

## The Illusion of Transparency: Biased Assessments of Others' Ability to Read One's Emotional States

Thomas Gilovich  
Cornell University

Kenneth Savitsky  
Williams College

Victoria Husted Medvec  
Northwestern University

Three sets of studies provide evidence for an *illusion of transparency*, or a tendency for people to overestimate the extent to which others can discern their internal states. People often mistakenly believe that their internal states "leak out" more than they really do. The authors attribute this bias to a tendency for people to adjust insufficiently from the "anchor" of their own phenomenological experience when attempting to take another's perspective. Evidence for this illusion is provided by showing that liars overestimate the detectability of their lies (Studies 1a, 1b, and 1c) and that people believe their feelings of disgust are more apparent than they actually are (Studies 2a and 2b). A final pair of experiments (Studies 3a and 3b) explores the implications of the illusion of transparency for people's reluctance to intervene in emergencies. All 3 sets of studies also provide evidence consistent with the proposed anchoring and adjustment interpretation.

Fans of Edgar Allan Poe will recall that the key passage in *The Tell-Tale Heart* is one in which the protagonist does his best to play it cool during a conversation with three police officers. It is a performance made more difficult by the fact that the officers happen to be standing directly above the hidden body of the protagonist's murder victim. As he becomes increasingly anxious that the officers suspect his guilt, he begins to hear what he takes to be his victim's heart beating underneath the floorboards. He becomes convinced that the sound, which in reality is the beating of his own heart, can be heard by the officers as well. Eventually, his emotions get the best of him and he gives himself away:

Was it possible they heard it not? . . . no, no! They heard!—they suspected!—They *knew*!—they were making a *mockery* of my

horror! . . . I could bear those hypocritical smiles no longer! . . . "Villains!" I shrieked, "dissemble no more! I admit the deed!—tear up the planks!—here, here!—it is the beating of his hideous heart!" (Poe, 1976, p. 262)

One element of the central character's reaction doubtless belongs more to the world of fiction than to everyday life: his conviction that the heartbeat was his victim's rather than his own. Although people sometimes project their mental states onto others, they typically recognize that their own strong physiological reactions belong to themselves and not to someone else.

In contrast, a second element of the protagonist's behavior, namely his exaggerated view of the officers' ability to read his internal reactions, may be quite common. In particular, we contend that, like Poe's character, people often overestimate the extent to which their thoughts, feelings, and sensations "leak out" and are available to others. A dinner guest may feel that her distaste over her host's atrocious cooking is more apparent than it really is, a secret admirer may believe his infatuation with a colleague is more obvious than is actually the case, or (closer to our own telltale hearts) a social psychologist conducting a deceptive experiment may overestimate the extent to which her participants can sense her apprehension and see through the cover story. Borrowing a term from Miller and McFarland (1987, 1991), we refer to this tendency to overestimate the extent to which others can read one's internal states as the *illusion of transparency*.

Why might people be susceptible to such an illusion? We contend that the bias stems primarily from the powerful impact of an individual's own phenomenology. People are typically quite aware of their own internal states and tend to focus on them rather intently when they are strong. To be sure, people recognize that others are not privy to the same information as they are, and they attempt to adjust for this fact when trying to anticipate another's perspective. Nevertheless, it can be hard to get beyond one's own perspective even when one knows that it

---

This research was supported by Research Grants MH45531 from the National Institute of Mental Health and SBR9319558 from the National Science Foundation. Portions of this research were presented at the annual meetings of the Eastern Psychological Association, 1995, Boston; the American Psychological Society, 1996, San Francisco, and 1997, Washington, DC; the American Psychological Association, 1996, Toronto, Ontario, Canada; and at the Empire State Social Psychology Conference, 1996, Blue Mountain Lake, New York.

We thank Antonio Annunziato, Todd Bickford, Theresa Buckley, Ann Charlton, Deborah Fidler, Talia Korenbrot, Sara Lederman, Shane Steele, and Tyler Story for their help in collecting the data, and Sara Lederman for serving as the confederate in Studies 3a and 3b.

Correspondence concerning this article should be addressed to Thomas Gilovich, Department of Psychology, Cornell University, Uris Hall, Ithaca, New York 14853-7601; Kenneth Savitsky, Department of Psychology, Williams College, Bronfman Science Center, Williamstown, Massachusetts 01267; or Victoria Husted Medvec, Department of Organization Behavior, J. L. Kellogg Graduate School of Management, Northwestern University, 2001 Sheridan Road, Evanston, Illinois 60208. Electronic mail may be sent to tdg1@cornell.edu, kenneth.k.savitsky@williams.edu, or vhm@nwu.edu.

is necessary to do so: The "adjustment" that one makes from the "anchor" of one's own internal experience is likely to be insufficient (Jacowitz & Kahneman, 1995; Quattrone, Lawrence, Finkel, & Andrus, 1984; Quattrone, 1982; Tversky & Kahneman, 1974). The net result is a residual effect of one's own phenomenology and a sense that one's internal states are leaking out and are more available to others than is actually the case.

In this regard, the illusion of transparency has much in common with the *spotlight effect*, a phenomenon we have demonstrated and discussed at length elsewhere (Gilovich, Medvec, & Savitsky, 1998; Gilovich, Kruger, Medvec, & Savitsky, 1998; Savitsky & Gilovich, 1998). The spotlight effect refers to the tendency to overestimate the extent to which others notice and attend to one's external appearance and behavior. In one study, for example, undergraduate participants who were asked to wear a T-shirt depicting the singer Barry Manilow (a figure of dubious renown among college students) drastically overestimated the likelihood that observers would take notice of their attire and recall who was pictured on the shirt (Gilovich, Medvec, et al., 1998). People appear to believe that the social spotlight shines more brightly on them than it actually does.

Like the proposed illusion of transparency, the spotlight effect appears to stem from the difficulty of getting beyond one's own phenomenological experience. Here too, people typically recognize that others are less focused on them than they are themselves, but it can be difficult to properly adjust for that realization. The adjustment people make from their own experience tends to be insufficient, causing them to overestimate the extent to which they are the object of others' scrutiny. Support for this anchoring and adjustment interpretation was obtained by showing that the spotlight effect was diminished considerably when participants made their estimates after they had had time to acclimate to wearing the embarrassing shirt. The delay led the participants to be less absorbed with their own appearance and thus to begin their estimates from a lower anchor (Gilovich, Medvec, et al., 1998). The present investigation of the illusion of transparency expands on this work by suggesting that people believe the social spotlight shines through to their internal states as well.

Indirect evidence for an illusion of transparency comes from several recent sources. In one study, participants watched a humorous videotape while their facial expressions were covertly recorded. Across a variety of experimental conditions, the participants thought they had been more expressive than observers rated them as being. Indeed, when the participants were allowed to view their own videotapes, they generally expressed surprise at how inexpressive they had been, relative to how expressive they had felt. As the authors put it, people may not know just how little they show (Barr & Kleck, 1995).

Additional support comes from research by Vorauer and Ross (1998), who found that people tend to overestimate the extent to which others can make accurate trait inferences about them from a sample of their behavior. Participants whose behavior was completely determined by the constraints of the situation nevertheless felt as though observers could discern important elements of their true nature by observing their actions. Vorauer and Ross suggest that this error occurs when individuals fail to appreciate the extent to which they are uniquely "in the know"

about themselves (see also Vorauer, in press; Vorauer & Claude, 1998).

Finally, the marked discrepancy between one's own knowledge of one's internal states and one's sense of what others are able to discern is nicely illustrated by an experiment (cited by Griffin & Ross, 1991) in which participants were asked to tap a well-known melody on a tabletop and then to estimate the proportion of listeners who would be able to identify the song they had tapped. As predicted, tappers grossly overestimated the listeners' abilities.

To appreciate these results, Griffin and Ross (1991) invited the reader to consider the different subjective experiences of the participants in each of the two roles.

First, imagine yourself as the tapper. As you tap rhythmically on the table in communicating the opening bars of the catchy tune you have chosen (let's say, "Yankee Doodle" or "Auld Lang Syne") you inevitably experience much more than your own tapping. Rather than impoverished knocks on the table, you "hear" the tune and the words to the song; indeed you are apt to hear . . . a full orchestration, complete with rich harmonies between strings, winds, brass, and human voice. Now imagine you are the listener. For you, there are no notes, words, chords, or instruments; you hear only an aperiodic series of taps. Indeed, you are unable even to tell how the brief, irregular moments of silence between taps should be construed—that is, whether each is a sustained note, a musical "rest" between notes, or a simple pause as the tapper contemplates the "music" to come next. (p. 335)

What is required, then, is for tappers to realize that the rich information they possess is radically different from the impoverished stimulus available to listeners. What is more, tappers must then adjust adequately to capture the perspectives of their listeners. Note, however, that these adjustments tended to be insufficient, resulting in biased intuitions of listeners' abilities. In the present research, we demonstrate that this phenomenon applies not just to musical melodies heard inside one's head, but also to internal states generally, including nervousness, disgust, alarm, and potentially many others.

The illusion of transparency also bears some resemblance to several additional established phenomena, including the hindsight bias (Fischhoff, 1975, 1982; Fischhoff & Beyth, 1975), the curse of knowledge (Camerer, Lowenstein, & Weber, 1989; Keysar & Bly, 1995; Keysar, Ginzler, & Bazerman, 1995), the self-as-target bias (Fenigstein, 1984; Zuckerman, Kernis, Guarniera, Murphy, & Rappoport, 1983), and various other accounts of egocentrism and biased perspective-taking (e.g., Fenigstein & Abrams, 1993; Griffin & Ross, 1991; Piaget, 1928/1959a, 1926/1959b; L. Ross & Ward, 1995, 1996; Stephenson & Wicklund, 1983; M. Ross & Sicoly, 1979). We discuss the most pertinent examples of this work in the context of specific experiments described below and defer a treatment of more general connections to earlier research to the General Discussion.

### Overview of Present Experiments

We report three sets of studies in this paper. In the first, we demonstrate that participants induced to lie overestimate the detectability of their lies. Second, we show that participants asked to sample a foul-tasting drink while trying to maintain a neutral facial expression exaggerate the extent to which their

disgust leaks out and can be detected by observers. Finally, we explore the speculation originally offered by Miller and McFarland (1987, 1991) that an illusion of transparency can help explain an enduring mystery in the literature on bystander nonintervention.

