of the self. Many types of impression management are therefore crucial when people create profiles in professional (e.g., trying to get a job) and personal online settings (e.g., trying to get a date).

An early study evaluated the frequency of deception in online LinkedIn résumés (public and private profiles) compared to traditional paper résumés (Guillory & Hancock, 2012). Lying rates were consistent across settings, though when participants coded their résumés for the types of lies (e.g., lies related to responsibility, ability, involvement, or interests), lying rates differed as a function of the résumé's publicness. People lied less about responsibilities in their public LinkedIn résumés compared to a traditional résumé or private LinkedIn résumé, presumably because other people in their professional online network could validate whether they had a specific job or duty while employed. People lied more on their public LinkedIn résumés when characteristics were unverifiable, such as personal interests. For example, in a limited-cue online professional environment such as LinkedIn, it is difficult to judge whether someone is interested in travel, photography, or learning a new language. Therefore, people often lied about content that was less verifiable in online settings compared to offline settings, and this type of impression management is a purposeful strategy to appear interesting, likable, and competent (see also, Markowitz & Hancock, 2018).

People expect that some deception exists in settings where people try to appear attractive and likable, such as online dating (Drouin, Miller, Wehle, & Hernandez, 2016). Prior work suggests that men tend to overstate their height and women tend to understate their weight in dating profiles (see Markowitz et al., 2018), but core characteristics of each person are often truthful (e.g., the size of his or her family, whether he or she is divorced; Ellison et al., 2012). When deception occurs, how is it reflected in the language patterns of the profile?

Lies in the profile (e.g., one's online height, weight, and age relative to offline height, weight, and age) tend to correlate with fewer self-references in the profile text (Toma & Hancock, 2012). This pattern of self-references is consistent with traditional deception research on verbal cues (Hauch et al., 2015; Markowitz & Griffin, 2020; Newman et al., 2003). Daters also used fewer negative emotion words, however, contrary to prior work suggesting that liars leak more negative emotion when they lie (Ekman, 2001). Considering the role of context and pragmatic goals helps to clarify the relationship among deception, language, and online dating profiles. Daters who lie in their profile may use less negative affect in their "About Me" section to appear dateable or to amplify their attractiveness. This self-presentation strategy is reasonable, because online daters have access to more people than do traditional daters (Finkel, Eastwick, Karney, Reis, & Sprecher, 2012); the cost of a negative first

impression may be much higher when another person is just one swipe away.

HIGH-STAKES AND LOW-STAKES LIES

Deception researchers are often concerned with *stakes,* or the degree that lies may cause "serious harm to targets of deceit and adverse consequences to deceivers if caught" (Burgoon et al., 2016, p. 124). Evidence suggests that high-stakes lies may contain some language signals that are consistent across settings as well. For example, Markowitz and Hancock (2016) evaluated the rate of linguistic obfuscation in fraudulent compared to genuine academic biomedical publications. The authors compared just over 250 academic publications retracted for fraud to just over 250 genuine publications and observed that the fraudulent papers contained more obfuscation, including less readable writing (e.g., more words per sentence, more syllables per word; Flesch, 1948), than genuine papers. In another investigation, Humpherys, Moffitt, Burns, Burgoon, and Felix (2011) evaluated over 200 corporate documents submitted to the U.S. Securities and Exchange Commission and observed that fraudulent reports contained significantly more words, more sentences, and more complex words than did nonfraudulent reports, evidence supporting the linguistic obfuscation hypothesis as well (Bloomfield, 2002; Courtis, 1998, 2004; Li, 2008; Markowitz & Hancock, 2016). Other evidence suggests that obfuscation, as indicated by rates of linguistic abstraction and jargon (Markowitz & Hancock, 2016), is also found in the codes of conduct from companies that commit ethics infractions such as anticompetitive activity, fraud, or environmental violations (Markowitz, Kouchaki, Hancock, & Gino, 2021). These language patterns can also affect how people form perceptions about a target (e.g., people rate a company as being less warm, less moral, and less trustworthy if its corporate writing is obfuscated compared to nonobfuscated) and cheating behavior (e.g., people cheat more on problem-solving tasks after reading a corporate document that contains high- compared to low levels of obfuscation). Together, these consistencies across studies likely emerged because of contextual similarities: When people write about fraud, try to appear credible, and the consequences of deception detection are high, patterns of linguistic obfuscation tend to differentiate lies from truths.

