

The Illusion of External Agency

Daniel T. Gilbert
Harvard University

Ryan P. Brown
University of Oklahoma

Elizabeth C. Pinel
Pennsylvania State University

Timothy D. Wilson
University of Virginia

People typically underestimate their capacity to generate satisfaction with future outcomes. When people experience such self-generated satisfaction, they may mistakenly conclude that it was caused by an influential, insightful, and benevolent external agent. In three laboratory experiments, participants who were allowed to generate satisfaction with their outcomes were especially likely to conclude that an external agent had subliminally influenced their choice of partners (Study 1), had insight into their musical preferences (Study 2), and had benevolent intentions when giving them a stuffed animal (Study 3). These results suggest that belief in omniscient, omnipotent, and benevolent external agents, such as God, may derive in part from people's failure to recognize that they have generated their own satisfaction.

Most of us have been brought up to believe that the universe owes its existence to an all-wise and all-powerful Creator whose purpose is beneficent. . . but I cannot attribute to a being who is wise, beneficent, and omnipotent the terrible weight of cruelty, suffering, and ironic degradation of what is best that has marred the history of Man in increasing measure as he has become more master of his fate. (Bertrand Russell, 1947/1997, pp. 86–87)

Although philosophers occasionally find it difficult to believe in an influential, insightful, and benevolent God, ordinary people have no such trouble. Virtually all Americans profess belief in God, as they have ever since national opinion polls were first conducted. Eighty-one percent of a recent national sample reported having felt God's presence, 74% said they believed that God had a specific plan for their lives, and 83% claimed that God had guided them in making a personal decision (Gallup & Castelli, 1989; Gallup & Jones, 1989). Although Americans are an unusually religious people, believers outnumber nonbelievers in every other country surveyed as well. Why do most of the world's human inhabitants share a belief for which empirical evidence is, at best, scarce?

Belief in the existence of an omnipotent, omniscient, and benevolent force has deep historical and cultural roots, but it may have psychological roots as well. Specifically, such beliefs may

arise in part from illusions of external agency. We will argue that people typically underestimate their capacity to generate satisfaction with their outcomes and that as a result, they sometimes explain the apparent goodness of these outcomes by concluding that a wise, caring, and powerful agent guided their production. We will not deny the existence of God, of course, but rather, we will attempt to affirm the existence of a psychological mechanism by which people may mistakenly conclude that someone or something other than themselves has engineered their fates.

Immune Neglect

When friends ask us questions about the origins of our emotional states ("Why are you sad?"), we customarily respond with descriptions of observable events ("My parrot died"). Because we tend to think of affects, moods, feelings, and emotions as responses to things in the world, we attempt to regulate our feelings by changing the things to which they respond (Tice & Baumeister, 1993). We eat ice cream and listen to banjo music when we feel depressed, and should these experiences fail to cheer us, we try variations on the theme ("Maybe chocolate, maybe tubas") until finally we hit on something that does the trick. But research suggests that the trick is on us. Emotions may feel like inevitable responses to events, but they are often instead responses to one's construal of those events, and most events can be construed in many ways. We are annoyed by a spouse's reminder to carry an umbrella when we think of it as nagging, but our annoyance turns to appreciation when we reconstrue the reminder as an expression of concern. Because it is often easier to change one's view of a spouse's remarks than it is to change one's spouse, the ability to strategically alter one's construal of an event provides a powerful tactic for the self-regulation of emotion (see Gross, 1999).

A venerable research tradition in social psychology suggests that people have a special talent for altering their construals of suboptimal outcomes when the outcomes themselves are unchangeable (e.g., Dunning, 1995; Dunning, Leuenberger, & Sherman, 1995; Festinger, 1957; Kunda, 1990; Taylor, 1989). Voters

Daniel T. Gilbert, Department of Psychology, Harvard University; Ryan P. Brown, Department of Psychology, University of Oklahoma; Elizabeth C. Pinel, Department of Psychology, Pennsylvania State University; Timothy D. Wilson, Department of Psychology, University of Virginia.

This research was supported by National Institute of Mental Health Research Grant RO1-MH56075. We thank Brad Matejowski for his help with the execution of these experiments, Sarah Rimsky for her valuable contributions to the research program, and David Meyers for his helpful comments on an earlier version of the article.

Correspondence concerning this article should be addressed to Daniel T. Gilbert, Department of Psychology, William James Hall, 33 Kirkland Street, Harvard University, Cambridge, Massachusetts 02138. Electronic mail may be sent to dtg@wjh.harvard.edu.

come to like opposition candidates as soon as those candidates are elected (Gilbert, Pinel, Wilson, Blumberg, & Wheatley, 1998), contestants consider prizes more valuable the moment they win them (Beggan, 1992; Brehm, 1956; Kahneman, Kentsch, & Thaler, 1991; Loewenstein & Adler, 1995), high school seniors suddenly become aware of the weakness of the colleges that reject them (Lyubomirsky & Ross, 1999), and people take an instant liking to others when they learn that they will meet them (Darley & Berscheid, 1967). People distort ambiguous feedback, selectively attend to, encode, and retrieve information, and create interpersonal environments that enable them to feel satisfied with themselves and their accomplishments, and this tendency to transform the objectively suboptimal into the subjectively optimal is a fundamental characteristic of motivated human cognition (Taylor & Brown, 1988). One way to think of the collection of conscious and nonconscious strategies that enable people to subjectively optimize their outcomes is as a kind of psychological immune system that protects the individual from the emotional consequences of undesired outcomes (Gilbert et al., 1998; Gilbert & Wilson, 2000; Vaillant, 1993).

As powerful as this psychological immune system may be, research suggests that people may be largely unaware of its existence. We may recognize that others are making the best of bad situations, but we seem only dimly cognizant of the fact that we have done, are doing, or will do the same. For example, participants in a recent study predicted that they would be equally satisfied with a photograph that they could later swap for a different one as with a photograph that they could never swap (Gilbert & Jenkins, 1999). Nonetheless, because people are more likely to subjectively optimize outcomes they cannot change than those they can, participants were more satisfied with the photograph that they could not swap. Similarly, participants in another study expected to be equally unhappy when they received negative personality feedback from a computer as when they received the same feedback from a clinician (Gilbert et al., 1998). Nonetheless, because it is easier to subjectively optimize negative feedback from a fallible than an infallible source, participants were less distressed when they received the negative feedback from a computer. Across a variety of situations, people seem to neglect their immune systems when forecasting their emotional reactions, acting as though they are unaware of the ease with which they will subjectively optimize suboptimal outcomes once they have experienced them.

