Communication Yearbook 34

Edited by
Charles T. Salmon

2010

Published Annually for the
International Communication Association

Routledge
Taylor & Francis Group
NEW YORK AND LONDON

First published 2010
by Routledge
270 Madison Avenue, New York, NY 10016

Simultaneously published in the UK
by Routledge
2 Park Square, Milton Park, Abingdon, Oxon OX14 4RN

Routledge is an imprint of the Taylor & Francis Group, an informa business

Typeset in Times by EvS Communication Networx, Inc.
Printed and bound in the United States of America on acid-free paper by Sheridan Books, Inc

ISSN: 0147-4642

ISBN10: 0-415-87857-8 (hbk)
ISBN10: 0-203-84627-3 (ebk)

ISBN13: 978-0-415-87857-9 (hbk)
ISBN13: 978-0-203-84627-8 (ebk)

Contents

The International Communication Association

The International Communication Association (ICA) was formed in 1950, bringing together academics and other professionals whose interests focus on human communication. The Association maintains an active membership of more than 4,000 individuals, of whom some two-thirds teach and conduct research in colleges, universities, and schools around the world. Other members are in government, law, medicine, and other professions. The wide professional and geographic distribution of the membership provides the basic strength of the ICA. The Association serves as a meeting ground for sharing research and useful dialogue about communication interests.

Through its divisions and interest groups, publications, annual conferences, and relations with other associations around the world, ICA promotes the systemic study of communication theories, processes, and skills. In addition to *Communication Yearbook*, ICA publishes the *Journal of Communication*, *Human Communication Research*, *Communication Theory*, *Journal of Computer-Mediated Communication*, *Communication, Culture & Critique*, *A Guide to Publishing in Scholarly Communication Journals*, and the *ICA Newsletter*.

For additional information about the ICA and its activities, visit online at www.icahdq.org or contact Michael L. Haley, Executive Director, International Communication Association, 1500 21st Ave. NW, Washington, DC 20036 USA; phone (202) 202-955-1444; fax 202-955-1448; email ica@icahdq.org.

Editors of the *Communication Yearbook* series:

Volumes 1 and 2, Brent D. Ruben
Volumes 3 and 4, Dan Nimmo
Volumes 5 and 6, Michael Burgoon
Volumes 7 and 8, Robert N. Bostrom
Volumes 9 and 10, Margaret L. McLaughlin
Volumes 11, 12, 13, and 14, James A. Anderson
Volumes 15, 16, and 17, Stanley A. Deetz
Volumes 18, 19, and 20, Brant R. Burleson
Volumes 21, 22, and 23, Michael E. Roloff
Volumes 24, 25, and 26, William B. Gudykunst
Volumes 27, 28, and 29, Pamela J. Kalbfleisch
Volumes 30, 31, 32, and 33, Christina S. Beck
Volume 34, Charles T. Salmon

CHAPTER CONTENTS

2 A Few Transparent Liars

Explaining 54% Accuracy in Deception Detection Experiments

Timothy R. Levine

Michigan State University

Deception detection experiments consistently find that people are statistically significantly, but only slightly, better than chance. The stability of this finding and the lack of variance in judge ability are at odds with current and classic deception theory that explains accuracy in terms of message recipient's ability to spot leaked deception cues. An alternative explanation based on limited variance in message source transparency provides a more coherent account of deception detection findings and has important implications for past and future deception theory.

Introduction

Research consistently finds that people are only slightly better than chance at detecting deception. Meta-analysis of more than 200 experiments finds that people are, on average, 54% accurate when they have a 50–50 chance of being right (Bond & DePaulo, 2006). This level of accuracy is statistically better than could be obtained by chance alone, but it also suggests that people are not much better than a random coin flip at correctly distinguishing honest communication from outright lies.

There are a number of intriguing aspects of the deception detection accuracy literature. First, one of the more curious facets of the literature is the across-study stability of accuracy results. The slightly-better-than-chance accuracy finding appears to be among the most reliable, consistent, and robust finding in all of social science. Ninety-eight percent of all accuracy results fall between 39% and 67%, and more than 90% of studies produce results within 10% of the across-study mean. Figure 2.1, adapted from the results of Bond and DePaulo (2006), visually depicts the across-study consistency of this finding with a frequency distribution of study results. In Figure 2.1, the percent accuracy reported in each study is plotted on the horizontal axis and the number of studies obtaining a given result is presented on the vertical axis. A normal curve is superimposed on the distribution. As the reader can see, previous findings in the literatures are neatly and normally distributed around the across-study grand mean. As those familiar with meta-analysis know, the distributions of

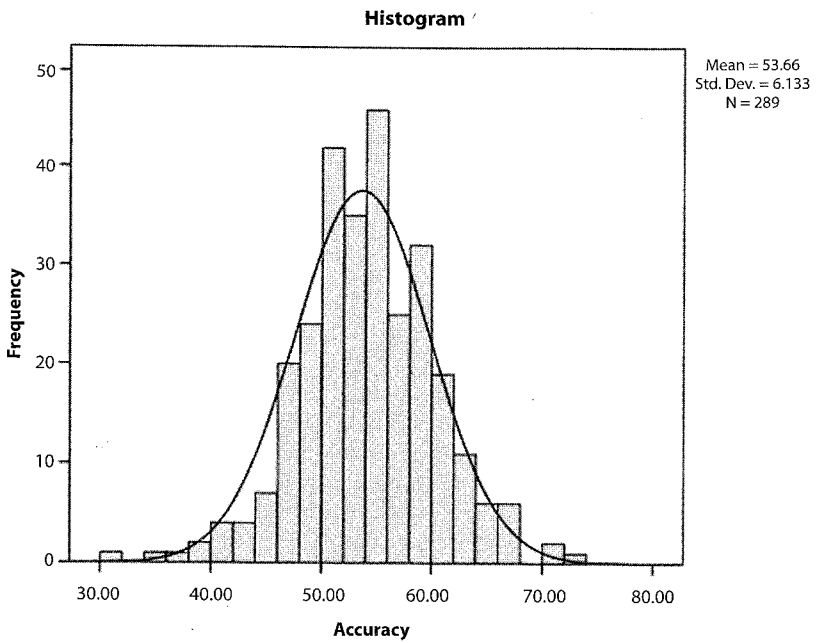

Figure 2.1 Percent accuracy observed in deception detection experiments (adapted from Bond & DePaulo, 2006).

study findings in social scientific literatures are rarely so orderly. Instead, "messy" literatures with multiple moderators and much unresolved heterogeneity are much more typical.