Across all three sets of studies, our approach is the same: Participants' intuitions about how they will be judged by others are compared with how participants are actually rated by observers. The illusion of transparency exists, then, when participants' estimates of the extent to which observers can discern their internal states exceed observers' actual ability to do so.

### Study 1a

As an initial investigation of the illusion of transparency, we put Poe's account of the telltale heart to empirical test. Do people who lie overestimate the detectability of their deception? Of course, we could not simulate the high-stakes nature of the deception in Poe's tale. Nevertheless, we suspected that participants would tend to overestimate the detectability of their lies even for fairly innocuous falsehoods told in the confines of the laboratory. Accordingly, we had groups of participants play a round-robin lie detection game in which each of them told lies and truths to the rest of the group. Within each session, each participant served as a "liar" in one round and as a truth-telling foil in all other rounds. Furthermore, when not called on to lie, each player served as an observer whose job it was to detect which player was lying. Of key interest were the intuitions of participants when they were the liars regarding the detectability of their own lies. According to our illusion of transparency hypothesis, liars should feel as though their lies are more obvious and detectable than they really are.

### Method

**Participants.** Thirty-nine Cornell University undergraduates participated in one of seven groups of 5 participants each, or one group of 4. Participants in this and all experiments reported in this article were recruited from a variety of courses in psychology and human development and earned extra credit for their participation.

**Materials.** We created 25 personal information questions that could unambiguously be answered truthfully or deceptively (e.g., "Name a foreign country you have visited," "Name a famous person you have met," "What brand of shampoo do you typically use?"). Each question was typed onto two separate index cards—one labeled "truth" and the other labeled "lie." In addition to the particular question, each card contained a sentence fragment for the participant to complete (e.g., "The brand of shampoo I typically use is . . .").

**Procedure.** After being screened to ensure that they were unacquainted with one another, the participants were escorted into a large laboratory room where they were randomly assigned to one of five chairs facing a podium. Participants each donned a name tag marked with a number from 1 to 5. The experimenter explained that the study was designed to investigate people's ability to detect lies and that they would be asked to play five rounds of a round-robin lie detection game. In each round, the players would be asked, one by one, to walk to the front of the room and receive a card from the experimenter. This card would contain a question, which they were to answer aloud in front of the assembled participants. Some of the cards, they were told, would require a true answer whereas others would require them to tell a lie.

The experimenter went on to explain that participants' main job was to be astute lie detectors. They were informed that the player correctly

identifying the greatest number of liars would be awarded a prize—a coupon for a free ice cream at a campus ice-cream shop. When lying, on the other hand, their job was to appear as if they were telling the truth.

Participants were informed that there would be one liar per round of the game and that each individual would be the liar once and only once in the course of the experiment. Prior to the experimental session, the experimenter had sorted the 25 cards such that there was only one lie card (and four truth cards) in each round, and such that each participant was to be the liar in one and only one round. The particular five questions for which participants were asked to lie were counterbalanced across experimental sessions.

In each of five rounds, then, each participant received a card from the experimenter, was given a brief moment to compose his or her answer, and made a single statement to the rest of the group from the podium at the front of the room. Each round thus consisted of five statements, delivered by the participants one after another (e.g., "I spent a summer in Kamchatka," "I have met David Letterman," "I usually use extra strength Head & Shoulders shampoo").

Participants completed a brief questionnaire after each round. On their questionnaire, liars were asked to estimate the number of participants, besides themselves, who would guess correctly that they had been the liar in that round. Answers could range from 0 (*none of the other players would peg them as the liar*) to 4 (*all of the other participants would guess that they had been the liar*). Special care was taken to inform participants that one person out of four, on average, would be expected to guess the identity of the liar in each round by chance alone.<sup>1</sup>

At the same time as the liar was estimating his or her detectability, truth-tellers were asked to guess the identity of the liar.<sup>2</sup> They did this by circling the player number of the participant they suspected had lied in that round. In addition, all truth-tellers were asked "How many people do you think will guess (incorrectly) that *you* were the liar in this round?" This question was included to address an alternative interpretation for the hypothesized results, discussed below.

When all questionnaires had been completed, the experimenter collected them and began the next round. After the final round, each player "came clean" and revealed to the group what he or she had lied about. The participant who had correctly identified the greatest number of liars was then awarded his or her prize, and everyone was debriefed. The experimental session lasted approximately 30 min.

### Results

Because the data within each session are highly interdependent, all analyses were conducted at the level of the experimental session rather than the individual participant.

As anticipated, liars overestimated the likelihood that their fellow participants would be able to identify them as the liar. Across the eight sessions, liars estimated that an average of 48.8% of the participants would correctly peg them as the liar when, in fact, only 25.6% did so—an accuracy rate indistinguishable from chance. The difference between liars' intuitions

<sup>1</sup> In the single group of 4, participants donned name tags numbered 1–4, played four rounds of the lie detection game, and were told that 1 participant out of 3 could be expected to guess the identity of the liar by chance alone. All other procedural details were consistent with the other sessions.

<sup>2</sup> To eliminate any ancillary cues to the identity of the liar, all participants received identical questionnaires but were instructed to answer different questions, depending on whether they had been a liar or a truth-teller in that particular round.

and the actual accuracy rate was highly statistically significant, paired  $t(7) = 5.41, p < .001$ .

We attribute this finding to an illusion of transparency: Liars presumably felt as if their feelings of nervousness about lying leaked out, or that others could "see right through them." There are, however, several alternative interpretations of this result. One derives from the *self-as-target bias*, or people's exaggerated judgments of the extent to which others' thoughts and actions are directed at them (Fenigstein, 1984; Zuckerman et al., 1983). Applied to the present experiment, liars' overestimations of the number of observers who would correctly identify them might not have stemmed from any feeling of transparency, but from a simple conviction that they would tend to be the target of others' guesses and suspicions more than one would expect by chance.

But note that if this alternative interpretation is true, and participants felt like the generalized targets of others' suspicions, they would have felt this way even when they were telling the truth. The illusion of transparency hypothesis, in contrast, entails elevated estimates of detectability only when an individual is lying. Our data permitted a test of this issue. Recall that in addition to asking participants to estimate their detectability when they were lying, we also asked them to estimate the number of their fellow participants who would guess that they were lying when they were telling the truth. We averaged across each participant's four truth-telling rounds and compared this average with the estimate each participant made when he or she was lying. Consistent with our expectations, but in contrast to what would be expected if our results stemmed exclusively from a self-as-target bias, participants expected significantly fewer of their fellow participants to pick them as the liar when they were telling the truth ( $M = 34\%$ ) than when they were lying ( $M = 49\%$ , reported above), paired  $t(7) = 6.11, p < .0005$ .

### Study 1b

We conducted two follow-up experiments to explore additional alternative interpretations of the results of Study 1a. First, recall that in the previous study, each participant was the liar once and only once and, moreover, that the participants were aware of this aspect of the design. One might argue that this was partly responsible for observers' low accuracy rates. Specifically, if a participant guessed a particular player as the liar in Round 1, he or she may have been reluctant to guess that player again as the liar in any subsequent round. If the original guess had been in error, the participant would thereby decrease his or her chances of correctly identifying the liar in these latter rounds. Thus, the participants' awareness that each of them was the liar once and only once might have artifactually lowered their accuracy rates.

In truth, we do not think this represents much of a challenge to our results. Because participants were given no feedback about the accuracy of their guesses until the end of the experiment, nor were they forbidden to guess the same player as the liar more than once, their overall accuracy should not have been affected. Indeed, the logic presented above could also be used to argue that our design artifactually *increased* observers' accuracy rates, making Study 1a a conservative test of our hypothesis. Still, to be certain that this aspect of our design in no way

affected the observers' ability to detect deception, we conducted a replication in which participants were informed that although there was only one liar per round, repetitions were allowed. Specifically, participants were informed that a computer had randomly selected which participant was to be the liar in each round, and so any individual could be the liar once, more than once, or not at all during the course of the experiment.

In actuality, however, each player was the liar once and only once across the five rounds of the experiment. All that changed were participants' perceptions of the frequency with which each could receive a lie card.<sup>3</sup> Forty Cornell University students served as participants in eight groups of 5. As before, liars substantially overestimated the detectability of their lies: They estimated that an average of 50% would correctly select them as the liar; in fact, only 27% did so, paired  $t(7) = 6.56, p < .0005$ . Once again, people's lies were less detectable than they suspected.

### Study 1c

Study 1c was designed to investigate two more plausible alternative interpretations of the results observed in Studies 1a and 1b. First, liars' exaggerated fear that they would be detected may have stemmed not from an illusion of transparency but from an abstract theory that lies are easy to detect. Empirical research has repeatedly shown that people's ability to detect lies is quite modest (DePaulo, Zuckerman, & Rosenthal, 1980; Ekman, 1985; Knapp & Comadena, 1979; Kraut, 1980; Zuckerman, DePaulo, & Rosenthal, 1981)—a result that most people find surprising. Before being exposed to this evidence, people tend to believe that lies are readily detectable.

Of course, any such belief that lies are readily detectable would be aided and abetted by the illusion of transparency that we have proposed. Still, our account emphasizes the feelings of detectability that arise, *in vivo*, in the particular situation, rather than intuitions based on an abstract theory about human cognitive and perceptual abilities. If the liars in our experiment simply thought that lies were readily detectable, their estimates of the number of observers who would correctly select them as the liar in a particular round might have followed rather dispassionately from this abstract theory rather than from any personal feelings of transparency.

A second plausible alternative interpretation comes from research suggesting that whenever individuals possess some knowledge, they can have difficulty assuming the perspective of another individual who is not in the know. Instead, they mistakenly attribute to the other person some degree of awareness of their privileged information. This tendency has been referred to as the *curse of knowledge*—one is "cursed" by one's own knowledge in the sense that it can be difficult to set that knowledge aside when imagining how things appear to someone else (Camerer et al., 1989; Keysar & Bly, 1995; Keysar et al., 1995).

How might the curse of knowledge account for our results? Liars were (of course) well aware that they were lying and may

<sup>3</sup> We did not include the self-as-target control question on the questionnaires used in Studies 1b or 1c.

have had difficulty getting beyond that fact when anticipating the perspective of others. As a consequence, they may have felt that others would share their knowledge, resulting in an overestimation of their own detectability. As we discuss later, this interpretation resembles our own theorizing to some degree. Even so, the curse of knowledge explanation does not require that participants perceived anything akin to the feelings of leakage that we believe underlie the illusion of transparency. Indeed, the curse of knowledge hypothesis, unlike our own speculations, makes no distinction between an individual who tells a lie and another person who merely knows for certain that a particular individual has done so. Both should be equally cursed by that knowledge.

