New Agendas in Communication
A Series from Routledge and
the College of Communication
at The University of Texas at Austin

Roderick Hart and Stephen Reese, Series Editors

This series brings together groups of emerging scholars to tackle important interdisciplinary themes that demand new scholarly attention and reach broadly across the communication field's existing courses. Each volume stakes out a key area, presents original findings, and considers the long-range implications of its "new agenda."

The Interplay of Truth and Deception
edited by Matthew S. McGlone and Mark L. Knapp

Journalism and Citizenship
edited by Zizi Papacharissi

Understanding Science
edited by LeeAnn Kahlor and Patricia Stout

Political Emotions
edited by Janet Staiger

Media Literacy
edited by Kathleen Tyner

The Interplay of Truth and Deception

New Agendas in Communication

Edited by
Matthew S. McGlone and
Mark L. Knapp

(2010)

Routledge
Taylor & Francis Group
NEW YORK AND LONDON

Chapter 2

Some Considerations for a New Theory of Deceptive Communication

Timothy R. Levine and Rachel K. Kim

Scholarly books, chapters, and articles on the topic of deceptive communication often begin with the assertion that deception is frequent, ubiquitous, and commonplace. For example, readers of the deception literature encounter claims such as "communicators frequently decide that honesty is not the best policy" and "deception and suspected deception arise in at least one quarter of all conversation" (Buller & Burgoon, 1996, p. 203). Careful consideration of the probable prevalence of deception, however, requires considering the question of relative frequency in comparison to some standard.

Perfect estimates of the prevalence of deception do not exist, and the overall frequency of deceptive communication in a society is probably not the sort of fact that is discoverable through the scientific method or otherwise. Just how does one create a representative sample of communication and then reliably ascertain truth and deception? Nevertheless, there is good reason to believe that (a) successful deception is substantially more prevalent than discovered deception, but (b) overall, while a daily occurrence, deceptive communication probably occurs with much less frequency than nondeceptive communication. Explaining the reasons behind these educated guesses requires defining deception and reviewing what is known about deception.

Consideration of definition and reliable research findings will lead to the main thesis of this chapter, namely that current theoretical approaches to deception are lacking and new theoretical directions are needed to understand deceptive communication and its detection. This chapter selectively reviews the research on deception and deception detection in an effort to suggest new and alternative perspectives on the topic. The chapter ends with a rough sketch of what a new theoretical perspective might look like.

Defining Deception

Deception involves knowingly or intentionally misleading another person.

Although this definition is simple, many important implications follow from defining deception in this way.

For a start, truth and deception are not polar opposites and deception and falsehoods are far from synonymous. For example, a so-called "honest mistake," that is, saying something that one incorrectly believes to be true, is not deception. Or, saying something known to be false is not deceptive if it is said in such a way that the hearer should know it is false. Sarcasm is an obvious example. None of these cases involves an intent to mislead. However, saying something that is literally true in a sarcastic way so that the listener infers something false can be deceptive. In short, what is literally true can be deceptive and saying something false need not be.

Following this line of thought, useful distinctions can be made between actual deception, deceptive attempts, messages perceived as deceptive, and messages that are functionally deceptive. Actual deception is meant to deceive, and achieves this end; the target person is misled by design. In deceptive attempts, someone tries to deceive, and there is deceptive intent, but the target is not actually misled. This situation may be thought of as failed deception. In perceived deception, the target person attributes deceptive intent and thinks that someone has or is lying regardless of the actual intent. Finally, messages that are functionally deceptive mislead others regardless of the actual or perceived intent. Functionally deceptive messages lead to the same outcome as deception without getting into the message source's head.

A related implication is that message features, message intent, and message function or impact need to be distinguished because these things do not map perfectly onto one another. So, someone can say something that is objectively false, omit information, change the subject, and so forth, in a manner that is either likely to, and intended to, deceive or not. The objective truth or falsity of messages may or may not actually function as deception, and such messages may or may not be perceived as deception. In short, speaker intent and message consequence in conjunction define deception, not the objective qualities of messages or information dimensions.

Deception Prevalence

Because deception, by definition, involves deceptive intent, many if not most messages that depart from the whole truth and nothing but the truth fall well short of deception. Most communication, for example, is by necessity truncated. If someone asks how you are, a simple "fine" will usually suffice, and a fully disclosive answer is typically inappropriate. After all, Grice's (1989) maxim of quantity also applies to refraining from providing too much information in addition to providing sufficient

information. As a consequence, studies assessing the prevalence of "information control" (e.g., Turner, Edgley, & Olmstead, 1975) are not necessarily informative about the frequency of deception. Turner et al. find that omission, evasion, and inaccuracies are commonplace in conversation. While this is no doubt true, these forms of information control are not equivalent to deception because deceptive intent is unknown.

The impossibility of sampling deception has already been mentioned, but at least two studies provide some idea of how often people lie. DePaulo, Kashy, Kirkendol, Wyer, and Epstein (1996) had 70 college students and 70 nonstudents keep a lying diary for 1 week. During that time, the vast majority of respondents (95%) reported at least one lie, and on average, college students reported two lies per day while nonstudents reported a single lie per day. For the students, a lie was told in 30% of all interactions, and 20% of conversations for the nonstudents. These findings suggest that lying is an everyday occurrence, but also that it is relatively infrequent compared to nondeceptive communication. If one considers the sheer amount of communication we engage in during the course of a day, one to two deceptive messages is proportionally small.