A second, related feature of the literature is that individual-study deviation from the across-study 54% average appears to be a simple function of the number of judgments made in the study (Bond & DePaulo, 2006). Studies involving larger numbers of judgments lawfully produce accuracy levels closer to the across-study mean. This means that study-to-study differences are more a function of mere reliability than substantive factors such as the population that was sampled, the theoretical orientation of the authors, the hypotheses being testing, or the specific independent variables under study. Usually, in social science, such things as hypotheses, theories, participants, and independent variables matter, so it is most curious that the deception detection literature is different from its parent fields of inquiry in these regards.

A third unusual feature of the literature is that small differences routinely produce statistically significant results even with small samples. The difference between the observed 54% accuracy and the 50% accuracy expected by mere chance is not only statistically significant in meta-analysis, but also at the level of the individual study. In social science, it usually takes large sample

sizes to produce small, but statistically significant findings because the standard errors in social data tend to large. Standard errors, in turn, tend to be large when sample sizes are small because meaningful individual differences almost always exist in how humans respond in social situations. In contrast, the tight confidence intervals typical in the deception detection literature result from unusually small standard errors which, in turn, are a function of an unusual degree of homogeneity across judges.

Perhaps what makes the deception detection literature most intriguing, however, is the theoretical puzzle of adequately explaining these findings. If currently accepted theory is correct, there should be more variance in the findings from study to study, findings should not be just a mere function of reliability, there should be more variance in judge ability, and the accuracy ceiling should be higher. Decades of accumulated data systemically and persistently defy theoretical expectations, and this lack of theory-data correspondence begs explanation.

The difference between 54% observed accuracy and the 50% chance accuracy is highly statistically significant and thus clearly reflects some systematic causal mechanism. Veracity judgments are not solely attributable to chance, and consequently there must be a reason for the pattern of results. Simply finding that people do better than chance is not surprising. Most people, social scientists and otherwise, believe that some people are astute social observers and most of the rest of us have at least modest skill at reading others. What is perplexing is why people systematically do better than chance, but seldom much better than chance. Why not higher accuracy? And why is the just over chance accuracy finding so persistent and stable? In short, a viable explanation for this robust finding must not only account for why people systematically exceed chance, but also why people do not exceed chance by very much.[1]

Current theory interprets and explains accuracy results in terms of the individual message receiver's skill, competence, and ability to ferret out lies based on observing the behavior of the message source. The presumption is that these experiments provide a reasonable test of people's ability to distinguish truths from lies and that it is possible for the participants in these experiments to do well at their task. Given this, accuracy findings are usually interpreted as informative about people's ability (or lack there of) to detect deception. The message recipients are, after all, the subjects in these experiments, and consequently their performance is the focus.

Accuracy findings, however, may tell us as much or more about the message sources being judged and the conditions under which they are judged than about the people doing the judging. Along these lines, the current paper explains accuracy in deception detection findings in terms of variance in message source performance rather than variance in judge ability. It is argued that in the context presented by typical deception detection experiments, a "few transparent liars" model provides a better account of detection accuracy

findings than a judge ability model or a source-receiver interaction model. Simply stated, stable and slightly better than chance accuracy is a function of a few transparent liars.

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 standard 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%, systematic increases above chance levels reflect increasing mastery of course content. There is no reason to expect scores systematically below 50% because students do not intentionally seek a low grade by purposely missing questions that they know. Simply put, test scores are most often thought of as telling us about student performance: how much they studied, how effectively they studied, how smart they are, their testing taking ability, and the like.

Such true-false tests provide a reasonable analogy for deception detection experiments. In such experiments, participants are exposed to a series of truths and lies and are asked to distinguish which are which by making dichotomous truth-lie judgments. In these experiments there is an equal probability of a message being honest or deceptive and accuracy is scored as the percent of judgments correct across truths and lies, so the chance rate in these experiments is 50% just like a true-false test (Levine, Park, & McCornack, 1999). Honest sources have no reasons to fool judges in believing that they are really lying, and judges have nothing to gain by intentionally making incorrect judgments, so there is no reason to expect a systematic mechanism that would produce worse than chance performance. Thus, 50% minus chance variation should provide a basement for scores, and systematic variance above a 50% level would reflect the extent to which people accurately distinguish truths from lies in the context of these experiments.

Back to the true-false test analogy: 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 a similar distribution of scores. The average student performs significantly better than chance, but in an absolute sense, fails miserably. Also, the standard deviations are consistently small. Over the four tests, all students tend to hover around the mean that varies little from test to test. What kind of explanation might best explain this pattern of test scores?

Again, this analogy is reasonably comparable to the deception detection literature. Across more than 200 separate experiments spanning several decades, average accuracy hovers around 54% (Bond & DePaulo, 2006). Findings from study to study are remarkably homogeneous and standard errors both between and across studies are small. A number of statistically significant moderators

exist, but the impact of known moderators is small in an absolute sense and the results of most studies fall within plus or minus 10% regardless of the particular independent variables included in the study, the communication medium involved, who the research participants are, and the characteristics of the deceptive messages judged (Bond & DePaulo, 2006).