Low-stakes, everyday lies tend to have a similar verbal composition as well. Research by Markowitz and Hancock (2018) had mobile daters record their lies to a recent match, and coders assessed the impression management strategies associated with each lie. The majority of the deceptions could be categorized as self-presentation lies (e.g., lies to appear attractive, interesting, likable, dateable) and availability management lies (e.g., lies that use the ambiguities of computer-mediated communication to control how available or eager the dater appears to the match). People who communicated self-presentation lies would describe how much they enjoyed going to the gym, how much they loved dogs, or how much they loved reading. People who communicated availability management lies tried to avoid activities with another person and often provided excuses for being unavailable, such as their phone dying, being busy with work, or not seeing a message appear on their screen. Self-presentation and availability management lies are considered low-stakes impression management deceptions, because they help to move a conversation forward and on average are unlikely to terminate a relationship if detected.

The same impression management lies occur between people who are nonromantic partners as well. Early work by Hancock et al. (2009) had participants communicate in an online instant message chat with another person, and availability management lies represented nearly one-fifth of the lies. These effects have been replicated with similar deception frequencies across text messaging with friends (Birnholtz, Guillory, Hancock, & Bazarova, 2010; French, Smith, Birnholtz, & Hancock, 2015; Reynolds, Smith, Birnholtz, & Hancock, 2013). Together, these studies indicate that the message content of lies tends to be similar across platforms (e.g., instant messaging chat, texting with friends, texting a romantic stranger) when people have consistent impression

management goals (e.g., self-presentation, availability management).

CONTEXT MATTERS

In this chapter, we have reviewed evidence in support of the idea that deception affects language, but not uniformly. The COLD model (Markowitz & Hancock, 2019) argues that psychological dynamics (e.g., emotional and cognitive aspects of a deception), pragmatic goals (e.g., what the liar is trying to accomplish by communicating falsely), and genre conventions (e.g., message-related characteristics relevant to a discourse community) can modify the relationship between deception and language, and a comprehensive view of this relationship can help us understand mixed empirical evidence while guiding new predictions about dishonesty. Given this wealth of evidence, what conclusions can we draw about how deception affects language, and what can future researchers learn from taking this context-contingent view?

Deception research has largely concluded that Pinocchio's nose does not exist, suggesting that there is no single cue to identify false from truthful behavior across settings (DePaulo et al., 2003; Markowitz, 2020; Vrij, 2008). Prior research has also argued that current studies are too underpowered to detect any reliable effect of deception on behavior (Luke, 2019). This leads to at least three possible outcomes for the relationship among deception, behavioral cues (e.g., language), and context (Levine, personal communication, June 9, 2019).

First, there may be a subset of cues from language that distinguishes lies and truths at the 5% level across individuals, messages, situations, and settings. Consistent with Luke (2019), however, the extant literature cannot detect these effects, because primary studies are underpowered. A second possibility is that null effects described in primary studies are indeed genuine, and cues do not extend across individuals, messages, situations, and settings. While meta-analyses provide substantial counterevidence to this claim (e.g., Hartwig & Bond, 2014), there are noteworthy concerns in deception research such as small sample sizes, selective reporting, and publication bias in primary studies (Thornton & Lee, 2000). Third, we have not accounted for complexities of the relationship between deception and language and therefore, failed to capture moderators or individual differences that may help to reveal their genuine relationship. As Hauch et al. (2015, p. 330) acknowledge, considering "alternative theoretical approaches may find other cues or moderators to be important" and in this chapter, we propose that a context-contingent model (Markowitz & Hancock, 2019) may be helpful to at least identify more aspects of a deception that matter for language.

Recent work has used the COLD model to test the reliability of cue-based approaches across lies, truths, and genre-related speech. In a preregistered, a priori powered experiment, Markowitz and Griffin (2020) gave participants the following instructions: lie (e.g., "Your task is to write about your attitudes on abortion. When discussing this topic, however, we would like you to lie about how you feel"), tell the truth (e.g., "Your task is to write about your attitudes on abortion. You can tell us anything about your attitudes on the topic and please write in detail. It is important that you tell us your truthful and honest views, as people may attempt to guess your true views"), or write within a particular genre (e.g., "Your task is to write about your attitudes on abortion. You can tell us anything about your attitudes on the topic and please write in detail"). Note that participants were randomly assigned to either give their attitudes on abortion or friends, but topic differences did not emerge.

