



Why so confident? The influence of outcome desirability on selective exposure and likelihood judgment [☆]

Paul D. Windschitl ^{a,*}, Aaron M. Scherer ^a, Andrew R. Smith ^c, Jason P. Rose ^b

^a Department of Psychology, University of Iowa, Iowa City, IA 52242, United States

^b Department of Psychology, University of Toledo, Toledo, OH 43606, United States

^c Department of Psychology, Appalachian State University, Boone, NC 28608, United States

ARTICLE INFO

Article history:

Received 10 June 2011

Accepted 22 October 2012

Keywords:

Selective exposure
Confirmation bias
Desirability bias
Wishful thinking
Overconfidence
Likelihood judgment

ABSTRACT

Previous studies that have directly manipulated outcome desirability have often found little effect on likelihood judgments (i.e., no desirability bias or wishful thinking). The present studies tested whether selections of new information about outcomes would be impacted by outcome desirability, thereby biasing likelihood judgments. In Study 1, participants made predictions about novel outcomes and then selected additional information to read from a buffet. They favored information supporting their prediction, and this fueled an increase in confidence. Studies 2 and 3 directly manipulated outcome desirability through monetary means. If a target outcome (randomly preselected) was made especially desirable, then participants tended to select information that supported the outcome. If made undesirable, less supporting information was selected. Selection bias was again linked to subsequent likelihood judgments. These results constitute novel evidence for the role of selective exposure in cases of overconfidence and desirability bias in likelihood judgments.

© 2012 Elsevier Inc. All rights reserved.

Introduction

People routinely face uncertainty and grapple with questions such as “Is it true?” and “Will it happen?” In this Information Age, when people ponder such questions, they can often readily access relevant information. However, the available information can be heterogeneous in its implications, and the sheer amount of it can be daunting. Therefore, the act of selecting some information to consider further, while leaving other information neglected, becomes critical. The potential for bias is substantial. It is easy to imagine how fund managers, policy makers, medical patients, and others who seek only selective types of information could develop distorted expectations and confidence about target outcomes, leading to bad decisions and consequences.

The present paper addresses the influence that people’s motives for a particular conclusion can have on information selection and resulting confidence levels (i.e., likelihood judgment). We had three main research questions. First, does the desirability of an outcome have a causal impact on information selection. Second, what is the direction of the effect?—Does high desirability fuel the seeking of supporting evidence? Third, what role does a selection bias have in shaping confidence/optimism about the outcome? As a

concrete example, imagine that Alex learns from her financial advisor that she will earn more from her stock holdings if Company A and B merge. Naturally, Alex now hopes these two companies will merge. If she becomes curious about the prospects of the merger, would Alex’s desire for the merger bias her interest in reading information that appears to support or cast doubt on the merger? Does the desire ultimately bias her perception of the likelihood of the merger?

To test our research questions, we developed a paradigm that involves experimental manipulations of outcome desirability, as well as measures of both information selection and likelihood judgment. We know of no other published study that includes all these features. There are, however, two literatures that include studies relevant to various parts of our research—the literature on motivated reasoning and the more narrowly defined literature on the desirability bias. In the following sections, we first discuss how our work relates to—and is distinguishable from—existing research on motivated reasoning. Then we discuss how our research extends the current literature on the desirability bias.

Motivated reasoning

The literature defined by the term *motivated reasoning* is vast. As many review papers attest, people are often prone to arrive at conclusions they find desirable or comforting (Balcetis, 2008; Kunda, 1990; Pyszczynski & Greenberg, 1987; Roese & Olson, 2007; Taylor & Brown, 1988; Trope & Liberman, 1996). Many cognitive processes

[☆] This work was supported by Grant SES 07-20349 to Paul D. Windschitl from the National Science Foundation.

* Corresponding author. Fax: +1 319 335 0191.

E-mail address: paul-windschitl@uiowa.edu (P.D. Windschitl).

are flexibly dependent on directional motives—including attention, visual perception, memory processes, depth of processing, and logical reasoning (e.g., Balcetis & Dunning, 2006; Clark & Wegener, 2008; Dawson, Gilovich, & Regan, 2002; McDonald & Hirt, 1997). Most pertinent to the present paper would be research showing that motivations influence information selection. For example, Holton and Pyszczynski (1989) found that receiving harsh feedback from a confederate increased participants' interest in seeing negative information about the confederate. And work using selective-exposure paradigms reveals that people's tendency to view and process information depends on whether it is expected to fit with current attitudes and recent choices (for reviews, see Hart et al., 2009; Jonas, Schulz-Hardt, Fischer, & Frey, 2006).

This motivated-reasoning research provides general fodder for expecting that participants in our studies would tend to select information favoring an outcome they desire. However, there are two features of our research that inject some healthy skepticism as to whether findings from previous work can be presumed to provide answers to our research questions (i.e., with no need for an empirical test).

