
When Social Worlds Collide: Overconfidence in the Multiple Audience Problem

Leaf Van Boven

Justin Kruger

Cornell University

Kenneth Savitsky

Williams College

Thomas Gilovich

Cornell University

Individuals sometimes try to convey different identities to different people simultaneously or to convey certain information to one individual while simultaneously concealing it from another. How successfully can people solve these multiple audience problems and how successfully do they think they can? The research presented here corroborates previous findings that people are rather adept at such tasks. In Study 1, participants who adopted different identities in preliminary interactions with two other participants (acting the part of a studious nerd with one and a fun-loving party animal with the other) were able to preserve these identities when they interacted subsequently with both individuals at the same time. In Study 2, participants were able to communicate a secret word to one audience while simultaneously concealing it from another. Despite their skill at these tasks, however, participants in both studies were overconfident in their abilities, believing that they were better able to solve these multiple audience problems than they actually were.

Social interactions are often strategic. Depending on our goals, we may attempt to convey different information—even different personalities altogether—to different people in different situations (Cooley, 1902; Fleming, 1994; Goffman, 1959, 1963; James, 1892; Jones & Pittman, 1982; Mead, 1934; Schlenker, 1980; Swann, 1987). As James (1892) put it,

Many a youth who is demure enough before his parents and teachers, swears and swaggers like a pirate among his “tough” young friends. We do not show ourselves to our children as to our club-companions, to our customers as to the laborers we employ, to our own masters and employers as to our intimate friends. (p. 42)

Although switching from one alternate self to another as the situation demands can be a formidable task, it is not an impossible one. Most of us remember that it is our children who get the baby talk and our “masters” the respectful reverence—not the other way around. But what happens when we face the boss and our children at the same time? What happens, in other words, when our social worlds collide? A wedding reception, college reunion, or company picnic may bring one into simultaneous contact with individuals with whom one has cultivated very different identities. One’s grandparents, college chums, coworkers, boss, and children may all be assembled in the same location. What is one to do? Can people successfully maintain their multiple identities in such situations, or is the inconsistency in who they are apparent to others?

Only recently have psychologists begun to address these issues empirically and determine whether people are able to handle such multiple audience problems successfully (Clark & Schaefer, 1987; Fleming & Darley, 1991; Fleming, Darley, Hilton, & Kojeten, 1990). The predominant finding in these studies is that people are rather successful in such predicaments, conveying one

Authors’ Note: This research was supported by a research grant (No. SBR9319558) from the National Science Foundation. Portions of this research were presented at the 1998 annual meeting of the American Psychological Society in Washington, DC. Justin Kruger is now at the University of Illinois at Urbana-Champaign. Correspondence concerning this article should be addressed to Leaf Van Boven, Cornell University, Department of Psychology, Uris Hall, Ithaca, NY 14853-7601; e-mail: leaf.van.boven@cornell.edu.

PSPB, Vol. 26 No. 5, May 2000 619-628

© 2000 by the Society for Personality and Social Psychology, Inc.

impression or message to one person (or group of people) while simultaneously communicating a different impression or message to others. In one experiment, for example, participants were able to convey the location of a secret map to a group of their friends while simultaneously concealing its location from a group of strangers (Fleming et al., 1990, Study 2). In another study, participants were able to persuade their friends to select certain popular songs as secret songs while leading others to select decoys (Fleming et al., 1990, Study 3).

Because multiple audience problems are common, and because the decision to enter or avoid such predicaments depends in part on one's confidence that one can navigate them successfully, it is important to ascertain not only how skilled people are in these dilemmas but also how skilled they *think* they are. Are individuals aware of their ability to pull off such complex social maneuvers? Anecdotal evidence from the studies of Fleming and colleagues (1990) suggests that the answer might be "no." Despite participants' successes in those studies, most described the task as daunting. Following one of Fleming et al.'s studies, for example, "all 8 [participants] reported that the task was extremely difficult. Nonetheless, all senders were able to complete the task" (p. 603). In another, "all 10 senders reported that the task seemed extremely difficult. . . . Nonetheless, all senders were again able to complete the task" (p. 606).

This discrepancy between confidence and success suggests that individuals may tend to be underconfident in their ability to solve multiple audience problems—that they are more skilled than they realize. Thus, when we bring our college friends (who know us as party animals) home to meet our high school friends (who view us as bookish), we may underestimate our ability to preserve our discrepant identities.

We believe, however, that this issue warrants further scrutiny. First, Fleming and colleagues' (1990) reports are based on anecdotal observations in studies designed to explore people's actual ability to solve multiple audience problems, not people's perceptions of their ability. Second, believing a task to be difficult need not imply underconfidence. One can believe that a task is difficult and still have an accurate assessment of one's competence, or be overconfident.

Indeed, there may be good reasons to expect overconfidence, rather than underconfidence, in multiple audience problems. First, there is a large literature demonstrating that people tend to be overconfident in gauging whether they have been successful at a task (e.g., Dunning, Griffin, Milojkovic, & Ross, 1990; Keren, 1987; Lichtenstein, Fischhoff, & Phillips, 1982; Oskamp, 1965; Vallone, Griffin, Lin, & Ross, 1990; Wright, Rowe, Boger, & Gammack, 1994).

