

Machines and Mindlessness: Social Responses to Computers

Clifford Nass*

Stanford University

Youngme Moon*

Harvard University

Following Langer (1992), this article reviews a series of experimental studies that demonstrate that individuals mindlessly apply social rules and expectations to computers. The first set of studies illustrates how individuals overuse human social categories, applying gender stereotypes to computers and ethnically identifying with computer agents. The second set demonstrates that people exhibit overlearned social behaviors such as politeness and reciprocity toward computers. In the third set of studies, premature cognitive commitments are demonstrated: A specialist television set is perceived as providing better content than a generalist television set. A final series of studies demonstrates the depth of social responses with respect to computer "personality." Alternative explanations for these findings, such as anthropomorphism and intentional social responses, cannot explain the results. We conclude with an agenda for future research.

Computer users approach the personal computer in many different ways. Experienced word processors move smoothly from keyboard to mouse to menu, mixing prose and commands to the computer automatically; the distinction between the hand and the tool blurs (Heidegger, 1977; Winograd & Flores, 1987). Novices cautiously strike each key, fearing that one false move will initiate an uncontrollable series of unwanted events. Game players view computers as

*Correspondence concerning this article should be addressed to Clifford Nass, Department of Communication, Stanford University, Stanford, CA 94305-2050 [e-mail: nass@leland.stanford.edu]; or to Youngme Moon, Graduate School of Business Administration, Harvard University, Boston, MA 02163 [e-mail: ymoon@hbs.edu].

windows into other worlds, whereas e-mail users treat the computer as a mere conduit, little different than a telephone. In short, orientations to computers are derived from a host of individual, situational, and technological variables.

Despite this variance, all computer users know a fundamental truth: *The computer is not a person and does not warrant human treatment or attribution.* It is hard to imagine how one could reach any other conclusion. Unlike dolls or robots, which have faces and bodies, a personal computer¹ looks nothing like a person. It is much less suggestive of a human form than a car, for example, which has headlights for eyes, a hood line for a mouth, and turn signals for facial expressions (McCloud, 1993; Norman, 1992). A computer is unaware of a user's emotions, and it never expresses emotions of its own. It doesn't even refer to itself as "I." Perhaps the clearest list of deficiencies comes from Shylock in *The Merchant of Venice* (Shakespeare, 1971, p. 215): A computer has "no . . . eyes, . . . hands, organs, dimensions, senses, affections, passions. . . . [It does not] bleed . . . laugh . . . [or] die."

With such clear and compelling evidence of the differences between computers and people, we² have not been surprised that of the thousands of adults who have been involved in our studies, not a single participant has ever said that a computer should be understood in human terms or should be treated as a person. That is, *anthropomorphism*, the assignment of human traits and characteristics to computers, will be ruled out by the adult participants in our studies (the issue of anthropomorphism and children is more complex; see, for example, Turkle, 1984).

This rejection of anthropomorphism stands in stark contrast to people's actual *behavior* in our labs and in more naturalistic settings. In this article, we will argue that there is clear evidence that individuals mindlessly (Langer, 1989) apply social rules and expectations to computers. Following Langer's (1989, 1992) explication of mindlessness, we present numerous studies that demonstrate the breadth of individuals' mindless responses to computers. In the first set of experiments, we describe how people tend to *overuse human social categories*, such as gender and ethnicity, by applying them to computers. In the second set of experiments, we provide evidence that people engage in *overlearned social behaviors*, such as politeness and reciprocity, toward computers. In the third set of experiments, we describe how people exhibit *premature cognitive commitments* with respect to computers, as demonstrated by responding to a medium's labeling as a

¹ We use the term "computer" to refer to the standard hardware and software that is encountered by the vast majority of users. Although there are certainly computers that do provide first steps toward emulating a wide range of human capabilities and behaviors, these machines are extremely rare, perform very poorly as compared to humans, and do not seem to have influenced the vast majority of people's thinking about computers.

² Throughout the article, the term "we" refers to Byron Reeves as well as the authors.

“specialist.” To demonstrate the depth of social responses, we describe a variety of consequences that occur, consistent with the human personality psychology literature, when individuals assign “personality” to a computer. We then argue that alternative explanations for these findings, such as anthropomorphism, intentional social responses, and demand characteristics, cannot explain our observations. We conclude with an agenda for further research.

Mindless Responses to Computers

Mindless behavior, which has been observed in a wide variety of social situations (see Langer, 1989, for a review), occurs as a result of conscious attention to a subset of contextual cues (Langer, 1992). These cues trigger various scripts, labels, and expectations, which in turn focus attention on certain information while diverting attention away from other information. Rather than actively constructing categories and distinctions based on all relevant features of the situation, individuals responding mindlessly prematurely commit to overly simplistic scripts drawn in the past.

This specification of mindlessness provides both the standard for demonstrating that social responses to computers are mindless and the means for eliciting these mindless responses. We can conclude that individuals are responding mindlessly to computers to the extent that they apply social scripts—scripts for human-human interaction—that are inappropriate for human-computer interaction, essentially ignoring the cues that reveal the essential asocial nature of a computer. As a practical matter, we have found that the best means for identifying these scripts is to turn to the literature in experimental social psychology (and to a lesser extent, sociology and anthropology; see Reeves & Nass, 1996). We have then replicated, as closely as possible, the same experimental situation, stimuli, measures, and statistics used by the original experimenters to establish the “rule” in human-human interaction, with one important change: We have replaced the human target with a computer. Thus, the extant literature not only has specified the range of social expectations and behaviors that might be (surprisingly) directed toward computers, but it has also provided the means for testing these behaviors in the context of human-computer interaction (Reeves & Nass, 1996).

To elicit *mindless* social responses in this context, individuals must be presented with an object that has *enough* cues to lead the person to categorize it as worthy of social responses, while also permitting individuals who are sensitive to the entire situation to note that social behaviors were clearly *not* appropriate. On the one hand, this suggests that a rock would likely not trigger many social responses, as there would not be enough humanlike cues to encourage a social categorization. On the other hand, it would not be characteristic of mindless behavior to behave as if one was talking with a person when a friendly voice said, “Hi, come on in,” even

if it turned out the message was prerecorded. (Of course, if there were indications that the voice was prerecorded, such as an exact repetition of the message whenever anyone appeared, then a continued social response to the voice would be indicative of mindlessness.)