A test-taker explanation for this outcome might hold that the students studied only a little. They learned just enough to do better than mere chance but not nearly enough to do well. For example, the students may have learned and remembered 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 producing a 55% on the exam (10 + 45 = 55). So, a "little bit of knowledge" model can explain the outcome. But, it seems very odd that all the students would only know a little on all four tests. Usually, some students do better than other students. The lack of variance is odd indeed.

An alternative explanation is that the scores tell us more about the test than the test takers. Imagine a true-false test with 90 questions that are impossibly difficult necessitating guessing by everyone and ten questions that are so easy that everyone gets them correct whether or not they took the class, studied, etc. This test tells us nothing at all about what was learned, and it does not matter who takes the test. Regardless of knowledge, motivation, and ability, everyone does pretty much the same. People consistently do better than 50%, but no one does much better. Such a test always produces an average of 55% plus or minus chance variation. 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 so that little variance in ability is observed.

Explaining 54% Accuracy

Conventional Wisdom Regarding Reasons for Accuracy and Inaccuracy

The current thinking about deception detection has largely evolved from Ekman and Friesen's (1969) idea of "leakage." The original idea was that there are emotional correlates of deception, emotions are conveyed nonverbally, and emotional expression is not entirely under conscious control. Compared to the honest message source, deceivers are apt to experience guilt, fear of detection, and perhaps other emotions created by the act of lying. Deceivers try to control behavioral displays so as not to give themselves away, but cues associated with deception-linked emotion leak out anyway, often through nonverbal channels that are more difficult to control. Consequently, a message recipient who is actively looking for the right leaked cues should be fairly adept at distinguishing truth from lie.

The leakage idea was expanded upon by Zuckerman, DePaulo, and Rosenthal's (1981) four-factor theory. The four factor framework specifies four

internal psychological states that differentiate truths and lies. These four include emotions, arousal, cognitive effort, and over control. Relative to truth-tellers, liars are more likely to experience greater levels of arousal, emotions like fear and guilt, greater cognitive effort, and more effort to control of non-verbal displays. Because each of these internal states is thought to be associated with specific nonverbal behaviors, clues to deception are leaked nonverbally. For example, the increased cognitive effort associated with lying is thought to lead to an increased number of speech errors and longer response latencies that signal deceit.

A contemporary iteration of this thinking is reflected in Interpersonal Deception Theory (IDT) (Burgoon & Buller, 1996). Liars strategically present themselves as honest, but non-strategically leak deception cues. Message receivers pick up on these cues and become suspicious. Liars, however, pick up on leaked suspicion, and strategically adapt. So do receivers. Net accuracy depends on the liar's encoding skill relative to the receiver's decoding skill and how the interaction progresses dynamically over time. IDT, therefore, is a source-receiver interaction model that rests on variance in both sender and receiver ability. Both senders and receivers are leaky, both senders and receivers are adept at spotting and correctly interpreting leakage, and both senders and receivers strategically use the information gained from leakage to adjust their own behavioral performance.

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

The accepted reasons why people are far from perfect is that (a) there is an imperfect link between any given behavioral display and veracity, and (b) lay people rely on cues that lack diagnostic utility in addition to, or instead of, authentic cues. That is, there are no perfectly reliable deception cues and people often look for the wrong things. Consistent with this latter point, research finds that the most common belief about deception is that liars avoid eye contact (Bond & The Global Deception Research Team, 2006) whereas meta-analysis finds no link between gaze and honesty (DePaulo et al., 2003; Sporer & Schwandt, 2007). IDT further adds that liars engage in strategic countermoves, constantly adapting their performance to appear more honest and misleading less skilled judges (Burgoon & Buller, 1996). The net result is above chance accuracy that is far from perfect.

There are many 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 that does so perfectly (DePaulo et al., 2003). Further, there is a less than perfect but non-zero correspondence

between what people look for when detecting deception and the behaviors that have actual diagnostic utility (Zuckerman et al., 1981). Thus, above chance but less than ideal accuracy makes much sense from this perspective, and there is a wealth of supportive findings that can be cited as evidence in favor of this stance.

Not all findings in the literature, however, fall neatly in line with the ability-to-spot-leakage account. First, efforts to enhance accuracy through non-verbal training have failed to document much in the way of improvement. Meta-analysis finds only marginal improvements from non-verbal training (Frank & Feeley, 2003), and studies offering additional controls find even more meager results (Levine, Feeley, McCornack, Harms, & Hughes, 2005). If low accuracy findings stemmed simply from looking for the wrong cues, one would expect much better results from training studies. Further, the training approaches that appear most effective are not based on nonverbal leakage (Blair, Levine, & Shaw, in press; Hartwig, Granhag, Stromwall, & Kronkvist, 2006).

A 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. People also tend to self-select into professions where they have aptitude. Thus, police, military interrogators, customs officials, etc. should be better at detecting deception than the average college sophomore. Yet, meta-analysis suggests that this is not the case. Neither age, nor expertise, nor specific occupation meaningfully impact accuracy (Aamodt & Custer, 2006; Bond & DePaulo, 2006).

Further, both judge ability accounts and source-receiver interaction explanations suggest substantial within and between study variance in accuracy. Ability and relative ability should vary from person to person and situation to situation. Such variance should not only stem from training and experience, but also from other individual differences such as self monitoring, perspective taking, emotional intelligence and the like. If some substantial proportion of message sources leak subtle cues to deception, more socially adroit people should be more accurate than their socially oblivious counterparts producing variance, increasing standard errors, and making findings generally more variable than they are. Yet, meta-analysis again shows that cognitive and social abilities have little impact on accuracy (Aamodt & Custer, 2006). The within- and across-study consistency of findings is difficult to reconcile with the leakage-based, perceptive judge perspective. In short, ability-based explanations can be stretched to explain the average effects but not the lack of variance apparent in the existing data.