To investigate these two alternative interpretations—the abstract theory that lies are detectable and the curse of knowledge—we ran a version of our experiment in which each of the 5 participants was yoked to his or her own personal observer. Each yoked observer received a card from the experimenter identical to the card given to the actual participant. Thus, each observer knew the question posed to his or her partner in each round at the exact moment the partner did and was also aware of whether the card called for the truth or a lie. Finally, the observers completed questionnaires analogous to those completed by their partners after every round. Thus, the observer yoked to the liar in each round also estimated the number of players who would correctly identify the liar. These estimates could then be compared with the intuitions of the liars themselves.

The yoked observers allowed us to examine simultaneously both alternative interpretations discussed above. First, the yoked observers should hold the same abstract theory about lying and lie detection as the participants, but should be immune to any in vivo feelings of transparency, as they themselves do not make statements to the rest of the group. Thus, if the results of Studies 1a and 1b stemmed from an abstract theory about the detectability of lies, the observers should overestimate the detectability of their partners' lies to the same extent as the partners themselves. If, on the other hand, our results stem from an illusion of transparency, only the liars should overestimate the detectability of their lies.

A curse of knowledge interpretation likewise predicts that the yoked observers, who are "cursed" by the same information as the liars, will overestimate the detectability of the liars' lies every bit as much as the liars themselves. According to the illusion of transparency hypothesis, in contrast, the yoked observers should not experience the sensation of leakage and, hence, should not overestimate the detectability of their partner's lies. In summary, then, to the extent that the yoked observers do not overestimate the detectability of their partner's lies, both alternative interpretations can be ruled out. Finally, the yoked observers also control for any alternative interpretation involving response bias or the particularities of the response mode (e.g., that participants ignored our statements about the level of identification accuracy that could be expected by chance).

## Method

Eighty Cornell University students participated in one of eight sessions of this experiment, 40 as players, replicating our earlier design, and an

additional 40 as yoked observers. The experiment followed the same basic procedure as the earlier studies, but with observers completing questionnaires after each round as well. Also, one additional item was added to these questionnaires: Liars and their partners were both asked to rate the "obviousness" of the lie on a 7-point scale ranging from *not at all obvious* (1) to *very obvious* (7).

Finally, we also included an individual difference scale, in part to explore the mechanism hypothesized to give rise to the illusion of transparency. Recall our thesis that the illusion derives from the difficulty of putting aside one's own phenomenological experience when attempting to view oneself from the perspective of another. If individuals base their estimates of leakage on their own phenomenological experience, then those for whom these internal experiences are more available should be particularly prone to believe that their internal states are leaking out.

To test this possibility, we had each participant complete an inventory of dispositional self-consciousness (Fenigstein, Scheier, & Buss, 1975; see also Carver & Glass, 1976; Fenigstein, 1987). Participants responded to each of 23 items on 5-point scales with endpoints *extremely uncharacteristic of me* (0) and *extremely characteristic of me* (4). The Self-Consciousness Scale consists of three subscales: Private Self-Consciousness, Public Self-Consciousness, and Social Anxiety. It was the first of these that was of particular interest to us. Private self-consciousness refers to an individual's tendency to focus internally, reflecting on his or her inner thoughts and feelings (e.g., "I reflect about myself a lot"). We reasoned that individuals scoring high on this subscale, who were likely to have been keenly aware of their own internal states during the course of the experiment, would be especially likely to believe that those internal states had leaked out. In our terms, when attempting to capture others' perspectives on themselves, these individuals may have begun the inferential chain from a more pronounced anchor.

## Results

As before, liars overestimated the detectability of their lies: On average, they predicted that 44.3% of their fellow players would detect their deception; in fact, only 32.4% did so, paired  $t(7) = 2.91, p < .05$ . Of greater interest to the present investigation, however, is that liars' estimates diverged from those made by the yoked observers: Observers estimated that only 25.3% of the players would detect the liars' deception, far fewer than the 44.3% estimate made by the liars themselves, paired  $t(7) = 4.12, p < .005$ . Additionally, liars rated their own lies as significantly more obvious than did their yoked observers,  $M_s = 3.0$  and  $2.0$ , respectively, paired  $t(7) = 3.90, p < .01$ .

These results indicate that the discrepancy between liars' estimates and observers' actual accuracy cannot be attributed to participants' abstract theory that lies are easy to detect. These results also cast doubt on the curse of knowledge alternative interpretation of our results. It seems, then, that rather than deriving dispassionately from a theory of human psychology or from an inability to put one's knowledge aside when considering another person's perspective, liars' heightened estimates of the detectability of their lies are the result of an illusion of transparency.

Finally, was dispositional self-consciousness related to participants' feelings of transparency? To address this question, we correlated each of the three self-consciousness subscales with liars' estimates of the number of their fellow participants who would catch them in their lie as well as their ratings of how obvious their lie had been. As expected, neither public self-consciousness nor social anxiety was significantly correlated

with liars' estimates of their own leakage,  $r_s = -.03$  and  $-.18$ , respectively, or with their obviousness ratings, both  $r_s = .10$ .

In contrast, private self-consciousness was significantly related to liars' feelings of transparency,  $r = .41$ ,  $p < .01$ , as well as their ratings of how obvious their lie had been,  $r = .40$ ,  $p < .02$ . Multiple regression analyses predicting liars' detectability estimates and obviousness ratings from all three self-consciousness subscales simultaneously also revealed significant effects only for private self-consciousness:  $b = .85$ ,  $t(34) = 2.66$ ,  $p < .02$ , for detectability;  $b = 1.38$ ,  $t(34) = 2.70$ ,  $p = .01$ , for obviousness. This relationship between estimates of leakage and individuals' habitual focus on their own thoughts and feelings suggests that liars based their estimates of how detectable their lies were on their own phenomenological experiences. Those who tended to be highly aware of their own internal states apparently began from a more pronounced anchor, resulting in greater feelings of transparency.

### Discussion of Studies 1a, 1b, and 1c

Across three studies, we found consistent support for an illusion of transparency. Participants induced to lie overestimated the detectability of their lies in all three experiments. Like the protagonist in *The Tell-Tale Heart*, our participants' deception was not as obvious as they thought. The three studies also diffuse the concern that the observed discrepancy between liars' estimates and the actual detectability of their lies was caused by some factor other than the hypothesized illusion of transparency. In particular, we have provided empirical evidence indicating that our findings do not derive from the self-as-target bias (Study 1a), an abstract theory that lies are easy to detect (Study 1c), the curse of knowledge (Study 1c), or a feature of our design that could have artifactually lowered actual accuracy rates (Studies 1b and 1c).

Finally, although our lie-detection studies were not designed as explicit tests of the anchoring and adjustment mechanism that we believe underlies the illusion of transparency, they nevertheless provided evidence consistent with that interpretation. First, as already discussed, the illusion was most pronounced among participants in Study 1c who had the highest private self-consciousness scores. This is consistent with the proposed mechanism in that these individuals are the ones who are most focused on their internal states and thus felt them most keenly. They doubtless recognized that the observers were not privy to the fullness and intensity of their internal experience, but the adjustments they made in light of this realization nevertheless began from a higher anchor—that is, from a richer emotional experience. The net result is that these participants gave higher estimates of the number of observers who would be able to detect their deception.

Additional evidence in support of the proposed anchoring and adjustment mechanism comes from an ancillary finding in Study 1a. Recall that participants in that study were asked to estimate how many observers would think they were lying both when they were lying and when they were telling the truth. The complement of the latter estimates, of course, represents their estimates of how many people could discern that they were telling the truth, or how much leakage they felt when they were being honest. Would these estimates yield an illusion of transparency

as well? We would expect not, because when telling the truth there is no strong internal experience on which to anchor one's judgments of leakage. Lying typically generates a host of emotions that can potentially leak out (Ekman, 1985) and from which—according to our anchoring and adjustment model—the process of judgment begins. Not so for telling the truth. When there is no strong internal sensation from which to adjust, there is no illusion of transparency.

We tested this prediction by averaging each participant's implicit estimates of the number of observers in Study 1a who would be able to discern that they were telling the truth when, in fact, they were being truthful. Each participant made four such estimates. These estimates were then compared with the actual number of observers who indicated that the participant was telling the truth. As expected, this analysis yielded no evidence of an illusion of transparency. On average, truth tellers estimated that 63% of observers would say they were telling the truth, when, in fact, 73% did so. We emphasize that these data should be interpreted with caution: They were obtained by reverse scoring participants' estimates of the number of observers who would think they were lying, and this feature of the study was not designed as an explicit test of the anchoring and adjustment mechanism. Nevertheless, the data are consistent with that mechanism and nicely complement the private self-consciousness data as support for the anchoring and adjustment interpretation.

The three studies reported thus far indicate that the illusion of transparency is robust across a variety of procedural changes—at least in the domain of lie detection. But what about other domains? Does it apply to emotional states other than those a person experiences when telling lies? Our remaining studies were designed to find out.

### Study 2a

Consider the following awkward situation: You are a dinner guest at the home of a friend who takes pride in his culinary talents. You wait with the other guests in eager anticipation as your host toils away in the kitchen, preparing what you expect will be some succulent entrée. But when dinner arrives, you discover to your dismay that the meal contains a generous amount of an ingredient you find absolutely unpalatable. The very thought of it prompts revulsion.

What to do? One option is to admit the aversion up front, perhaps apologizing for your unsophisticated palate. We suspect, however, that most people would reject this course of action in favor of downing a few well-timed bites of the offensive entrée and doing their best to conceal their feelings of disgust from the host. We further suspect that people elect this course of action despite strong doubts about whether they will be able to execute it effectively. That is, people may believe that despite their best efforts to conceal their true feelings, their distaste will leak out and be apparent to others.

Are such doubts justified? Judging from the first three experiments, perhaps not. The illusion of transparency implies that a person's feelings of disgust may not be as obvious as he or she believes. There is a substantial discrepancy between the phenomenological experience of the individual, whose feelings of disgust may be quite pronounced, and the cues available to

outside observers—a discrepancy that the individual may fail to appreciate fully and adjust for adequately.