Earlier, we discussed the readily available information (including appearance-related cues) that would demonstrate to an observer fully aware of context that social scripts should not be applied to computers. What cues, then, might encourage a categorization of computers as social actors? Although there has been no systematic investigation of this point, there are a few characteristics that distinguish computers from most other technologies and are closely associated with the human “prototype” (Nass, Steuer, Henriksen, & Dryer, 1994; Smith, 1989): (1) words for output (Brown, 1988; Turkle, 1984); (2) interactivity, that is, responses based on multiple prior inputs (Rafaeli, 1990); and (3) the filling of roles traditionally filled by humans (Cooley, 1966; Mead, 1934). These characteristics, which may provide sufficient bases for individuals to cue “humanness,” are incorporated into virtually every application on every personal computer. Hence, if these cues encourage social responses, the treatment of computers as human is ubiquitous.

Experimental Evidence Supporting Mindless Responses to Computers

Overuse of Categories

Our first series of studies examined whether individuals would carry over uniquely human social categories to the computer realm. We began with one of the most psychologically powerful social categories: gender (Bem, 1981). Specifically, we focused on three well-established gender stereotypes. First, research has shown that dominant behavior by males tends to be well-received (dominant men tend to be perceived as “assertive” and “independent”), whereas dominant behavior by females tends to be poorly received (dominant women tend to be perceived as “pushy” or “bossy”; Costrich, Feinstein, Kidder, Maracek, & Pascale, 1975). Second, when people are evaluated, they tend to consider the evaluation to be more valid if it comes from a male than if it comes from a female (Eagly & Wood, 1982). Finally, people not only tend to categorize certain topics as being “masculine” or “feminine,” but they tend to assume that women know more about “feminine” topics and men know more about “masculine” topics (e.g., Heilman, 1979).

To determine whether computers would trigger the same scripts, expectations, and attributions associated with these gender stereotypes, we designed an experiment involving computers with voice output³ (Nass, Moon, & Green, 1997). During the experiment, participants were told that they would use computers for three

³ Voice might be another cue that encourages social interaction (Amalberti, 1993).

separate sessions: tutoring, testing, and evaluation. During the tutoring session, the tutor computer verbally presented (via a prerecorded female or male voice) a series of facts on each of two topics, "computers and technology" (a stereotypically "male" topic) and "love and relationships" (a stereotypically "female" topic). After the tutoring session, the participant was directed by the tutoring computer voice to move to a "tester" computer for the testing session. The tester computer, which had no voice, administered a text-based, multiple-choice test. Each question ostensibly had a "correct" answer. Upon completing the testing session, the tester computer told the participant to go to a third computer, the "evaluator" computer, for the evaluation session. The evaluator computer, using a *different* prerecorded female or male voice, reviewed each question, indicated whether the participant had given a correct answer, and evaluated the performance of the tutor computer. The evaluations were generally positive, for example: "Your answer to this question was correct. The tutor computer chose useful facts for answering this question. Therefore, the tutor computer performed well." The evaluator computer thus played two dominant roles: It evaluated the performance of the participant, and it evaluated the performance of the tutor computer.

After the on-line evaluation session, the evaluator computer asked the participant to complete a pencil-and-paper questionnaire that consisted of two sets of questions. The first set asked participants for their assessment of the tutor computer's performance during the tutoring session. The second set of questions asked for an assessment of the evaluator computer during the evaluation session.

The results supported the hypothesis that individuals would mindlessly gender-stereotype computers. Both male and female participants found the female-voiced evaluator computer to be significantly *less friendly* than the male-voiced evaluator, even though the content of their comments was identical. In addition, the generally positive praise from a male-voiced computer was more compelling than the same comments from a female-voiced computer: Participants thought the tutor computer was significantly *more competent* (and friendlier) when it was praised by a male-voiced computer, compared to when it was praised by a female-voiced computer. And finally, the female-voiced tutor computer was rated as significantly more informative about love and relationships compared to the male-voiced tutor, while the male-voiced tutor was rated as significantly more informative about computers (Nass et al., 1997).

It is important to note that in this experiment, participants were wholly aware that the voice was a narrator for the computer (as opposed to some human "behind" the computer). In addition, they knew that the voice did not necessarily reflect the gender of the computer programmer. In fact, when asked (in postexperimental debriefs), participants indicated that they believed the various computer programs were probably written by the same person, even though different voices were played in the different software programs. And despite their behavior, they also uniformly agreed that male-voiced computers were no

different than female-voiced computers and that to engage in gender-stereotyping with respect to computers would be ludicrous.

Ethnicity. The second social category examined was ethnicity (Nass, Isbister, & Lee, in press). In this experiment, an interactive video manipulation⁴ was used to provide participants with an “ethnicity cue.” And because minority group individuals tend to exhibit much higher levels of ethnic identification (Wilder & Shapiro, 1984), we involved Korean rather than Caucasian participants.

Participants were given a series of hypothetical choice-dilemma situations in which an individual had to decide between two courses of action. In general, one option tended to be more rewarding and attractive but less likely to be attained. (For example, one of the situations described the dilemma of a college football player who could either go for a risky play that could win the game, or go for a cautious play that could tie.)

The participants were instructed to read a description of the situation, make a decision, and then ask the computer agent—represented by a Korean or Caucasian video face—what decision it would make. After being presented with the agent’s decision and its arguments in favor of that decision, the participants answered a series of questions (using a paper-and-pencil questionnaire) concerning their perception of the agent’s decision, the quality of the agent’s argument, and their own decision. This procedure was repeated for eight different choice-dilemma situations.

The results in this experiment were consistent with mindless stereotyping: The social category cue (ethnicity) triggered a series of expectations, assumptions, and attributions, regardless of the context in which the cue appeared. Specifically, participants in the same-ethnicity condition perceived the agent to be more attractive, trustworthy, persuasive, and intelligent, and as making a decision more similar to their own, compared to those in the different-ethnicity condition.

As a direct test of whether individuals mindlessly categorize the computer as a social actor, we repeated this experiment, telling the second group of participants that they were having a videoconference with another person. Ethnicity in humans can predict fundamental values (Wallace, 1997); hence, it was not unreasonable for participants to strongly rely on a real person’s ethnicity in assessing advice on these dilemmas. Remarkably, there were no interactions between ethnicity and human versus computer for any dependent variables; that is, the ethnicity of the image had as strong an effect as a computer agent as it did as an actual person! Once categorized as an ethnically marked social actor, human or nonhuman was no longer an issue.

⁴ Faces have been argued to encourage social responses (Ekman, Friesen, & Ellsworth, 1972; Reeves & Nass, 1996, chap. 3; Shepard, 1990).

Ingroup versus outgroup. Our third study in this line of research provided the most stringent test of the mindless use of social categories (Nass, Fogg, & Moon, 1996). Specifically, we wanted to determine whether people would rely on an arbitrarily assigned social category, that is, ingroup versus outgroup (Tajfel, 1982), when they interacted with computers. Anyone who has been part of a “color war,” in which half of the people are assigned to the blue team and the other half are assigned to the competing green team, knows that the mere act of being labeled and made dependent on others leads to feelings of loyalty and a perception that one’s teammates are superior to those on the other team.