More recently, Serota and Levine (2008) tried to access the prevalence of deception in American life with a different methodology. They asked a nationally representative sample of 1,000 individuals (stratified by age, sex, education, income, and region of the country) if they had lied in the last 24 hours. The mean number of lies per day was slightly less than two, a value similar to that reported above. The distribution, however, was highly skewed. Sixty percent of respondents reported zero lies in the past 24 hours, but a few of those who did lie reported as many as 54 lies. These findings suggest that many people may not lie each and every day, although we strongly suspect that almost everyone lies sometimes. More interestingly, if these findings are accurate, the findings suggest that most lies may be told by a relatively few very prolific liars. Together with the earlier findings, the conclusion seems to be that lying is prevalent in that we are likely to encounter lies on a daily basis, but infrequent in comparison to everyday honest communication. In other words, most people do lie at times, but most people are much more honest than not.

Deception Motives

Despite widespread social and moral prohibitions against deception, deception nevertheless occurs. The question that naturally arises is why do people deceive others?

The answer is that people lie for a reason. More precisely, we believe that people often deceive others when the truth poses some obstacle to goal obtainment. Absent some psychopathology, people do not deceive when the truth works just fine. In short, most people follow the maxim "Do not lie if you do not have to" most of the time.

This maxim is consistent with what noted philosopher and ethicist Sissela Bok (1999) has labeled the "principle of veracity." According to Bok, there exists a moral asymmetry between truth and deception. The telling of truth requires no justification, deceit does. Honesty and trust provide a necessary foundation for human relations and symbolic exchange. Violating these requires ethical justification whereas adherence does not.

Previous work has sought to classify deception motives in a variety of ways. For example, Turner et al. (1975) list five motivations including (a) to save face, (b) to manage relationships, (c) to exploit, (d) to avoid tension/conflict, and (e) to control situations. More commonly, motives are classified by the locus of benefit; that is, whether the deception primarily benefits self, other, or the relationship (e.g., Hample, 1980; Metts, 1989).

Examining existing typologies, it appears that none of the goals achieved through deception are at all unique to deception. That is, the various category systems delineating various motives for deception do not differ from more general social motivations guiding nondeceptive behavior. For example, consider face goals. The goal of a face-maintaining message is not to deceive per se, but to manage self and other's face needs, and these ends can be accomplished through both honest and deceptive means. Similarly, virtually all instrumental and relational goals can, depending on the situation, be achieved through both honest and deceptive actions. Thus, deception is best thought of as a possible tactic, strategy, or means for goal attainment rather than a desired end in itself.

The probability of using deceptive rather than honest means for goal attainment is likely conditional on situation features and constraints, and not on the nature or type of the goal pursued. According to Bok's (1999) principle of veracity, the moral culpability associated with deception creates an initial imbalance in the assessment of deceptive and truthful alternatives, and adopting deceptive means requires justification that is not necessary for truthful means. So, although deception is in almost everyone's social repertoire, it is generally employed as a tactical or strategic option of last resort or path of least resistance. People will not be deceptive when the truth is sufficient, efficient, and effective for goal attainment. It is only when the truth poses an obstacle to goal obtainment, regardless of what that goal might be, that people entertain the possibility of being deceptive. That is, people are deceptive only when truthful alternatives are more effortful or less efficacious.

We tested this thinking in a series of three studies (Levine, Kim, & Hamel, 2008). In Studies 1 and 2, participants considered what they would say in a set of common situations where motive to deceive was varied by making the truth problematic or not for goal obtainment. In Study 1, participants ($N = 66$) selected honest or deceptive messages in response to situations, while Study 2 ($N = 68$) replicated the first with

open-ended responses coded for deceptive content. Consistent with predictions, the selection or generation of deceptive messages was virtually nonexistent in situations where deception motive was absent (Study 1 = 1.6%, Study 2 = 0.0%) but substantial when such a motive was present (Study 1 = 62.5%, Study 2 = 64.3%).

The self-report findings of the first two studies were supplemented with behavioral observations in Study 3, where 126 participants were given an opportunity to cheat for monetary gain and subsequently interrogated about cheating. For noncheaters, sincere honesty was enough for attaining goals, but being truthful was problematic for cheaters. Once again as predicted, a substantial number of cheaters (60%) chose to lie about cheating while not a single noncheater did so. These three studies provide evidence consistent with Bok's (1999) principle of veracity. The findings indicate that people lie for a reason and generally avoid deception when truth will suffice for goal attainment. On the other hand, those with a motive to deceive do not always choose to lie. The consistency of the findings across our three studies can be seen in Figure 2.1.

We also tested a projected motive model explaining deception detection upon the idea that people not only often act in accordance with the veracity principle, but they likely believe that others adhere to it too (Levine, Kim, & Blair, 2008). That is, when people consider whether a

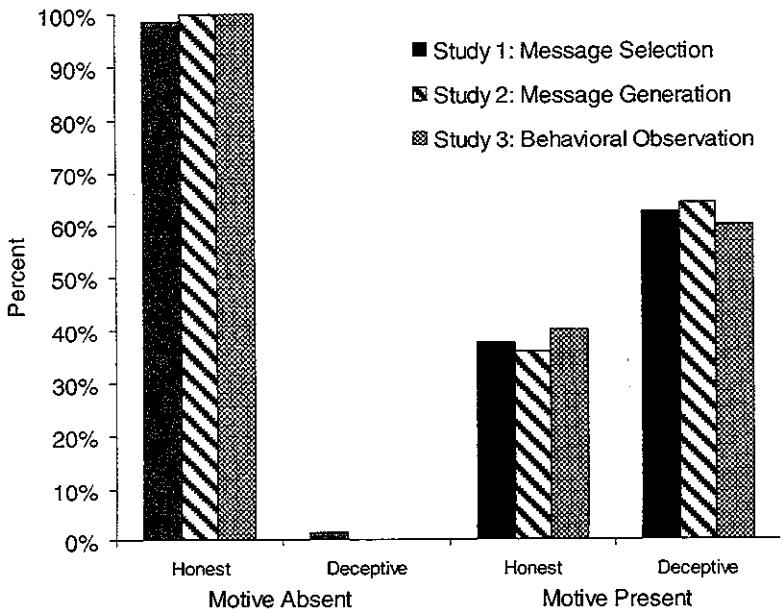

Figure 2.1 Honest and deceptive message production in Levine, Kim, and Hamel (2008).