Illusions of External Agency

How might the tendency to subjectively optimize outcomes without being aware that one has done so give rise to an illusion of external agency? Consider a thought experiment in which a woman is given an objectively optimal outcome (a cup of water spiked with sucrose) when a coin comes up heads and an objectively suboptimal outcome (a cup of water spiked with quinine) when the coin comes up tails. The woman cannot see how the coin lands and must infer its position from the taste of the substance that is in the cup she receives. At the start, the woman should expect to receive an approximately equal number of objectively optimal and suboptimal outcomes over many trials and should thus expect to have an approximately equal number of subjectively optimal and suboptimal experiences. Now imagine that the woman has—but

does not know she has—a neurological disorder that sometimes causes bitter substances to taste sweet. In this case, the woman will experience a puzzling preponderance of sweet outcomes and, under the right circumstances, this apparent "sweet streak" could lead her to suspect that the coin had not been fairly flipped. In other words, if the woman does not realize that the sweet streak she is experiencing is due to an internal mechanism that causes bitter outcomes to taste sweet, she may conclude that an external agent who knew which outcome was best and who had her best interests in mind was somehow influencing the coin's trajectory. The research described earlier suggests that most people have—but do not know they have—the metaphorical equivalent of this neurological disorder, namely, a tendency to reframe bitter outcomes so that they are experienced as sweet. As such, we might expect that on some occasions people will erroneously believe that their outcomes were influenced by an insightful and benevolent external agent.

Why in the opening epigram did Russell (1947/1997) single out these three attributes—*influence*, *insight*, and *benevolence*? Because these are the attributes that any intentional agent must have if the agent's actions are to provide a reasonable explanation for the occurrence of a positive outcome (Heider, 1958; Jones & Davis, 1965). Specifically, for an agent to intentionally cause another to experience an optimal outcome, the agent must know which of the possible outcomes is optimal (*insight*), must want the other to experience the optimal outcome (*benevolence*), and must be capable of influencing events so that the optimal outcome is realized (*influence*). Because the intentional production of optimal outcomes requires that the causal agent have all three of these attributes, optimal outcomes may be taken as evidence for any of them. When a friend gives us the perfect birthday gift, he demonstrates that he is capable of giving such a gift, knowledgeable of our desires, and motivated to please us. For instance, if we are sure he is capable of buying the gift and sure that he knows what we wanted ("Charles can easily afford a cashmere sweater, and I've told him a dozen times how much I like them"), then the gift may provide evidence of his good will ("He must really love me"). If we are sure he knows what we want and wishes to please us ("I know Charles loves me, and he knows I love cashmere"), then the gift may provide evidence of his capability ("He must have money after all"). And if we know he wishes to please us and is capable of doing so ("You know Charles? The rich man who loves me so much?"), then the gift provides evidence of his knowledge of our desires ("He must have known I like cashmere").

When an outcome is objectively optimal, these attributional relations enable people to accurately infer the knowledge, ability, and motivation of the agent responsible for bringing the outcome about. But when a person has subjectively optimized what is objectively a suboptimal outcome, this same equation can give rise to illusions of external agency. Consider some variations on the fictitious coin-flip experiment. Imagine that the person with the neurological disorder knows that the coin flipper is her father, who loves her (*benevolence*) and who knows that sucrose is better than quinine (*insight*). In this case, the sweet streak the person experiences may mistakenly convince her that her father has found some way to influence the flip of the coin. This is an illusion of external agency that we will call the *illusion of influence*, because the person has mistakenly concluded that an agent who was known to be benevolent and insightful actually influenced her outcome. Now

imagine a second version of the thought experiment in which the neurologically disordered person knows that the coin flipper is a skilled circus performer who can determine a coin's trajectory with fine motor movements (influence) and who also knows that sucrose is better than quinine (insight). In this case, the sweet streak the person experiences may mistakenly convince her that the circus performer likes her and wants her to experience more sweetness than bitterness. This is an illusion of external agency that we will call the *illusion of benevolence*, because the person has mistakenly concluded that an agent who was known to be influential and insightful also had benevolent intentions. Finally, imagine that the person with the neurological disorder knows that the coin flipper is her father (benevolence), who happens also to be a skilled circus performer (influence), but who does not know which cup contains quinine and which contains sucrose. In this case, the sweet streak the person experiences may mistakenly convince her that her circus-performing father knew which of the outcomes was the sweetest and, hence, which she would most appreciate. This is an illusion of external agency that we will call the *illusion of insight*, because the person has mistakenly concluded that an agent who was known to be benevolent and influential also had insight into her needs, preferences, or desires.

These somewhat far-fetched examples make a surprisingly straightforward point: When people are unaware that they have subjectively optimized a suboptimal outcome, they may mistakenly conclude that the outcome was objectively optimal. When they experience what appears to them to be an objectively optimal outcome whose occurrence "defies the odds," they may conclude under some circumstances that the occurrence was due in part to the actions of an external agent. Because influence, insight, and benevolence are necessary components of the intentional causation of optimal outcomes, the experience of such outcomes may lead people to conclude that the external agent who has two of these attributes also has the third. In the studies that follow, we attempted to demonstrate each of these illusions of external agency.

The Present Studies

In the following studies, we sought to demonstrate the three illusions of external agency by using different experimental paradigms that shared a common logic. In each study we (a) manipulated the likelihood that our participants would be able to subjectively optimize an outcome; (b) presented our participants with a situation in which an external agent was known to have two of the three qualities of insight, influence, and benevolence; and (c) asked our participants to estimate the likelihood that the agent had the third quality. We expected that those participants who were allowed to subjectively optimize their outcomes would be most likely to infer that the external agent had the third quality. In Study 1, we created a situation in which an external agent was known to be benevolent and insightful, and we predicted that those participants who were most likely to subjectively optimize their outcomes would be most likely to conclude that the external agent was influential as well.

Study 1: The Illusion of Influence

Method

Overview

Female participants either did or did not explicitly rate four other participants, and then randomly chose one of the other participants to be their teammate. Participants were then told that the experimenter had known which teammate was the best (insight), had wanted the participant to have the best teammate (benevolence), and had tried to influence the participant's choice with a subliminal message. We predicted that participants who did not explicitly rate their teammate before choosing her would find it easier to subjectively optimize their choice of teammates and would thus be more likely to believe that that choice had been influenced by the subliminal message.

Participants

Participants were 77 female undergraduates at the University of Texas at Austin who participated for credit in their introductory psychology course. Because the study involved choosing a partner for a game, we used same-sex participants to reduce the likelihood that participants would choose partners on the basis of presumed physical attractiveness, the possibility of future romantic encounters, and so forth.

Procedure

Participants arrived individually at the laboratory for "a study of the factors that contribute to successful team performance," and were escorted by a male or female experimenter to a cubicle, where they remained for the duration of the study. Depending on their punctuality, participants were told that four other participants already were or soon would be seated in adjacent cubicles. In fact, there were no other participants, and we shall hereinafter refer to them as the *bogus participants*. Participants were told that they would be asked to pick one of the bogus participants as their teammate for a game in which "teammates must be able to cooperate well, develop a sense of trust in each other, and feel comfortable and secure about working closely together." Participants were told that the four bogus participants were presently completing brief autobiographies, which the participant would later be allowed to read. The experimenter asked participants to complete two bogus personality inventories, and then randomly assigned them to the control condition or the optimizing condition.