In the first condition, we attempted to create a feeling of shared identity between the person and the computer by (1) reminding the participants that they were dependent on the computer, (2) giving the participants a blue armband, (3) putting a blue border around the computer’s monitor, and (4) referring to the participant and the computer as the “blue team.” In the second condition, each participant was referred to as the “blue person” (armband) working with the “green computer” (border), and the participant was encouraged to focus on individual responsibility. This study, like most of the studies in our research program, used simple text-based software.

The results showed that even when confronted with such a minimal manipulation, and an understanding that the computer could not return the participant’s loyalty, participants in the “team” condition were significantly more likely to cooperate with the computer, to conform to the computer’s suggestions, to assess the computer as more friendly and more intelligent, and to perceive the computer as being similar to themselves, compared to participants in the “nonteam” condition. Further research suggested that the mere use of a matching armband and border could mindlessly induce social responses (Reeves & Nass, 1996, chap. 13). Of course, participants in these experiments claimed (in postexperimental debriefs) that the labeling was irrelevant to their behaviors and attitudes.

Together, the results from these three studies confirm that people tend to rely on social categories when interacting with computers, even when the cues associated with those categories do not have the same meaning or even make “sense” in the human-computer context.

Overlearning

Mindless behavior can also emerge from “overlearning,” that is, from deeply ingrained habits and behaviors. Individuals are so facile at performing some behaviors that once the script is initiated, they stop searching for additional context cues and simply respond according to the script. A straightforward demonstration of this is provided by an experiment we conducted on politeness (Nass, Moon, & Carney, 1999).

Politeness. Research has indicated that when an individual is asked to directly evaluate another person in a face-to-face setting (e.g., “How do you like my new haircut?”), the resulting evaluation tends to be positively biased (e.g., “It looks wonderful.”). That is, people tend to give “polite” evaluations in these situations—even when those evaluations involve some dishonesty—because they are reluctant to hurt the feelings of another person. In our experiment, we replicated this scenario using text-based computers. Participants worked with Computer A, and were then interviewed about Computer A’s performance. In the first condition, the interview was conducted by Computer A. In the second condition, the interview was conducted by Computer B (an identical second computer). In the final condition, the interview was conducted via paper-and-pencil questionnaire. Consistent with the politeness hypothesis, evaluations were significantly more positive when the computer asked about itself as compared to the other two conditions (these two conditions obtained identical, and likely truthful, responses). In other words, people were polite to a computer!

These polite responses occurred despite the fact that in postexperimental debriefs, participants uniformly denied believing that computers have “feelings” or warrant polite treatment. Thus, this was a classic case of overlearning: The social rule that dictates insincere responses (the “politeness” rule) automatically came into play as soon as the computer asked about itself. The participants mindlessly failed to consider that the basis for the rule—emotional harm to the questioner—did not apply in this context.

Reciprocity. A second domain in which overlearning was examined was reciprocity. All societies train their members to observe the rule of reciprocity: “One should provide help, favors, and benefits to those who have previously helped them” (Cialdini, 1993; Fogg, 1997; Gouldner, 1960). The rule of reciprocity has been shown to be extremely powerful (Cialdini, 1993); indeed, anthropologists Leaky and Lewin (cited in Cialdini, Green, & Rusch, 1992) assert that reciprocity is *the* central characteristic of being human. Thus, if the overlearning of social scripts extends to the computer domain, then individuals will reciprocate when a computer is helpful to them.

The first experiment (Fogg & Nass, 1997) involved two tasks: a task in which a computer “helped” a user and a task in which the user was asked to “help” a computer. In Task 1, participants conducted a series of web searches with a computer; the results of the searches were either extremely useful or not at all useful. In Task 2, participants worked with a computer that was trying to create a color palette to match human perception. They were told that by making accurate comparisons of sets of presented colors, they could help the computer create this palette. Participants could choose how many comparisons to do: the more comparisons, the more the participant helped the computer. In one condition, participants performed Task 2 on the same computer that they performed Task 1; in the other condition, participants used different (but identical) computers for Tasks 1 and 2.

The results were consistent with reciprocity norms: Participants who worked with a helpful computer in Task 1 and then returned to the same computer in Task 2 performed significantly more “work” for the computer in Task 2, compared to participants who used two different computers for the two tasks. Same-computer participants even performed Task 2 with greater accuracy, a different measure of effort. There was also evidence of a “retaliation” effect: When participants worked with a computer in Task 1 that was not very helpful and then returned to the same computer in Task 2, they made significantly fewer comparisons than participants who used different computers.

Reciprocal self-disclosure. We decided to follow up this study with an investigation into reciprocal self-disclosure (Moon, in press). Research has shown that people are typically reluctant to divulge intimate information about themselves to anyone but their closest friends and relatives (e.g., Kelly & McKillop, 1996). One notable exception to this rule involves reciprocity: There is substantial evidence that people will engage in intimate self-disclosure—even with relative strangers—if they first become the recipients of such disclosures from their conversational partners (see Moon, in press, for a review). In short, disclosure “begets” disclosure, such that people who receive intimate disclosure feel obligated to respond with a personal disclosure of equal intimacy.

In this experiment, we were interested in whether people would engage in reciprocal self-disclosure with a computer, providing that the computer initiated the disclosure process by divulging information first. Participants were interviewed by a computer on a variety of topics. In the no-reciprocity condition, the computer asked the interview questions in a relatively straightforward manner, for example: “What has been your biggest disappointment in life?” or “What have you done in your life that you feel most guilty about?”

In the reciprocity condition, the computer preceded each interview question with some seemingly parallel information about itself: “This computer has been configured to run at speeds up to 266 MHz. But 90% of computer users don’t use applications that require these speeds. So this computer rarely gets used to its full potential. What has been your biggest disappointment in life?” or “There are times when this computer crashes for reasons that are not apparent to its user. It usually does this at the most inopportune time, causing great inconvenience to the user. What have you done in your life that you feel most guilty about?” The information disclosed by the computer was descriptive in nature and always referred to factual matters. The computer never made a statement that implied that it had emotions, feelings, attitudes, or motivations, and it never referred to itself as “I.”

Because the interview questions in the reciprocity condition were much lengthier than those in the no-reciprocity condition, there was also a third condition in this experiment. In this control condition, the number of words for each question equaled the number of words in the reciprocity condition. But unlike the

reciprocity questions, the control questions did not involve computer “disclosures,” for example: “You are now ready for the next question in this interview. The next question is about disappointment. In this question, you will be asked about the biggest disappointments in your life. The specific question is as follows . . . What has been your biggest disappointment in life?” or “The next question in this interview is about guilt. More specifically, you will be asked what you have done in your life that you feel most guilty about. The question is: What have you done in your life that you feel most guilty about?”