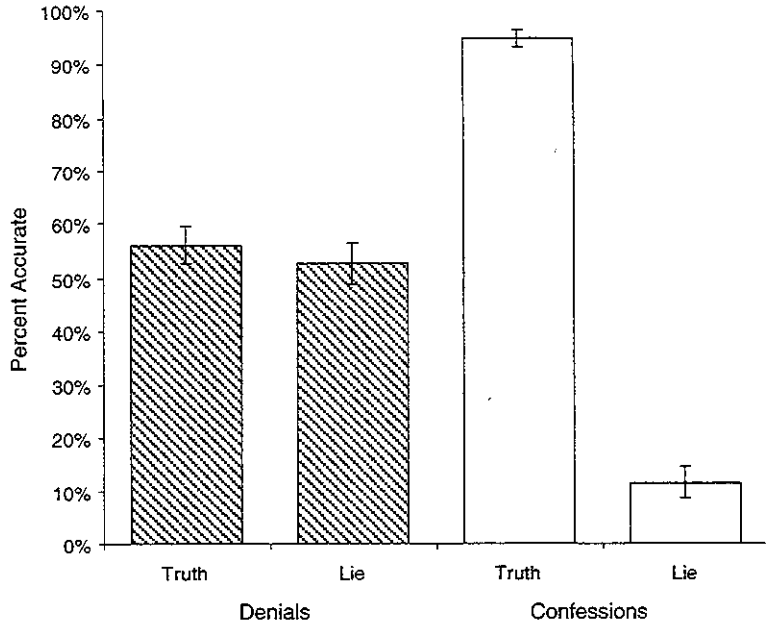

Figure 2.2 Detection accuracy for confessions and denials in Levine, Kim, and Blair (2008), Study 1.

message might be deceptive, they consider if the message source has reason to lie. If there is no perceived motive for deception, then honesty is likely presumed. Consistent with this reasoning, Levine Kim, and Blair (2008) found confessions were judged as honest more frequently than denials (η^2 = .64), and while detection accuracy was just above 50% for denials regardless of message veracity, it was 95% for true confessions, and only 12% for false confessions (see Figure 2.2). These findings also highlight the importance of the content-in-context in judgmental processes, which is generally lacking in extant deception research and theory.

Taken together, these findings have important implications for theory and research in deception. For one, motive plays a crucial role in deception, yet most studies of deception typically lack proper context reflecting the centrality of motive. For example, in many cue studies the only motive for deception is following instructions and held constant, and potential liars in detection studies often have no apparent motive at all. Another implication is that situational factors may have a stronger impact than individual differences in the probability of producing deceptive messages. Aside from perhaps a few prolific liars, people generally eschew deception unless situation features and constraints make being truthful problematic for attaining goals.

Deception Detection Accuracy

Perhaps one of the most reliable and well-documented findings in all of social science is that people are statistically significantly better than chance, but only slightly more so, at detecting deception. Meta-analysis of more than 200 separate lie detection experiments finds that people are, on average, about 54% accurate in assessing message veracity when they have a 50–50 chance of being right (Bond & DePaulo, 2006). The results of most individual studies fall within 10% of this across-study average (i.e., approximately between 45% and 65%). Thus, the only-slightly-better-than-50% finding is stable and consistent. Not surprisingly, it has become very widely accepted among deception researchers.

More recently, Bond and DePaulo (2008) looked at the variance in accuracy judgments rather than just average accuracy levels. They decomposed accuracy scores into four components: demeanor, truth-bias, transparency, and ability. *Demeanor* is the tendency of a person being judged to appear honest (or deceptive) independent of whether or not they are lying. Variance in demeanor means that some people are just more (or less) believable than others. *Truth-bias* is the tendency to believe others whether or not they are telling the truth. Variance in truth-bias means some people are more gullible than others; others are more skeptical. *Transparency* refers to how good or bad people are at lying. That is, people who are transparent tend to leak the fact that they are lying and are therefore relatively easier to distinguish when they are lying or telling the truth than those who are less transparent. Finally, *ability* is an individual difference in skill at telling if someone is lying or not. Thus, demeanor and transparency reflect variance in the message source, while truth-bias and ability reflect variance in the message judge. Further, demeanor and truth-bias reflect variance in *bias*; i.e., are tendencies to believe (or not) that are independent of actual veracity, whereas transparency and ability reflect variance in *skill* in presenting or discriminating between honest and deceptive messages.

Bond and DePaulo (2008) found that variance in demeanor is huge, both in an absolute sense and relative to the other sources of variation. Some people are just more believable than others, and this aura of believability has a large impact on judges. There are also substantial individual differences in truth-bias, with these differences being much smaller than the variance in demeanor, but much larger than the two. Finally, the variance in transparency is much larger than the variance in ability. Individual differences in ability contribute little (maybe only plus or minus 1% or 2% in overall accuracy). Thus, variance in believability and accuracy stems more from the message source than the person judging the message, and the variance in bias swamps variance in accuracy. This explains why accuracy values across studies are so stable. The lack of individual differences in judge ability leads to small standard errors and stable findings.