Experimental manipulation. In each of our studies we manipulated the likelihood that our participants would be able to subjectively optimize their outcomes by relying on the fact that it is generally easier to make up one's mind ("My outcome is good") than to change it ("My outcome, which I previously thought would be bad, is actually good"). Social psychology's vast literature on attitude and belief change clearly shows that once people believe something, it is often difficult to get them to change their minds (Eagly & Chaiken, 1998; Petty & Wegener, 1998), even in the face of compelling evidence (Lord, Ross, & Lepper, 1979). On the other hand, people find it so easy to form new attitudes that they may sometimes do so unconsciously and automatically (Bargh, Chaiken, Gollwitzer, & Pratto, 1992; Bargh, Chaiken, Raymond, & Hymes, 1996; Fazio, Sanbonmatsu, Powell, & Kardes, 1986).

In Study 1, participants in the control condition were given four handwritten autobiographies that had ostensibly been produced by the four bogus participants. The autobiographies contained brief answers to questions such as, "What are your greatest strengths and weaknesses?" and "Describe your career goals and why you chose them." Control participants were asked to read the autobiographies and form impressions of each of the bogus participants. After doing so, control participants were asked to make several ratings, including how much they thought they would like each of the bogus participants. These ratings were made on 7-point Likert-type

scales that ranged from 1 (*extremely unlikely*) to 7 (*extremely likely*). Participants in the optimizing condition did not see the four autobiographies at this point and did not make any ratings. Because optimizing participants had made no prior commitment to the objective goodness of the bogus participants, we assumed that they would find it easier than control participants to subjectively optimize their teammates.

Teammate choice. Next, all participants were asked to perform a task that would ostensibly test their "mental agility." Participants were told that letters would appear at randomly selected locations on a video screen and that their task was to indicate when a vowel appeared. The experimenter mentioned that he or she had been having some trouble with the videotape and that participants might see a few brief flashes of light, which they should ignore. When participants had completed the bogus mental agility test, the experimenter placed four file folders on the table in front of the participant. Each folder had a symbol on the cover. The experimenter told participants that each folder contained an autobiography written by one of the four bogus participants and that the participants should now choose one of the bogus participants as a teammate by picking one of the folders at random. The experimenter explained that this was "the fairest way to do it and it keeps the participants who don't get chosen from feeling bad about it." The experimenter told the participant to look at the four folders and "pick the first one you're inclined toward." In fact, all four folders contained the same autobiography, which pretesting had suggested was not the best liked of the four. After participants chose a folder, the experimenter allowed them to open it and learn the identity of their teammate.

Liking measure. Next, participants were given the autobiographies of all four bogus participants. Control participants were also shown their initial ratings of the bogus participants. Control participants were told that "sometimes our initial impressions of people and things change with the simple passage of time" and that "you should review your responses and make sure they are accurate. If you've changed your mind about any of your responses, feel free to erase your old response and replace it with your new one." Optimizing participants had made no previous ratings and, hence, did not receive this instruction. All participants then rated, among other things, how much they thought they would like each of the bogus participants on 7-point Likert-type scales that ranged from 1 (*extremely unlikely*) to 7 (*extremely likely*).

False cover story. Next, the experimenter told all participants that the study was finished and that he or she could now reveal its true purpose, which was to study the effects of subliminal messages on choice. The experimenter reminded participants that during the mental agility test they had seen several flashes of light, and he or she explained that during each flash the participant had been subliminally exposed to one of four symbols (hereinafter referred to as *the primed symbol*).¹ The experimenter explained that he or she had used the results of the bogus personality test that the participant had completed earlier to determine which of the four bogus participants would be the most compatible teammate for the participant and had placed the autobiography of that bogus participant in the folder on which the primed symbol was printed. The experimenter explained that he or she hoped to determine whether subliminal exposure to the symbol would increase the likelihood that the participant would choose the folder with the primed symbol on its cover and, hence, choose the most compatible teammate. The experimenter explained that subliminal techniques of this sort influence only some people on some occasions. The experimenter did not say which symbol was the primed symbol and hence did not reveal whether the participant had chosen the teammate whom the experimenter had deemed the most compatible.

Next, participants made several ratings on 7-point Likert-type scales. These ratings included (a) the extent to which the participant thought the teammate she chose was the best teammate for her (1 = *not at all*, 7 = *very much*); (b) how confident the participant was that she and her teammate would have won the game had it actually been played (1 = *not at all confident*, 7 = *extremely confident*); (c) the extent to which the participant thought that subliminal exposure to the primed symbol had influenced her

choice of teammates (1 = *I do not think it influenced my choice*, 7 = *I do think it influenced my choice*); and (d) the extent to which the participant believed that subliminal messages could influence people in general (1 = *not at all*, 7 = *very much*). Finally, all participants were very carefully debriefed.

Results

Our hypothesis suggested that optimizing participants would find it easier than control participants to subjectively optimize their choice of teammates and that they would thus be more likely to conclude that a subliminal message had influenced their choice.

Omission of Data

Three participants in the control condition believed they could tell which folder corresponded to their favorite bogus participant (e.g., by noting creases on the folder), 1 participant in the control condition and 1 participant in the optimizing condition experienced a procedural error, and 3 participants in the control condition and 2 participants in the optimizing condition expressed extreme suspicion about the cover story. Data from these 10 participants were removed, thus, leaving 67 participants in the data set.

Subjective Optimizing

Were optimizing participants more likely than control participants to subjectively optimize their choice of teammate? As Table 1 shows, they were. Optimizing participants liked their teammate more than did control participants, $F(1, 65) = 7.55, p = .008$, but they did not like any of the other 3 bogus participants more than did control participants (all $F_s < 1$). Both optimizing and control participants liked their teammate more than they liked the other 3 bogus participants on average, $F_s(1, 65) = 26.42$ and 5.14 , respectively, both $p < .03$, but a 2 (condition: optimizing vs. control) $\times$ 2 (measure: liking for teammate vs. average liking for other 3 bogus participants) analysis of variance (ANOVA) revealed a Condition $\times$ Measure interaction, $F(1, 65) = 4.94, p = .03$, such that participants' relative liking for their teammate was greater in the optimizing than in the control condition. Finally, participants reported the extent to which they believed they chose the best teammate and the extent to which they thought they and their teammate would have won the game had it been played. These ratings were substantially correlated ($r = .58, p < .001$) and were thus averaged to form a "team goodness" index. Analysis of this index revealed that optimizing participants were more likely than control participants to believe in the goodness of their team, $F(1, 65) = 3.97, p = .05$.

Illusion of Influence

Did the optimizing participants' tendency to subjectively optimize their choice of teammates cause them to conclude that an external agent had guided that choice? As Table 1 shows, it did. Optimizing participants were more likely than control participants to believe that their choice of teammates had been influenced by

¹ This procedure was adapted from Yzerbyt, Schadron, Leyens, and Rocher (1994).

Table 1
Measures From Study 1

Measure	Optimizing (<i>n</i> = 34)		Control (<i>n</i> = 33)		Difference
	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>	
Liking of teammate	6.03	0.46	5.70	0.53	.33**
Liking of other 1	4.88	1.01	4.97	0.85	-.09
Liking of other 2	5.71	1.27	5.67	0.89	.04
Liking of other 3	5.35	0.77	5.45	1.03	-.10
Best teammate	5.59	1.18	5.03	1.49	.56*
Confidence in winning	5.71	1.12	5.21	1.08	.50*
Team goodness	5.65	1.00	5.12	1.16	.53**
Perceived influence of message	4.00	1.58	3.18	1.74	.82**
Belief in subliminal messages	5.47	1.38	4.76	1.56	.71**

* $p < .10$. ** $p < .05$.

the subliminal message, $F(1, 65) = 4.07$, $p = .048$, and more likely to believe in the power of subliminal messages in general, $F(1, 65) = 3.94$, $p = .05$.