When we looked at the results of this experiment, we found self-disclosure tendencies to be consistent with the norms of reciprocity: Responses in the reciprocity condition were higher in intimacy (measured in terms of depth and breadth) than responses in the other two conditions. Thus, overlearned social scripts can be activated not only in a context in which they do not make “sense” but also in a context in which the “trigger” for the script makes explicit the nonhuman source of the information.⁵

Premature Cognitive Commitment With Single Exposure

As Langer (1992) points out, mindlessness is distinct from mere overlearning because the former may result from a single exposure to a stimulus, as opposed to repeated exposures. This can happen, for example, when information is given by an authority figure. In these cases, information is accepted uncritically, without attention to other aspects of the situation. To determine whether this type of “premature cognitive commitment” occurs when people interact with machines, we decided to conduct an experiment that manipulated the labeling of machine roles. Because computers might naturally be perceived as authoritative in the content they produce, we focused on a technology that does *not* produce content and is never thought of as an expert: a television set.

Specialist versus generalist. Would the mere labeling of a television as a “specialist” influence individuals’ perception of the content it presented? To conduct this study, participants were brought into the laboratory and watched segments from news shows and situation comedies. Those who were assigned to watch the “generalist” set were told they would watch an ordinary TV that we used to show both news and entertainment shows. On top of the TV was a sign that read “News and Entertainment Television.” Those who were assigned to the “specialist” condition were told that they would watch programs on two different televisions: They would watch news on a television set that we happened to use only to show news

⁵ In its use of placebic information to generate compliance, this study is related to Langer, Blank, and Chanowitz (1978).

programs, and entertainment on a television set (on the other side of the room) that we happened to use only to watch entertainment programs. On top of each of these two TVs were signs that read “News Television” and “Entertainment Television,” respectively.

After viewing a series of news and entertainment segments, we asked participants to evaluate what they had seen. The results demonstrated a premature cognitive commitment to the notion of expertise: Participants in the “specialist” condition thought the news segments were significantly higher in quality, more informative, interesting, and serious than did participants in the “generalist” condition, even though everyone saw identical news segments. Similarly, though everyone saw the same programs, participants in the “specialist” condition thought the entertainment segments were significantly funnier and more relaxing than participants in the “generalist” condition. Thus, even meaningless assignments of “expertise” can result in mindless acceptance of content.

Breadth and Depth of Social Responses

The previous studies (see Reeves & Nass, 1996, for additional studies) establish the breadth of mindless responses *across* areas of social psychology. We have also been interested in establishing a rich set of results *within* a domain of psychology: personality.

Personality. In the computer science literature, “personality” has traditionally been one of the “holy grails” of artificial intelligence. The assumption has been that the creation of even crude computer “personalities” necessarily requires tremendous computing power and realistic humanlike representations. In contrast, we decided to rely on the tendency of users to make premature cognitive commitments to generate strong effects from a computer personality via the simple scripting of text.

In these studies, participants worked with a computer that displayed either a “dominant” or “submissive” personality style. Although the core content was kept identical, the manipulation of dominance and submissiveness in the computers was rather simple: (1) the dominant computer used strong, assertive language during the task (e.g., “you should definitely do this”), whereas the submissive computer used more equivocal language (e.g., “perhaps you should do this”); and (2) the dominant computer expressed high confidence in its actions during the task (an average confidence level of 80%), whereas the submissive computer expressed low confidence (an average confidence level of 20%).

In our first study (Nass, Moon, Fogg, Reeves, & Dryer, 1995), we used a standard personality test to categorize participants according to whether they had “dominant” or “submissive” personalities. We then paired them with a computer that either matched or mismatched their personality. We hypothesized that

participants would respond to the computers making a premature cognitive commitment to the textual cues of personality. Consistent with the principle of “similarity-attraction,” which posits that individuals are attracted to other *people* who are similar to themselves, we found that dominant participants were more attracted to, assigned greater intelligence to, and conformed more with the dominant computer, compared to the submissive computer. Submissive participants reacted the same way to the submissive computer compared to the dominant computer, despite the essentially identical content (Nass et al., 1995).

We have replicated and extended this result, using the same manifestation of personality. For example, we have found that when people use a “matching” computer, they are more willing to purchase items via the computer (Moon, 1998). We have also established “gain” effects: When a computer’s personality cues change from being dissimilar to the user to being similar to the user, individuals are more positively disposed to the computer than when its personality cues are consistently similar (Moon & Nass, 1996). Cued personality similarity can even overcome the “self-serving bias”: When a computer’s “personality” matches that of the user, individuals are more likely to give the computer credit for success and less likely to blame the computer for failure, compared to when there is a personality mismatch (Moon & Nass, 1998). And finally, personality cues influence the perception of neutral content. In this study, when the personality of a computer-based “News and Entertainment Guide” matched users’ personalities, users found the music, humor, and health advice (which was identical for all participants) to be significantly better (Moon, 1998). Thus, even minimal cues can mindlessly evoke a wide range of scripts, with strong attitudinal and behavioral consequences.

Addressing Alternative Explanations

The previous sections of this article have shown that mindlessness provides a robust explanation for the wide variety of social behaviors we have observed in human-computer interaction. However, are there explanations for social responses toward computers that allow for *mindful* responses on the part of individuals? Certainly, as we have presented these results over the years, alternative explanations for these phenomena have been proposed. In this section, we argue that these alternative explanations do not bear the weight of the evidence.

For individuals to thoughtfully apply social rules to computers, at least one of three factors must be involved: (1) individuals erroneously believe that computers warrant human treatment, (2) individuals orient their responses to some human “behind” the computer, or (3) individuals determine that the experimenter wants participants to exhibit social responses, and the participants comply. We address each alternative in turn.

Anthropomorphism

If individuals have a belief that computers are essentially human, that is, if individuals *anthropomorphize* computers, human-appropriate responses to computers reflect a reasonable application of social rules and behaviors. Indeed, prior to the current research, anthropomorphism was the standard explanation for social responses to computers (e.g., Barley, 1988; Turkle, 1984; Winograd & Flores, 1987). The argument was that social responses to computers emerged from ignorance concerning the ontological status of computers qua people.

We reject the assumption of anthropomorphism, largely based on the fact that the participants in our experiment were adult, experienced computer users. When debriefed, they insisted that they would never respond socially to a computer, and vehemently denied the specific behaviors they had in fact exhibited during the experiments. Because these behaviors extended over many minutes and relied on intentional responses on the part of the user (as distinct from physiological or other automatic responses, such as shouting at the computer), we believe that the participants were sincere in their protestations.