As just noted, looking across individual studies, there is relatively little between-study variance, with almost all studies reporting accuracy rates between 45% and 65%. The 4% improvement-over-chance finding is highly statistically significant and thus clearly reflects some systematic mechanism. An interesting question is why this 54% ±10% finding occurs so reliably. Why not higher accuracy? Why not lower? Why is accuracy so stable if it is not just due to mere chance?

The finding that people detect deception better than chance is not especially surprising. Most of us believe that people generally have at least some skill at reading others. What is perplexing, however, is why people systematically do better than chance, but never much better than chance. That is, a viable explanation for this reliable finding must not only account for why people systematically exceed chance, but also why people never, ever, seem to exceed chance by very much. A close look at current interpretations of these findings suggests that, while intuitive, currently accepted views really do not mesh very well with extant data.

In line with cultural inclinations associated with individualism, there is a tendency to interpret and explain detection accuracy in terms of the individual's skill, competence, and ability to ferret out lies. That is, accuracy findings are most frequently interpreted as informative about people's ability to detect deception, with a myopic emphasis on the message recipient's ability. The message recipients, after all, are the subjects in these experiences, and consequently their performance is the focus. However, all communication, including deception, requires at least two people, and there is no reason why the basis for the accuracy findings might not reside more in the message source or message itself than in the message recipient. In fact, this is just what the Bond and DePaulo (2008) findings suggest. In short, accuracy findings tell us more about the source and nature of messages being judged than the people doing the judging.

A Test-Taking Analogy

Many readers are likely familiar with giving and taking tests. Imagine that a class of 100 students is given four 100-question true-false exams over a semester. The typical interpretation of scores on tests like these is that scores of 50% reflect mere chance performance and no knowledge of class material, and that as scores approach 100%, nonrandom increases above chance levels reflect increasing mastery of course content. There is no reason to expect scores systematically below 50% because surely students would not intentionally seek a low grade by purposely missing a question they know. That is, test scores tell us about the student; how

much they studied, how effectively they studied, how smart they are, their testing-taking ability, and so forth.

Now, imagine that students take a series of tests like the one described above, and the average score on each of the tests is around 55%, with all four tests producing approximately the same outcome and distribution score. The average student performs significantly better than chance, but in an absolute sense, fails miserably. Also, the standard deviations are small, so that over the four tests, all students tend to hover around the mean. Given that students tend not to do systemically worse than chance, there are few outliers on the low end to pull the mean down. What kind of explanation might best explain this pattern of test scores?

The first explanation that may occur to many readers is that the students studied a little but not very much. They learned something, but only a little. For example, the students may have learned and remembered only about 10% of the class content. So, if students knew the answers to 10 questions and guessed on the other 90 questions, then they would get the 10 they knew for sure right and get 50% of the 90 they guessed on and thus receive a 55% on the exam (10 + 45 = 55). So, a "little bit of knowledge" model can explain the outcome. But, it seems odd that *all* the students would only know a little.

An alternative explanation is that the scores tell us more about the test then the test takers. Imagine a true-false test with 90 questions that are impossibly difficult (thus necessitating guessing) and 10 questions that are so easy that everyone gets them right whether or not they took the class, studied, and so forth. This test tells us nothing at all about what was learned. Such a test, however, produces a reliable average of 55%. Further, now the lack of test-taker variance is not at all odd. It is expected. The nature of the test makes test-taker ability irrelevant. Further, to the extent that little test-taker ability variance is observed, the "test is the explanation" account is more plausible than the "little bit of knowledge" model. The test-based explanation is predicated on the idea of variance in test-item difficulty rather than in test-taker ability and is nicely consistent with the new Bond and DePaulo (2008) results.

We propose that most deception detection experiments involve an ecology very different from how people really detect lies. In most laboratory deception detection experiments, there is no motive for deception, and if there is, it is held constant. Also, the type of information people usually use to detect lies outside the lab (consistency with prior knowledge, confession, information from third parties, physical evidence, etc., see Park, Levine, McCornack, Morrison, & Ferrara, 2002) is withheld by research design. It is little wonder that people mostly guess, except when there is the occasional totally transparent liar. The 54% findings reflect the ecology of the experimental design; motive is constant, there is little external information, and there is an equal number of truths and lies.

Reasons for (In)accuracy

Much of the current thinking about deception detection has evolved from Ekman's (2001) idea of *leakage*. The idea is that there are emotional correlates of deception, that emotions are conveyed nonverbally, and that emotional expression is not entirely under conscious control. Compared to the honest message source, deceivers are apt to experience guilt, fear of detection, and so forth. Deceivers try to control behavioral displays so as not to give themselves away, but nonverbal cues to deception leak out anyway, often through nonverbal channels that are more difficult to control. Consequently, the message recipient who actively looks for the right leaked cues should be adept at distinguishing truths from lies.

The leakage idea was expanded upon by Zuckerman, DePaulo, and Rosenthal's (1981) four-factor theory of deception. The four-factor framework specifies four internal states that differentiate truths and lies: arousal, emotions, cognitive effort, and over control. Because relative to truth-tellers, liars are more likely to experience heightened internal states of arousal, emotions like fear and guilt, cognitive effort, and over control of nonverbal displays, and because each of these internal states is thought to be associated with specific nonverbal behaviors (e.g., increased cognitive effort leads to long response latencies), clues to deception are leaked nonverbally.