Consistent with prior evidence from primary studies (Newman et al., 2003), false speech contained fewer self-references and more negative emotion terms than truths. Truthful speech was not significantly different than genre-related speech across these language features, but differences in analytic thinking, auxiliary verbs, and adjectives indeed emerged. These data provide evidence that explicit lies and truths contain signals that are consistent with prior evidence. Speech that is not explicitly false or truthful (genre-related speech) can differ from lies and truths, but established cue-based models may not pick up on these differences. Addressing the "third bar problem" (Markowitz & Griffin, 2020), specifically

considering the language effects of (1) lies, (2) truths, and (3) the genre, may provide a more comprehensive view of a deception setting through language patterns. Exploratory techniques, which remain consistent with principles of open and transparent science, should complement theoretically and empirically grounded investigations.

In our prior work (Markowitz & Hancock, 2019), we acknowledged that psychological dynamics, pragmatic goals, and genre conventions are three important aspects of the context that matter in deception research related to language, but there are likely more. If researchers are interested in deception detection and using language to discriminate between lies and truths, it is possible that each deception setting needs to be modeled uniquely. For example, prior evidence suggests that papers with fraudulent data tend to have more linguistic obfuscation than do papers with genuine data, but these texts originated from biomedical publications (Markowitz & Hancock, 2016). It is unclear whether papers from physics would exhibit the same pattern, since the writing conventions of these papers may differ from those in the original sample. Therefore, in order for researchers to investigate how deception affects language and how context influences this relationship, frameworks for each deception type may be required. Contextual factors should be considered in the design of a study, not just as a post hoc rationalization of mixed findings.

Together, we believe that the effect of deception on language is real, but small effect sizes and contextual constraints impact the interpretability of the extant literature. Future research should consider how psychological dynamics, pragmatic goals, and genre conventions affect deceptions but continue to investigate how other aspects of context (e.g., the cultural setting, interpersonal dynamics; Burleson, 2009) may impact the relationship between deception and language. We also believe future research might benefit from trying to automatically situate deceptive discourse production in context. Some social media companies are already attempting to perform this action. For example, in May 2020, Twitter added a warning to President Donald Trump's Tweet about the harms of mail-in voting, signaling people to "get the facts" by clicking on resources for more truthful information. A similar approach could be useful for low-stakes, everyday lies as well. Perhaps a person's social media record (e.g., his or her text data, social connections) could be assessed with other, available online content to evaluate a message's veracity. Triangulating such information would allow message receivers to understand how a statement's veracity was determined and indicate the deceptive tendencies of a speaker.

REFERENCES

Akehurst, L., Köhnken, G., Vrij, A., & Bull, R. (1996). Lay persons' and police officers' beliefs regarding deceptive behaviour. *Applied Cognitive Psychology, 10*(6), 461–471.

Ariely, D. (2012). *The (honest) truth about dishonesty: How we lie to everyone—especially ourselves*. HarperCollins.

Baayen, R. H., Piepenbrock, R., & Bulickers, L. (1995). *The CELEX Lexical Database*. Linguistic Data Consortium, University of Pennsylvania.

Back, M. D., Stopfer, J. M., Vazire, S., Gaddis, S., Schmukle, S. C., Egloff, B., & Gosling, S. D. (2010). Facebook profiles reflect actual personality, not self-idealization. *Psychological Science, 21*(3), 372–374.

Biber, D., Connor, U., & Upton, T. A. (2007). *Discourse on the move: Using corpus analysis to describe discourse structure*. John Benjamins.

Birnholtz, J., Guillory, J., Hancock, J., & Bazarova, N. (2010). "On my way": Deceptive texting and interpersonal awareness narratives. *Proceedings of the 2010 ACM Conference on Computer Supported Cooperative Work*. Association for Computing Machinery.

Blackburn, K. G., Yilmaz, G., & Boyd, R. L. (2018). Food for thought: Exploring how people think and talk about food online. *Appetite, 123,* 390–401.

Blair, J. P., Levine, T. R., & Shaw, A. S. (2010). Content in context improves deception detection accuracy. *Human Communication Research, 36,* 423–442.

Blandón-Gitlin, I., Fenn, E., Masip, J., & Yoo, A. H. (2014). Cognitive-load approaches to detect deception: Searching for cognitive mechanisms. *Trends in Cognitive Sciences, 18*(9), 441–444.

Bloomfield, R. J. (2002). The "incomplete revelation hypothesis" and financial reporting. *Accounting Horizons, 16*(3), 233–243.