A second reason to expect communicators to be overconfident in their ability to solve multiple audience problems derives from the strategy they most often employ. Although people faced with multiple audience problems may adopt a variety of approaches and techniques, research has shown that the most popular and successful strategy is for a communicator to refer to information that is "mutually known, believed, or assumed based on personal experiences" (Clark & Schaefer, 1987, p. 211). That is, a communicator can capitalize on information shared between himself or herself and one audience but not another (Clark & Schaefer, 1987; Fleming & Darley, 1991; Fleming et al., 1990). A woman may attempt to feign agreement with her boss's praise of conservative talk-show host Rush Limbaugh, for example, by stating in front of her liberal friends that Limbaugh's show has joined opera—an art form she and her friends passionately despise—atop her list of great entertainment.

Given that individuals use shared information to solve multiple audience problems, it stands to reason that the more unique knowledge they share with their intended audience, the more successful they will be. Likewise, the more unique knowledge they think they share with their intended audience, the more confident they will be. To the extent that actual shared knowledge aligns with perceived shared knowledge, an individual's confidence will be well placed. If, however, there is a misalignment between actual and perceived shared knowledge, one's confidence may be off the mark.

Evidence from two lines of research suggests that there may indeed be a systematic misalignment between actual and perceived shared knowledge that would lead to overconfidence in the multiple-audience problem. First, people tend to overestimate the knowability of the things they themselves know (Camerer, Loewenstein, & Weber, 1989; Fussell & Krauss, 1991, 1992; Keysar, Ginzler, & Bazerman, 1995). The process of determining whether another individual shares some bit of knowledge with oneself can be a difficult one, and one's own knowledge can function as something of an anchor. Although people realize that others do not share their entire knowledge base, the adjustments they make to capture others' perspectives, similar to adjustments to anchors generally, tend to be insufficient (Jacowitz & Kahneman, 1995; Quattrone, 1982; Quattrone, Lawrence, Finkel, & Andrus, 1984; Tversky & Kahneman, 1974). As a result, people may overestimate the extent to which the things they know are also known by others. A communicator's own knowledge about the varieties and peculiarities of wines from the Loire valley, for example, may lead to inflated estimates of the extent to which a particular audience is similarly informed.

A second reason people may be miscalibrated in their perceptions of shared knowledge is suggested by recent research on the spotlight effect, an egocentric bias whereby people overestimate the extent to which others notice and remember details about them (Gilovich, Kruger, & Medvec, 2000; Gilovich, Medvec, & Savitsky, 2000). Other people simply do not notice and commit to memory as many details about an individual as that individual suspects. As a result, people may overestimate the amount of information about themselves that others remember. A person may be surprised, for instance, to learn that her enthusiastic comments about the genius of Iggy Pop or the Sex Pistols were little noted nor long remembered by her friends.

People's perceptions of the information they share with others are thus likely to be inflated, either because they have anchored on their own repertoire of general knowledge (Fussell & Krauss, 1992) or because they expect others to have paid more attention to them in past interactions than they actually have (Gilovich, Medvec, & Savitsky, 2000). Whether referring to Pavarotti, Pinot Noir, or punk rock, they are likely to think they have pulled off a multiple audience problem more successfully than they actually have. The present research examines this hypothesis. In two experiments, we attempt to determine (a) whether people are indeed overconfident in their ability to solve multiple audience problems and (b) whether perceptions of shared knowledge play a role in this overconfidence.

STUDY 1: WHEN SOCIAL WORLDS COLLIDE

Multiple audience problems come in many forms. Prior research has focused almost exclusively on the variety in which communicators are asked to convey secret messages to one audience while concealing them from another (Clark & Schaefer, 1987; Fleming & Darley, 1991; Fleming et al., 1990). In Study 1, we sought to extend this research to include a multiple audience problem of a different sort—one in which communicators attempt simultaneously to maintain discrepant social identities (cf. James, 1982).

Toward this aim, we asked target participants to adopt two radically different identities with two different participants. With one, targets were instructed to come across as a studious nerd, and with the other, they were instructed to present themselves as a fun-loving party animal. We then brought the targets together with both interaction partners for a third discussion and asked them to do their best to maintain the discrepant identities they had established previously.

Based on prior research (Clark & Schaefer, 1987; Fleming & Darley, 1991; Fleming et al., 1990), we predicted that targets would be somewhat successful in

maintaining their dual identities in the final round of discussion. More important for our purposes, however, we predicted that targets would be overconfident about their ability to maintain distinct identities, believing that they had navigated the multiple audience situation more successfully than they actually had.

METHOD

Participants

Twenty-four Cornell University undergraduates participated in triads in exchange for course credit in various introductory psychology courses.

Procedure

After being screened to ensure that they were unacquainted with one another, the three participants were randomly assigned to the roles of Discussants 1, 2, and 3. For ease of explication, we refer here to Discussants 1 and 2 as Observers 1 and 2 and Discussant 3 as the target.

Participants were informed that the study would consist of three short discussions, two of which would be between only two of them, and a third that would be between all three. Participants were then escorted to separate cubicles where two of them received instructions for the first discussion and the third completed unrelated filler questionnaires.

Discussion 1. The first discussion was between the target and Observer 1. Both participants read that the topic of the discussion would be “the best way to spend one’s time in college” and that they would each deliver an informal 2-minute statement of their position on the topic and then engage in 2 minutes of interactive discussion. Unbeknownst to Observer 1, the target was given some additional, secret instructions. Specifically, targets were instructed to adopt the identity of an extremely studious individual, someone who “prefers studying to socializing,” who “virtually lives in the library,” and who could be described as a “nerd.” Targets were asked to do their best to convince their interaction partner that they were this kind of person without revealing that they had been instructed to do so.