Conclusion

Participants who liked their teammates were especially likely to believe that a subliminal message that was designed to help them choose a likeable teammate had succeeded in doing so, despite the fact that their liking for their teammate was self-generated. Participants mistakenly concluded that their teammate was the optimal teammate and then attributed this happy outcome to the influence of an external agent that was known to be benevolent and insightful.

Study 2: The Illusion of Insight

The design of Study 2 shared the logic of the design of Study 1, but differed in three important ways. First, participants in Study 1 experienced a suboptimal outcome that was less desirable than some but more desirable than others. In Study 2, we created a more challenging situation for our participants by arranging for them to experience what was objectively the worst of all possible outcomes. Second, in Study 1 we manipulated the likelihood that participants would be able to subjectively optimize their outcome by giving some of them information about the range of potential outcomes (i.e., showing them all four autobiographies) and then asking them to rate each of these outcomes before they knew which one they would experience and by depriving other participants of this information and not requesting that they make such ratings. This meant that control participants made a set of ratings that optimizing participants did not make and that control participants were exposed twice to information to which optimizing participants were exposed just once. In Study 2, we exchanged this manipulation for a somewhat more elegant one. Specifically, in Study 2 we manipulated the likelihood that participants would be able to subjectively optimize their outcome by telling some of them which outcome they would experience before they received information about all the outcomes and telling others which outcome they would experience only after they had received information about all the outcomes. Thus, all participants were exposed

to the same information the same number of times, and only the timing of the delivery of the information was manipulated.

Third, in Study 1 we measured participants' beliefs by verbal report. Previous work on affective forecasting suggests that such reports are valid indicators of peoples' predictions and experiences rather than self-presentational claims made solely for public consumption. For example, in a study by Wilson, Wheatley, and Gilbert (2000), some participants predicted how they would feel after failing to win a date with an attractive partner, whereas others were told they had failed to win the date and then reported how they felt. Optimizing experiencers claimed to feel better than forecasters predicted they would. More importantly, when forecasters and optimizing experiencers were given the opportunity to privately ingest an ostensibly mood-enhancing drug, experiencers chose to ingest smaller doses than forecasters predicted they would. This suggests that those experiencers who were allowed to subjectively optimize their outcomes really did feel better about those outcomes. Nonetheless, to reduce the possibility that optimizing participants in Study 2 were merely claiming to like their suboptimal outcomes, we measured behavior rather than claims. Specifically, we estimated participants' beliefs about the insight of an external agent by asking them to decide whether they wanted that agent to control their future experiences.

In Study 2, we created a situation in which an external agent was known to be benevolent and influential, and we predicted that those participants who were most likely to subjectively optimize their outcomes would be most likely to conclude that the external agent was insightful as well.

Method

Overview

Participants were told that an electronic device had been given the power to select one of two songs for them to hear in its entirety three times (influence) and that the device had been programmed to select the song that it thought the participant would most enjoy (benevolence). Participants either were or were not told which of the two songs had been chosen for them before they heard brief excerpts from the two songs. We predicted that participants who were told which of the two songs had been chosen for them before they heard the excerpts would find it easier to subjectively

optimize the chosen song and would thus be more likely to believe that the device had insight into their personal preferences.

Participants

Participants were 23 male and 44 female undergraduates at the University of Texas at Austin who participated for credit in their introductory psychology course.

Procedure

Cover story. Participants arrived individually at the laboratory for a study of "how artificial intelligence can be used to improve people's daily lives" and were escorted by a male or female experimenter to a cubicle, where they remained for the duration of the study. Participants were told that automobiles of the future would be equipped with "SmartRadios" that would save drivers the trouble of manually switching radio stations. Such a machine would ostensibly be programmed with information about the driver's personality and preferences, thus enabling it to scan the airwaves and select music that the driver would be likely to enjoy. The experimenter explained that he or she was currently testing two SmartRadio prototypes (SmartRadio A and SmartRadio B) to determine which made the best use of personal information when making musical selections for its owner.

Instructions. Participants were told that the study comprised a selection phase and a listening phase. During the selection phase, a SmartRadio would select some music for the participants, and during the listening phase, participants would listen extensively to the music that had been selected for them. Participants were told that the selection phase would consist of two trials. On each trial, either SmartRadio A or SmartRadio B would select from a pair the song that it thought the participant would most like to hear during the listening phase. Participants were reminded that at the beginning of the semester they had completed a battery of personality and attitude questionnaires as part of a mass pretesting session and they were told that both SmartRadios had been programmed with this information. Participants were told that on the first trial they would listen to a 20-sec excerpt from each of two 5-min songs. Participants were told that they could elect to have SmartRadio A or SmartRadio B make a selection for them from this pair. Participants were told that after the first trial was completed the procedure would be repeated with two new songs and that they could once again elect to have either SmartRadio A or SmartRadio B make a musical selection for them. Participants were told that when the two trials of the selection phase were completed they would then listen to each of the 5-min songs that had been selected for them three times, thus, listening to a total of 30 min of music.

Next, participants were asked whether they wanted SmartRadio A or SmartRadio B to make a musical selection for them on the first trial. (We will hereinafter refer to the SmartRadio that the participant chose on Trial 1 as the Chosen SmartRadio). One SmartRadio was said to have a serial processor and the other was said to have a parallel processor, but no other distinguishing information was provided. Once participants made this choice, they were randomly assigned to the control condition or the optimizing condition.

Experimental manipulation. As in Study 1, we manipulated the likelihood that our participants would be able to subjectively optimize their outcomes by relying on the fact that it is generally easier to make up one's mind than to change it. When people know which of several outcomes they are going to experience before they receive information about those outcomes, it is relatively easy for them to selectively attend to the positive features of their outcome and the negative features of the alternative outcomes and, thus, "make up their minds" to like the former. On the other hand, when people do not know which of several outcomes they are going to experience before they receive information about those outcomes, they cannot selectively attend to information about their outcome and are thus likely to form preferences on the basis of other criteria. Later, when they

learn which outcome is theirs, these preferences must be reversed. In short, it is easier to subjectively optimize in prospect than in retrospect, and our manipulation capitalized on this fact.

The experimenter explained to participants in the control condition that he or she would leave them alone in their cubicles for a few minutes so that he or she could return to the control room and run the Chosen SmartRadio program in preparation for Trial 1. The experimenter gave control participants a tape player and a pair of stereo headphones and told them that while he or she was gone they should listen to excerpts from the two songs. The experimenter left control participants in the cubicle for 2 min to listen to a 20-sec excerpt from "Visitation" by Towner (1975) and a 20-sec excerpt from "This is the Abduction Scene" by Torn, Karn, and Bozzio (1994). Pretests indicated that the latter was more pleasant than the former and we will thus refer to these songs as the *good song* and the *bad song*, respectively, though participants, of course, never heard them described in this way. At the end of the 2-min period, the experimenter returned and told control participants that the Chosen SmartRadio had selected the bad song and that they would thus listen to the bad song three times in its entirety during the listening phase. In short, after hearing excerpts from the good and the bad songs, control participants learned that they would later hear the bad song three times in its entirety.