This is not to say that individuals of all ages cannot or do not develop very strong relationships with computers and other objects (e.g., Martin, 1997; Schaffer, 1991; Sherman, 1991; Turkle, 1984). One can observe many individuals who cherish a computer or other object because they have become emotionally attached to it, who give computers (and other technologies, most notably automobiles) a name, and who carry on running dialogues with machines that cannot listen. These responses are not evidence for anthropomorphism, because anthropomorphism, as defined here, involves the thoughtful, sincere belief that the object *has* human characteristics. If the adults who exhibit the previous behaviors were asked whether the object they were orienting to actually had human traits and characteristics, the evidence suggests that adults (the case of children is much more complex; see Turkle, 1984) would say “no.” Indeed, social and emotional behavior directed toward cherished objects seems to be related to the evocation of memories and emotion management, rather than a direct response to the object itself (Sherman, 1991; Wapner & Redondo, 1990). In the same way, people can cry while reading a book; the tears are not directed to the book as medium.

These rich relationships with computers and other technologies also should not be confused with the social responses we describe here. For the former, individuals are very often *quite mindful* of their emotional attachments to objects. However, for the social responses we observe, individuals are *unaware* of their behaviors, and apply a wide range of social rules mindlessly. That is, rather than an object for reflection, the computer seems to be a peer in a social interaction.

Thus, there are key differences between anthropomorphism (a sincere belief that the computer warrants human treatment), cherished objects (in which the object is oriented to with a focus on its ability to evoke or manage feelings and

attitudes), and the responses described in this paper, termed *ethopoeia*. Ethopoeia (from the Greek) involves a *direct* response to an entity as human while knowing that the entity does not warrant human treatment or attribution. Models of thoughtful human attribution and behavior or evocation of memories and feelings cannot explain the processes that elicit stereotyping, politeness, reciprocity, and the like toward a computer, but an obliviousness to the unique characteristics of a computer as an interactant certainly can.

Orientation to the Programmer

Some have argued that social responses to the computer are not actually social responses to the *computer* at all; instead, they are social responses to an unseen human “behind” the computer, usually a *programmer*. The basic argument here is that individuals frame interactions with computers as interactions with imagined programmers; since programmers are people, it is not surprising that individuals display social responses. A related explanation is the “intentional stance” explanation (Dennett, 1987): When confronted with a complex entity obviously designed by a person, humans will ascribe humanlike goals and characteristics as a heuristic for understanding (Johnson-Laird, 1983, pp. 475–476). It then seems a small (albeit dangerous) leap from this ascription to a fuller endowment of human traits and characteristics.

Three categories of evidence challenge this explanation of “the computer as programmer.” First, in the dozens of experiments we have conducted, the overwhelming majority of participants have indicated, both spontaneously and in direct questioning, that they did not have the programmer or any other human “in mind” during the interaction. Although it is possible that these thoughts are not conscious, this certainly challenges the “intentional stance” explanation.

Second, in all of the experiments involving multiple computers (e.g., the politeness study, the gender study, the reciprocity study), participants unanimously agreed that the programs on the two computers were written by the same person. This is not a coincidence: Not only were all of the software tasks in fact written by the same person, but we did everything we could to make the interfaces as uniform as possible with respect to interaction style, language, font, and so forth. If individuals were thinking of the programmer when they were participating in these experiments, the fact that there was a single programmer would have *eliminated* the difference between the boxes, just as having a conversation with the same person from two different telephones does not make one feel that one is speaking with two different people. The significant differences in these experiments clearly undermine the orientation to the programmer argument.

Finally, Nass and Sundar (1996; see Reeves & Nass, 1996, chap. 16) performed a critical test of the “programmer explanation.” Participants worked on tutoring tasks with two different computers that had “different styles of teaching.”

For half of the participants, the computer was consistently referred to as a “programmer”; for the other half of the participants, the computer was called a “computer.” For example, “you will be working with this [computer] [programmer] to learn about the United States” or “this [computer] [programmer] will provide you with 15 facts.” If people orient to the programmer when they interact with a computer, this minimal manipulation should have no effect, as the words “computer” and “programmer” would be synonyms. Conversely, if individuals respond directly (and mindlessly) to the computer, then there would be differences between computer-oriented responses and programmer-oriented responses. Consistent with the mindlessness explanation, participants who did not have to think of the programmer had significantly more positive feelings about the experience than did “programmer” participants.⁶

Together, these three strands of evidence suggest that mindlessness is a better explanation for social responses than is an orientation to the programmer.

Demand Characteristics

The final challenge to mindless responses focuses on the extent to which (1) the experimental situations or (2) the questionnaires *encourage* users to demonstrate social responses, responses that would not be exhibited in more normal situations. We address each in turn.

One might argue that many of the experiments described above involve behaviors of the computer that are so unusually social that individuals believe that they are being asked to *pretend* that they are in a social situation. That is, participants exhibit a “willing suspension of disbelief” (Coleridge, 1889); they assume that in order to be engaged in the experimental task, they are expected to “forget” that they are only dealing with a computer. This view is different from the “orientation to the programmer” in that in the present perspective, the social orientation is *directly* to the computer, rather than *indirectly through* the computer to an unseen programmer.

There are a number of flaws in this argument. First, in more than half of the experiments described above, individuals were interacting with simple text on a screen. None of the experiments included sophisticated input modalities, such as speech recognition or visual processing. The computers never referred to themselves as “I” (all self-references were to “this computer”), never referred to the user by name, gender, or other identifying characteristic, and never indicated or implied that they had feelings, emotions, or other human traits or characteristics. All of the experiments included cover stories that focused on traditional computing tasks,

⁶ Of course, any difference between the programmer and computer conditions would lead us to reject the “computer as programmer” explanation.

such as teaching, evaluation, or presentation and processing of information according to fixed criteria, rather than richly social situations. And not only did the experimenter follow a fixed script, but participants encountered the experimenter only before their work with the computer and after they completed the experiment; thus, it would have been virtually impossible for the experimenter to signal the desired responses.

Regarding the questionnaires, every effort was made to create measures that did not suggest human traits or characteristics. Questions concerning the performance of the computer and questions about the user's affect or performance do nothing to encourage a social response; these measures nonetheless indicated social responses. And in many of the studies (e.g., the reciprocity experiments), we moved beyond questionnaire measures to richer behaviors that could not be susceptible to the argument that the questions cued social responses.