Following the interaction, participants returned to their separate rooms and completed questionnaires. Observers were asked to rate the targets on each of four traits (serious, outgoing, bookish, and sociable) on a 10-point scale anchored at *doesn't describe the person at all* (1) and *describes the person perfectly* (10). They also rated how important they thought each of three activities (doing schoolwork, attending parties, and attending course lectures) was for targets on a 10-point scale anchored at *low priority* (1) and *high priority* (10). Finally, they rated the relative importance to the target of doing schoolwork

and attending parties on a scale anchored at *schoolwork is more important* (1) and *attending parties is more important* (10). Because our key interest pertains to targets' beliefs about observers' impressions of them, targets were asked to estimate, as accurately as possible, how their interaction partner would rate them on each of the scales.

Discussion 2. The second discussion was between the target and Observer 2 and was identical to the first discussion in every way (including topic) except that targets were now asked to portray the opposite identity. That is, targets were asked secretly to convince their interaction partner that they were "a reckless party animal" who prefers to "live life for the moment" and "have a good time" at all costs. Following the interaction, both participants returned to their separate rooms and completed the same set of dependent measures described above. The order of the identities that targets were instructed to convey was counterbalanced across sessions.

Discussion 3. The purpose of the first two rounds of discussion was for the target to establish a different social identity with each observer. Would the target be able to maintain these discrepant personalities when confronted with both observers simultaneously? To find out, we staged a third and final interaction between all three participants. This time, the target was asked to do his or her best to maintain both identities. We chose a discussion topic for the third round that was pertinent to the previous interactions: "What advice would you give to incoming Cornell freshmen about the proper balance between academics and social life?"

The procedure for the final round of discussion was similar to the earlier rounds: Each discussant delivered a 2-minute position statement followed by 2 minutes of open discussion. For this third round of discussion, the target was always "randomly" selected to speak first. After the discussion, participants returned to their rooms and completed the same measures as before. Observers judged the targets on the same measures used after the previous discussions and targets predicted how each observer in turn would rate them. Finally, all participants were probed for suspicion and debriefed.

RESULTS

Because the data within each session were interdependent, all analyses were conducted at the level of the experimental session ($n = 8$) rather than the individual participant ($n = 24$). After appropriate reverse scoring, we averaged observers' ratings and targets' predictions of observers' ratings across the seven items to create overall indexes.¹ For all measures reported below, lower values correspond to rating the target as more of a nerd and higher values to rating the target as more of a party animal.

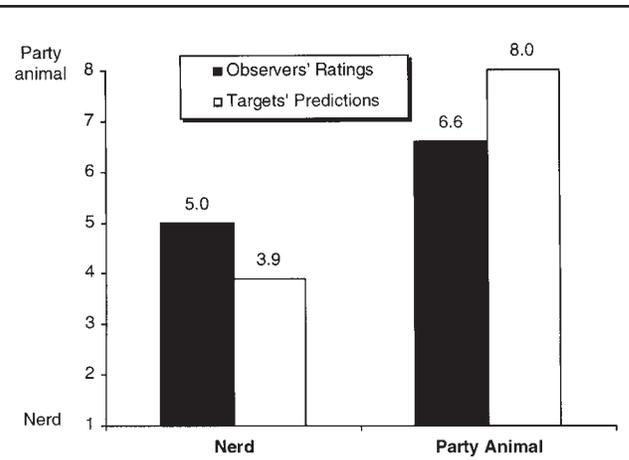


Figure 1 Observers' final impressions of targets and targets' predictions of observers' final impressions.

Constructing dual identities. To provide a fair test of targets' actual and perceived ability to solve the multiple audience problem, it is essential that targets conveyed different identities in the first two discussions. They did. Targets who presented themselves as a party animal were given higher ratings by observers ($M = 8.3$) than were targets who portrayed a nerd identity ($M = 2.8$), paired $t(7) = 10.13, p < .001$. Moreover, targets were well aware of the fact that they had established divergent identities, expecting observers to rate them as either a party animal ($M = 8.9$) or a nerd ($M = 3.3$), depending on the impressions they attempted to convey, paired $t(7) = 5.43, p < .001$.

Maintaining dual identities. How did targets fare in the multiple audience situation in the final round of discussion? As expected, they were able to maintain some semblance of their initial identities in the third discussion. As the dark bars in Figure 1 depict, observers to whom targets had initially conveyed the impression of a party animal rated the target as more of a party animal ($M = 6.6$) than did observers to whom targets had initially conveyed the impression of a nerd ($M = 5.0$), paired $t(7) = 2.18, p < .07$. To be sure, these impressions were reliably less polarized than in the previous rounds, $F(1, 7) = 16.03, p < .01$, but the fact that observers were left with systematically different impressions suggests that targets were at least somewhat successful in handling their multiple audience problem.

Perceived ability to maintain dual identities. Our primary interest was whether targets were overconfident in their ability to maintain divergent social identities. As can be seen from the white bars in Figure 1, the targets' ratings indicate that they did indeed believe they had maintained the discrepant identities of party animal and nerd through the third discussion ($M_s = 8.0$ and 3.9 , respectively), paired $t(7) = 4.55, p < .01$. Were they overconfident

in their assessments of the extent to which they had done so? A 2 (intended identity: party animal vs. nerd) \times 2 (role: observer vs. target) repeated-measures ANOVA revealed a significant interaction, $F(1, 7) = 10.20, p < .05$, indicating that they were. Targets anticipated that the observers to whom they had initially conveyed the impression of a party animal would rate them as more of a party animal than those observers actually did, $t(7) = 2.68, p < .05$. Similarly, targets predicted that the observers to whom they had initially conveyed the impression of a nerd would rate them as more of a nerd ($M = 3.9$) than those observers actually did, although this difference was not statistically significant, $t(7) = 1.41, ns$.