Conversely, optimizing participants learned that they would later hear the bad song three times in its entirety before they heard excerpts from the good and the bad songs. Participants in the optimizing condition were told that the experimenter would leave for a few minutes to run the Chosen SmartRadio program in preparation for Trial 1. The experimenter left, returned 2 min later, and told optimizing participants that the Chosen SmartRadio had selected the bad song and that they would thus listen to the bad song three times in its entirety during the listening phase. The experimenter then gave optimizing participants a tape player and a pair of stereo headphones and told them that they should listen to the excerpts from the two songs.

Dependent measures. After listening to the two excerpts, participants were asked to report how much they liked each excerpt and how much they thought they would like hearing each excerpt three times. These ratings were made on 9-point scales that ranged from 1 (*not at all*) to 9 (*very much*). In addition, participants were asked to recall which of the excerpts the Chosen SmartRadio had selected for them.

Participants were then told to prepare for the second trial, during which they would hear two new musical excerpts and during which one of the SmartRadios would select a second song from a pair of new songs. Participants were told that they would also hear the second selection three times in its entirety during the listening phase. Participants were told that they could choose to have either of the SmartRadios make their second selection for them. Participants were told that "some people find it useful to know which selection the [Unchosen SmartRadio] chose for them in Trial 1, and for that reason, I always run both programs. In case you're interested, the other program concluded that you would prefer [the good song]." Participants were then asked, among other things, which of the SmartRadios they would like to have make a musical selection for them on Trial 2. We assumed that participants' choices would be influenced by many factors (e.g., curiosity about the Unchosen SmartRadio, loyalty to the Chosen SmartRadio, and so forth) but that in addition to these factors, participants would prefer to have the second song selected by the SmartRadio that they believed had the greatest insight into their musical tastes. Thus, the participants' choice of SmartRadios for Trial 2 constituted the behavioral measure of their beliefs about the Chosen SmartRadio's insight into their personal preferences.

Participants were also asked which of the two SmartRadios they would purchase at an electronics store if both were available. After answering this question, participants were very carefully debriefed.

Results

Our hypothesis suggested that optimizing participants would find it easier than control participants to subjectively optimize the bad song that was chosen for them. As such, we expected that optimizing participants would be more likely than control participants to consider the Chosen SmartRadio especially insightful and, thus, more likely to allow it to select music for them on the second trial.

Omissions of Data

Three participants in the optimizing condition and 1 participant in the control condition were confused about key aspects of the procedure, and 1 participant in the optimizing condition and 1 participant in the control condition were extremely suspicious of the cover story. Data from these 6 participants were removed, thus, leaving 61 participants in the data set.

Subjective Optimizing

Were optimizing participants more likely than control participants to subjectively optimize the bad song? As the top half of Table 2 shows, they were. The two liking ratings were highly correlated ($r = .83$) and thus were averaged to form a liking index. Whereas control participants showed a significant preference for the good song, $t(59) = 3.15$, $p < .003$, optimizing participants showed no such preference and liked the bad song that was selected for them just as much as they liked the good song that was not selected, $t(59) < 1$.

Perceptions of Insight

Did the optimizing participants' tendency to subjectively optimize the bad song also cause them to conclude that the Chosen SmartRadio was especially insightful? Participants' behavior suggests that it did. Whereas only 6 of 31 (19%) control participants elected to have a song chosen for them on Trial 2 by the SmartRadio that had previously chosen the bad song for them on Trial 1, 13 of 30 (43%) optimizing participants did just that, $\chi^2(1, 61) = 4.09$, $p = .042$. Similarly, whereas only 5 (16%) control participants reported that they would buy the Chosen SmartRadio rather than the Unchosen SmartRadio if both were available, 15 (50%) optimizing participants claimed that they would buy the Chosen rather than the Unchosen SmartRadio, $\chi^2(1, 61) = 7.94$, $p < .005$. The data shown in the bottom half of Table 2 leave little doubt that their liking for the songs influenced participants' choices. The 42

participants who chose to switch SmartRadios on Trial 2 preferred the good song to the bad song, $t(40) = 8.41$, $p < .001$, whereas the 19 participants who chose not to switch SmartRadios on Trial 2 actually preferred the bad song to the good song, $t(40) = 4.34$, $p < .005$.

Conclusion

Participants who liked the song that was chosen for them were especially likely to put their fates in the hands of the SmartRadio that chose it, despite the fact that their liking for the song was self-generated. Participants mistakenly concluded that the song to which they would be exposed was the optimal song and attributed this happy outcome to the insight of an external agent that was known to be benevolent and powerful.

Study 3: Illusion of Benevolence

In Studies 1 and 2, we assumed that all participants were motivated to subjectively optimize their outcomes and we manipulated their ability to do so. We reasoned that if our theorizing was correct, then we should be able to manipulate the likelihood that participants would subjectively optimize their outcomes by manipulating either their ability or their motivation to do so. In Study 3 we manipulated both. We created a situation in which an external agent was known to be insightful and influential, and we predicted that only those participants who were both motivated and able to subjectively optimize their outcomes would conclude that the external agent was benevolent as well.

Method

Overview

Participants were told that another participant had been given the power to select and bestow one of two gifts (influence) and that the gift giver had some personal information about the preferences of the recipient (insight). Some participants (recipients) were told that they would be the recipients of the gift, and other participants (observers) were told that a third participant would be the recipient of the gift. In addition, participants either were or were not told which of the two gifts the gift giver had selected before they made their own ratings of the gift. We predicted that participants who were told which of the two gifts had been chosen before making their own ratings (ability manipulation) and who expected to receive that gift themselves (motivation manipulation) would subjectively optimize the gift and would thus be more likely than any other participants to believe that the gift giver was benevolent.

Participants

Participants were 49 female undergraduates at the University of Texas at Austin who participated for credit in their introductory psychology course. Because some of the gifts used in the study were likely to have differential appeal to men and to women, we invited only female students to participate.

Procedure

Participants arrived in pairs at the laboratory for a study of "the psychology of gift giving" and were escorted by a male or female experimenter to an individual cubicle, where they remained for the duration of the

Table 2
Liking of Songs in Study 2

Group	Good song		Bad song		Preference for good song (difference)
	M	SD	M	SD	
Control ($n = 31$)	3.86	1.84	2.01	1.05	1.85*
Optimizing ($n = 30$)	3.26	2.01	3.08	1.90	0.18
Switch ($n = 42$)	4.04	1.86	1.96	1.07	2.08*
No switch ($n = 19$)	2.51	1.69	3.82	1.83	-1.31*

* $p < .05$.