At a more basic level, one of the best arguments against demand characteristics is that for many of the experiments, participants were *unaware* of the responses they were supposed to provide. For example, in the personality experiments, participants did not know the principles of similarity-attraction that manifested in so many different domains. Similarly, it is unlikely that participants were aware of gain theory and the extent of ingroup biases caused by simple labeling. Moreover, in some cases, the norms *against* exhibiting behaviors such as gender stereotyping would have resulted in demand effects that *countered* the actual results we obtained. In sum, the limitations of experimental demands suggest that the results we obtained are found in normal use of the computer.

Agenda for Future Research

The idea that individuals apply social rules when interacting with computers—the “Computers are Social Actors” paradigm—has been well-established (Reeves & Nass, 1996). This paradigm has generated both a wide range of predictions about human-computer interaction and a method for demonstrating the validity of those predictions as well. The present article presents a process—mindlessness—that accounts for these seemingly bizarre responses to computers.⁷

Although mindlessness provides a general explanation for the cognitive processes underlying this phenomenon, it fails to pinpoint *precisely* when and why mindless behavior will occur and when individuals will respond to computers merely as tools. To develop this deeper understanding, we must develop a much more detailed specification of the nature of machines and mindlessness. We propose some critical issues in the sections that follow.

⁷An alternative, but related, process explanation based on evolutionary psychology can be found in *The Media Equation* (Reeves & Nass, 1996, chap. 1).

Characteristics of Computers

The term “computer” (along with the more general term “media”) is theoretically limited, if not vacuous, because the experience of a user can vary enormously depending on the particular characteristics of the computer’s hardware and software. In this regard, some key questions include the following: What are the particular *dimensions* of computers (Nass & Mason, 1990) that encourage or discourage mindless social responses? Are there some dimensions that are more powerful than others? How do the various dimensions interact?

One possibility is that the more computers present characteristics that are associated with humans, the more likely they are to elicit social behavior. According to this hypothesis, computers with voice input and output should elicit more social behavior than their text-based counterparts, software agents with human features should be more likely to encourage social responses than software agents with animal features, and computers that express “emotions” should be socially compelling.

And yet although these predictions seem reasonable enough, it is in fact unclear whether the relationship between humanlike characteristics and social responses is linear. *Perfect* implementations of technologies mimicking human characteristics (such as voice input/output) may generate powerful social responses, but it is not at all clear whether ersatz versions of these technologies are “close enough” to elicit more social responses than would have occurred in their absence (Reeves & Nass, 1996). Indeed, it is equally possible that a lack of verisimilitude increases the saliency of the computer’s “nonhumanness.”

For example, speech recognition systems often make errors that are unlike any that a human would ever make (Kurzweil, 1997); similarly, even the best speech synthesis systems exhibit intonations and cadences that would never come from a person (Nass & Gong, 1999; Olive, 1997). How do such technologies affect users? Do they make them more mindful of the fact that they are not working with a social entity? Unfortunately, the research with respect to this question is rather limited. There is some evidence that users are less likely to respond socially to a poor implementation of a humanlike software character than to a good implementation of a doglike character (Kiesler, Sproull, & Waters, 1996). Other researchers have speculated that false expressions of emotion may backfire, since they remind users of the nonsocial nature of the interaction (Picard, 1997; although people do have a bias toward acceptance of virtually all information; see, e.g., Gilbert, 1991). Clearly, more work is needed in this area.

A related set of issues involves *aggregations* and *asymmetries* with respect to humanlike characteristics. In human adults, the various characteristics and capabilities associated with “being human” (such as the ability to talk, listen, express emotion, look like a human, etc.) are usually inseparable and immutable. In computers, however, not only can particular capabilities be added or removed at will,

but these capabilities are constantly being upgraded as technology improves. This raises a number of questions: In computers, do different combinations of characteristics have additive or synergistic effects with respect to social responses? What happens when highly sophisticated technologies are combined with relatively crude ones in a single machine? Is highly mechanistic language coupled with full-motion video perceived by users to be “somewhat social,” or does the mismatch only cause users to be more mindful of the computer’s essential nonhumanness (see Isbister & Nass, in press)? In short, are there interactions among the various dimensions, or simply main effects? Further compounding all of these questions is the possibility that different dimensions of technology might cue different sets of social scripts; for example, voice technologies might encourage the use of social scripts that involve behavior but have no effect on attribution scripts.

And finally, there is the issue of whether social rules are elicited solely by cues associated with *humans* or whether convincing approximations of reality—that is, virtual realities (Steuer, 1992)—are also sufficient to cue social responses. In other words, a computer that presents ambient sound likely creates a more “realistic” representation of the physical world; does it also evoke more social responses? Conversely, do cues uniquely associated with computers, such as extremely rapid calculation or “crashing,” remind users that social responses are inappropriate?

Individual Differences

In human-computer interaction research, the individual differences that researchers have traditionally focused on are expertise and gender. We have either varied or controlled for both of these factors in a number of our experiments and have found no evidence that social responses to computers are confined to a certain category of people, nor do expertise and gender seem to interact with other factors. This holds true even in our personality studies: Although individual differences such as personality have led to different responses to interfaces, they have not led to differences in the *particular* social rules that are elicited, or the *strength* of those responses.

For present purposes, the most significant individual difference may be the extent to which users, because of either disposition or the demands of the task, focus on the task. In previous research, states of “flow” (Csikszentmihalyi & Csikszentmihalyi, 1988) have been found to lead to intensified processing; however, it could be that these flow states also lead users to ignore seemingly task-irrelevant dimensions that would suggest nonhumanness. Though there is evidence that people can be trained to be more mindful of context cues (see Langer, 1989, 1992), it is unclear how broadly this ability generalizes and the circumstances under which users can overcome cognitive limitations.

Which Social Rules?

Although there may be categories of social rules that are more likely to be mindlessly triggered by computers, no formal typology exists. We therefore propose a few hypotheses. Social attitudes and behaviors that are controlled by more primitive or automatic processes (e.g., multiple voices automatically representing multiple social actors; see Reeves & Nass, 1996, chap. 16) are more likely to be mindlessly elicited than more socially constructed attitudes and behaviors (e.g., which jokes are funny; see Morkes, Kernal, & Nass, 2000). Rules that are used frequently (e.g., conversational politeness norms; see Grice, 1967) are more likely to be mindlessly elicited than rules that are used rarely (e.g., what to say when someone is exhibiting strong emotions). Social behaviors that are uniquely directed at members of a person's culture (e.g., intracultural rituals) may be more difficult to elicit via computers, since interaction breakdowns (Winograd & Flores, 1987) that result from certain characteristics of the computer, including the computer's limited vocabulary and its sometimes inexplicable behavior, may remind the user of a "foreigner" or a person from a different culture. As noted earlier, there may also be linkages between particular types of social rules and particular characteristics of media.