DISCUSSION

The results of Study 1 are consistent with past research in that targets were able to solve a multiple audience problem with some success by maintaining simultaneously discrepant identities (e.g., Clark & Schaefer, 1987; Fleming & Darley, 1991; Fleming et al., 1990). More important, the results also confirmed our hypothesis that participants would be overconfident in their ability to do so: Targets thought they had maintained the two discrepant personalities more successfully than they actually had.

How did participants go about solving their multiple audience problem? Although the data from Study 1 do not address this issue directly, our informal observations suggest two common strategies. First, some targets made ambiguous, “middle-of-the-road” statements about the proper balance between academics and social life that did not explicitly contradict what they had said before (e.g., “On one hand, you’ve got to study, but on the other hand, you’ve got to have fun”). Such diluted presentations were bound to be perceived as quite a contrast to the target’s original statements, which explains why the observers’ ratings at the end of Discussion 3 were not as extreme as those made after Discussions 1 and 2. Still, these diluted presentations do not contradict anything said in the earlier sessions, and, given the extremity of the target’s initial persona, such innocuous statements would presumably leave observers with remnants of their initial impressions intact—thereby solving the target’s multiple audience problem. Indeed, virtually any presentation, innocuous or not, likely would have met with similar success: To the extent that observers were influenced at all by their initial impressions, they would have come away with different final impressions. In other words, the effect may have been the result of the perseverance of the strong first impressions created during the first two rounds of discussion (Ross, Lepper, & Hubbard, 1975). Note, however, that although this account might explain observers’ different impressions,

it provides no ready explanation for the results of primary interest in this study—targets’ overconfidence.

Overconfidence would arise, however, from a second strategy that many targets appeared to employ. In particular, many participants referred strategically to specific comments they had made in the earlier rounds of discussion in the hopes that the observers might recall them and understand their significance; for example, “I think incoming students should spend their fall break the way I did” and “It’s like I said before, Saturdays are good for one thing and one thing only.” Note that this strategy is likely to lead to overconfidence in participants’ assessments of their ability to solve their multiple audience problem if, as we maintain, people tend to overestimate the extent to which others notice and remember the details of their earlier utterances. Such overconfidence, of course, is precisely what we observed.

STUDY 2: SECRETS AND LIES

We conducted Study 2 to examine more directly whether people’s excessive confidence in their ability to solve multiple audience problems stems in part from their exaggerated assumptions of shared knowledge. In particular, we manipulated participants’ perceptions of the amount of knowledge they shared with their intended audience. If the tendency to exaggerate shared knowledge contributes to overconfidence in multiple audience problems, then the more knowledge communicators believe they share with their intended audience, the more overconfident they are likely to be.

Study 2 also sought to generalize the findings from the first study to a multiple audience situation more akin to that studied by Fleming and colleagues (Fleming & Darley, 1991; Fleming et al., 1990). In particular, we sought to verify that communicators would be overconfident in their attempts to convey information to one audience while simultaneously concealing it from another.

With these goals in mind, we asked participants in Study 2 to convey a secret message to one group of fellow participants while simultaneously concealing that message from another group. Half of the participants did so after having engaged in a familiarization exercise with their intended audience. We expected that this manipulation would increase participants’ perceptions of the amount of knowledge they shared with their intended audience and thus increase their confidence in their ability to solve their multiple audience problem.

Of course, by manipulating perceived familiarity, we manipulated actual familiarity as well. To the extent that success in multiple audience problems is achieved by referring to shared knowledge, familiar participants ought to enjoy more success than their unfamiliar counterparts. Given the reasons outlined in the Introduction, however, we doubted that the familiarity manipulation

would affect actual ability as strongly as perceived ability. Thus, we expected “familiar” participants to be more overconfident than unfamiliar participants in their ability to solve the multiple audience problem.

METHOD

Participants

One hundred twelve Cornell University undergraduates participated in groups of 10 or 12 in exchange for course credit in various introductory psychology courses.

Procedure

After ensuring that they did not know one another, participants were randomly divided into two groups (A or B) and given nametags that indicated their group affiliation. Depending on whether there were 10 or 12 participants, both groups had either 5 or 6 members. Group A constituted the familiar group and Group B the unfamiliar group, although they were not identified as such to the participants.

Familiarity manipulation. The two groups were separated for some preliminary exercises. The members of the familiar group assembled in a circle and engaged in a 5-minute familiarization procedure. Specifically, each participant took turns answering several questions in front of the group: “What is your major and why?” “What did you do last summer?” “What is your favorite city and why?” and “What is your most embarrassing moment?” Meanwhile, participants in the unfamiliar group answered these same questions, but on questionnaires and in separate cubicles. Because the two groups were separated during this time, neither group knew the activities of the other.

Secret messages task. All participants then reconvened, sat together with their own group, and were introduced to the secret messages task. It was explained that each of them would take a turn reading aloud a short, prewritten essay. The essays consisted of about 300 words and were selected from a National Teacher Examination study guide (Bobrow et al., 1989; see the appendix for a sample essay). For each essay, we underlined five keywords; participants were to select one of them to be their secret word. Their job was to convey the identity of the secret word to their own group without revealing it to the other group. Speakers were told they were free to use any strategy they wished to accomplish this goal, with two restrictions. First, speakers were required to read the essays verbatim; they were not allowed to ad lib or omit words or phrases. Second, speakers were forbidden from making bodily gestures that could be seen by the members of one group but not the other.² Two participants (one from

each group) were randomly chosen for each of the five or six rounds of the experiment (depending on whether there were 10 or 12 participants in that session). The two selected participants then each chose an essay to read from a list of 12 and were escorted to private cubicles to familiarize themselves with the essay and plan what strategy to use. Meanwhile, members of the audience (i.e., the remaining members of both groups) were told the title of the essay each speaker had chosen and the five keywords contained in each essay (i.e., the words from which each speaker would select his or her secret word).