Boyd, R. L., & Pennebaker, J. W. (2015). Did Shakespeare write *Double Falsehood*?: Identifying individuals by creating psychological signatures with text analysis. *Psychological Science, 26*(5), 570–582.

Burgoon, J., Mayew, W. J., Giboney, J. S., Elkins, A. C., Moffitt, K., Dorn, B., . . . Spitzley, L. (2016). Which spoken language markers identify deception in high-stakes settings?: Evidence From earnings conference calls. *Journal of Language and Social Psychology, 35*(2), 123–157.

Burleson, B. R. (2009). The nature of interpersonal communication: A message-centered approach. In C. R. Berger, M. E. Roloff, & D. R. Roskos-Ewoldsen (Eds.), *The handbook of communication science* (pp. 145–165). Sage.

Burns, M., & Moffitt, K. (2014). Automated deception detection of 911 call transcripts. *Security Informatics, 3*(1), 1–9.

Chung, C. K., & Pennebaker, J. W. (2007). The psychological functions of function words. In K. Fiedler (Ed.), *Social communication* (pp. 343–359). Psychology Press.

Courtis, J. K. (1998). Annual report readability variability: Tests of the obfuscation hypothesis. *Accounting, Auditing and Accountability Journal, 11*(4), 45–9472.

Courtis, J. K. (2004). Corporate report obfuscation: Artefact or phenomenon? *British Accounting Review, 36*(3), 291–312.

DePaulo, B. M., Malone, B. E., Lindsay, J. J., Muhlenbruck, L., Charlton, K., & Cooper, H. (2003). Cues to deception. *Psychological Bulletin, 129*(1), 74–118.

Drouin, M., Miller, D., Wehle, S. M. J., & Hernandez, E. (2016). Why do people lie online?: "Because everyone lies on the internet." *Computers in Human Behavior, 64,* 134–142.

Dwoskin, E., & Timberg, C. (2018). How merchants use Facebook to flood Amazon with fake reviews. Retrieved August 23, 2019, from *www.washingtonpost.com/business/economy/how-merchants-secretly-use-facebook-to-flood-amazon-with-fake-reviews/2018/04/23/5dad1e30-4392-11e8-8569-26fda6b404c7_story.html?noredirect=on*.

Ekman, P. (2001). *Telling lies: Clues to deceit in the marketplace, politics, and marriage.* Norton.

Ellison, N. B., Hancock, J. T., & Toma, C. L. (2012). Profile as promise: A framework for conceptualizing veracity in online dating self-presentations. *New Media and Society, 14*(1), 45–62.

Fang, F. C., Steen, R. G., & Casadevall, A. (2012). Misconduct accounts for the majority of retracted scientific publications. *Proceedings of the National Academy of Sciences USA, 109*(42), 17028–17033.

Finkel, E. J., Eastwick, P. W., Karney, B. R., Reis, H. T., & Sprecher, S. (2012). Online dating: A critical analysis from the persepctive of psychological science. *Psychological Science in the Public Interest, 13*(1), 3–66.

Flesch, R. (1948). A new readability yardstick. *Journal of Applied Psychology, 32*(3), 221–233.

French, M., Smith, M. E., Birnholtz, J., & Hancock, J. T. (2015). Is this how we (all) do it?: Butler lies and ambiguity through a broader lens. *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. Association for Computing Machinery.

Funk, C., & Kennedy, B. (2019). Public confidence in scientists has remained stable for decades. Retrieved August 26, 2019, from *www.pewresearch.org/fact-tank/2019/03/22/public-confidence-in-scientists-has-remained-stable-for-decades*.

Global Deception Research Team. (2006). A world of lies. *Journal of Cross-Cultural Psychology, 37*(1), 60–74.

Guillory, J., & Hancock, J. T. (2012). The effect of Linkedin on deception in resumes. *Cyber-Psychology, Behavior and Social Networking, 15*(3), 135–140.

Guillory, J. E., Hancock, J. T., Woodruff, C., & Keilman, J. (2015). Text messaging reduces analgesic requirements during surgery. *Pain Medicine, 16*(4), 667–672.

Hancock, J. T., Birnholtz, J. P., Bazarova, N. N., Guillory, J., Amos, B., & Perlin, J. (2009). Butler lies: Awareness, deception and design. *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)*. Association for Computing Machinery.

Hancock, J. T., Curry, L. E., Goorha, S., & Woodworth, M. (2007). On lying and being lied to: A linguistic analysis of deception in computer-mediated communication. *Discourse Processes, 45*(1), 1–23.