The most recent iteration of this thinking is reflected in Interpersonal Deception Theory (IDT; Buller & Burgoon, 1996). Liars strategically present themselves as honest, but nonstrategically leak deception cues. Message recipients pick up on these cues and become suspicious. Liars, however, pick up on leaked suspicion, and strategically adapt. In turn, so do receivers. Net accuracy depends on the liar's encoding skill (i.e., transparency) relative to the receiver's decoding skill (i.e., ability) and how the interaction progresses over time.

Thus, according to these dominant theoretical perspectives, the reason people are systematically better than chance in accuracy is that verbal and nonverbal cues indicative of deceit are inadvertently leaked, these leaked behaviors signaling deceit are then perceived, and consequently truth is correctly distinguished from lie. Accuracy is generally a function of a message judge's ability to recognize valid leakage.

From this perspective, the reason people are far from perfect is that (a) there is an imperfect link between any given behavioral display and veracity, and (b) people rely on cues that lack diagnostic utility in addition to valid cues. That is, there are no perfect cues to deception and people often look for the wrong indicators. For example, research finds that the most common belief about deception is that liars avoid eye contact whereas meta-analysis finds no link between gaze and honesty (Bond & The Global Deception Research Team, 2006; Depaulo et al., 2003).

IDT further adds that liars engage in strategic countermoves, constantly adapting their performance to appear more honest and mislead weary judges. The net result is above chance accuracy that is far from perfect.

There are good findings consistent with this view. Meta-analysis of deception cues finds that there are behaviors that probabilistically distinguish truths from lies, but no behavior or set of behaviors does so perfectly (DePaulo et al., 2003). Further, there is less than perfect correspondence between what people look for when detecting deception and the behaviors that have actual utility. Thus, above chance but less than ideal accuracy makes much sense from this perspective, and there are many supportive findings that can be cited as evidence in favor of this stance.

Not all findings in the literature, however, are in line with the ability-to-spot-leakage account. First, the Bond and DePaulo (2008) meta-analysis suggests relatively little variance in transparency and especially ability. At least on the surface, these new findings seem at odds with the skill- and leakage-based accounts. Important here is Bond and DePaulo's finding that variance in transparency in detection accuracy studies is massively larger (by a factor of 10!) than variance in ability. This suggests that accuracy is not so much a function of a person's ability to recognize leakage, but instead much more a function of individual differences in what is leaked independent of who is doing the judging. In terms of our previous test-taking analogy, the variance in accuracy is attributable to differences in test question difficulty rather than student knowledge. Back to deception, most people can lie seamlessly, but a few people's lies get detected by everyone. Accuracy findings tell us more about the liar than the detector. Clearly, this conclusion is at odds with current theory, especially IDT's cat and mouse account.

Second, efforts to enhance accuracy through nonverbal training have failed to document much in the way of improvement. For example, meta-analysis finds only marginal improvements from nonverbal training, and studies offering additional controls find even more meager results (Frank & Feeley, 2003; Levine, Feeley, McCornack, Harms, & Hughes, 2005). If low accuracy simply involved looking for the wrong cues, one would expect much better results from training studies.

Finally, a leakage and judge ability account would also predict that professional expertise would be a strong determinant of accuracy. If accuracy is a skill, people should get better with practice and experience. Police, military interrogators, customs officials, and so forth, should be better at detecting deception than the average college sophomore. Yet, meta-analysis suggests that this is not the case (Bond & DePaulo, 2006).

In short, while there is considerable evidence consistent with leakage-based accounts, there are also findings that are at odds with these accounts. Ideally, good theory should keep analogous findings to a minimum. Thus, there seems to be room for theoretical improvement.

Truth-Bias

Another reliable finding from the accuracy literature is the prevalence of truth-bias. Truth-bias is the tendency to judge messages as honest independent of actual message veracity (Levine, Park, & McCornack, 1999; Levine, Kim, Park, & Hughes, 2006). So far, truth-bias has been discussed as a source of individual variance in judgments, and Bond and DePaulo (2008) report substantial individual differences in truth-bias. Truth-bias, however, is not just (or even mostly) an individual difference. It also has a strong situational component, and it varies considerably based on situational priming (Kim & Levine, 2008; McCornack & Levine, 1990). But most importantly, truth-bias is reliably observable and has substantial impact across both individuals and situations. Bond and DePaulo (2006) found evidence for the existence of a general truth-bias across studies, and a truth-bias has been observed in every experiment we have conducted. That is, truth-bias varies in degree from person to person and situation to situation, but despite these differences, most people are truth-biased most of the time.

Consistent evidence for these claims can be seen in Table 2.1 where truth-bias rates in our research are presented in the first column. As the

Table 2.1 Truth-Bias and Veracity Effect Results

Study	*Truth-Bias*	*Truth and Lie Accuracy*	
McCornack & Levine (1990)	72%	81.8%	31.3%
Levine et al. (1999), Study 4	68%	68.5%	37.5%
Levine & McCornack (2001)	72%	75.0%	31.0%
Study 2	69%	76.7%	39.2%
Study 3	56%	56.8%	44.1%
E. Park, Levine et al. (2002)	66%	67.0%	37.0%
Levine et al. (2005), Study 1	63%	65.3%	38.6%
Study 2	62%	66.4%	43.0%
Study 4	62%	66.4%	43.2%
Levine et al. (2006)	66%	67.1%	34.3%
Levine et al. (2008), Study 1	72%	74.5%	32.1%
Study 2	68%	74.2%	37.7%
Study 3	70%	62.9%	22.5%
Kim & Levine (2008)	71%	73.8%	32.3%

reader can see, truth-bias scores range from a low of 56% to a high of 72%. In every case, the percent of messages judged as honest is greater then the 50% that is expected from nonbiased processing. We have even found a truth-bias when 100% of the messages judged were lies (Levine et al., 2006) and when judges were made highly suspicious (Kim & Levine, 2008; McCornack and Levine, 1990).