One important feature is that we focus on cases in which people are tasked with judging likelihood, and they are aware that there will be a *moment of truth*. That is, we are interested in cases when people know that they will be learning whether the outcome about which they provided a judgment did or did not happen (was or was not true). This characteristic distinguishes our studies from many studies within the literature on motivated reasoning. In most studies of motivated reasoning, people do not need to worry about their conclusions being invalidated or checked for accuracy. For example, when people change their attitudes to avoid dissonance (Festinger & Carlsmith, 1959), change their self-perceived traits after learning what traits bode well for a successful life (Dunning, 2003; Kunda & Sanitioso, 1989), or change how they rate the validity of a test because they failed it (Wyer & Frey, 1983), they do not need to worry that their motivated conclusions will be invalidated soon (or perhaps ever). There is no impending moment of truth.

There are reasons to suspect that optimistic distortions in information search and subsequent judgments might be dampened or absent (possibly even reversed) when there is a moment of truth in sight. When a moment of truth is relevant, accuracy motivations might be enhanced, leading people to attend to evidence more carefully and avoid letting motivated biases influence their information gathering and processing (Gilovich, Kerr, & Medvec, 1993; see also Armor & Sackett, 2006; Tyler & Rosier, 2009). Also, in contexts when a person desires an outcome and will soon learn the true outcome, that person knows he/she will either be pleased or disappointed. Because unexpected bad news is worse than expected bad news, people might brace for bad news by becoming increasingly pessimistic (see Shepperd & McNulty, 2002; Shepperd, Ouellette, & Fernandez, 1996; Sweeny & Krizan, in press; Weber, 1994). They might even become more curious about whether bad news is coming, so they seek out and check information consistent with an undesirable outcome, which could provide evidence for a pessimistic likelihood judgment.

A second important feature is that we designed our paradigm to test for the effect of desirability when it is clearly unconfounded with other factors. In our main studies, which are described later, we used random assignment and experimentally created different levels of outcome desirability (the desirability was newly established), thereby ensuring that outcome desirability varied independently of other outcome characteristics or associations. This strategy differed from previous studies that have harnessed existing differences in desirability rather than directly manipulating it. The strategy of using existing differences leaves these previous studies open to alternative interpretations. For example, several studies have shown correlations between the extent to which

respondents rated political or sports outcomes as desirable and the extent to which they expected those outcomes to occur (e.g., Babad, 1997; Granberg & Brent, 1983). Whereas one interpretation of these correlations is that desires drove expectations, the opposite causal path is equally plausible (see Kay, Jimenez, & Jost, 2002), and third-variable interpretations are also viable (for discussions, see Krizan, Miller, & Johar, 2010; Krizan & Windschitl, 2007, 2009; Massey, Simmons, & Armor, 2011).

Even among studies that have avoided obvious confounds associated with not experimentally manipulating desirability, problems relating to preexisting differences still persist. Consider, for example, a clever paradigm used by Ditto, Munro, Apanovitch, Scepanisky, and Lockhart (2003) in which participants had to interpret the results for a saliva test. They scrutinized the test results to different degrees as a function of whether they thought the result suggested good health outcomes or bad health outcomes. This is an important and fascinating result. However, as Ditto and his colleagues documented, the college-student participants had an *a priori* expectation that the test results would be favorable—leading to greater scrutiny of an unfavorable result. Ditto et al. noted that the *a priori* expectation might be due, quite rationally, to the fact that participants tended to have a history of good health (or motivational processes that operate over time to bolster an expectation of good health).¹ These unresolved possibilities do not provide an answer to whether a newly established desire that is unconfounded with other factors can have immediate consequences on information selection and optimism.

Previous studies on the desirability bias

The most direct way of testing the influence of desirability on optimism is to experimentally manipulate outcome desirability independently of other outcome characteristics or associations, and then solicit forecasts about the outcomes. This is precisely what many studies on the *desirability bias* (aka *wishful thinking*) have done (e.g., Bar-Hillel & Budescu, 1995; Bar-Hillel, Budescu, & Amar, 2008; Irwin, 1953; Lench & Ditto, 2008; Marks, 1951; Windschitl, Smith, Rose, & Krizan, 2010). In a typical version of these studies, participants learn about two possible outcomes and are given a monetary reason—manipulated independently of all other factors—for hoping that one outcome is the true outcome.