Hartwig, M., & Bond, C. F. (2014). Lie detection from multiple cues: A meta-analysis. *Applied Cognitive Psychology, 28*(5), 661–676.

Hauch, V., Blandón-Gitlin, I., Masip, J., & Sporer, S. L. (2015). Are computers effective lie detectors?: A meta-analysis of linguistic cues to deception. *Personality and Social Psychology Review, 19*(4), 307–342.

Humpherys, S. L., Moffitt, K. C., Burns, M. B., Burgoon, J. K., & Felix, W. F. (2011). Identification of fraudulent financial statements using linguistic credibility analysis. *Decision Support Systems, 50*(3), 585–594.

Hyland, K. (2001). Humble servants of the disci-

pline?: Self-mention in research articles. *English for Specific Purposes, 20*(3), 207–226.

Hyland, K. (2003). Self-citation and self-reference: Credibility and promotion in academic publication. *Journal of the American Society for Information Science and Technology, 54*(3), 251–259.

Ireland, M. E., & Mehl, M. (2014). Natural language use as a marker of personality. In T. M. Holtgraves (Ed.), *The Oxford handbook of language and social psychology* (pp. 201–218). Oxford University Press.

Ireland, M. E., Slatcher, R. B., Eastwick, P. W., Scissors, L. E., Finkel, E. J., & Pennebaker, J. W. (2011). Language style matching predicts relationship initiation and stability. *Psychological Science, 22*(1), 39–44.

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

Kacewicz, E., Pennebaker, J. W., Davis, M., Jeon, M., & Graesser, A. C. (2014). Pronoun use reflects standings in social hierarchies. *Journal of Language and Social Psychology, 33*(2), 125–143.

LaFollette, M. C. (1992). *Stealing into print: Fraud, plagiarism, and misconduct in scientific publishing.* University of California Press.

Larrimore, L., Jiang, L., Larrimore, J., Markowitz, D. M., & Gorski, S. (2011). Peer to peer lending: The relationship between language features, trustworthiness, and persuasion success. *Journal of Applied Communication Research, 39*(1), 19–37.

Levelt, W. J. M., Drenth, P. J. D., & Noort, E. (2012). Flawed science: The fraudulent research practices of social psychologist Diederik Stapel. Retrieved from *https://pure.mpg.de/pubman/faces/viewitemfullpage.jsp?itemid=item_1569964*.

Levine, T. R. (2018). Ecological validity and deception detection research design. *Communication Methods and Measures, 12*(1), 45–54.

Levine, T. R. (2020). *Duped: Truth-Default Theory and the social science of lying and deception.* University of Alabama Press.

Levine, T. R., Ali, M. V., Dean, M., Abdulla, R. A., & Garcia-Ruano, K. (2016). Toward a pan-cultural typology of deception motives. *Journal of Intercultural Communication Research, 45*(1), 1–12.

Levine, T. R., Kim, R. K., & Hamel, L. M. (2010). People lie for a reason: Three experiments documenting the principle of veracity. *Communication Research Reports, 27*(4), 271–285.

Li, F. (2008). Annual report readability, current earnings, and earnings persistence. *Journal of Accounting and Economics, 45*(2–3), 221–247.

Lu, S. F., Jin, G. Z., Uzzi, B., & Jones, B. (2013). The retraction penalty: Evidence from the web of science. *Scientific Reports, 3,* 1–4.

Luke, T. J. (2019). Lessons from Pinocchio: Cues to deception may be highly exaggerated. *Perspectives on Psychological Science, 14*(4), 646–671.

Margolin, D., & Markowitz, D. M. (2018). A multitheoretical approach to big text data: Comparing expressive and rhetorical logics in Yelp reviews. *Communication Research, 45*(5), 688–718.

Markowitz, D. M. (2018). Academy Awards speeches reflect social status, cinematic roles, and winning expectations. *Journal of Language and Social Psychology, 37*(3), 376–387.

Markowitz, D. M. (2020). The deception faucet: A metaphor to conceptualize deception and its detection. *New Ideas in Psychology, 59,* Article 100816.

Markowitz, D. M., & Griffin, D. J. (2020). When context matters: How false, truthful, and genre-related communication styles are revealed in language. *Psychology, Crime and Law, 26*(3), 287–310.

Markowitz, D. M., & Hancock, J. T. (2014). Linguistic traces of a scientific fraud: The case of Diederik Stapel. *PLOS ONE, 9*(8), e105937.