Once the two speakers indicated they were ready, they returned to the room, one at a time, and read their essays to the assembled participants from behind a podium. Members of the audience—both the speaker’s own group and the other group—were instructed to circle the keyword they believed to be the speaker’s secret word. When speakers were finished reading their essay, they estimated how many members of each group would correctly guess their secret word. (The experimenter took special care to explain exactly how many of each group could be expected to guess the secret word by chance alone.)

When both speakers had read their essays and all questionnaires had been completed, two new speakers were chosen and the procedure was repeated with the remaining essays. This continued until all members of both groups had read essays. All participants were then probed for suspicion and debriefed.

RESULTS

Because the data within each session were interdependent, all analyses were conducted at the level of the experimental session ($n = 10$) rather than the individual participant ($n = 112$).

Actual Ability

Overall, communicators successfully communicated their secret word to a greater proportion of their intended audience than their unintended audience. Across conditions, speakers communicated the secret word to 48% of the members of their intended audience and to 41% of the members of their unintended audience, $F(1, 9) = 7.11$, $p < .05$. As the dark bars in Figure 2 reveal, speakers in the familiar groups were no more successful in this endeavor than were speakers in the unfamiliar groups. Indeed, a 2 (familiar vs. unfamiliar groups) \times 2 (intended vs. unintended audience) ANOVA yielded no hint of an interaction, $F < 1$, or of a main effect for condition, $F(1, 9) = 1.19$, *ns*. Whatever effects the familiarity manipulation might have had, it did not increase speakers’ actual ability to solve their multiple audience problem.

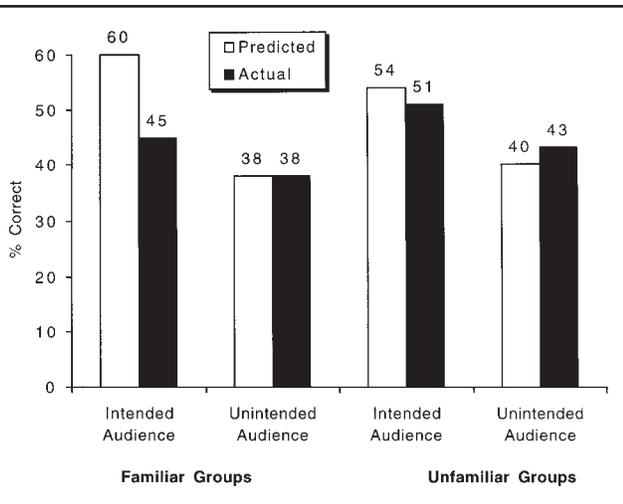


Figure 2 Speakers' predictions of the percentage of their own group and the other group that would be able to identify the secret word and the actual percentage of each group that was able to do so.

Perceived Ability

Although our familiarity manipulation did not affect participants' actual ability to solve their multiple audience problem, it had a significant impact on their perceived ability to do so. The white bars in Figure 2 depict speakers' predictions of the percentages of each group that would be able to identify correctly the secret word. As expected, speakers from the familiar groups were more confident in the ability to solve the multiple audience problem than were speakers from the unfamiliar groups. Familiar speakers believed they had communicated the secret word to 22% more members of their own group than the other group, whereas unfamiliar speakers believed they had communicated the secret word to only 14% more members of their own group, $F(1, 9) = 4.62, p < .06$.³

Did this increase in confidence lead to overconfidence? A 2 (intended vs. unintended audience) \times 2 (perceived vs. actual percentage) ANOVA among participants from familiar groups yielded a significant interaction, $F(1, 9) = 10.24, p < .02$, indicating that it did. Speakers who had engaged in a familiarization exercise were overconfident in their ability to convey their secret message to their intended audience while concealing it from their unintended audience. Speakers from the unfamiliar groups, in contrast, were not, $F(1, 9) = 1.46, ns$.

Interestingly, the overconfidence demonstrated among participants from familiar groups came not from their estimates of their ability to conceal the secret message from their unintended audience ($t < 1$) but from their estimates of their ability to convey it to their intended audience, $t(9) = 3.56, p < .01$. (Participants from unfamiliar groups were not overconfident in their

ability to conceal or convey the secret message, both $t < 1$.) Stated differently, participants from familiar groups were more overconfident in their ability to convey the secret message to their intended audience than were participants from the unfamiliar groups, as indicated by a 2 (familiar vs. unfamiliar group) \times 2 (predicted vs. actual percentage) interaction, $F(1, 9) = 4.62, p < .06$. They were not, however, more overconfident in their ability to conceal the secret message from their unintended audience ($F < 1$). The three-way interaction in a 2 (familiar vs. unfamiliar group) \times 2 (intended vs. unintended audience) \times 2 (predicted vs. actual percentage) ANOVA was not significant, $F(1, 9) = 1.42, ns$. One explanation for this pattern of results may be that participants' attention was more focused on the intended audience than the unintended audience. Communication, after all, usually is focused on conveying information, not concealing it. Although both conveying and concealing information were required to solve the multiple audience problem, to the extent that participants were focused more on conveying information, their strategies (and thus their judgments of confidence) may have been focused there as well.