To investigate this possibility, we simulated the essential components of the dinner-guest scenario. Participants were asked to conceal their feelings of disgust over a foul-tasting drink and then to estimate how successfully they had done so. These estimates were then compared with the participants' actual success at concealment. Our prediction was that participants would overestimate the extent of their leakage, believing that their disgust was more apparent than it actually was.

## Method

**Participants.** Twenty-five Cornell University undergraduates volunteered to participate in an experiment entitled "Neutral Expressions." Fifteen were run individually as tasters, and the remaining 10 were recruited later and served as observers.

**Procedure.** In a first phase of the experiment, participants arrived individually and were greeted by an experimenter. They were then invited to sit at a table on which were placed 15 small plastic cups, each containing a small amount of red liquid. On the side of the table opposite their seat was a video camera mounted on a tripod. The experimenter informed the participant that the study was concerned with facial expressions, specifically *neutral* facial expressions, or "expressions that show no thought or emotion." To explore people's ability to maintain such blank expressions, the study called for participants to sample a variety of "special drinks" while their reactions were videotaped. In particular, the participant was to taste each drink, concealing his or her reactions, good or bad.

Participants were assured that the drinks were harmless, but were informed that 5 of the 15 cups contained a drink concocted by the experimenter to have an unpleasant taste. The remaining 10 cups were said to contain a good-tasting drink. This information was accurate: Five cups contained a small amount (approximately 5 ml) of a mixture of water, red food coloring, and the vinegar brine solution in which pickled grape leaves are packed. Pretesting revealed that it was indeed regarded as an unpleasant taste by all. The remaining 10 cups contained cherry-flavored Kool Aid, prepared as directed. The two different kinds of drinks were visually indistinguishable. The location of the foul-tasting drinks was counterbalanced across participants with two restrictions: The foul-tasting drink never appeared in the first position and two cups containing the foul-tasting drink never appeared side by side.

Before proceeding with the study, the experimenter noted that the videotape made of the taster would later be shown to a group of 10 observers, who would be told that five of the drinks had an unpleasant taste but would not be told which ones. The observers' job, then, would be to watch the videotape and try to determine which drinks had the unpleasant taste, based only on the taster's facial expressions. Tasters were told, in turn, that their job was to make the observers' task as difficult as possible.

The tasters then proceeded to taste each drink, one at a time, while looking at the camera and being videotaped. After each drink, the video-camera was paused and the taster completed a brief questionnaire. The questionnaire elicited the tasters' predictions of the accuracy of the 10 observers. In particular, the tasters were asked to estimate the number of observers who would correctly identify whether the drink was one with a pleasant or unpleasant taste. On the unpleasant-tasting trials, it was made clear to the tasters that an average of 3.33 observers could be expected to guess correctly by chance alone. On pleasant-tasting trials, it was made clear that the corresponding chance accuracy rate was 6.67.

In the second phase of the experiment, the observers were assembled in small groups and watched the videotapes made in the first phase of

the study. Because the tasters were asked to predict the number of observers out of 10 who would correctly identify the foul-tasting drink, the videotape of each taster was seen by a total of 10 observers. They were given a detailed description of the first phase of the study, including the fact that five drinks had an unpleasant taste. Their job was to guess whether each drink had a pleasant or unpleasant taste, based on the taster's expressions. If they felt uncertain about the taste of a given drink, they were encouraged to guess. They were given no feedback regarding the accuracy of their responses.

## Results

We predicted that tasters would fall victim to the illusion of transparency and overestimate the number of observers who could tell if they were sipping a foul-tasting drink. Because the pleasant-tasting drink did not constitute as strong an internal experience as the unpleasant drink, we did not expect participants to exhibit a similar illusion of transparency for these trials. To examine these questions, we collapsed separately across pleasant and unpleasant trials for each taster and computed the average number of observers that tasters predicted would guess correctly for each type of drink. We then computed the average number of observers who actually did guess correctly.

Both hypotheses were supported. Tasters estimated that an average of 4.91 observers would correctly identify the foul-tasting drinks, which was significantly higher than the average of 3.56 observers who actually did make correct identifications on these trials,  $t(14) = 3.63, p < .005$ . Furthermore, the tasters' estimates were significantly greater than the chance accuracy rate of 3.33,  $t(14) = 3.97, p < .005$ , whereas the observers' actual accuracy was not,  $t < 1$ . In contrast, tasters estimated that an average of 6.20 observers would correctly identify the good-tasting drinks, a figure that did not differ significantly from the actual accuracy rate of 6.83. For the good-tasting drinks, neither the tasters' estimates nor the accuracy of the observers was significantly different from the chance accuracy rate of 6.67.

As with the lie detection studies, we attribute these results to the illusion of transparency. Participants felt that cues to their disgust over the foul-tasting drink had leaked out and were noticed by more observers than was actually the case. As with the first lie detection study, however, there are a couple of alternative interpretations of these data. First, the results might merely reflect inaccurate abstract theories about the detectability of disgust, not any illusions about excessive leakage. Alternatively, they could reflect the curse of knowledge: Participants knew when they were tasting pleasant or unpleasant drinks and it may have been difficult for them to get beyond that knowledge when estimating what the observers were likely to know.

To test each of these alternative interpretations, we replicated the taste test study using the same yoked-partner procedure as in Study 1c. As in that study, the yoked partners had the same abstract theories as the tasters and were "cursed" by the same knowledge of whether a given drink had a pleasant or unpleasant taste. If the results of study 2a were due to faulty abstract theories or the curse of knowledge, the estimates made by the yoked partners should not differ systematically from those of the tasters themselves. If the previous results were due to the illusion of transparency, in contrast, the estimates of the tasters and yoked partners should diverge.

## Study 2b

### Method

**Participants.** Fifty-two Cornell University undergraduates served as participants. Thirty-two arrived in pairs and were randomly assigned to be either a taster or a partner; the remaining 20 were recruited later and served as observers.

**Procedure.** The basic procedure was the same as that in Study 2a, with two changes. One change was the addition of the yoked partner. Participants in the first phase of the experiment arrived in pairs, and one was randomly chosen by a coin flip to be the taster. The taster was seated across the table from the video camera, with the partner seated on an adjacent side, facing the taster. The partner's job, it was explained, was to observe as the taster sampled each drink and determine how well (or how poorly) the taster was able to maintain a neutral expression. Critically, the partner did so while knowing whether each drink had a pleasant or unpleasant taste. The experimenter informed the partner whether each drink was a pleasant or unpleasant drink at the exact moment the taster tasted it. To prevent this information from being recorded on the videotape, and to minimize the possibility of distracting the partner from the task of scrutinizing the taster's facial expressions, the experimenter alerted the partner as to whether each drink was pleasant or unpleasant via a subtle tap on the shoulder (one tap for a pleasant drink; two taps for unpleasant). Partners indicated that this mode of communication was sufficient and nonintrusive.

The second change involved the number of drinks the taster sampled and the way in which the estimates of both the taster and partner were elicited. Only 10 drinks were sampled. They were arranged in front of the taster in two groups, labeled "Set 1" and "Set 2," and were numbered 1–5 within each set. One of the drinks in each set contained the foul-tasting mixture of water, food color, and brine, and the other 4 contained Kool Aid. The location of the foul-tasting drink in each set was counterbalanced across participants with two restrictions: The foul-tasting drinks were never first in a set, and they never appeared in the same position in both sets.

Tasters and partners were informed that the taster would sample all the drinks in a set, at which point the experimenter would pause the video camera and administer a brief questionnaire. The questionnaire asked tasters and partners to estimate the number of observers who would correctly identify the single foul-tasting drink in the set. The experimenter explained that an average of 20%, or 2 observers out of 10, could be expected to answer correctly for each set by chance alone. Tasters and partners also rated the extent to which they thought the taster's feelings about the unpleasant drink had leaked out. These judgments were made on a 10-point scale with endpoints labeled *I [he/she] kept a perfect neutral facial expression* (1) and *I [he/she] leaked out to a great extent* (10). The dependent measures in this study were altered from those used in Study 2a to examine whether the same results could be obtained with different methods of elicitation, and thus further the generalizability of the illusion of transparency.

As before, the observers were run in a second phase of the study. Because tasters and partners were asked to predict the number of observers out of 10 who would correctly identify the foul-tasting drink, the videotape of each taster was seen by 10 observers. Observers were given a detailed description of the first phase of the study, including the fact that one drink (and only one drink) out of each set of five had an unpleasant taste. Their job was to guess the location of the foul-tasting drink in each set.

### Results

We predicted that tasters would fall victim to the illusion of transparency and overestimate the number of observers who would identify the single foul-tasting drink in each set. To exam-

ine this question, we collapsed across the two sets of drinks for each taster and computed the average number of observers that tasters predicted would guess correctly and the average number who actually did guess correctly.<sup>4</sup> These results are presented in the two leftmost bars of Figure 1.

As can be seen from Figure 1, our prediction was supported: On average, tasters estimated that substantially more observers would correctly identify which drink had been unpleasant than actually did so,  $t(15) = 5.01$ ,  $p < .0005$ . Figure 1 also shows that partners likewise tended to overestimate the observers' accuracy,  $t(15) = 2.09$ ,  $p = .05$ , but not to the same extent. Indeed, the estimates made by the partners were significantly lower than those made by the tasters themselves,  $t(15) = 2.11$ ,  $p = .05$ . This suggests that, above and beyond the curse of knowledge or the effects of any abstract theory, the tasters demonstrated an illusion of transparency. Consistent with this interpretation, tasters tended to believe that they leaked out marginally more than their partners thought they did ( $M_s = 3.41$  and 2.53, respectively),  $t(15) = 1.88$ ,  $p < .10$ .

### Discussion of Studies 2a and 2b

The findings of Studies 2a and 2b provide further support for the existence of an illusion of transparency. In both studies, participants believed that they failed to hide their reactions to an unpleasant taste. In reality, tasters were remarkably successful at concealing their distaste: Observers performed no better than chance and fell far short of tasters' estimates. We can conclude that the illusion of transparency effect is not just a manifestation of abstract theories of leakage or of the curse of knowledge because the yoked partners in Study 2b—who would have possessed all the same theories, and who were provided with factual knowledge equivalent to that possessed by the tasters themselves—did not overestimate observers' accuracy as much. As in the lie detection studies reported earlier, people can conceal their internal states better than they suspect.