Advocates of sender leakage and judge ability perspectives frequently blame methodological scapegoats for the failure to achieve theoretically predicted variance in accuracy. The three most frequently blamed culprits are the use of low stakes lies, sanctioned lies, and dichotomous (rather than continuous) deception judgments. Although it is axiomatic from within the logic

of the leakage perspective that sanctioning and stakes make a substantial difference, the one published study directly testing cues from unsanctioned lies against the behaviors observed during truth and sanctioned lying found little evidence that sanctioning made a meaningful difference (Feeley & deTurck, 1998). In fact, the behavior for which the largest differences were observed produced differences that were in the direction opposite to that predicted by leakage and four-factor theories. Sanctioned lies were characterized by more speech errors and unsanctioned lies by fewer speech errors than the truthful control. Furthermore, detection accuracy studies using unsanctioned, higher stakes lies report results that are no different from the literature as a whole (e.g., Levine, Kim, Park, & Hughes, 2006). Studies that produce results falling farther for the 54% average are not those using better, presumably more valid stimulus materials, but are instead those studies involving few judgments, producing less stable results (Bond & DePaulo, 2006). As the number of judgments made in a study increases, estimates approach the across-study average. If substantial moderators existed, methodological or otherwise, this finding simply would not be the case. Finally, once converted to a common metric, meta-analysis shows that studies using continuous scaling find comparable results to dichotomously scored accuracy studies (Bond & DePaulo, 2006). Thus, not only do the usual methodological suspects appear innocent of any substantial sabotage, their failure to make the predicted differences offers further evidence inconsistent with the prevailing theoretical views.

The Psychometrics of Veracity Judgments

Finally, and most important, findings from the most recent meta-analysis are decisively incongruent with the judge ability perspective. Bond and DePaulo (2008) looked at the variance in accuracy judgments rather than just average accuracy levels. They decomposed veracity judgments into four components: demeanor, truth-bias, transparency, and ability.[2] *Demeanor* is the tendency of a person being judged to appear honest (or deceptive) independent of whether or not the person is lying. Variance in demeanor indicates that some people are 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 leaky people are when lying and how sincere they are when they are telling the truth. People who are transparent leak the fact that they are lying and it is relatively easier to distinguish when they are lying from when they are 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, whereas truth-bias and ability reflect variance in the message judge. Further, demeanor and truth-bias reflect variance in *bias*; that is, they are tendencies to believe (or not) that

are independent of actual veracity, whereas transparency and ability reflect variance in *openness or skill* in presenting or discriminating between honest and deceptive messages. Because deception detection experiments average across an equal number of truths and lies, demeanor and truth-bias do not affect overall accuracy (Levine et al., 2006) whereas transparency and ability do impact overall accuracy. Thus, systematically above-chance accuracy stems from transparency, ability, or both.

Bond and DePaulo (2008) found that variance in demeanor is large, both in an absolute sense and relative to the other three 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 individual differences in truth-bias and transparency, with these differences being much smaller than the variance in demeanor, but much larger than the variance in ability. Individual differences in ability contribute very little to overall accuracy. Comparatively, the variance in demeanor was 200 times as large as the variance in judge ability. 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 ability. This explains why accuracy values within studies are so stable. The lack of individual differences in judge ability leads to small standard errors and even small differences can be statistically significant. This finding also renders theoretical accounts of deception detection based on either judge ability or source-receiver interactions impotent. Variables that do not vary cannot co-vary, and variables to that do not co-vary do not explain anything. The finding of only trivial variance in judge ability means that deception detection accuracy cannot be explained in terms of judge ability or source-receiver interactions.

A Few Transparent Liars

The thesis of the current essay is that deception detection accuracy findings are the result of a few transparent senders. Important to the argument is Bond and DePaulo's (2008) finding that the variance in sender transparency in detection accuracy studies is massively larger than the variance in judge ability, although variance in transparency is not large in an absolute sense. This suggests that above-chance accuracy is not a function of individual differences in judge's ability to recognize leakage, but instead a function of individual differences in how much is leaked independent of who is doing the judging. In terms of the previous test taking analogy, the variance in test results is attributable to test question difficulty rather than student knowledge or competence. Returning to deception, the current perspective holds that most people can lie seamlessly without diagnostically useful leakage, but a few people tend to give themselves away and consequently are systemically detected by most observers. There are enough transparent senders to produce accuracy rates that statistically exceed chance level, but too few (under the

conditions in most deception detection experiments) to allow for accuracy rates that exceed chance by much. Those few transparent senders are seen as transparent by almost everyone so there is much more variance in transparency than ability. Hence, accuracy findings tell us more about the sender than the detector. The slightly-better-than-chance finding is a function of a relatively few transparent liars.

The Few Transparent Liars explanation fits the literature more neatly and cleanly than theories predicated on the ability to spot leakage. A Few Transparent Liars accounts both for the mean levels of accuracy observed in deception detection experiments and for the lack of variability observed around mean levels (see Bond & DePaulo, 2006, 2008). It explains the lack of large effects for professional experience (see Bond & DePaulo, 2006) and nonverbal training (see Frank & Feeley, 2003). It explains why cue studies fail to find evidence of substantial leakage (see DePaulo et al., 2003; Sporer & Schwandt, 2006, 2007). In short, the few transparent liars explanation makes previously puzzling anomalous findings cohere.

The current view does not deny the existence of leakage. Instead, leakage in some form or another is what makes some liars leaky. However, slightly above chance accuracy means that although leakage sometimes happens, it does not characterize the lies told by most of the people most of the time, at least in quantities that would be diagnostically useful even to a trained and experienced eye. Thus, one key difference between the current view and traditional leakage-based theory is the prevalence and centrality of leakage. Traditional leakage perspectives portray leakage as typical of lies, or at least unsanctioned, high-stakes lies. Leakage is presumed to characterize deceptive communication such that it can be used by a skilled and knowledgeable judge to distinguish truths from lies. This traditional view has been recently re-articulated by O'Sullivan (2009):

> In the previous chapter, Mark Frank reviewed how the disruptions in feeling and thinking caused by lying can result in observable clues that lie detectors could use to detect deceit. Quite surprisingly, few people seem to use these clues. (p. 74)

Leakage perspectives hold that deception cues exist, but low accuracy stems from a failure of judges to correctly use the information available to them. In contrast to leakage theory where leakage is seen as a given and judge ability as variable, the Few Transparent Liars Model views leakage as variable and judge ability as near constant.