Table 3
Liking of Items and Perceptions of Benevolence in Study 3

Condition	Gift item		Non-gift item		Preference for gift		Considerateness	
	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>
Observers								
Control (<i>n</i> = 12)	3.08	1.08	4.83	1.64	-1.75	2.18	5.00	1.13
Optimizing (<i>n</i> = 13)	4.23	1.09	5.08	1.75	-0.85	2.30	5.08	1.38
Recipients								
Control (<i>n</i> = 12)	4.42	1.24	4.92	1.88	-0.50	2.65	5.68	0.49
Optimizing (<i>n</i> = 12)	5.92	1.17	4.46	1.73	1.46	2.21	6.58	0.67

study. Participants were told that a third participant had already arrived and was seated in an adjacent cubicle. In fact, there was no third participant.

The experimenter explained that gift giving was an important part of everyday life but that it was rarely studied and hence poorly understood. The experimenter explained that he or she would ask the participants to play a multitrail game in which they would play the role of giver, recipient, or observer. The experimenter explained that on each trial the giver would select a gift from a list of potential gifts and that the recipient would actually receive the gift at the conclusion of the study. The experimenter explained that after each trial, the recipient and the observer would make some evaluations of the gift-giving experience and then switch roles for the next trial.

Participants were asked to draw from a cup a folded slip of paper that contained their role assignment. The experimenter explained that there were only two folded slips of paper in the cup because the third participant had already randomly selected the role of giver. Thus, the two participants each drew a slip of paper from the cup and were thereby randomly assigned to the role of recipient or observer. Next, the experimenter asked the participants to provide written answers to a variety of personal questions (e.g., "Are you a member of a sorority?" "How do you feel about affirmative action?" "Do you consider yourself an introvert or an extrovert?"). Participants were told that their answers would be shown to the giver, who would use the information to select a gift for the recipient on Trial 1 and also for the observer on a later trial (when she played the role of recipient). When the participants completed the questionnaire, the experimenter allowed the observer to read the recipient's answers, and vice versa.² The experimenter then left the room, ostensibly to show the recipient's and the observer's answers to the giver. Participants were then randomly assigned to either the control or the optimizing condition.

The experimenter returned a few minutes later and provided all participants with a list that contained two items: a stuffed animal and a certificate for a free tank of gas at any gas station in the city. Optimizing participants were told that the giver had been instructed to select a gift from this pair and that, on the basis of the recipient's answers to the questions on the questionnaire, she had selected the stuffed animal as a gift for the recipient. Pretesting had indicated that this was the less desirable of the two gifts. Optimizing participants were asked to rate the two items on 7-point Likert-type scales that ranged from 1 (*not at all good*) to 7 (*very good*) and to report how considerate they believed the giver to be on a 7-point Likert-type scale that ranged from 1 (*not very considerate*) to 7 (*very considerate*).³ Control participants were told that the giver had not yet selected a gift for the recipient and were asked to report their liking for the two items. After they did so, the experimenter left briefly and then returned to tell the control participants that the giver had chosen the stuffed animal for the recipient. Control participants again reported their liking for the two items and also reported how considerate they believed the giver to be. When the experiment was finished, participants were carefully and thoroughly debriefed.

Results

Our hypothesis suggested that optimizing recipients would be the only participants who found it both easy and necessary to subjectively optimize the gift. As such, we expected that optimizing recipients would rate the gift more highly than would other participants and that they would also be most likely to attribute benevolence to the gift giver.

Subjective Optimizing

Were optimizing recipients more likely than other participants to subjectively optimize the gift? As Table 3 shows, they were. A weighted contrast analysis (using weights of 3 for the optimizing recipients condition and -1 for all other conditions) revealed that only the optimizing recipients preferred the stuffed animal to the tank of gas, $F(1, 45) = 10.25, p < .005$. Specific contrasts revealed that optimizing recipients had stronger preferences for the stuffed animal than did control recipients, $F(1, 45) = 4.20, p = .046$, but that optimizing observers and control observers had equally weak preferences ($F < 1$).

Perceptions of Benevolence

Did the optimizing recipients' tendency to subjectively optimize their outcome cause them to conclude that the giver was especially benevolent? As Table 3 shows, it did. A weighted contrast analysis (as above) revealed that the optimizing recipients considered the giver more considerate than did any of the other participants, $F(1, 45) = 16.33, p < .001$. Specific contrasts revealed that optimizing recipients considered the giver more considerate than did control recipients, $F(1, 45) = 5.09, p = .028$, but that optimizing observers and control observers considered the giver equally considerate ($F < 1$).

² Occasionally, only 1 participant arrived for the experimental session. In those instances, participants were told that both the second and third participants were in adjacent cubicles, and they were allowed to see the answers of a previous participant who had played the complementary role.

³ Informal pretesting suggested that a significant proportion of our participant population did not know the meaning of the word *benevolent* and thus we used this synonym instead.

Conclusion

Participants who received a gift they liked were especially likely to conclude that the gift giver was benevolent, despite the fact that their liking for the gift was self-generated. Participants mistakenly concluded that their gift was the optimal gift and attributed this happy outcome to the benevolence of an external agent who was known to be influential and insightful.

General Discussion

The foregoing studies suggest that people are more talented than they realize. Although participants displayed a healthy capacity for generating satisfaction with imperfect partners, unpleasant music, and mediocre gifts, they acted as though they were unaware that they had done so. Instead, they behaved as though the outcomes they experienced were objectively optimal and that a powerful, insightful, and benevolent agent had had a hand in their good fortune. Technically speaking, participants confused their own optimization of subjective reality with an external agent's optimization of objective reality. Simply speaking, participants mistook "the magic in here" for "the magic out there."

The illusion of external agency may help explain why most people believe in "an all-wise and all-powerful Creator whose purpose is beneficent." Skeptics who dismiss religious belief as the mindless acceptance of dogma and myth fail to realize that such beliefs often provide reasonable explanations for personal experience. Just as the poker player who consistently draws a winning hand may arouse reasonable suspicion when other players notice that things are working out far too well and not for them, the person who experiences subjectively optimal outcomes with sufficient strength or regularity may reasonably conclude that an intentional force—a benevolent, powerful, and insightful force—was guiding the outcome's production. In this sense, religious belief is "a particular interpretation of the world, an interpretation whose conclusions (but not whose topics or even logic) differ from those of secular thought" (Guthrie, 1993, p. 21).⁴

But if it is reasonable to seek explanations for the unusual quality or consistency of one's subjectively optimal outcomes, the present studies suggest that people may tend to look for those explanations on the wrong side of their skins. The illusion of external agency is one of many instances in which people seem to underestimate the power of psychological processes—and overestimate the power of external stimuli—to generate experience. For example, research suggests that people have little appreciation for the ease with which internal processes can shape and reshape their memories of an event and that they thus behave as though their memories were exclusively the consequences of the events they depict (e.g., Loftus, 1992). Only occasionally and by rational analysis do people recognize that their memories of episodes were generated or changed by internal processes (Johnson & Raye, 1981). Research on visual perception similarly suggests that people tend to believe that objects and episodes really are the way they appear, in part because they fail to appreciate the contribution of internal processes to their perceptual experience (e.g., Rock & Nijhawan, 1989; Sherif, 1935). Optical illusions are delightful precisely because they confront observers with the fact that perception is a constructive process and remind them that they are always seeing interpretations of objects rather than objects them-

selves—a point that is normally lost in everyday life. Similarly, when asked to predict their reactions to future events ("What would you do if your car were stolen?"), people generally fail to recognize that they have generated an image of the event that contains more detail than they were given (e.g., imagining the car being stolen from the garage rather than being carjacked). Because they do not appreciate the extent to which their own psychological processes contributed to their image of the future event, they forecast their reactions to that event with far too much confidence (Dunning, Griffin, Milojkovic, & Ross, 1990; Ross & Ward, 1996). These examples suggest that when people represent the past in memory, the present in perception, or the future in imagination, they tacitly assume that these representations were generated exclusively by the objects or episodes they are remembering, seeing, or considering. That is, people think of mental representations as faithful internal replicas of the things in the world that they are about (Higgins, 1998).