DISCUSSION

As in Study 1 and in previous research (e.g., Fleming & Darley, 1991; Fleming et al., 1990), participants in this study were modestly successful in solving their multiple audience problem, conveying a secret message more to their intended audience than to their unintended audience. Although participants who had undergone a familiarization exercise with their intended audience were no more successful in this task than were participants who had not (a finding we address in Note 4), they were more confident. As a result, familiar participants were overconfident in their ability to communicate their secret message to their intended audience but unfamiliar participants were not.

These findings highlight the potential for discrepancy between actual and perceived familiarity. Participants who had undergone a familiarization exercise believed they gained a repertoire of shared knowledge to which they could refer in solving their multiple audience problem. For example, one participant, a former marching band member, told his group about an embarrassing incident in which he lost control of his baton during a performance, sending it careening into a crowd of onlookers. When reading his essay, he attempted to use his group's knowledge of this event to help them identify his secret word. Specifically, he pantomimed the event, twirling baton and all, as he read it. To confuse his unintended audience, then, he acted out a variety of other marching band activities as he read each of the other key

words (including some enthusiastic drumming on a nearby metal wastebasket).

However, because people tend to exaggerate the amount of knowledge they share with others, such strategies, although clever, were less effective than the speakers believe. A closer look at the incident described above illustrates this rift. Although the speaker felt confident that his intended audience would remember his embarrassing moment and thus get the secret message he was sending, a postexperimental interview revealed that his presumption was ill-founded. "Oh," asked one of his fellow group members, "were you the one who was in a band?"⁴

GENERAL DISCUSSION

Everyday social interactions can be challenging. Especially formidable are those interactions in which one must simultaneously convey different information to different people. Echoing the findings of past research, both experiments reported here suggest that people are somewhat successful in these multiple audience situations. By referring to bits of knowledge known by both oneself and a particular audience, one may convey information to that audience—whether an individual or a group—without others catching on.

At the same time, our results reveal that for all of people's success in such delicate situations, they believe they are more successful than they actually are. Studies 1 and 2 both revealed significant overconfidence in communicators' assessments of how successfully they had solved their multiple audience problem. Interestingly, this overconfidence occurred despite the fact that participants in our studies, similar to those in the studies reported by Fleming and his colleagues (Fleming & Darley, 1991; Fleming et al., 1990), found the task rather daunting. It appears that overconfidence does not depend on a feeling of ease in performing a task.

The results of Study 2 further suggest that this overconfidence derives from exaggerated perceptions of shared knowledge. Participants who had undergone a familiarization exercise attempted to make use of this familiarity in solving their multiple audience problem. However, they routinely overestimated the extent to which their intended audience would be able to pick up on these references—a mistake participants in the unfamiliar groups did not make.

We suggest that there are at least two reasons for people's miscalibration of the extent to which others share their knowledge. For one, individuals may anchor on their own knowledge base and adjust insufficiently for the fact that others do not necessarily share all of their knowledge (Tversky & Kahneman, 1974). Indeed, a long history of research has noted that people can have pro-

found difficulties when it comes to taking the perspective of others (e.g., Ichheiser, 1949; Inhelder & Piaget, 1958; Keysar et al., 1995; Piaget, 1962; Ross & Ward, 1996; Stephenson & Wicklund, 1983). An instructive example is provided by Newton (cited in Griffin & Ross, 1991), who found that participants asked to tap a well-known melody on a tabletop grossly overestimated listeners' ability to recognize the melody. Although tappers themselves may have heard a rich representation of the song they tapped, "complete with rich harmonies between strings, winds, brass, and human voice" (Griffin & Ross, 1991, p. 335), listeners heard no such orchestration. To capture listeners' perspectives, then, tappers had to put aside their own representation of the melody and intuit others' relatively impoverished perspectives. Tappers' erroneous estimates indicate that they did so insufficiently. Applied to the present experiments, communicators' subtle hints and arcane references, although perfectly clear to them, were undoubtedly less clear to their audience. Communicators' overconfidence likely stemmed in part from their insufficient correction for this fact.

A second (and related) potential cause of communicators' misguided estimates of the extent to which others can decipher their clues derives from the recently documented spotlight effect, or people's tendency to overestimate the extent to which others notice them and remember details about their behavior and appearance (Gilovich, Medvec, & Savitsky, 2000). How might this bias contribute to overconfidence in the multiple audience dilemma? Many of the hints relayed by communicators in our experiments were references to details about themselves that they had revealed during the preliminary portion of the experiment (the first two discussions in Study 1 or the familiarization exercise in Study 2). To the extent that communicators overestimated the salience of these statements and actions in the eyes of others, references to them may have been less intelligible than the communicators expected. In this research, we made no attempt to determine the frequency with which individuals strategically referred to knowledge about themselves, for which the spotlight effect would be most relevant (e.g., "I was a drum major") as opposed to more general knowledge about things other than the self (e.g., "Thursday nights are when Colletown rocks"). An attempt to do so in future research might help to isolate that portion of the present finding that derives from the spotlight effect and the portion that derives from the overestimation of more general shared knowledge.