Consistent with this perspective, recent meta-analyses of nonverbal deception cues find small and inconsistent effects for objective nonverbal deception cues (DePaulo et al., 2003; Sporer & Schwandt, 2006; 2007). Sporer and Schwandt (2007), for example, note that of the 154 tests of specific nonverbal behaviors predicted by four-factor theory that are reported in the literature,

only 28 (18%) are statistically significant. Leakage findings thus occur at rates better than chance, but non-supportive results out-number supportive findings more than 5 to 1. Thus, leakage can very plausibly account for the statistically significant difference between 50% and 54% in accuracy, but not much more than that. Forty years of accumulated evidence shows that leakage-based accounts have little explanatory power, and consequently research findings do not support the prominent place nonverbal leakage has in deception theory.

The Few Transparent Liars Model is notably inconsistent with source-receiver interaction models such as IDT as accounts of deception detection accuracy. IDT views deception detection as a function of sender *and* receiver skill. Senders and receivers are theoretically specified to be sensitive to each other's leakage, and deception success or failure depends on the skills of each relative to the other. In the current perspective, judge ability does not meaningfully vary. Since it does not vary, it cannot co-vary, and it cannot *statistically interact* with receiver skill. This view does not deny that senders and receivers interact in a more general communication sense, but it does specify that deception detection accuracy is typically a function of limited sender variance and not variance in message judge ability or the statistical interaction sender-judge ability.

Several important qualifications and clarifications to the current argument need to be made explicit. First, in the test-taking analogy and so far in this essay, transparency was described as an all-or-nothing phenomenon for the sake of simplicity. Few, if any, people, however, are likely to be perfectly transparent. The point of the model is that within the typical deception detection experiment, some small proportion of people is at least somewhat transparent while the majority of people are not very transparent. The model need not, and does not, literally presume all or nothing transparency. Instead, it is based on more variance in sender transparency than judge ability.

A second qualification relates to the generality or boundary conditions of the model. The few transparent liars model applies to the context and ecology of typical deception detection experiment in the leakage tradition. Most previous deception detection experiments were designed with the ability to spot leakage perspective firmly ingrained, and hence these experimental designs test the ability to spot deception-produced leakage. Other types of information are typically precluded by design (Park, Levine, McCornack, Morrison, & Ferrara, 2002). For example, one way people are likely to assess the probability of deception is by considering if a person has reason to lie. That is, people project motive (Levine, Kim, & Blair, 2010). In deception detection experiments, however, no motive for lying is usually apparent. Or if there is a motive, it is constant across all sources. Thus, consideration of motive is of no use in determining truth from fabrication. Deception detection experiments also preclude useful prior factual knowledge, diagnostically useful information from informants, access to physical evidence,

use of context information, use of communication content in context, the strategic use of evidence, and knowledge of confessions. Research indicates that outside the lab, lies are usually detected well after the fact on the basis of information other than leaked behavioral cues, such as those listed in the previous sentence (Park et al., 2002). Thus, the ecology of the deception detection experiment makes sense from an ability to spot leakage orientation, but makes less sense if the ability to spot leakage is largely irrelevant to how lies are actually detected. Consistent with this, recent experiments have reported higher levels of accuracy, but these elevated accuracy findings are not obtained from passive observations of deception generated nonverbal leakage. Instead, recent experiments reporting impressive accuracy findings are based on the strategic use of evidence (Granhag, Stromwal, & Hartwig, 2007; Hartwig et al., 2006), strategic questioning designed to increase transparency (Levine & Blair, 2010; Levine, Shaw, & Shulman, 2010), or content in context (Blair et al., in press) approaches.

It does not follow, however, that the leakage-based deception detection literature is therefore uninformative. To the contrary, these experiments provide a very consistent picture of the extent to which people can distinguish truths from lies in near real time based only on passive observation of sender performance. The contention here is that, except for a few inept liars, people cannot and do not detect deception under such conditions. These findings are theoretically important because accepted theory clearly predicts otherwise. Thus, alternative theory is needed to which the current perspective will contribute. These findings also have important practical implications. Knowing what does not work is highly valuable in preventing wasted effort, discouraging a false sense confidence, and in prompting new directions.

A third qualification is that while "a few transparent liars" provides a catchy phrase for the current model, the reader is reminded that transparency applies to truths and lies. Consequently, slightly-above-chance accuracy is specified to result from a few transparent *senders*. Traditional leakage is something liars do, but the idea of transparency applies to honest and deceptive senders alike. Both leaky lies and exculpatory statements enhance transparency.

The reader is cautioned that although accuracy in deception detection experiments is typically better than chance for honest sources and below chance for liars (the "veracity effect," Levine et al., 1999), it does not follow that above chance accuracy results only from transparent honest senders. The veracity effect stems from the tendency to believe a source independent of message veracity (Levine et al., 1999). The veracity effect and base-rate effects can be explained well by "mere chance" models (Levine et al., 2006). The current transparency model, in contrast, was created to explain systematic improvements over and above chance. Transparent sources are believed when they are honest and doubted when they are lying. Thus, judgments based on sender transparency, unlike sender demeanor and judge truth-bias, are contingent on message veracity. So, the more a given judge is truth-biased, the more that judge will get honest messages correct and lies wrong creating the well documented difference between truth and lie accuracy. As long as judges view an equal number of truths and lies, however, truth-bias does not affect average accuracy (Park & Levine, 2001; Levine et al., 2006). Gains in truth accuracy and degradations in lie accuracy average out. Transparency, on the other hand, leads to improvement over chance. In the Park-Levine model, truth-bias affects the slope of the regression of truth-lie base-rates onto total accuracy. Transparency is hypothesized to affect the y-intercept in Park-Levine model. A complete lack of transparency would produce 50% accuracy at the 50–50 truth-lie base-rate regardless of truth-bias. Because accuracy is 54% at the 50–50 base-rate, a little (but not zero) transparency is posited here. Applied to Park-Levine, the current thesis is that it is sender transparency, not judge ability, which accounts for the y-intercept in the Park-Levine model. Thus, the current model is consistent with Park and Levine (2001) but the focus is different. A few transparent liars can be thought of a modular add-on to the veracity effect and Park-Levine models, all of which will fit under the more general logic of the forthcoming Truth-Bias Theory.