Direct Comparisons to Human-Human Interaction

Validation of the "Computers are Social Actors" paradigm has traditionally involved investigating whether humans exhibit the same basic patterns of behavior toward computers that are found in human-human interaction. Direct, meta-analytic comparisons between the human-human and human-computer contexts have not been conducted, in part because our adherence to well-established methodologies has constrained our ability to construct experimental designs equally appropriate for human-computer and human-human testing.

Recently, however, we have begun to perform experiments that allow for these types of metacomparisons. In these experiments, half of the participants are led to believe they are working with a computer, whereas the other half are led to believe that they are *using* a computer to communicate with a *person* in another room (e.g., Morkes, Kernal, & Nass, 2000; Nass, Isbister, & Lee, in press). In reality, *all* participants experience identical interactions with the computer, regardless of their beliefs about "who" they are communicating with; we are thus able to manipulate participants' perceptions of their interaction partners while controlling all other elements of the experimental situation. With a few exceptions (see Morkes, Kernal, & Nass, 2000), the "human" conditions in these experiments have *not* elicited stronger social responses than the "computer" conditions.

Of course, more research is needed before any general conclusions can be reached. This research should include other comparison groups that provide

a more rigorous challenge to the human-computer condition, such as conditions that involve a meeting with the ostensible interaction partner before the computer-mediated interaction begins, face-to-face conditions, and audio- and video-conferencing.

Final Words

We have spent the past 10 years surprising ourselves (and others) with the breadth and depth of people's social responses to computers. This article represents the second generation of research in our paradigm in that it focuses on an explication of the process by which social responses occur. The second generation will be complete when theory and experiments answer the following question: Which characteristics of computers (and other media) lead which individuals to follow which social rules how similarly to human-human interaction, and why?

References

- Amalberti, R. (1993). User representations of computer systems in human-speech interaction. *International Journal of Man-Machine Studies*, 38(4), 547–566.
- Barley, S. R. (1988). The social construction of a machine: Ritual, superstition, magical thinking, and other pragmatic responses to running a CT Scanner. In M. Lock & D. Gordon (Eds.), *Knowledge and practice in medicine: Social, cultural, and historical approaches* (pp. 497–540). Hingham, MA: Reidel.
- Bem, S. L. (1981). Gender schema theory: A cognitive account of sex typing. *Psychological Review*, 88(4), 354–364.
- Brown, B. (1988). *The human-machine distinction as predicted by children's para-social interaction with toys*. Unpublished doctoral dissertation, Stanford University, Stanford, CA.
- Cialdini, R. B. (1993). *Influence: Science and practice* (3rd ed.). New York: Harper Collins.
- Cialdini, R. B., Green, B. L., & Rusch, A. J. (1992). When tactical pronouncements of change become real change: The case of reciprocal persuasion. *Journal of Personality and Social Psychology*, 63, 30–40.
- Coleridge, S. T. (1889). *Biographia literaria, or, biographical sketches of my literary life and opinions*. London: George Bell.
- Cooley, C. H. (1966). *Social process*. Carbondale, IL: Southern Illinois University Press.
- Costrich, N., Feinstein, J., Kidder, L., Maracek, J., & Pascale, L. (1975). When stereotypes hurt: Three studies of penalties in sex-role reversals. *Journal of Experimental Social Psychology*, 11, 520–530.
- Csikszentmihalyi, M., & Csikszentmihalyi, I. S. (Eds.). (1988). *Optimal experience: Psychological studies of flow in consciousness*. New York: Cambridge University Press.
- Dennett, D. C. (1987). *The intentional stance*. Cambridge, MA: MIT Press.
- Eagly, A. H., & Wood, W. (1982). Inferred sex differences in status as a determinant of gender stereotypes about social influence. *Journal of Personality and Social Psychology*, 43, 915–928.
- Ekman, P., Friesen, W. V., & Ellsworth, P. (1972). *Emotion in the human face: Guidelines for research and an integration of findings*. New York: Pergamon Press.
- Fogg, B. J. (1997). *Charismatic computers: Creating more likable and persuasive interactive technologies by leveraging principles from social psychology*. Unpublished doctoral dissertation, Stanford University, Stanford, CA.
- Fogg, B. J., & Nass, C. (1997). Do users reciprocate to computers? *Proceedings of the CHI Conference* (Atlanta, GA). New York: Association of Computing Machinery.
- Gilbert, D. (1991). How mental systems believe. *American Psychologist*, 46(2), 107–119.

- Gouldner, A. W. (1960). The norm of reciprocity: A preliminary statement. *American Sociological Review*, 25, 161–178.
- Grice, H. P. (1967). Logic and conversation. In P. Cole & J. Morgan (Eds.), *Syntax and semantics 3: Speech acts* (pp. 41–58). New York: Academic Press.
- Heidegger, M. (1977). *The question concerning technology, and other essays*. New York: Harper & Row.
- Heilman, M. E. (1979). High school students' occupational interest as a function of projected sex ratios in male-dominated occupations. *Journal of Applied Psychology*, 64, 275–279.
- Isbister, K., & Nass, C. (in press). Consistency of personality in interactive characters: Verbal cues, non-verbal cues, and user characteristics. *International Journal of Human-Computer Interaction*.
- Johnson-Laird, P. N. (1983). *Mental models: Towards a cognitive science of language, inference, and consciousness*. Cambridge, MA: Harvard University Press.
- Kelly, A. E., & McKillop, K. J. (1996). Consequences of revealing personal secrets. *Psychological Bulletin*, 120(3), 450–465.
- Kiesler, S., Sproull, L., & Waters, K. (1996). A prisoner's dilemma experiment on cooperation with people and human-like computers. *Journal of Personality and Social Psychology*, 70(1), 47–65.
- Kurzweil, R. (1997). When will HAL understand what we are saying? Computer speech recognition and understanding. In D. Stork (Ed.), *HAL's legacy: 2001's computer as dream and reality* (pp. 131–170). Cambridge, MA: MIT Press.
- Langer, E. J. (1989). *Mindfulness*. Reading, MA: Addison-Wesley.
- Langer, E. J. (1992). Matters of mind: Mindfulness/mindlessness in perspective. *Consciousness and Cognition*, 1, 289–305.
- Langer, E. J., Blank, A., & Chanowitz, B. (1978). The mindlessness of ostensibly thoughtful action: The role of placebic information in interpersonal interaction. *Journal of Personality and Social Psychology*, 36, 635–642.
- Martin, S. L. (1997). *The meaning and significance of women's experience of their cherished personal possessions*. Unpublished doctoral dissertation, Stanford University, Stanford, CA.
- McCloud, S. (1993). *Understanding comics—The invisible art*. New York: Harper Perennial.
- Mead, G. H. (1934). *Mind, self, and society, from the standpoint of a social behaviorist* (C. W. Morris, Ed.). Berkeley, CA: University of California Press.
- Moon, Y. (1998). *When the computer is the "salesperson": Computer responses to computer "personalities" in interactive marketing situations* (Working Paper No. 99-041). Boston, MA: Harvard Business School.
- Moon, Y. (in press). Intimate exchanges: Using computers to elicit self-disclosure from consumers. *Journal of Consumer Research*.
- Moon, Y., & Nass, C. (1996). How "real" are computer personalities? Psychological responses to personality types in human-computer interaction. *Communication Research*, 23(6), 651–674.
- Moon, Y., & Nass, C. (1998). Are computers scapegoats? Attributions of responsibility in human-computer interaction. *International Journal of Human-Computer Interaction*, 49(1), 79–94.
- Morkes, J., Kernal, H., & Nass, C. (2000). Effects of humor in task-oriented human-computer interaction and computer-mediated communication: A direct test of SRCT theory. *Human-Computer Interaction*, 14(4), 395–435.
- Nass, C., Fogg, B. J., & Moon, Y. (1996). Can computers be teammates? *International Journal of Human-Computer Studies*, 45(6), 669–678.
- Nass, C., & Gong, L. (1999). Maximized modality or constrained consistency? *Proceedings of the AVSP 99 Conference* (Santa Cruz, CA). New York: Association of Computing Machinery.
- Nass, C., Isbister, K., & Lee, E.-J. (in press). Truth is beauty: Researching conversational agents. In J. Castells (Ed.), *Embodied conversational agents*. Cambridge, MA: MIT Press.
- Nass, C., & Mason, L. (1990). On the study of technology and task: A variable-based approach. In J. Fulk & C. Steinfeld (Eds.), *Organizations and communication technology* (pp. 46–67). Newbury Park, CA: Sage.
- Nass, C., Moon, Y., & Carney, P. (1999). Are respondents polite to computers? Social desirability and direct responses to computers. *Journal of Applied Social Psychology*, 29(5), 1093–1110.
- Nass, C., Moon, Y., Fogg, B. J., Reeves, B., & Dryer, D. C. (1995). Can computer personalities be human personalities? *International Journal of Human-Computer Studies*, 43, 223–239.