We believe there are at least three important reasons behind the strength and persistence of truth-bias. First, truth-bias stems in part from innate, hard-wired cognitive systems that govern how we process incoming information. Gilbert (Gilbert, 1991; Gilbert, Krull, & Malone, 1990) makes a convincing case that belief is a mental default, and while people can reject information as false, doing so requires additional cognitive resources and processes subsequent to comprehension. Doing otherwise would require a less efficient cognitive system, and thus there is likely an evolutionary basis for truth-bias. Second, communication requires truth-bias. If one questioned the veracity of everything others told them, communication could not operate. As Grice (1989) demonstrates, making sense out of what others say requires a presumption that they are communicating cooperatively. Finally, humans are social, and getting along with others requires some degree of trust, coordination, and consideration. Consistent with Goffman's (1955) observations, people give others considerable leeway so that social interaction is not disrupted.

Now it becomes necessary to make a brief digression. The widely accepted and incredibly reliable 54% accuracy plus or minus 10% finding has already been noted. The deception detection experiments leading to this conclusion, however, share some important elements that impact these findings. Accuracy is calculated by averaging across an equal number of truths and lies. The pervasive existence of truth-bias has important implications for the generality of the accuracy conclusions given these common research design features. These implications stem from a set of findings we call the "veracity effect" (Levine et al.,1999).

The *veracity effect* holds that, because people are truth-biased, people are more accurate in judging truths than lies, and therefore source veracity affects detection accuracy. Recall that meta-analysis suggests that (a) people are significantly, but only slightly better than chance, and (b) people tend to be truth-biased. If both are true, since people in deception experiments usually judge messages of which half are truths and half are lies, then people must be getting the truths right more frequently than the lies and truth accuracy must be greater than lie accuracy. Levine et al. (1999) found that the veracity accounted for between 30% and 60% of the variance in several detection accuracy rates in previous studies. Based on these findings, Levine et al. (1999) speculate that in many situations a priori message veracity may be the single biggest factor of deception detection accuracy. Examples of the veracity effect in our findings

can be seen by comparing the second and third columns in Table 2.1. To date, the veracity effect has been evident in every deception detection experiment conducted by Levine and colleagues. More than a dozen replications to date means this is a nicely reliable finding!

Because of the veracity effect, a very different understanding of the literature is obtained when truth accuracy and lie accuracy are scored and reported separately. Just because average accuracy is better than chance does not mean that people can pick out lies more than 50% of the time. In short, different conclusions are reached when considering truth and lie accuracy separately, or when the ratio of truths and lies are unequal.

The findings consistent with the veracity effect led us to the issues of truth-lie base-rates and the Park-Levine probability model (Levine et al., 1999; Levine et al., 2006; Park & Levine, 2001). The idea is that if a priori message veracity predicts subsequent accuracy as suggested by the veracity effect findings, then the ratio of truths to lies is crucial. If the veracity effect is true, then the 54% accuracy findings from the meta-analyses reviewed at the beginning of the chapter are strictly limited to averaged rates under conditions of exactly half truths and half lies. Of course, outside the lab the probability of deception is seldom precisely 50–50. Therefore, a common design feature of 40 years of deception detection experiments has led to meta-analysis results with very limited generality.

The Park-Levine (2001) model was designed to model average detection accuracy rates at different truth-lie base-rates. That is, it seeks to answer the question, what if the base-rate is something other than exactly 50–50? According to the model, detection accuracy at any specific base-rate can be determined by the following formula:

$$\text{Average accuracy across truths and lies} = P(H \mid T) \times P(T) + P(\sim H \mid \sim T) \times P(\sim T)$$

In words, the model predicts that observed total accuracy will be the product of truth accuracy times the proportion of messages that are true plus the product of lie accuracy times the proportion of messages that are lies where the proportion of true messages equals one minus the proportion of messages that are lies. This formula predicts linear effects and, so long as people are truth-biased, a positive slope. The more truth-biased people are, the steeper the slope. If judges are lie-biased, the slope would be negative. The y-intercept is a direct positive function of truth-lie transparency.

An initial test of the model yielded promising results. Levine et al. (2006) had one set of randomly selected subjects in a control group view a series of truths and lies with a 50–50 truth-lie base-rate. This condition

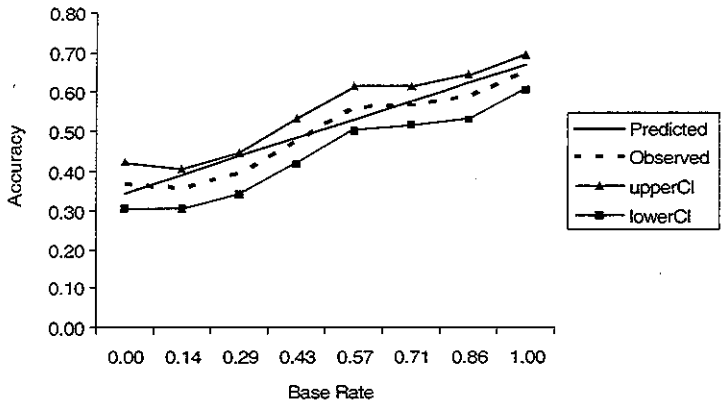

Figure 2.3 Detection accuracy as a function of truth-lie base-rate in Levine et al. (2006).