Methodological Implications

Testing the few transparent liars logic in future research requires, at minimum, a research design that allows for sender and judge variance to be partitioned. This, in turn, requires that multiple judges assess multiple senders with senders as a repeated factor in the design. In such designs, accuracy can be scored both for individual judges and for individual senders, and judge variance can be compared to sender variance. In a first study based on the current logic, Levine, Shaw, & Shulman (2010) used this sort of design with 128 participants each judging 44 different senders who either were deceptive or honest. Consistent with predictions, variance in sender transparency was ten times larger than variance in judge ability, and a strategic questioning induction designed to alter detection accuracy increased variance in transparency but not variance in judge ability.

A stronger test of the few transparent liars idea, however, would not only partition sender and judge variance, but also fully distinguish between transparency, demeanor, ability, and truth-bias variance. Research accomplishing this would require a design like that described above except that additionally each sender would need to produce both honest and deceptive messages. Such a design would allow research to partition variance into sender effect, judge effect, and sender-judge interaction effects.

Some Explanatory Speculation

Presuming that the Few Transparent Liars Model is accurate, one interesting question is why so few people are transparent. Although admittedly speculative,

perhaps the most plausible explanation is simply practice. Humans develop the cognitive abilities needed for deception in early childhood, between the ages of 3 and 5 (Peskin, 1992). The lies young children tell are likely to be highly transparent, and those transparent lies are likely to be punished or at least discouraged. Rather than abstaining from lies, however, children gradually learn to avoid detection through trial and error, and improve (i.e., lower transparency) with repeated practice over the years. By early adulthood, most people have had considerable practice in lying so it makes sense that they would be good at telling lies.

Why then are a few people transparent? One possibility is that since some people lie more than others (Serota, Levine, & Boster, 2010), some people are more practiced than others. A second explanation is that similar to most other skill sets, there are individual differences in natural ability and thus some people benefit more from practice than others. Further, these two explanations may interact so that people with less aptitude practice less and this exacerbates individual differences in transparency over time. Third, transparency is likely not limited to a stable, trait-like individual difference. Instead, transparency also probably varies within individuals as a result of situational and chance factors. Even the most seamless liars may slip up and reveal their lies on occasion. The net result of these three forces produces some small percentage of lies that are transparent and a large percentage that are not.

The Few Transparent Liars Model is not only predicated on some small variance in transparency, but also on a corresponding lack of variance in judge ability. Evolutionary perspectives on deception often presume that because human ability to deceive is highly evolved, the ability to detect must also have evolved. As Smith (1987) puts it, there must exist "a coevolutionary struggle between the deceiver and the deceived. There is an evolutionary arms race to develop better deception tactics and subsequently the pressure to develop better deception detection devices" (p. 59). Such a perspective makes the few transparent liar model seemingly at odds with evolutionary perspectives. Yet, this need not be the case. Whether the human ability to deceive evolved due to providing some unique evolutionary advantage, or as "a parasitism of the preexisting system for correct communication" (Smith, p. 59), it may well be the case that human susceptibility to being duped reflects a highly adaptive trade-off that has and does serve our species well.

As Gilbert (1991) argued, a cognitive system that evaluated all incoming information for veracity would be radically less efficient than a system that presumed truth as a default, then subsequently "unbelieved" information that was later discovered to be false. From an evolution of communication perspective, "it is well-known that group living, that has characterized the course of development of genus *Homo*, guarantees a series of noteworthy advantages" (Adenzato & Adrito, 1999, p. 10), and the development of language and communication greatly facilitates group functioning (Dunbar, 1988). It

is the contention here that efficient communication requires a presumption of honesty. Although this presumption of honesty enables social exploitation and deception, it is most plausible that the net evolutionary advantage of efficient communication and social functioning greatly outweigh the occasional costs of being duped. Further it may be vastly more efficient to have social systems that discourage deceit than to either evolve real time deception detection ability or forsake efficient communication and cognitive function.

To the extent that this current evolutionary take on deception and deception detection has merit, it is predicted that parents everywhere will teach their children that lying is wrong and that all human cultures and religions will develop prohibitions against deception, at least within the salient in-group. Perhaps this is why the "a liar won't look you in the eye" belief transcends culture (Bond & The Global Deception Research Team, 2006). A lack of eye contact may not be associated with actual deception and consequently has no utility as a deception detection tool, but gaze aversion is associated with shame. All human cultures benefit if deception is discouraged, at least within the in-group, and propagating the belief that deception is a shameful practice serves this end. In any case, evolving a finely toned cognitive system adept at spotting leakage need not be an evolutionary mandate just because we get duped once in a while. Readers skeptical of the profound gullibility that characterizes the human species are directed to Farquhar (2005). Readers preferring experimental evidence may prefer Levine et al. (1999, 2006) for evidence of the strong impact of truth-bias in deception detection research findings.