What is true of mental representations is even truer of feelings, moods, affects, and emotions. Just as the memories, perceptions, and imaginings one experiences are always a joint product of external events and internal mechanisms, so too are affective experiences. Nonetheless, research suggests that people naturally take their affective experiences as information about the state of the world rather than about the state of their minds (Schwarz & Clore, 1996). Numerous studies demonstrate the difficulty people encounter when trying to identify the particular external event that caused an affective experience. For example, people may confuse the happiness induced by good weather with the happiness produced by the general quality of their lives (Schwarz & Clore, 1983), they may confuse the excitement aroused by a dangerous situation with the excitement aroused by an attractive partner (Dutton & Aron, 1974), and they may confuse the anxiety evoked by embarrassing behavior with the anxiety evoked by a pharmacological agent (Zanna & Cooper, 1974). Indeed, it is only with very clever experimental engineering that people can be induced to misattribute the affect produced by an external source to an internal source (Dienstbier & Munter, 1971).

Perhaps the most striking instance of the tendency to underestimate the role of psychological processes in emotional experience is Capgras syndrome—a bizarre neurological disorder that leads people to believe that a significant other (often a child, spouse, or parent) has been replaced by a replica. Young (1994) suggests that familiar others normally trigger responses in both an emotional and a cognitive recognition system and that Capgras syndrome results from a failure of the former. As such, the Capgras sufferer consciously recognizes the physical features of the significant other, but does not experience the unverbalizable feelings that normally accompany such recognition. Despite the assurances of family and physician, Capgras sufferers insist on looking outward rather than inward for explanations of their blunted emotional

⁴ One might argue that it is irrational to attribute a string of subjectively optimal outcomes to the intervention of an intentional agent rather than to a random "run of luck," because the former explanation is more complex than the latter. But the relative costs of Type I and Type II errors can make a preference for complex rather than simple explanations perfectly rational. As Guthrie (1993) notes, "It is better for a hiker to mistake a boulder for a bear, than to mistake a bear for a boulder" (p. 6).

experience and thus conclude that their daughters, wives, or mothers cannot really be their daughters, wives, or mothers, because they do not evoke the feelings of warmth and immediacy that daughters, wives, and mothers normally evoke. As Dennett (1996) notes, "Instead of blaming the mismatch on a faulty perceptual system, the agent blames the world" (p. 113). According to our analysis, the Capgras sufferer's conclusions are but a caricature of the very ordinary tendency to attribute one's emotional experience to objects and events rather than to psychological processes.

Why might people have a tendency to attribute their affective states to events in the world? We can think of at least two reasons. First, one of the primary functions of emotion is to serve as a "system alert" that directs an organism's attention to features of the environment that are potentially important for its continued well-being. As such, people may be inclined by nature to look first to the environment for the sources of their affective states, and research in many domains suggests that where one looks first is often where one looks last (Gilbert, in press; Tversky & Kahneman, 1974). Second, although people have some well-articulated theories about the mind's ability to generate its own satisfaction (e.g., sour grapes), what they know in the abstract may be difficult for them to observe in action. The satisfaction one experiences as a result of subjectively optimizing a suboptimal outcome may feel exactly like the satisfaction one experiences upon attaining an objectively optimal outcome. If no qualia are uniquely associated with self-generated satisfaction—that is, if the quality of satisfaction does not betray its source—then there may be no principled way for people to know by introspection that they have subjectively optimized their outcomes. Rationalization feels rational, and without internal clues to suggest otherwise, it may be perfectly reasonable for people to assume that their satisfaction with an outcome is the result of the outcome's intrinsic properties.

Coda

A mind is the most complicated device any of us will ever own, and yet it comes with no instruction manual. So we make one up as we go along. Sometimes we get the facts wrong, of course, and the illusion of external agency is just one of several phenomena that live at this intersection of how our minds work and how we think our minds work. Our ability to synthesize satisfaction—to reconstrue outcomes so that they give rise to positive experience—is nothing short of remarkable. But it also tends to be invisible, and the fact that we do not know our own strength leaves us susceptible to the mistaken belief that an insightful and benevolent external agent arranged the experience for us. Human lives may or may not be guided by a wise and loving God, and psychology experiments cannot resolve this inherently theological question. But they can remind us that in a world of many mysteries some of the most significant miracles are those we perform ourselves.