simulated previous experiments. The results were entered into the model as estimates of P (H | T) and P (~H | ~T). Then another set of subjects participated in an experiment in which the base-rate P(T) was experimentally manipulated. Participants made veracity judgments in 1 of 8 different base-rate conditions ranging from all honest to all deceptive messages. The accuracy in the eight base-rate conditions was predicted based on the formula, which condition participants were in, and the estimates obtained from the control group data. The base-rate induction substantially impacted accuracy explaining 24% of the total variance in accuracy (it explained 21% of the variance in the earlier Levine et al., 1999, Study 4). The effects were linear, with the linear contrast accounting for 95% of the base-rate effect (98% in Levine et al., 1999). The correlation between predicted values and those obtained was $r = .97$. The model predicted the slope and the intercept of the line best fitting the data to within two decimal places. Raw accuracy was predicted to within a couple percentage points. The results are portrayed in Figure 2.3. These results are exciting because they reflect a degree of prediction and precision not often seen in experimental communication research. An interesting next step will be to see if the model generalizes to interactive deception experiments.

How People Really Detect Deception

A final important consideration into deception detection comes from our "How People Really Detect Lies" study (Park et al., 2002). The idea is that deception detection experiments are restrictive in another even more fundamental way than just a common P = .50 base-rate. Whether the messages are taped or interactive, in most of these experiments all

there is for judges to rely on is at-the-time verbal and nonverbal behavior of message sources. Outside the lab, however, people can fact check, talk to others, and so forth. Consequently, we expected that most lies outside the deception lab probably get detected well after the fact and by discovery methods other than utilizing verbal and nonverbal source behaviors at the time of deception. Park et al. (2002) simply asked participants to recall a lie that they had detected, to describe what happened, how they found out that the person was lying, and how much time had elapsed between the telling of the lie and its detection. Only 2% of the recalled lies were caught at the time of the telling based on source verbal and nonverbal behaviors. Most were detected after the fact, often much later, and discovery methods often included information from others, physical evidence, and subsequent confessions.

Conclusion: A Preliminary Sketch of Truth-Bias Theory

In this chapter, we offer a view of deception that departs considerably from current theory and conventional takes on the literature. Little emphasis is placed on nonverbal behaviors, leakage, and the mechanisms that produce leakage such as arousal, cognitive effort, guilt, and the like. We even will go so far as to suggest that the 54% detection accuracy finding is limited to only a very specific experimental paradigm that departs in important ways from the ecology of deception in nonresearch settings.

Our view begins with motives. Most people deceive for a reason, and would not deceive if a nondeceptive message would work just as well. The reasons people deceive are generally the same that guide nondeceptive behavior. Deception and truth are both means to goals. Deception is seldom an end in itself.

People presume others deceive for a reason. People project the motives of others, and if there is not an obvious motive for deception, people tend to believe. In contrast, many deception detection experiments do not provide a context for people to project motives, or if they do, motive is a constant so it cannot be used to distinguish between messages.

Deception, as opposed to information management, and relative to truth, is infrequent. We would ask the reader "How many times have you lied today?" We asked a nationwide sample this question, and 60% answered that they had not lied in the last 24 hours. Perhaps there is a reporting bias, but we believe that true deception is very infrequent relative to honest messages, but by no means infrequent in an absolute sense. If only 1% of messages were deceptive, given the sheer quantity of messages we receive, we would likely encounter lies daily. Nevertheless, most of what we encounter is probably honest and lacks deceptive intent.

Everyday deception is not especially arousing, emotionally charged, or cognitively effortful, and people worry little about detection unless the stakes are high. People learn to lie in childhood, and by the time they are teenagers most people can lie seamlessly most of the time. In fact, in the situations where people do lie, deception may often be an easier option than honesty (see McCornack, 1997). Further, at least for everyday deception, between- and within-individual variation in nonverbal behaviors swamps any deception-specific behaviors (see Levine et al., 2005). As a consequence, there are few diagnostically useful behaviors that reliably distinguish everyday deception from nondeceptive behavior. This may or may not be true of high-stakes situations, but those are unusual.

Consequently, additional information (e.g., confessions, prior knowledge inconsistent with a statement) is needed for deception to be detected accurately. And, this is the type of information people most often use outside the lab. As a further consequence, most deception is detected well after the fact by means other than at-the-time communicative performance of the source.

Message recipients do not actively consider the possibility of deception unless suspicion is aroused. A state of nonsuspicion is the default or status quo. Suspicion is most often a function of projecting a motive, information from an outside source, or message inconsistency with prior receiver knowledge. Absent this, there is little suspicion.

Even if suspicious, truth-bias nevertheless generally predominates. Truth-bias is likely to be especially strong in communication with close others, face-to-face (or other interactive) communication, in nonresearch settings, when there is no external information suggestive of deceit, and when an obvious motive for deception is absent.

As a consequence of deception being a low frequency event and message recipients seldom considering it as even a possibility, most people are correct in their assumption of honesty most of the time. This is adaptive and allows communication to function. However, deceptive messages slip through undetected. As a further consequence, at-the-time accuracy is a function of truth-bias and truth-lie base-rate; average accuracy approximates 50%, truth accuracy is higher, and lie accuracy is lower.

References

Bok, S. (1999). *Lying: Moral choice in public and private life*. New York: Vintage Books.

Bond, C. F., & The Global Deception Research Team. (2006). A world of lies. *Journal of Cross-Cultural Psychology, 37*, 60–74.

Bond, C. F., Jr., & DePaulo, B. M. (2006). Accuracy of deception judgments. *Review of Personality and Social Psychology, 10*, 214–234.

Bond, C. F., Jr., & DePaulo, B. M. (2008). Individual differences in judging deception: Accuracy and bias. *Psychological Bulletin, 134*, 477–492.