Some Further Implications

To the extent to which better-than-chance accuracy rests on sender transparency, an interesting question is if sender transparency is a stable, trait-like individual difference, a fleeting behavioral response that is time and situation specific, or both. While either of these possibilities is consistent with the current thinking, it is perhaps most likely that transparency has both stable and situational properties. Some people may be chronically transparent while others may exhibit transparency in one instance but not another. Because there is limited variance transparency, it would likely to difficult to track longitudinally. But, the possibility that transparency might, in part, be situational suggests that transparency might be something that could be prompted or triggered. Consistent with this speculation, Levine, Shaw, & Shulman (2010) report success in increasing transparency with strategic questioning.

To the extent that transparency is also, in part, attributable to a few chronically leaky liars, then an implication of a few leaky liars is that chronically leaky liars probably realize that they are poor liars and consequently try not to lie. Whereas research indicates that people, on average, tell one to two lies per day (DePaulo, Kashy, Kirkendol, Wyer, & Epstein, 1996), the frequency of lie telling is not normally distributed across the population (Serota et al., 2010).

Because prevalence data are highly positively skewed, most lies are told by a few prolific liars and the modal number of lies per day is zero (Serota et al., 2010). It is likely that prolific liars are low on transparency and high on honest demeanor making them not only prolific but also highly successful. Thus, the type of people who create the better-than-chance accuracy findings in the lab (were random assignment of participations to truth-lie experimental condition predominates) may self-select out outside the lab and simply avoid putting themselves in situations where they need to lie. Alternatively, the people who keep accuracy scores near chance and inflate truth bias may be those most likely to lie outside the lab in everyday life.

This prediction was recently tested by Levine and Blair (2010) and Levine, Shaw, and Shulman (2010). Senders self-selected into deception or honest conditions, and under indirect questioning judges rated the demeanors of the dishonest sources as significantly more honest than the honest senders. This suggests that unlike most deception detection experiments, if liar and truth-tellers decide for themselves rather or not to lie, the ability to distinguish truth for lie may drop to below chance levels because the few transparent liars don't lie, and those with the most honest demeanors do lie.

Second, if the current thinking is correct, one way to obtain high accuracy using a traditional deception detection experimental design is for the researcher to purposely select leaky liars for inclusion in the experiment. If a few transparent liars typically exist, they could be selected though pre-screening and the researcher could stack the deck to obtain any specific accuracy result as a function of the ratio to transparent to non-transparent liars used. If a large number of sources were screened, and if only transparent liars were selected, very high accuracy will result.

Further, if mean accuracy could be experimentally altered by increasing transparency, the thinking presented in this paper predicts that individual differences in ability might emerge. In the task-taking analogy, reducing test item difficulty would allow variance in student ability to emerge. That is, if the test questions were not impossibly difficult, then students who studied well would do better than other students. Applied to deception, a substantial increase in transparency would allow expertise effects to immerge. Preliminary support for this reasoning was obtained by Levine and Blair (2010).

A third implication is that if poor accuracy is a result of a few transparent liars, then deception accuracy might be enhanced by prompting additional variance in transparency. Levine, Shaw, and Shulman (2010) attempted to increase sender transparency with strategic and direct interrogative questioning designed to challenge liars. Detection accuracy was 68% when sources were subjected to direct, strategic questioning compared to 44% accuracy under indirect background questioning. More importantly, the improvement in accuracy under direct questioning as associated with a 50% gain in variance in transparency but less than a 10% gain in judge ability variance. Further still, in the direct questioning condition were accuracy exceeded chance by a substantial margin, the variance in sender transparency was 15 times larger than the variance in judge ability. These findings suggest that it is transparency variance that explains accuracy and those strategies that increase transparency increase accuracy independent of judge ability.

Finally, although the few transparent liars model was created to explain the results of deception detection experiments, the model has implications for deception detection outside the research context. The model suggests that the passive observation of nonverbal behavior has little merit as a deception detection technique. If most people are not very leaky, watching for leakage makes little sense. This does not mean, however, that deception detection efforts are pointless. Instead, it suggests that deception detection practitioners need to rely on techniques other than the passive observation of nonverbal leakage. For example, approaches based on the strategic use of evidence (Granhag et al., 2007) and content in context (Blair et al., in press) have been producing accuracy rates above 70% in recent studies.

Summary and Conclusion

Past and current deception theory explains deception detection accuracy in terms of people's ability to recognize and interpret leaked behavioral cues indicative of deceit. Accepted theory specifies that people leak diagnostically useful behavioral clues that can used to distinguish truths from lies. From this traditional perspective, people attentive and sensitive to other's behavior can spot a liar. This view is predicated on some constancy in sender transparency but variance in judge ability.

Research on deception detection suggests people are poor lie detectors, at least when all they have to go on is the nonverbal and linguistic behavior of potential liars. Whereas people typically do better than chance, they seldom do much better than chance. Across studies, average accuracy hovers around 54%, and most results fall neatly within 10% of this average regardless of the features of the truths and lies judged, who the research participants are, and the specific independent variables under consideration. These findings present a perplexing puzzle for current theory because if current theories have verisimilitude, then research findings should be more variable and have a higher ceiling. Simply put, if clues to deception are leaked, people should, at least under favorable conditions, be able to spot leakage cues and make accurate judgments.

One often invoked explanation for the poorer than expected performance in accuracy experiments is that while authentic cues exist, people simply rely on the wrong cues. If this were the case, however, training would substantially improve accuracy. It does not. Another explanation (e.g., O'Sullivan, 2008) blames the use of student samples, arguing that adults in certain professions can spot a liar. Again, this is simply not the case. Age, education, and profession have little impact. At this point, the failure of the data to coincide with

theoretical expectations is attributed to methodological limitation. Yet, once again, studies with these limitations removed are no more supportive of leakage theory than the literature as a whole. All this suggests that it is time to question the theory rather than the data.

This chapter explains slightly above chance accuracy in deception detection experiments by the existence of a few transparent senders. The existence of leakage is not denied, but leakage is presumed to be the exception rather than the rule. There are enough leaky liars so that people can do a little better than chance, but not enough to allow for much better than chance performance. If approximately 10% of liars were transparent, then 55% accuracy is both expected and understandable. Opposite from leakage-based theories, the current view is predicts more variance in sender transparency than judge ability.