- Nass, C., Moon, Y., & Green, N. (1997). Are computers gender-neutral? Gender stereotypic responses to computers. *Journal of Applied Social Psychology*, 27(10), 864–876.
- Nass, C., Steuer, J. S., Henriksen, L., & Dryer, D. C. (1994). Machines and social attributions: Performance assessments of computers subsequent to “self-” or “other-” evaluations. *International Journal of Human-Computer Studies*, 40, 543–559.
- Nass, C., & Sundar, S. (1994, May). Are programmers psychologically relevant to human-computer interaction? Paper presented at the annual meeting of the International Communication Association, San Francisco, CA.
- Norman, D. (1992). *Turn signals are the facial expressions of automobiles*. Reading, MA: Addison-Wesley.
- Olive, J. P. (1997). “The talking computer”: Text to speech synthesis. In D. Stork (Ed.), *HAL’s legacy: 2001’s computer as dream and reality* (pp. 101–130). Cambridge, MA: MIT Press.
- Picard, R. (1997). *Affective computing*. Cambridge, MA: MIT Press.
- Rafaeli, S. (1990). Interacting with media: Para-social interaction and real interaction. In B. D. Ruben & L. A. Lievrouw (Eds.), *Mediation, information, and communication: Information and behavior* (Vol. 3, pp. 125–181). New Brunswick, NJ: Transaction.
- Reeves, B., & Nass, C. (1996). *The media equation: How people treat computers, television, and new media like real people and places*. New York: Cambridge University Press.
- Schaffer, G. (1991). *College students’ cherished possessions: Their nature and role in adaptation*. Unpublished doctoral dissertation, Clark University, Worcester, MA.
- Shakespeare, W. (1971). *The complete works of William Shakespeare* (I. Ribner & G. L. Kittredge, Eds.). Waltham, MA: Xerox College Publishing.
- Shepard, R. N. (1990). *Mindsights: Original visual illusions, ambiguities, and other anomalies, with a commentary on the plan of mind in perception and art*. Ypsilanti, MI: W. H. Freeman.
- Sherman, E. (1991). Reminiscencia: Cherished objects as memorabilia in late-life reminiscence. *International Journal of Aging and Human Development*, 31(3), 219–235.
- Smith, E. E. (1989). Concepts and induction. In M. I. Posner (Ed.), *Foundations of cognitive science* (pp. 501–526). Cambridge, MA: MIT Press.
- Steuer, J. S. (1992). Defining virtual reality: Dimensions determining telepresence. *Journal of Communication*, 42(4), 73–93.
- Tajfel, H. (1982). *Social identity and intergroup behavior*. Cambridge, England: Cambridge University Press.
- Turkle, S. (1984). *The second self: Computers and the human spirit*. New York: Simon and Schuster.
- Wallace, W. L. (1997). *The future of ethnicity, race, and nationality*. Westport, CT: Praeger.
- Wapner, S., & Redondo, J. P. (1990). Cherished possessions and adaptation of older people to nursing homes. *International Journal of Aging and Human Development*, 31(3), 219–235.
- Weizenbaum, J. (1976). *Computer power and human reason: From judgment to calculation*. Ypsilanti, MI: W. H. Freeman.
- Wilder, D. A., & Shapiro, P. N. (1984). Role of out-group cues in determining social identity. *Journal of Personality and Social Psychology*, 47, 342–348.
- Winograd, T., & Flores, C. (1987). *Understanding computers and cognition: A new foundation for design*. Reading, MA: Addison-Wesley.

CLIFFORD NASS is an Associate Professor of Communication at Stanford University, with appointments by courtesy in Science, Technology, and Society, Sociology, and Symbolic Systems. He is Co-Director of the Interface Lab at the Center for the Study of Language and Information at Stanford University. He is author of *The Media Equation* (with Byron Reeves), as well as numerous papers on computers as social actors, human-computer interaction, organization theory, and statistical methodology. He has consulted on the design of over 100 software products and services. His present research focuses on voice-based and character-based interfaces.

YOUNGME MOON is an Assistant Professor of Business Administration in the Marketing unit at the Harvard Business School. Her research focuses on interactive marketing technologies and electronic commerce. In particular, she is interested in situations in which a given interactive technology—usually a computer interface—becomes a primary point of contact with consumers. She has published numerous papers in various journals, including *Journal of Consumer Research*, *Journal of Experimental Psychology: Applied*, *Journal of Applied Social Psychology*, and *Public Opinion Quarterly*.