Markowitz, D. M., & Hancock, J. T. (2016). Linguistic obfuscation in fraudulent science. *Journal of Language and Social Psychology, 35*(4), 435–445.

Markowitz, D. M., & Hancock, J. T. (2018). Deception in mobile dating conversations. *Journal of Communication, 68*(3), 547–569.

Markowitz, D. M., & Hancock, J. T. (2019). Deception and language: The Contextual Organization of Language and Deception (COLD) framework. In T. Docan-Morgan (Ed.), *The Palgrave handbook of deceptive communication* (pp. 193–212). Palgrave Macmillan.

Markowitz, D. M., Hancock, J. T., & Tong, S. T. (2018). Interpersonal dynamics in online dating: Profiles, matching, and discovery. In Z. Papacharissi (Ed.), *A networked self and love* (pp. 50–61). Routledge.

Markowitz, D. M., & Levine, T. R. (2021). It's the situation and your disposition: A test of two honesty hypotheses. *Social Psychological and Personality Science, 12*(2), 213–224.

Markowitz, D. M., Kouchaki, M., Hancock, J. T., & Gino, F. (2021). The deception spiral: Corporate obfuscation leads to perceptions of immorality and cheating behavior. *Journal of Language and Social Psychology, 40,* 277–296.

Markowitz, D. M., & Shulman, H. C. (2021).

The predictive utility of word familiarity for online engagements and funding. *Proceedings of the National Academy of Sciences, 118*(18), e2026045118.

Nahari, G., Ashkenazi, T., Fisher, R. P., Granhag, P.-A., Hershkowitz, I., Masip, J., . . . Vrij, A. (2019). "Language of lies": Urgent issues and prospects in verbal lie detection research. *Legal and Criminological Psychology, 24*(1), 1–23.

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*(5), 665–675.

Ott, M., Choi, Y., Cardie, C., & Hancock, J. T. (2011). Finding deceptive online spam by any stretch of the imagination. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics.

Pennebaker, J. W., Chung, C. K., Frazee, J., Lavergne, G. M., & Beaver, D. I. (2014). When small words foretell academic success: The case of college admissions essays. *PLOS ONE, 9*(12), e115844.

Reynolds, L., Smith, M. E., Birnholtz, J. P., & Hancock, J. T. (2013). Butler lies from both sides: Actions and perceptions of unavailability management in texting. In *Proceedings of the 2013 Conference on Computer Supported Cooperative Work*. Association for Computing Machinery.

Salen, K. (2015). The importance of online customer reviews. Retrieved August 23, 2019, from *www.invespcro.com/blog/the-importance-of-online-customer-reviews-infographic*.

Schober, M. F., & Glick, P. J. (2011). Self-deceptive speech: A psycholinguistic view. In C. Piers (Ed.), *Personality and psychopathology: Critical dialogues with David Shapiro* (pp. 183–200). Springer.

Tausczik, Y. R., & Pennebaker, J. W. (2010). The psychological meaning of words: LIWC and computerized text analysis methods. *Journal of Language and Social Psychology, 29*(1), 24–54.

ten Brinke, L., & Porter, S. (2012). Cry me a river: Identifying the behavioral consequences of extremely high-stakes interpersonal deception. *Law and Human Behavior, 36*(6), 469–477.

Thornton, A., & Lee, P. (2000). Publication bias in meta-analysis: Its causes and consequences. *Journal of Clinical Epidemiology, 53*(2), 207–216.

Toma, C., & Hancock, J. T. (2012). What lies beneath: The linguistic traces of deception in online dating profiles. *Journal of Communication, 62*(1), 78–97.

Turner, R. E., Edgley, C., & Olmstead, G. (1975). Information control in conversations: Honesty is not always the best policy. *Kansas Journal of Sociology, 11*, 69–89.

Vrij, A. (2008). *Detecting lies and deceit: Pitfalls and opportunities* (2nd ed.). Wiley.

Vrij, A., Leal, S., Granhag, P. A., Mann, S., Fisher, R. P., Hillman, J., & Sperry, K. (2009). Outsmarting the liars: The benefit of asking unanticipated questions. *Law and Human Behavior, 33*(2), 159–166.

Walther, J. B., & Jang, J. (2012). Communication processes in participatory websites. *Journal of Computer-Mediated Communication, 18*(1), 2–15.

Zuckerman, M., Koestner, R., & Driver, R. (1981). Beliefs about cues associated with deception. *Journal of Nonverbal Behavior, 6*(2), 105–114.