### Study 3a

As noted earlier, we took the term *illusion of transparency* from Miller and McFarland (1987, 1991), who proposed that such an illusion might help explain a puzzling inference on the part of participants in bystander intervention studies (Darley & Batson, 1973; Darley & Latané, 1968; Latané & Darley, 1970). Studies 3a and 3b were attempts to provide empirical support for their speculation.

When confronted with a potential emergency, people typically play it cool, adopt a look of nonchalance, and monitor the reactions of others to determine if a crisis is really at hand. No one wants to overreact, after all, if it might not be a true emergency. However, because each individual holds back, looks nonchalant, and monitors the reactions of others, sometimes everyone concludes (perhaps erroneously) that the situation is not an emergency and hence does not require intervention.

Note the inferential failure here: People witness the same behavior in others as they are engaging in themselves, but con-

<sup>4</sup> Examining each of the two sets separately produced results virtually identical to those reported above.

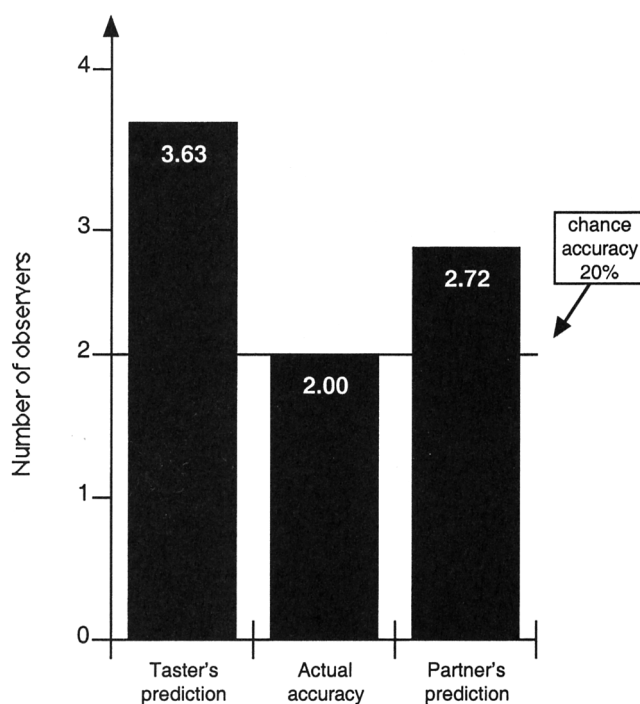


Figure 1. Mean number of observers predicted to guess the placement of the foul-tasting drink by tasters and partners, as well as observers' actual accuracy, Study 2b.

clude that the cause of others' behavior is different than their own. Why do they fail to conclude that the nonchalance of others has the same origin and the same meaning as their own—that everyone is simply playing it cool and looking to others to define the situation for them? This is not a terribly difficult inference to make, and people's failure to make it has loomed as something of a puzzle ever since the classic bystander intervention studies were first published.

Following Miller and McFarland (1987, 1991), we propose that the illusion of transparency can help resolve this puzzle. Simply put, people may not see others' behavior as similar to their own; indeed, they may be unaware that they appear as calm and unconcerned as everyone else. If people think that their alarm over the potential emergency is leaking out more than it is, they may conclude that others—who exhibit comparatively few signs of alarm—are genuinely much less concerned about the situation than they are. They may then use this information to conclude that they were mistaken and that no real emergency exists. Thus, a person's failure to draw the right conclusion about the apparent calm of others may not reflect an inferential failure about the meaning of others' behavior so much as a failure to recognize that how others look is precisely how one looks oneself.<sup>5</sup>

We examined this issue empirically in a modified version of the bystander intervention studies (Darley & Latané, 1968; Latané & Darley, 1970). Because this classic research is so widely known, we thought it would be difficult to conduct a study on a college campus involving a simulated physical or medical emergency without arousing suspicion.<sup>6</sup> We therefore

decided to create a novel kind of laboratory emergency, one that might provoke less suspicion while still retaining the essential features of an emergency situation.

In our study, participants found themselves confronted by an individual (a confederate) who appeared to be breaking the established rules of the experiment. Indeed, it appeared to participants that the problematic individual threatened to ruin the entire experiment—and this after the experimenter had solemnly and emphatically stressed the importance of the research and the necessity that all participants follow the rules of the study precisely. Participants were thus confronted with a situation analogous to that encountered in real-life emergencies: Should they intervene (and save the threatened experiment) or not? By embedding some crucial questions amidst a variety of filler items on a questionnaire completed by participants, we were able to investigate our illusion of transparency hypothesis. In particular, we predicted that participants would rate themselves as appearing more concerned over the transgressions than they actually appeared. That is, each would tend to believe that he or she appeared more alarmed at the behavior of the confederate than others rated him or her as appearing.

### Method

**Participants.** Forty Cornell University students were recruited for an experiment on group problem solving and participated in 1 of 10 groups (4 participants plus 1 female confederate).

**Procedure.** After being screened to ensure they were unacquainted with one another, participants were escorted into a laboratory room in which five chairs were arranged in an arc facing a chalkboard. All participants donned name tags with participant numbers (1–5).

The experimenter explained that the study was designed to investigate how various aspects of a group's working environment affect the group's problem-solving productivity. Their task in the experiment would be to unscramble as many anagrams as they could in 10 min, with the total number solved being the measure of the group's productivity. The experimenter explained that 1 participant would be randomly selected to be the "writer," and he or she would write each anagram on the chalkboard, one by one, for the group members to solve aloud. The writer would

<sup>5</sup> It should be noted that Miller and McFarland (1987, 1991) offer a different resolution to this puzzle—that people, on average, tend to believe that they are more embarrassed than others. A person's own inaction, then, is explained by something (fear of embarrassment) not believed to be as significant a determinant of others' behavior. This explanation for bystander nonintervention bears some similarity to the illusion of transparency explanation in that both suggest that conclusions about the self can diverge from conclusions about others in part because certain information is more available to an individual about him- or herself than about others.

Still, the explanations are distinct, and Miller and McFarland (1991) have noted that some of their experimental results that are consistent with their own approach are not open to reinterpretation in terms of an illusion of transparency. We agree, and we propose the illusion of transparency as one mechanism among several that jointly account for people's behavior in such complicated and multifaceted social situations.

<sup>6</sup> Indeed, as we were planning our experiment, several students, citing lessons learned about bystander nonintervention in their introductory psychology course, thwarted an actual suicide attempt, saving a student who had intended to leap into one of the gorges that border Cornell's campus (Carmona, 1993). The publicity surrounding this event made us even more dubious about pulling off a suspicion-free medical emergency.

have a list of all of the anagrams and their solutions and would keep track of the number solved by the group. The experimenter then handed around a cup containing slips of paper, on one of which was ostensibly printed the word "writer." In reality, all slips said "solver." Nevertheless, when the confederate drew a slip of paper, she announced that she had drawn the writer slip.

The experimenter stressed that it was crucial that all participants follow the rules of the experiment and work at their most efficient level. Furthermore, it was important that their performance on the anagram task accurately reflect their true ability level and amount of effort expended: "If for any reason your performance level is not an accurate indication of the effort and ability you bring to the task, then the experiment will be a failure and the data will be useless to us." Then, making eye contact with each participant, the experimenter remarked, "on a personal note," that the experiment was the final component of his dissertation research, that it was almost finished, and that the data so far looked highly promising. "We just need a few more pieces to fall into place, so it's even more important than usual that everyone takes the task seriously and follows all the rules."

One such rule, the experimenter emphasized, was that the writer should refrain from giving participants any hints or assistance in solving the anagrams. For their part, solvers were restricted to offering solutions to anagrams, or uttering the word "pass" if they wished the writer to move on to the next anagram. All other "table talk," including discussions of strategy or remarks of any kind to each other or to the writer was forbidden.

After a final review of the rules, the experimenter handed the confederate the list of 30 anagrams, set the dial on a kitchen timer for 10 min, and left the room. During this time, the confederate wrote the anagrams on the chalkboard for group members to solve, as instructed. Gradually, however, despite the experimenter's admonitions, she began to offer unsolicited assistance to the solvers. She began by providing them with small hints, such as informing them which letter a particular anagram started with, or suggesting a category to which the solution belonged (e.g., "Think of things you'd find at a circus"). Soon, however, her disregard for the rules became more blatant: She provided participants with words that rhymed with an anagram's solution, she gave participants credit for anagrams that, in her eyes, they had "almost" solved, and even turned the timer back to give them a few extra minutes. We carefully pretested this portion of the experiment so that we could create maximum alarm on the part of the participants without arousing their suspicions. Within each session, moreover, the confederate was instructed to continually adjust her behavior in an attempt to keep participants alarmed at her behavior and on the verge of intervention, but never actually to elicit intervention. Nevertheless, in a few cases participants did intervene or suspect the confederate's agenda. We return to this issue later.

When the timer finally rang, the experimenter reentered the room, expressed mild surprise at the number of anagrams the group had solved, and escorted all participants (including the confederate) to individual cubicles to complete the dependent measures. When all participants had finished, they were probed for suspicion, debriefed, and dismissed.

**Dependent measures.** We constructed the questionnaire to fit our cover story regarding the effects of a group's working environment on productivity. Thus, for example, participants were asked to rate the extent to which they believed the laboratory setting had inhibited their group's performance, compared with how they would have performed in a more naturalistic, real-world setting. Further, we asked participants to rate the extent to which they were concerned about the laboratory setting inhibiting their performance, and also how concerned they thought they appeared to the other group members. Finally, they were also asked to rate how concerned about the laboratory setting each of the other participants had appeared.

We included these questions, and several others of the same type, to

set the stage for the items of interest. Specifically, respondents were asked to indicate whether anyone had violated the rules of the experiment. If so, they were asked to circle the player number of the participant who had done so to the greatest extent. Following this, they were asked to rate how concerned they were about this individual's behavior, as well as how concerned they believed they appeared to others. These ratings were made on 7-point scales with endpoints labeled *not at all concerned* (1) and *very concerned* (7). Finally, they rated how concerned each of the other group members appeared on the same 7-point scale.

Our hypothesis involves a comparison between the participants' intuitions about how concerned they appeared about the potential emergency and how concerned they actually appeared to the other group members. By having each participant rate the appearance of every other participant, we were able to compute an average of how concerned each participant seemed to everyone else (not including the confederate). This average could then be compared with each player's self-rated appearance.