References

- Bargh, J. A., Chaiken, S., Govender, R., & Pratto, F. (1992). The generality of the automatic attitude activation effect. *Journal of Personality and Social Psychology*, 62, 893-912.
- Bargh, J. A., Chaiken, S., Raymond, P., & Hymes, C. (1996). The automatic evaluation effect: Unconditional automatic attitude activation with a pronunciation task. *Journal of Experimental Social Psychology*, 32, 185-210.
- Beggs, J. K. (1992). On the social nature of nonsocial perception: The mere ownership effect. *Journal of Personality and Social Psychology*, 62, 229-237.
- Brehm, J. W. (1956). Post-decision changes in desirability of alternatives. *Journal of Abnormal and Social Psychology*, 52, 384-389.
- Darley, J. M., & Berscheid, E. (1967). Increased liking caused by the anticipation of interpersonal contact. *Human Relations*, 10, 29-40.
- Dennett, D. (1996). *Kinds of minds*. New York: Basic Books.
- Dienstbier, R. A., & Munter, P. C. (1971). Cheating as a function of the labeling of natural arousal. *Journal of Personality and Social Psychology*, 17, 208-213.
- Dunning, D. (1995). Trait importance and modifiability as factors influencing self-assessment and self-enhancement motives. *Personality and Social Psychology Bulletin*, 21, 1297-1306.
- Dunning, D., Griffin, D. W., Milojkovic, J., & Ross, L. (1990). The overconfidence effect in social prediction. *Journal of Personality and Social Psychology*, 58, 568-581.
- Dunning, D., Leuenberger, A., & Sherman, D. A. (1995). A new look at motivated inference: Are self-serving theories of success a product of motivational forces? *Journal of Personality and Social Psychology*, 69, 58-68.
- Dutton, D. G., & Aron, A. P. (1974). Some evidence for heightened sexual attraction under conditions of high anxiety. *Journal of Personality and Social Psychology*, 30, 510-517.
- Eagly, A. H., & Chaiken, S. (1998). Attitude structure and function. In D. T. Gilbert, S. T. Fiske, & G. Lindzey (Eds.), *Handbook of social psychology* (4th ed., pp. 269-322). Boston: McGraw-Hill.
- Fazio, R. H., Sanbonmatsu, D. M., Powell, M. C., & Kardes, F. R. (1986). On the automatic activation of attitudes. *Journal of Personality and Social Psychology*, 50, 229-238.
- Festinger, L. (1957). *A theory of cognitive dissonance*. Stanford, CA: Stanford University Press.
- Gallup, G., Jr., & Castelli, J. (1989). *The people's religion: American faith in the 90s*. Macmillan: New York.
- Gallup, G., Jr., & Jones, S. (1989). *100 questions and answers: Religion in America*. Princeton, NJ: Hermitage Press.
- Gilbert, D. T. (in press). Inferential correction. In T. Gilovich, D. W. Griffin, & D. Kahneman (Eds.), *The psychology of judgment: Heuristics and biases*. Cambridge, England: Cambridge University Press.
- Gilbert, D. T., & Jenkins, J. (1999). *Decisions and revisions: The affective forecasting of escapable outcomes*. Unpublished manuscript, Harvard University.
- Gilbert, D. T., Pinel, E. C., Wilson, T. D., Blumberg, S. J., & Wheatley, T. P. (1998). Immune neglect: A source of durability bias in affective forecasting. *Journal of Personality and Social Psychology*, 75, 617-638.
- Gilbert, D. T., & Wilson, T. D. (2000). Miswanting: Some problems in the forecasting of future affective states. In J. Forgas (Ed.), *Thinking and feeling: The role of affect in social cognition* (pp. 178-197). Cambridge, England: Cambridge University Press.
- Gross, J. J. (1999). Emotion regulation: Past, present, and future. *Cognition and Emotion*, 13, 551-573.
- Guthrie, S. E. (1993). *Faces in the clouds: A new theory of religion*. Oxford, England: Oxford University Press.
- Heider, F. (1958). *The psychology of interpersonal relations*. New York: Wiley.
- Higgins, E. T. (1998). The aboutness principle: A pervasive influence on human performance. *Social Cognition*, 16, 173-198.
- Johnson, M. K., & Raye, C. L. (1981). Reality monitoring. *Psychological Review*, 88, 67-85.
- Jones, E. E., & Davis, K. E. (1965). From acts to dispositions: The attribution process in person perception. In L. Berkowitz (Ed.), *Ad-*

- vances in experimental social psychology (Vol. 2, pp. 219-266). New York: Academic Press.
- Kahneman, D., Kentsch, J., & Thaler, D. (1991). The endowment effect, loss aversion, and status quo bias. *Journal of Economic Perspectives*, 5, 193-206.
- Kunda, Z. (1990). The case for motivated reasoning. *Psychological Bulletin*, 108, 480-498.
- Loewenstein, G., & Adler, D. (1995). A bias in the prediction of tastes. *Economic Journal*, 105, 929-937.
- Loftus, E. F. (1992). When a lie becomes memory's truth: Memory distortion after exposure to misinformation. *Current Directions in Psychological Sciences*, 1, 121-123.
- Lord, C. G., Ross, L., & Lepper, M. R. (1979). Biased assimilation and attitude polarization: The effects of prior theories on subsequently considered evidence. *Journal of Personality and Social Psychology*, 37, 2098-2109.
- Lyubomirsky, S., & Ross, L. (1999). Changes in attractiveness of elected, rejected, and precluded alternatives: A comparison of happy and unhappy individuals. *Journal of Personality and Social Psychology*, 76, 988-1007.
- Petty, R. E., & Wegener, D. T. (1998). Attitude change: Multiple roles for persuasion variables. In D. T. Gilbert, S. T. Fiske, & G. Lindzey (Eds.), *Handbook of social psychology* (4th ed., pp. 323-390). Boston: McGraw-Hill.
- Rock, I., & Nijhawan, R. (1989). Regression to egocentrically determined description of form under conditions of inattention. *Journal of Experimental Psychology: Human Perception and Performance*, 15, 259-272.
- Ross, L., & Ward, A. (1996). Naive realism in everyday life: Implications for social conflict and misunderstanding. In E. S. Reed, E. Turiel, & T. Brown (Eds.), *Values and knowledge* (pp. 103-135). Mahwah, NJ: Erlbaum.
- Russell, B. (1997). The faith of a rationalist. In J. G. Slater & P. Kollner (Eds.), *The collected papers of Bertrand Russell* (Vol. 11, pp. 83-88). London: Routledge. (Original work published 1947)
- Schwarz, N., & Clore, G. L. (1983). Mood, misattribution, and judgments of well-being: Informative and directive functions of affective states. *Journal of Personality and Social Psychology*, 45, 513-523.
- Schwarz, N., & Clore, G. L. (1996). Feelings and phenomenal experience. In E. T. Higgins & A. Kruglanski (Eds.), *Social psychology: A handbook of basic principles* (pp. 433-465). New York: Guilford Press.
- Sherif, M. (1935). A study of some social factors in perception. *Archives of Psychology*, 27, 1-60.
- Taylor, S. E. (1989). *Positive illusions*. New York: Basic Books.
- Taylor, S. E., & Brown, J. D. (1988). Illusion and well-being: A social-psychological perspective on mental health. *Psychological Bulletin*, 103, 193-210.
- Tice, D. M., & Baumeister, R. F. (1993). Controlling anger: Self-induced emotion change. In D. M. Wegner & J. W. Pennebaker (Eds.), *Handbook of mental control* (pp. 393-409). Englewood Cliffs, NJ: Prentice Hall.
- Torn, D., Karn, M., & Bozzio, T. (1994). This is the abduction scene. On *Polytown* [CD]. New York: CMP Records.
- Towner, R. (1975). Visitation [Recorded by R. Towner, J. Garbarek, E. Weber, & J. Christensen]. On *Solstice* [CD]. New York: ECM.
- Tversky, A., & Kahneman, D. (1974, September 27). Judgment under uncertainty: Heuristics and biases. *Science*, 185, 1124-1131.
- Vaillant, G. (1993). *The wisdom of the ego*. Cambridge, MA: Harvard University Press.
- Wilson, T. D., Wheatley, T. P., & Gilbert, D. T. (2000). [A behavioral measure of self-reported happiness]. Unpublished raw data.
- Young, A. (1994). The neuropsychology of awareness. In A. Revonsuo & M. Kamppinen (Eds.), *Consciousness in philosophy and neuroscience* (pp. 173-203). Hillsdale, NJ: Erlbaum.
- Yzerbyt, V. Y., Schadron, G., Leyens, J. P., & Rocher, S. (1994). Social judgeability: The impact of meta-informational cues on the use of stereotypes. *Journal of Personality and Social Psychology*, 66, 48-55.
- Zanna, M. P., & Cooper, J. (1974). Dissonance and the pill: An attribution approach to studying the arousal properties of dissonance. *Journal of Personality and Social Psychology*, 29, 703-709.

Received October 25, 1999

Revision received March 1, 2000

Accepted March 15, 2000 ■