Part of the theoretical beauty of this argument lies in its ability to reconcile findings in the literature that otherwise seem odd and anomalous. Specifically, the few transparent liars explanation makes sense out of findings that are inconsistent with current theory based on individual differences in the ability to spot pervasive leakage. Not only does the few transparent liars explanation account for slightly better than chance accuracy, it is also accounts for the across-study stability of accuracy findings, the sender and receiver variance findings, the lack of individual differences in accuracy, the lack of effects for nonverbal training, sanctioning, and stakes, and the findings from meta-analyses of deception cue studies. In short, the few transparent liars explanation provides a more coherent account of deception findings than can be obtained with currently accepted theory.

The current model also leads to the derivation of interesting new predictions. For example, if a lack of naturally occurring transparency explains poor accuracy, and if transparency has a situational component, then deception detection accuracy might be improved by strategies designed to enhance sender transparency. This could be done with strategic questioning designed to make liars more leaky, with questioning that allows honest sources the opportunity to provide exculpatory answers, or, ideally, both. As a second example, if transparency also has a trait-like component, then below chance accuracy might be expected in certain situations where people have the choice to lie or not because transparent liars will not lie, and those who do chose to lie will be those with honest demeanors. Thus, the Few Transparent Liars model provides a new conceptual framework that both reconciles existing findings and generates exciting new avenues for future research.

Author's Note

This chapter was completed with support from the *National Science Foundation* (SBE0725685). Correspondence can be addressed to Tim Levine, Department of Communication, Michigan State University, East Lansing, MI, 48824, levinet@msu.edu.

Notes

1. Since the publication of the Bond and DePaulo (2006) meta-analysis, there have emerged a handful of recent studies which report levels of accuracy that are well above chance (e.g., Blair et al., in press; Levine & Blair, 2010; Granhag et al., 2007). These studies involved techniques that rest on strategic question asking, reliance on message content, and reliance on prior knowledge instead of traditional nonverbal leakage. The implications of these recent findings will be discussed later in this chapter.
2. Three of these four labels were adopted by the current author and not Bond and DePaulo (2008). These labels were chosen both because they seem more intuitively descriptive, and because they are more consistent with conventional usage in the previous literature.

References

Aamodt, M. G., & Custer, H. (2006). Who can best catch a liar? *Forensic Examiner, 15*, 6–11.

Adenzato, M., & Ardito, R. B. (1999). The role of theory of mind and deontic reasoning in the evolution of deception. In M. Hahn & S. C. Stoness (Eds.), *Proceedings of the Twenty First Conference of the Cognitive Science Society* (pp. 7–12). Mahwah, NJ: Erlbaum.

Blair, J. P., Levine, T. R., & Shaw, A. J. (in press). Content in context improves deception detection accuracy. *Human Communication Research.*

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.

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

Burgoon, J. K., & Buller, D. B. (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.

Dunbar, R. (1988). Theory of mind and the evolution of language. In J. R. Hurtford, M Studdert-Kennedy, & C. Knight (Eds.,) *Approaches to the evolution of language* (pp. 92–110). Cambridge: Cambridge University Press.

Ekman, P., & Friesen, W. V. (1969). Nonverbal leakage and clues to deception. *Psychiatry, 32*, 88–106.

Farquhar, M. (2005). *A treasury of deception: Liars, misleaders, hoodwinkers, and*

the extraordinary true stories of history's greatest hoaxes, fakes and frauds. New York: Penguin.

Feeley, T. H., & deTurck, M. A. (1998). The behavioral correlates of sanctioned and unsanctioned deceptive communication. *Journal of Nonverbal Behavior, 22*, 189–204.

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.

Granhag, P. A., Stromwal, L. A., & Hartwig, M. (2007). The Sue technique: The way to interview to detection deception. *Forensic Update, 88*, 25–29.

Hartwig, M., Granhag, P. A., Stromwall, L. A., & Kronkvist, O. (2006). Strategic use of evidence during police interviews: When training to detect deception works. *Law and Human Behavior, 30*, 603–619.

Levine, T. R., & Blair, J. P. (2010). *Questioning strategies, diagnostic utility, and expertise interactions in deception detection*. 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. (2010). (In)accuracy at detecting true and false confessions and denials: An initial test of a projected motive model of veracity judgments. *Human Communication Research, 36*, 81–101.

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.

Levine, T. R., Shaw, A., & Shulman, H. (2010). Increasing deception detection accuracy with strategic questioning. *Human Communication Research, 36*, 216–231.

O'Sullivan, M. (2008). Home runs and humbugs: Comments on Bond and DePaulo (2008). *Psychological Bulletin, 134*, 493–497.

O'Sullivan, M. (2009). Why most people parse palters, fibs, lies, whoppers, and other deception poorly. In B. Harrington (Ed.), *Deception: From ancient empires to Internet dating* (pp. 74–94). Stanford, CA: Stanford University Press.

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.

Peskin, J. (1992). Ruse and representation: On children's ability to conceal information. *Developmental Psychology, 28*, 84–89.

Serota, K. B., Levine, T. R., & Boster, F. J. (2010). The prevalence of lying in America: Three studies of reported deception. *Human Communication Research, 36*, 1–25.

Sporer, S. L., & Schwandt, B. (2006). Paraverbal indicators of deception: A meta-analytic synthesis. *Applied Cognitive Psychology, 20*, 421–446.

Sporer, S. L., & Schwandt, B. (2007). Moderators of nonverbal indicators of deception: A meta-analytic synthesis. *Psychology, Public Policy, and Law, 13*, 1–34.

Smith, E. O. (1987). Deception and evolutionary biology. *Cultural Anthropology, 2*, 50–64.

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