## Results

Because the responses within each group were highly interdependent, all analyses reported here were conducted at the level of the experimental session rather than the individual participant.

For each group, we first checked that all participants had in fact listed the confederate as the most rule-breaking group member; all had. We then computed participants' average ratings of how concerned they were and how concerned they believed they had appeared. Next, we derived a measure of how concerned each participant had actually appeared in the eyes of his or her group members by averaging, separately for each participant, the ratings that each of the other participants made about him or her. Finally, we averaged these "actual" appearance scores across all participants within each group. The three measures are depicted in Figure 2.

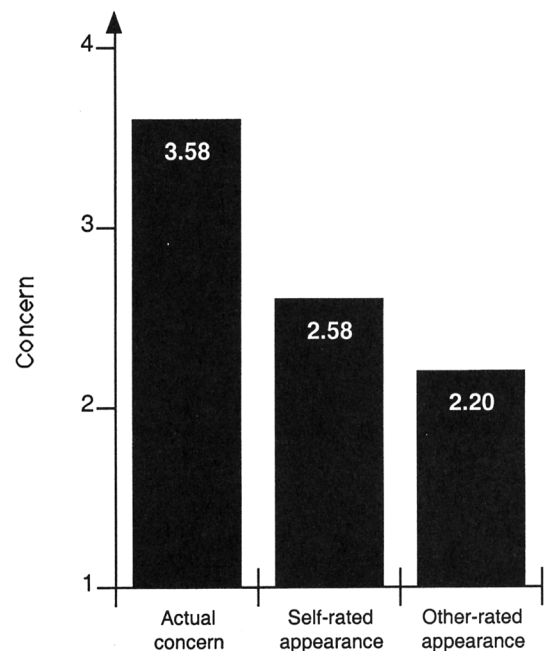


Figure 2. Mean self-ratings of actual concern, mean self-ratings of the appearance of concern, and mean ratings of how concerned each participant appeared to the other members of the group, Study 3a.

Two significant findings emerged. First, participants rated themselves, on average, as moderately concerned but as appearing somewhat less concerned than they actually were,  $t(9) = 4.09, p < .005$ . This finding could be interpreted in several ways, of course, but note that it is consistent with the anchoring and adjustment interpretation we have proposed. Participants doubtless recognized that the full nature of their emotional experience was not available to observers, and so they adjusted their ratings of how they appeared to everyone else downward from the anchor of their own experience. Their "appearance" ratings were therefore less extreme than their ratings of how they actually felt.

But were their adjustments insufficient, as such adjustments typically are? Did participants overestimate how concerned they looked to others? A comparison of how concerned participants thought they looked with how concerned others thought they looked indicates that they did indeed. As predicted, participants believed they appeared more concerned over the rule-breaking confederate than others rated them as appearing,  $t(9) = 2.84, p < .02$ .

Three participants, across three different experimental sessions, expressed some suspicion of either the procedures or the confederate. In addition, 2 participants engaged in some minimal form of intervention, as recorded by the confederate (e.g., "Are you supposed to be giving us hints?"). It is noteworthy that so few participants attempted to intervene and how mild the attempts at intervention were: No participant ever forcefully confronted the confederate, nor did anyone ever retrieve the experimenter to report the confederate's rule-breaking behavior.

Nevertheless, because a few participants did engage in some minimal forms of intervention or express some suspicion, we recomputed all analyses, eliminating all suspicious and interventionist individuals from the data set. Because our analyses were conducted at the group level, this could be accomplished simply by removing these individuals' data from the group averages while retaining data from all other participants in that session. These analyses again revealed support for our hypotheses: Participants believed they expressed less concern than they felt,  $t(9) = 5.00, p < .0005$ , but still believed that they appeared more concerned than they were rated by others,  $t(9) = 3.09, p < .02$ .

### Discussion

This experiment provides clear support for our hypothesis. Participants felt they appeared more concerned about the confederate's problematic behavior than they actually appeared, as rated by the other members of their group. Thus, the illusion of transparency may indeed help explain why people sometimes do not intervene in emergencies: Because people think their own alarm is more apparent than it really is, they assume that others are comparatively less alarmed, leading them to conclude, in turn, that the situation is not really an emergency and that no intervention is required.

### Study 3b

To ensure that the results of Study 3a were reliable, we conducted a replication that followed the same procedure, but with

one exception: In an effort to boost participants' feelings of concern over the confederate's rule-breaking behavior, we adopted a procedure intended to increase participants' feelings of personal involvement in the experiment and accountability for its outcome. Specifically, in place of the "free-for-all" group effort with which participants solved the anagrams in Study 3a, we adopted a "round-robin" procedure whereby each participant functioned as the exclusive solver, one by one, for 2.5 min. During each participant's allotted 2.5 min, they alone were allowed to solve the anagrams, and the confederate escalated her rule-breaking assistance gradually. Once again, we predicted a discrepancy between participants' intuitions about how concerned they appeared and their actual appearance of concern.

### Method

A total of 80 Cornell University students participated in 1 of 20 groups, each of which consisted of 4 participants and 1 confederate. Beyond the change noted above, all portions of this experiment, including the cover story and dependent measures, remained as before.

### Results

Figure 3 depicts participants' average self-ratings of actual concern, their average self-ratings of how concerned they

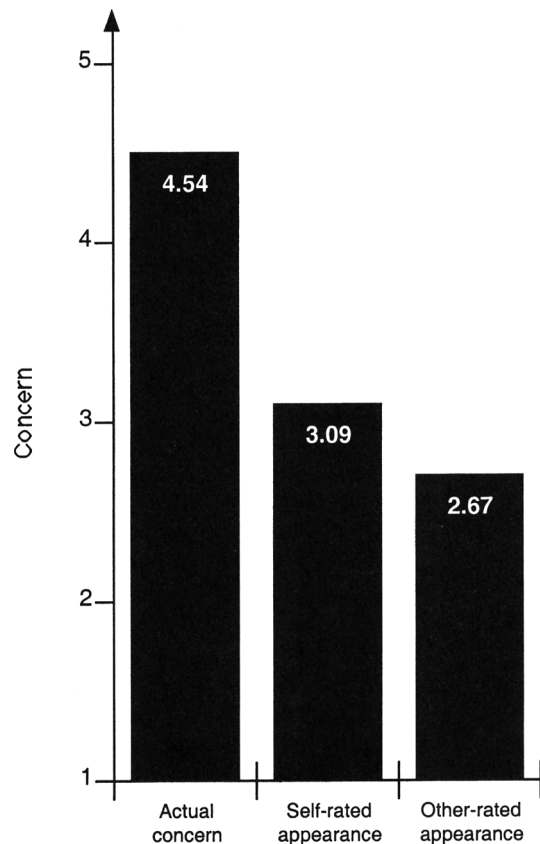


Figure 3. Mean self-ratings of actual concern, mean self-ratings of the appearance of concern, and mean ratings of how concerned each participant appeared to the other members of the group, Study 3b.

thought they appeared, and the average rating of how concerned each participant appeared to the other members of the group. Once again, the results conform to our predictions. First, participants rated themselves as appearing less concerned than they actually were,  $t(19) = 7.08, p < .0001$ . (Note that the new procedure appears to have been successful in increasing participant's self-reported concern by approximately one scale point over the levels observed in Study 3a.) This finding, once again, fits nicely with the anchoring and adjustment interpretation of the illusion of transparency. Participants recognized that the full extent of their alarm would not be picked up by the others in the group, and so they adjusted their estimates of how concerned they appeared to others downward from how concerned they actually felt. These adjustments were insufficient, however, resulting in a significant illusion of transparency: Participants rated themselves as appearing significantly more concerned than they were rated, on average, by their fellow participants,  $t(19) = 3.09, p < .01$ .

As before, a few individuals (3 participants in three different groups) expressed some suspicion about either the procedures or the confederate, and some (11 participants in eight different groups) did intervene. Though somewhat more frequent and direct than in Study 3a (e.g., "That's cheating, don't do that!"), the interventions were, on the whole, rather mild. Also as before, an analysis that excluded suspicious and interventionist participants revealed a statistically significant illusion of transparency: Individuals believed they appeared more concerned than they were rated by others,  $t(19) = 3.01, p < .01$ .

### Discussion

The results of Study 3b echo those observed in Study 3a: Participants confronted with a potential emergency situation were prone to an illusion of transparency, believing that their feelings of concern over a rule-breaking confederate were more observable than they actually were. In light of these results, once again, we believe that the illusion of transparency may provide one answer to the question of why people are reluctant to intervene in emergencies. Individuals may mistake the true source of others' calm exteriors not only because they fail to appreciate that others have attempted to conceal their feelings of alarm (Miller & McFarland, 1987, 1991) but also because they may underestimate their own ability to do so. Bystanders may tend to believe that their own alarm is more apparent than it really is, leading them to conclude that others are indeed comparatively less alarmed than they are. Tragically, this may lead them to infer that the situation is not an emergency and that no intervention is required.

### General Discussion

We obtained consistent support for the illusion of transparency across all studies reported in this paper. In Studies 1a, 1b, and 1c, participants who were asked to tell lies overestimated the proportion of observers who would detect their deception. Their lies were far less apparent than they suspected. We ruled out several alternative interpretations of this finding, including the self-as-target bias (Fenigstein, 1984; Zuckerman et al., 1983), the curse of knowledge (e.g., Keysar & Bly, 1995), and

an abstract (if mistaken) theory that lies are readily detected. In Studies 2a and 2b, participants who sampled a foul-tasting drink while attempting to maintain a blank facial expression nevertheless felt as though their feelings of disgust left telltale traces. Participants overestimated the number of observers who could identify which drinks were foul-tasting. Finally, in Studies 3a and 3b, participants who witnessed a confederate violate the rules of an experiment thought they exhibited more concern over the confederate's misbehavior than their fellow participants believed they did. That is, participants overestimated the extent to which their feelings of concern and alarm over the confederate's actions were apparent to others.

Elsewhere, we have demonstrated the illusion of transparency with respect to additional emotional states. In one set of studies (Savitsky, 1997), individuals who had to deliver public presentations overestimated how nervous they appeared. In another set (Van Boven, Medvec, & Gilovich, 1998), parties to a negotiation thought that their likes and dislikes—what they valued highly and what they could do without—were more apparent to the person with whom they wanted to strike a deal than was actually the case (see Vorauer & Claude, 1998).