Some readers may be tempted to attribute the present results to demand characteristics rather than to the mechanisms we have postulated. More specifically, solving the multiple audience problem was identified to par-

ticipants as their primary goal in the experiments. For a participant to indicate his or her belief that he or she had done so, then, was tantamount to indicating that he or she had indeed followed instructions and complied with the experimenters' wishes. We do not believe, however, that demand characteristics can explain all of the data we reported. In particular, participants in the unfamiliar groups of Study 2 received the very same instructions as did those in the familiar groups and thus should have experienced the same demand characteristics. Even so, they were not overconfident like their counterparts in the familiar group. Thus, experimental demand does not offer a viable alternative account of our findings.

Are multiple audience problems in everyday life characterized by the same degree of overconfidence that we observed in our laboratory experiments? One viewpoint might be that our studies overestimate overconfidence in the real world because participants were allowed to exchange only a small amount of information with each other. Thus, attempts to solve the multiple audience problem may have lacked the ammunition of a lifetime of shared information that everyday attempts enjoy. Indeed, the familiarization exercise in Study 2 was quite minimal compared to the rich interactions people have with their friends, colleagues, and romantic partners, and one might expect multiple audience problems that occur spontaneously, in more naturalistic settings, to be solved more successfully, attenuating overconfidence.

Despite the plausibility of this argument, we believe the opposite is a more likely possibility. In their research on confidence and accuracy in people's perceptions of their relationship partners, Swann and Gill (1997) have found that although people's confidence in their impressions of others increases as a function of relationship length, the accuracy of their impressions does not. As a result, people are often more, not less, overconfident in the knowledge they have about others as relationship length increases. It is no great leap to assume that an individual's assessments of what others know about him follow a similar path, with perceived knowledge (and hence confidence) increasing more steeply than actual knowledge.

Furthermore, note that participants in Study 2, unlike participants in everyday multiple audience problems, were allowed to see the situation from all relevant perspectives: as communicator, as intended audience, and as unintended audience. As a result, these participants had ample opportunity to recognize the fact that communicators' messages were not substantially clearer to them when they were the intended audience than when they were the unintended audience and to bring this realization to bear on their assessments of their own ability to solve the multiple audience problem (see Note 3).

This conservative feature of our experiment, coupled with the tendency for overconfidence to grow as relationship length increases, suggests that the present research may well underestimate overconfidence in the multiple audience problem.

APPENDIX

The following is a sample essay from Study 2. Participants were to select one of the underlined words to be their secret word, which they were to convey to their own group but conceal from the other group.

WAR GAMES

Thirty years ago, thousands of little boys were buying thousands of G. I. Joe dolls, and "war" was still a popular kid's game. Then came the Vietnam War, and we decided that war toys taught little boys to grow up to be killers. G. I. Joe disappeared from the shelves along with the platoons of military accessories: machine guns, grenades, and other replicas of modern weaponry. Well, the kids are battling it out again, here in the 1980s and 1990s. But this time their battlefield is not the sands of Iwo Jima or the rice fields of Cambodia; war has been transported to outer space, and G. I. Joe has been replaced by Luke Skywalker. Now the futuristic artillery of Star Wars fills toy stores and the merchants can never stock enough of it. It would seem that playing war will never disappear from the American scene, and it also seems that even if the kids should lose interest in shooting each other, their parents will not.

NOTES

1. The average interitem correlations for observers who had witnessed targets construct party animal and nerd identities, respectively, were .21 and .46 after their initial interactions and .16 and .61 after the final discussion. The average interitem correlations for targets' predictions of how they were rated were .50 and .69, respectively, after the initial interactions and were both .54 after the third interaction. All but one of these was greater than the .20 threshold (Briggs & Cheek, 1986) for sufficient reliability.

2. We instituted this second restriction after several speakers in pilot sessions adopted the strategy of concealing their faces from the unintended group and winking an eye after reading the secret word.

3. Participants' estimates of how well they had solved their multiple audience problem were unaffected by whether they went early or late in the session, that is, whether they were one of the first or one of the last speakers to read their essay. In other words, participants did not seem to learn from their experience as observers how difficult it would be as a speaker to solve the multiple audience problem.

4. As a testimony to the dubiousness of such a strategy, note that speakers from the familiar groups were no better able to convey the secret message than were speakers from the unfamiliar groups. This begs the question of just how speakers were able to navigate this tricky situation. Indeed, our informal observations suggest that many participants in the unfamiliar groups adopted no discernible strategy at all, and of those that did, their strategies seemed (to us, at least) of questionable efficacy. For example, one participant cleared his throat after reading his secret word, but not after any of the other words—a strategy (presumably) no more apparent to his intended audience than his unintended audience. To be fair, participants from the unfamiliar groups were in a very difficult situation; we cannot ourselves think of a

successful strategy one might adopt in their position. So, how were speakers able to solve the multiple audience problem? One possibility, admittedly speculative, is that audience members simply paid more attention to speakers from their own group than speakers from the opposing group. Any signals a speaker sent out, then, would be more likely to have been picked up by the intended than the unintended audience. If so, strategies such as clearing one's throat may have been effective because the unintended audience's attentions were elsewhere.