Buller, D. B., & Burgoon, J. K. (1996). Interpersonal deception theory. *Communication Theory, 6*, 203–242.

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*, 979–995.

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

Ekman, P. (2001). *Telling lies*. New York: Norton.

Frank, M. G., & Feeley, T. H. (2003). To catch a liar: Challenges for research in lie detection training. *Journal of Applied Communication Research, 31*, 58–75

Gilbert, D. T. (1991). How mental systems believe. *American Psychologist, 46*, 107–119.

Gilbert, D. T., Krull, D. S., & Malone, P. S. (1990). Unbelieving the unbelievable: Some problems in the rejection of false information. *Journal of Personality and Social Psychology, 59*, 601–613.

Goffman, E. (1955). On face-work: An analysis of ritual elements in social interaction. *Psychiatry, 18*, 213–231.

Grice, P. (1989). *Studies in the way of words*. Cambridge, MA: Harvard University Press.

Hample, D. (1980). Purpose and effects of lying. *The Southern Speech Communication Journal, 46*, 33–47.

Kim, R. K., & Levine, T. R. (2008). *The effect of suspicion on deception detection accuracy: A reconceptualization and replication of McCornack and Levine (1990)*. Manuscript submitted for publication.

Levine, T. R., Feeley, T., McCornack, S. A., Harms, C., & Hughes, M. (2005). Testing the effects of nonverbal training on deception detection accuracy with the inclusion of a bogus training control group. *Western Journal of Communication, 69*, 203–218.

Levine, T. R., Kim, R. K., & Blair, J. P. (2008). *(In)accuracy at detecting true and false confessions and denials: An initial test of a projected motive model of veracity judgments*. Manuscript submitted for publication.

Levine, T. R., Kim, R. K., & Hamel, L. M. (2008). *People lie for a reason: Three experiments documenting the principle of veracity*. Manuscript submitted for publication.

Levine, T. R., Kim, R. K., Park, H. S., & Hughes, M. (2006). Deception detection accuracy is a predictable linear function of message veracity base-rate: A formal test of Park and Levine's probability model. *Communication Monographs, 73*, 243–260.

Levine, T. R., Park, H. S., & McCornack, S. A. (1999). Accuracy in detecting truths and lies: Documenting the "veracity effect." *Communication Monographs, 66*, 125–144.

McCornack, S. A. (1997). The generation of deceptive messages: Laying the groundwork for a viable theory of interpersonal deception. In J. O. Greene (Ed.). *Message production* (pp. 91–126). Mahwah, NJ: Erlbaum.

McCornack, S. A., & Levine, T. R. (1990). When lovers become leery: The relationship between suspicion and accuracy in detecting deception. *Communication Monographs, 57*, 219–230.

Metts, S. (1989). An exploratory investigation of deception in close relationships. *Journal of Social and Personal Relationships, 6,* 159–179.

Park, E., Levine, T. R., Harms, C., & Ferrerra, M. (2002). Group and individual accuracy in deception detection. *Communication Research Reports, 19,* 99–106.

Park, H. S., & Levine, T. R. (2001). A probability model of accuracy in deception detection experiments. *Communication Monographs, 68,* 201–210.

Park, H. S., Levine, T. R., McCornack, S. A., Morrison, K., & Ferrara, M. (2002). How people really detect lies. *Communication Monographs, 69,* 144–157.

Serota, K., & Levine, T. R. (2008). *The prevalence of deception in American life.* Unpublished manuscript, Michigan State University.

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.

Zuckerman, M., DePaulo, B. M., & Rosenthal, R. (1981). Verbal and nonverbal communication of deception. In L. Berkowitz (Ed.), *Advances in experimental and social psychology* (Vol. 14, pp. 1–59). New York: Academic.

Chapter 3

Gray Area Messages

Gary D. Bond and Lassiter F. Speller

> Falsehoods committed by people who are mistaken or self-deceived are not lies, but literal truths designed to mislead are lies. (DePaulo et al., 2003, p. 74)

Adolescents tend to think about the world in absolute terms: true/false, right/wrong, good/evil (Perry, 1999). We believe that the idealism peculiar to youth springs from dualistic thinking. Like the wily detective in the interrogation room, we can extract the truth after sifting through layers of fabrication. Like Don Quixote, we can venture out and try to right the wrongs of the world. Like Michael the archangel, we can wage a winning battle against evil. When we cross the ambiguous threshold into adulthood, however, we realize that dualistic thinking is, well, sort of naïve. We begin to perceive the world as a conundrum filled more with shades of gray than with neat rows of absolute black and white. When we talk specifically about deception and truth in this chapter, we will assert that those endpoints are not black and white opposites, but are rather crude, arbitrarily-drawn communicative boundaries. Even the definitions of deception and truth carry with them several exceptions, as noted next.

Many definitions of deception include the deliberate intention to mislead others (DePaulo et al., 2003; Ekman, 2001; Ekman & O'Sullivan, 2006), but intentional misleading is sometimes not deception. For example, one might joke or use irony or sarcasm without intentionally misleading another if one believes that the listener knows one's intent (Taylor, Lussier, & Maring, 2003). If the listener has the necessary cognitive capability, coupled with an accurate shared mental model of the conversation, then the messenger may safely assume that the joke or sarcasm will be "got" by the listener. Ekman (2001) and Ekman and O'Sullivan (2006) define "lie" as the deliberate intention to mislead "without prior notification of the target of the lie" (Ekman & O'Sullivan, 2006, p. 674). This definition excludes jokes, irony, and sarcasm, but incorporates "lie" in the definition of the same construct. Transmitting information to a