It thus appears that the illusion of transparency is a robust phenomenon that applies to a host of different internal states. We attribute this illusion to an anchoring and adjustment bias. When individuals attempt to determine how apparent their internal states are to others, they begin the process of judgment from their own subjective experience. The adjustments they make from this anchor—adjustments that stem from the recognition that others are not as privy to their internal states as they are themselves—tend to be insufficient (e.g., Tversky & Kahneman, 1974). The net result, as we have shown, is a residual effect of one's own phenomenology and a feeling that one's internal states are more apparent to others than is actually the case.

Although we did not conduct an explicit test of this proposed mechanism, a variety of data that support the anchoring and adjustment interpretation was obtained from all three sets of experiments reported here. First, we found in Study 1c that participants who scored high on a measure of private self-consciousness were particularly prone to feelings of transparency. These individuals, because of their relatively inward self-focus, are likely to have a particularly keen sense of their own internal experience. Their processes of judgment should therefore begin from a higher anchor value and thus result in higher estimates, which they did. Second, we found in Studies 1a and 2a that the illusion exists only when a person is experiencing a palpable emotional state. In particular, participants exhibited an illusion of transparency when they were lying but not when telling the truth, and when they had sipped a foul-tasting drink but not a pleasant one. These findings likewise support the anchoring and adjustment interpretation because they show that when there is no pronounced internal experience to adjust from, there is no illusion of transparency. Finally, we found in Studies 3a and 3b that although participants thought they looked more alarmed over a troublesome situation than others thought they appeared, they nonetheless thought they looked less alarmed than they actually felt. This latter difference is consistent with the central premise that participants were adjusting downward from the anchor of their own experience.

Although together these findings offer substantial support for the anchoring and adjustment interpretation we have proposed, each is subject to alternative interpretation. More definitive support must thus await the outcome of experiments explicitly designed to test the role of the anchoring and adjustment processes we have proposed. There are at least two different ways to conduct such a test. One would involve manipulating the strength of the initial anchor. What would happen, for example, if one instituted a delay between, say, sipping a foul-tasting solution and estimating the number of observers who could tell it was unpleasant? By diminishing the internal representation of the foul taste, we would expect the participants' delayed estimates to begin from a lower anchor value and thus exhibit less of an illusion of transparency. The second way to test this underlying mechanism is to manipulate not the anchor but the adjustment (Gilbert, 1989; Gilbert, Pelham, & Krull, 1988). What would happen, for example, if participants were made cognitively busy while sipping the foul-tasting drink? The taste system being the way it is, we would expect the experience of the foul-tasting drink to be an automatic process but the adjustment from that experience to be more controlled. Busyness should then dampen the adjustment but not the anchor, and thus increase the magnitude of the illusion of transparency.

Of course, further studies such as these must be conducted to delineate the exact nature of the underlying causes of the illusion of transparency. Nevertheless, it is already clear from existing research that the potential consequences of this phenomenon are considerable. For example, Studies 3a and 3b lend support to the notion that the illusion of transparency may be one reason why people often fail to intervene in emergencies (Darley & Batson, 1973; Darley & Latané, 1968; Latané & Darley, 1970; see also Miller & McFarland, 1987, 1991). Each bystander may privately believe that he or she appears more concerned over the potential emergency than the others. This leads to the (sometimes tragic) conclusion that the situation is not actually an emergency and that no intervention is required.

There may be other negative consequences of the illusion of transparency as well. In particular, the illusion may play a significant role in interpersonal misunderstanding and conflict (cf. L. Ross & Ward, 1995). Among married couples, for example, the ability to read one another's nonverbal communications has been shown to be related to marital satisfaction (Gottman & Porterfield, 1981; Kahn, 1970; Noller, 1980). We suggest, further, that people's *beliefs* about how well they communicate their inner thoughts and emotions may also be important. An exaggerated view of how well one has conveyed one's inner state, or an unrealistic expectation that one's partner be able to "read one's mind," may be a source of significant interpersonal discord ("You should have known that I . . ."). Couples especially prone to an illusion of transparency may be especially prone to conflict.

Inherent in this speculation is the suggestion that the illusion of transparency applies not just when individuals attempt to conceal their inner states—as in all of the experiments reported here—but also when people attempt to communicate them. Here too, we suggest that people may succumb to an illusion of transparency, believing that they have conveyed their emotional states with greater fidelity than is actually the case. We have obtained preliminary evidence that people do indeed overesti-

mate how clearly they can communicate their emotional states to observers via their facial expressions (Savitsky, 1997). Although one may feel emotions such as love, envy, or disappointment with great intensity, such intensity is typically not mirrored fully in one's facial expressions—a discrepancy that is easy to lose sight of when imagining how one appears to others.

Are all internal states necessarily subject to the illusion of transparency? Probably not. In many cases, people may be well calibrated regarding how much their thoughts, feelings, and emotions are apparent to others. When might individuals succumb to the illusion and when not? We believe at least two variables are important. First, as we have suggested, an individual must experience the internal state with some intensity. Many states are inherently less intense, and we suspect that such states are less conducive to an illusion of transparency. Our data from Studies 1a and 2a provide some support for this notion. Because neither the experience of telling the truth nor sampling a pleasant-tasting drink evokes a compelling phenomenological experience, these conditions yielded no illusion of transparency.

A second precondition for the illusion of transparency may be that the individual believe there is some route by which the internal state can leak out and be detected by others. The internal states we have examined—detection apprehension, disgust, and alarm, as well as speech anxiety (Savitsky, 1997)—are all states for which people have theories about the cues one can use to detect their presence (e.g., the belief, true or not, that one can tell if an individual is lying by looking him or her "straight in the eye"). For those internal states for which people do not have such theories about leakage, one would anticipate less of an illusion of transparency. People may be less likely, for example, to believe that their test-taking anxiety can be detected than to believe their anxiety over delivering a speech can be detected.

Because the illusion of transparency stems in part from people's acute awareness of their internal states, the research on self-awareness and what it suggests about the ability to take another person's perspective might also be expected to have something to say about the question of when the illusion is likely to be more or less prevalent. Unfortunately, the existing literature is mixed on this issue. According to some theorists, self-consciousness enhances an individual's ability to take another person's perspective because it focuses an individual's attention on his or her status as a social object (Duval & Wicklund, 1972; Hass, 1984; Hass & Eisenstadt, 1990; Stephenson & Wicklund, 1983; Wicklund, 1975). As Stephenson and Wicklund (1983) put it, "The impact of self-awareness . . . is one of orienting the person toward the other's point of view, with a consequent heightened sense of the differences between one's own orientation . . . and that of the other" (p. 69). Vorauer and Ross (1998), however, present evidence that self-focus does not facilitate perspective taking, and even inhibits it, when the self is the object of judgment. In this case, "the information that people need to ignore to see things from their audience's perspective is precisely the same information that is particularly salient to them" (p. 7). It seems that a clearer understanding of the impact of self-focused attention on the illusion of transparency must await the results of future research.

As mentioned earlier, the illusion of transparency bears some resemblance to a variety of established phenomena. One phe-

nomenon that we were particularly concerned with disentangling from the illusion is the curse of knowledge. As we noted, one can be "cursed" by one's own knowledge in the sense that it can be difficult to set aside that knowledge when imagining how things appear to others (Camerer et al., 1989; Keysar, 1994; Keysar & Bly, 1995; Keysar et al., 1995). We argued that this phenomenon cannot fully account for our findings because yoked partners who were provided with all of the same knowledge as the liars in Study 1c and the tasters in Study 2b did not overestimate the leakage of those individuals' internal states as much as the individuals did themselves.<sup>7</sup> Note that at a deeper level, however, the two phenomena may be inextricably linked. In particular, the liars and tasters in our studies probably were cursed by their own knowledge, but it was their experiential knowledge of their own internal states, and not their abstract knowledge of which statements were lies or which drinks tasted bad, that prompted the illusion of transparency. Thus, the illusion of transparency may be something of an extreme case of the curse of knowledge, a case in which the attempt to take another person's perspective is hampered not by awareness of some pallid information but by the internal experience of one's own emotional states. In this way, our findings extend the curse of knowledge to encompass internal feelings and sensations as well as abstract, propositional knowledge.

The illusion of transparency also resembles, albeit more remotely, the self-as-target bias, or people's tendency to believe that they are the object of others' thoughts and actions more than is actually the case (Fenigstein, 1984; Zuckerman et al., 1983). Students "just know," for example, that the teacher is about to call on them—particularly if they are unprepared. The self-as-target bias appears to reflect the same difficulty in getting beyond one's own phenomenological experience that lies at the core of the illusion of transparency. In many situations, one's conscious attention is largely devoted to imagining what would happen or what one would do if another's actions were directed at the self. Because people are reluctant to view their thoughts as random or unwarranted, they use a bit of heuristic reasoning and conclude that "if I am so preoccupied with this concern, there must be a reason for it" (Savitsky, Medvec, Charlton, & Gilovich, 1998; Schwarz, 1990; Schwarz & Clore, 1983). Threats that one prepares for are threats that seem real.

The two phenomena are particularly similar when a person's thoughts are devoted to his or her shortcomings that the other person might exploit. Consider the student who is unprepared for class, or the Little Leaguer who is positioned in right field to minimize a fielding liability. Both are intensely focused on their own shortcomings, and both may believe that their fear is leaking out and can be picked up, either by a sadistic teacher or an opportunistic batter. In cases such as these, the self-as-target bias and illusion of transparency blend together considerably.

In summary, we have presented evidence that people are subject to an illusion of transparency. One (unsettling) conclusion that follows from this is that people are more skilled at dissembling than they suppose. We are better liars than we realize. So whether we are trying to get away with murder, like Poe's protagonist, or attempting to pull off the lesser feat of choking down a meal without our host detecting our feelings of disgust,

we may be wise to remember that our hearts are not as telltale as we think.

<sup>7</sup> Note that we observed no curse-of-knowledge effect in our lie-detection paradigm—that is, observers who were cursed with the same abstract knowledge as the liars did not overestimate the likelihood of the lies being detected. We did, however, observe a significant curse of knowledge in our tasting study: Observers overestimated the audience's ability to detect the foul-tasting drink, though not to the same extent as the tasters themselves did. This may be because the yoked partners in the tasting study knew everything the tasters did (i.e., they knew when a drink was foul-tasting and, because they sampled the drink themselves, they knew how foul-tasting it was). In the lie-detection paradigm, on the other hand, the yoked partners knew that a lie was being told, but not how much of a lie—whether it was a mild distortion or a real whopper.