REFERENCES

- Bobrow, J., Nathan, H., Fisher, S., Covino, W. A., Orton, P. Z., Bobrow, B., & Weber, L. (1989). *National teacher examinations: Core battery preparation guide*. Lincoln, NE: Cliffs Notes.
- Briggs, S. R., & Cheek, J. M. (1986). The role of factor analysis in the development and evaluation of personality scales. *Journal of Personality, 54*, 106-148.
- Camerer, C., Loewenstein, G., & Weber, M. (1989). The curse of knowledge in economic settings: An experimental analysis. *Journal of Political Economy, 97*, 1232-1253.
- Clark, H. H., & Schaefer, E. F. (1987). Concealing one's meaning from overhearers. *Journal of Memory and Language, 26*, 209-255.
- Cooley, C. H. (1902). *Human nature and the social order*. New York: Scribner.
- Dunning, D., Griffin, D. W., Milojkovic, J. D., & Ross, L. (1990). The overconfidence effect in social prediction. *Journal of Personality and Social Psychology, 58*, 568-581.
- Fleming, J. H. (1994). Multiple-audience problems, tactical communication, and social interaction: A relational-regulation perspective. In M. P. Zanna (Ed.), *Advances in experimental social psychology* (Vol. 26, pp. 215-291). New York: Academic Press.
- Fleming, J. H., & Darley, J. M. (1991). Mixed messages: The multiple audience problem and strategic communication. *Social Cognition, 9*, 25-46.
- Fleming, J. H., Darley, J. M., Hilton, J. L., & Kojeten, B. A. (1990). Multiple audience problem: A strategic communication perspective on social perception. *Journal of Personality and Social Psychology, 58*, 593-609.
- Fleming, J. H., & Scott, B. A. (1991). Costs of confession: An analysis of the Persian Gulf War POW videotapes in historical and theoretical perspective. *Contemporary Social Psychology, 15*, 127-138.
- Fussell, S. R., & Krauss, R. M. (1991). Accuracy and bias in estimates of others' knowledge. *European Journal of Social Psychology, 19*, 509-525.
- Fussell, S. R., & Krauss, R. M. (1992). Coordination of knowledge in communication: Effects of speakers' assumptions about what others know. *Journal of Personality and Social Psychology, 62*, 378-391.
- Gilovich, T., Kruger, J., & Medvec, V. H. (2000). *The spotlight effect revisited: Overestimating the manifest variability of our actions and appearance*. Manuscript submitted for publication.
- Gilovich, T., Medvec, V. H., & Savitsky, K. (2000). The spotlight effect in social judgment: An egocentric bias in estimates of the salience of one's own actions and appearance. *Journal of Personality and Social Psychology, 78*, 211-222.
- Goffman, E. (1959). *The presentation of self in everyday life*. Garden City, NY: Doubleday.
- Goffman, E. (1963). *Interaction ritual: Essays on face-to-face behavior*. New York: Pantheon.
- 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.
- Ichheiser, G. (1949). Misunderstandings in human relations: A study in false social perception. *American Journal of Sociology, 55*, 1-67.
- Inhelder, B., & Piaget, J. (1958). *The growth of logical thinking from childhood to adolescence*. New York: Basic Books.
- Jacowitz, K. E., & Kahneman, D. (1995). Measures of anchoring in estimation tasks. *Personality and Social Psychology Bulletin, 21*, 1161-1166.
- James, W. (1892). The self. In *Psychology: The briefer course* (pp. 41-49). New York: Henry Holt.
- Jones, E. E., & Pittman, T. S. (1982). Toward a general theory of self-presentation. In J. Suls (Ed.), *Psychological perspectives on the self* (Vol. 1, pp. 231-262). New York: Random House.
- Keren, G. (1987). Facing uncertainty in the game of bridge: A calibration study. *Organizational Behavior and Human Decision Processes, 39*, 98-114.
- Keysar, B., Ginzler, L. E., & Bazerman, M. (1995). States of affairs and states of mind: The effect of knowledge on beliefs. *Organization Behavior and Human Decision Processes, 64*, 283-293.
- Lichtenstein, S., Fischhoff, B., & Phillips, L. D. (1982). Calibration of probabilities: The state of the art in 1980. In D. Kahneman, P. Slovic, & A. Tversky (Eds.), *Judgment under uncertainty: Heuristics and biases*. New York: Cambridge University Press.
- Mead, G. H. (1934). *Mind, self and society*. Chicago: University of Chicago Press.
- Oskamp, S. (1965). Overconfidence in case-study judgments. *Journal of Consulting Psychology, 29*, 261-265.
- Piaget, J. (1962). Play, dreams, and imitation in childhood (C. Gattegno & F. M. Hodgson, Trans.). London: Routledge and Paul.
- 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., Lepper, M. R., & Hubbard, M. (1975). Perseverance in self-perception and social perception: Biased attributional processes in the debriefing paradigm. *Journal of Personality and Social Psychology, 32*, 880-892.
- 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: Lawrence Erlbaum.
- Schlenker, B. R. (1980). *Impression management: The self-concept, social identity, and interpersonal relations*. Monterey, CA: Brooks/Cole.
- Stephenson, B., & Wicklund, R. A. (1983). Self-directed attention and taking the other's perspective. *Journal of Experimental Social Psychology, 19*, 58-77.
- Swann, W. B. (1987). Identity negotiation: Where two roads meet. *Journal of Personality and Social Psychology, 53*, 1038-1051.
- Swann, W. B., Jr., & Gill, M. J. (1997). Confidence and accuracy in person perception: Do we know what we think we know about our relationship partners? *Journal of Personality and Social Psychology, 73*, 747-757.
- Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science, 185*, 1124-1131.
- Vallone, R. P., Griffin, D. W., Lin, S., & Ross, L. (1990). Overconfident prediction of future actions and outcomes by self and others. *Journal of Personality and Social Psychology, 58*, 582-592.
- Wright, G., Rowe, G., Boger, F., & Gammack, J. (1994). Coherence, calibration, and expertise in judgmental probability forecasting. *Organizational Behavior and Human Decision Processes, 57*, 1-25.

Received October 28, 1998

Revision accepted February 22, 1999