## References

- Barr, C. L., & Kleck, R. E. (1995). Self-other perception of the intensity of facial expressions of emotion: Do we know what we show? *Journal of Personality and Social Psychology*, 68, 608–618.
- Camerer, C., Lowenstein, G., & Weber, M. (1989). The curse of knowledge in economic settings: An experimental analysis. *Journal of Political Economy*, 97, 1232–1253.
- Carmona, J. (1993, September 7). Bystanders stop suicide attempt. *The Cornell Daily Sun*, pp. 1, 8.
- Carver, C. S., & Glass, D. C. (1976). The self-consciousness scale: A discriminant validity study. *Journal of Personality Assessment*, 40, 169–172.
- Darley, J. M., & Batson, C. D. (1973). "From Jerusalem to Jericho": A study of situational and dispositional variables in helping behavior. *Journal of Personality and Social Psychology*, 27, 100–108.
- Darley, J. M., & Latané, B. (1968). Bystander intervention in emergencies: Diffusion of responsibility. *Journal of Personality and Social Psychology*, 8, 377–383.
- DePaulo, B. M., Zuckerman, M., & Rosenthal, R. (1980). Humans as lie detectors. *Journal of Communication*, 30, 129–139.
- Duval, S., & Wicklund, R. A. (1972). *A theory of objective self-awareness*. New York: Academic Press.
- Ekman, P. (1985). *Telling lies: Clues to deceit in the marketplace, marriage, and politics*. New York: Norton.
- Fenigstein, A. (1984). Self-consciousness and the overperception of self as a target. *Journal of Personality and Social Psychology*, 47, 860–870.
- Fenigstein, A. (1987). On the nature of public and private self-consciousness. *Journal of Personality*, 55, 543–553.
- Fenigstein, A., & Abrams, D. (1993). Self-attention and the egocentric assumption of shared perspectives. *Journal of Experimental Social Psychology*, 29, 287–303.
- Fenigstein, A., Scheier, M. F., & Buss, A. H. (1975). Public and private self-consciousness: Assessment and theory. *Journal of Consulting and Clinical Psychology*, 43, 522–527.
- Fischhoff, B. (1975). Hindsight ≠ foresight: The effect of outcome knowledge on judgment under uncertainty. *Journal of Experimental Psychology: Human Perception and Performance*, 1, 288–299.
- Fischhoff, B. (1982). For those condemned to study the past: Heuristics and biases in hindsight. In D. Kahneman, P. Slovic, & A. Tversky (Eds.), *Judgment under uncertainty: Heuristics and biases* (pp. 335–354). New York: Cambridge University Press.
- Fischhoff, B., & Beyth, R. (1975). "I knew it would happen"—Remembered probabilities for once-future things. *Organizational Behavior and Human Performance*, 13, 1–16.
- Gilbert, D. T. (1989). Thinking lightly about others: Automatic compo-

- nents of the social inference process. In J. S. Uleman & J. A. Bargh (Eds.), *Unintended thought* (pp. 189–211). New York: Guilford Press.
- Gilbert, D. T., Pelham, B. W., & Krull, D. S. (1988). On cognitive busyness: When person perceivers meet persons perceived. *Journal of Personality and Social Psychology*, 54, 723–740.
- Gilovich, T., Kruger, J., Medvec, V. H., & Savitsky, K. (1998). *Those bad hair days: Biased estimates of how variable we look to others*. Unpublished manuscript, Cornell University.
- Gilovich, T., Medvec, V. H., & Savitsky, K. (1998). *The spotlight effect in social judgment: An egocentric bias in estimates of the salience of one's own action*. Unpublished manuscript, Cornell University.
- Gottman, J. M., & Porterfield, A. L. (1981). Communicative competence in the nonverbal behavior of married couples. *Journal of Marriage and the Family*, 43, 817–824.
- Griffin, D. W., & Ross, L. (1991). Subjective construal, social inference, and human misunderstanding. In M. P. Zanna (Ed.), *Advances in experimental social psychology* (Vol. 24, pp. 319–359). San Diego, CA: Academic Press.
- Hass, R. G. (1984). Perspective taking and self-awareness: Drawing an E on your forehead. *Journal of Personality and Social Psychology*, 46, 788–798.
- Hass, R. G., & Eisenstadt, D. (1990). The effects of self-focused attention on perspective-taking and anxiety. *Anxiety Research*, 2, 165–176.
- Jacowitz, K. E., & Kahneman, D. (1995). Measures of anchoring in estimation tasks. *Personality and Social Psychology Bulletin*, 21, 1161–1166.
- Kahn, M. (1970). Nonverbal communication and marital satisfaction. *Family Process*, 9, 449–456.
- Keysar, B. (1994). The illusory transparency of intention: Linguistic perspective taking in text. *Cognitive Psychology*, 26, 165–208.
- Keysar, B., & Bly, B. (1995). Intuitions of the transparency of idioms: Can one keep a secret by spilling the beans? *Journal of Memory and Language*, 34, 89–109.
- Keysar, B., Ginzler, L. E., & Bazerman, M. H. (1995). States of affairs and states of mind: The effect of knowledge on beliefs. *Organizational Behavior and Human Decision Processes*, 64, 283–293.
- Knapp, M. L., & Comadena, M. E. (1979). Telling it like it isn't: A review of theory and research on deceptive communications. *Human Communication Research*, 5, 270–285.
- Kraut, R. E. (1980). Humans as lie detectors: Some second thoughts. *Journal of Communications*, 30, 209–216.
- Latané, B., & Darley, J. (1970). *The unresponsive bystander: Why doesn't he help?* New York: Appleton-Century-Crofts.
- Miller, D. T., & McFarland, C. (1987). Pluralistic ignorance: When similarity is interpreted as dissimilarity. *Journal of Personality and Social Psychology*, 53, 298–305.
- Miller, D. T., & McFarland, C. (1991). When social comparison goes awry: The case of pluralistic ignorance. In J. Suls & T. Wills (Eds.), *Social comparison: Contemporary theory and research* (pp. 287–313). Hillsdale, NJ: Erlbaum.
- Noller, P. (1980). Misunderstandings in marital communication: A study of couples' nonverbal communication. *Journal of Personality and Social Psychology*, 39, 1135–1148.
- Piaget, J. (1959a). *Judgment and reasoning in the child*. Paterson, NJ: Littlefield, Adams, & Co. (Original work published 1928)
- Piaget, J. (1959b). *The language and thought of the child*. New York: Humanities Press. (Original work published 1926)
- Poe, E. A. (1976). The tell-tale heart. In S. Levine & S. Levine (Eds.), *The short fiction of Edgar Allan Poe* (pp. 259–262). Indianapolis, IN: Bobbs-Merrill.
- Quattrone, G. A. (1982). Overattribution and unit formation: When behavior engulfs the person. *Journal of Personality and Social Psychology*, 42, 593–607.
- Quattrone, G. A., Lawrence, C. P., Finkel, S. E., & Andrus, D. C. (1984). *Explorations in anchoring: The effects of prior range, anchor extremity, and suggestive hints*. Unpublished manuscript, Stanford University.
- Ross, L., & Ward, A. (1995). Psychological barriers to dispute resolution. In M. P. Zanna (Ed.), *Advances in experimental social psychology* (Vol. 27, pp. 255–304). San Diego, CA: Academic Press.
- Ross, L., & Ward, A. (1996). Naive realism in everyday life: Implications for social conflict and misunderstanding. In T. Brown, E. Reed, & E. Turiel (Eds.), *Values and knowledge* (pp. 103–135). Hillsdale, NJ: Erlbaum.
- Ross, M., & Sicoly, F. (1979). Egocentric biases in availability and attribution. *Journal of Personality and Social Psychology*, 37, 322–336.
- Savitsky, K. (1997). *Perceived transparency and the leakage of emotional states: Do we know how little we show?* Unpublished doctoral dissertation, Cornell University.
- Savitsky, K., & Gilovich, T. (1998). *Is our absence as conspicuous as we think?* Unpublished manuscript, Cornell University.
- Savitsky, K., Medvec, V. H., Charlton, A. E., & Gilovich, T. (1998). "What, me worry?": Arousal, misattribution, and the effect of temporal distance on confidence. *Personality and Social Psychology Bulletin*, 24, 529–536.
- Schwarz, N. (1990). Feelings as information: Informational and motivational functions of affective states. In E. T. Higgins & R. Sorrentino (Eds.), *Handbook of motivation and cognition: Foundations of social behavior* (Vol. 2, pp. 527–561). New York: Guilford Press.
- Schwarz, N., & Clore, G. L. (1983). Mood, misattribution, and judgments of well-being: Informative and directive functions of affective states. *Journal of Personality and Social Psychology*, 45, 513–523.
- Stephenson, B., & Wicklund, R. A. (1983). Self-directed attention and taking the other's perspective. *Journal of Experimental Social Psychology*, 19, 58–77.
- Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, 185, 1124–1131.
- Van Boven, L., Medvec, V. H., & Gilovich, T. (1998). *The illusion of transparency in negotiations*. Unpublished manuscript, Cornell University.
- Vorauer, J. D. (in press). The other side of the story: Transparency estimation in social interaction. In G. B. Moskowitz (Ed.), *Future directions in social cognition*. Mahwah, NJ: Erlbaum.
- Vorauer, J. D., & Claude, S. D. (1998). Perceived versus actual transparency of goals in negotiation. *Personality and Social Psychology Bulletin*, 24, 371–385.
- Vorauer, J. D., & Ross, M. (1998). *Self-awareness and feeling transparent: Failing to suppress one's self*. Unpublished manuscript, University of Manitoba, Winnipeg, Manitoba, Canada.
- Wicklund, R. A. (1975). Objective self-awareness. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 8, pp. 233–275). New York: Academic Press.
- 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.
- Zuckerman, M., Kernis, M. H., Guarniera, S. M., Murphy, J. F., & Rappoport, L. (1983). The egocentric bias: Seeing oneself as cause and target of others' behavior. *Journal of Personality*, 51, 621–630.

Received August 4, 1997

Revision received February 9, 1998

Accepted February 15, 1998 ■