One of the more surprising findings to emerge from this literature is that the nature of the forecast being solicited—a discrete prediction vs. a scaled judgment—has a strong impact on the whether a desirability bias is detected (for a meta-analysis, see Krizan & Windschitl, 2007). Studies using a classic marked card paradigm in which participants make a discrete outcome prediction about whether a marked card will be drawn from a deck show that participants are more likely to predict a marked card when it would be a desirable outcome rather than neutral (e.g., Irwin, 1953; Marks, 1951; Windschitl et al., 2010). However, the fact that discrete predictions are influenced by desirability can be explained without assuming that people alter their internal assessments of likelihood about the outcomes. For example, a differential-threshold account suggests that the desirability of an outcome doesn't change how evidence is sought or evaluated, but instead simply lowers the threshold for making an affirmative prediction (Bar-Hillel & Budescu, 1995; Krizan & Windschitl, 2007; Price & Marquez,

¹ In an attempt to isolate the role of motivation, Ditto et al. (2003) showed that observer participants, who did not share the same motivations as actor-participants, did not exhibit the same effects when making judgments about actor participants described in a vignette. However, we believe there are significant limitations with this approach (e.g., observer-participants would have not only lacked the same motivations as actor-participants, they would have also lacked any basis for strong *a priori* expectations about the actors).

2005; see also the biased-guessing account by Windschitl et al. (2010)).

Indeed, across different paradigms, researchers have very rarely found evidence that a manipulation of outcome desirability influences scaled judgments of how likely the outcome is (e.g., on a 0–100% scale; see Bar-Hillel & Budescu, 1995; Krizan & Windschitl, 2007). Although both Price and Marquez (2005) and Windschitl et al. (2010) were able to use card paradigms to show that outcome desirability affects discrete predictions, a switch to a scaled likelihood judgment led to nonsignificant desirability effects within otherwise identical paradigms. Furthermore, in a recent study that will be discussed in more detail later in this paper, Vosgerau (2010) found no evidence of an optimistic desirability bias in likelihood judgments about an outcome dependent on a series of coin flips. He argued that when people have a stake in an outcome (whether positive or negative), this may increase the estimated likelihood of the outcome, but positive outcomes are not given higher likelihood estimates than negative outcomes. In fact, in his critical study, people appeared to be pessimistic rather than optimistic in response to outcome desirability (vs. undesirability). In summary, there is an important dearth of evidence that outcome desirability does what perhaps many researchers assume: boost the perceived likelihood of an outcome.

The potential role of information selection in desirability biases

A notable characteristic of the studies cited above regarding the desirability bias, is that information-selection processes were not at play in those studies. Participants were simply given the information they needed to make their likelihood assessments. As noted by Krizan and Windschitl (2007), there are a variety of cognitive processes that might mediate effects of manipulated desirability on likelihood judgments (but most of these processes have received little research attention). One of those potential mediators is information selection or evidence search. In the studies we present here, participants were given a heterogeneous set of information from which they could make selections before providing likelihood judgments. We have already mentioned research on selective exposure that illustrates the substantial flexibility that people exhibit in their choosing of, and processing of, new information (Hart et al., 2009; Jonas et al., 2006). Another example is from research on information distortion, which demonstrates that decision making processes can involve the biased perception of new information—in the direction favoring a person's leading decision option (Carlson & Russo, 2001; DeKay, Patiño-Echeverri, & Fischbeck, 2009; Russo, Carlson, Meloy, & Yong, 2008; Russo, Medvec, & Meloy, 1996).

In short, we believed that, relative to the approach used in most previous studies on the desirability bias, our approach of presenting participants with a heterogeneous set of information—from which pieces of information needed to be selected and processed—might be more conducive to observing motivated effects tied to outcome desirability. Nevertheless, prior to conducting our research, there were three plausible possibilities regarding the results. First, with a moment of truth in sight and with an opportunity to deliberately select new information to read, people might be careful to select a balanced subset of information—i.e., information that both supported and challenged their preferred outcome (Armor & Sackett, 2006; Gilovich et al., 1993; Tyler & Rosier, 2009). Second, with a moment of truth in sight, they might become especially concerned about protecting themselves from disappointment and therefore become vigilant and select information that challenged their desired outcome or supported an undesired outcome (e.g., Shepperd et al., 1996; Sweeny & Krizan, *in press*; Weber, 1994). Third, despite the moment of truth, people

might prefer to read information that favored the desired outcome. This possibility is consistent with a bird's-eye view of the literatures on motivated reasoning and optimism. Although all three of the possible result patterns were plausible prior to our experiments, our expectations leaned toward the third possibility. We also expected that an information-selection bias would fuel a tendency to give higher likelihood judgments to outcomes that were desirable rather than not. That is, participants in our studies would produce clear evidence of what some researchers have aptly called “the illusory wishful thinking effect” (see Bar-Hillel & Budescu, 1995).

Preview of paradigm and studies

We used an information-buffet paradigm inspired by post-choice selective-exposure studies (Hart et al., 2009; Jonas et al., 2006). Our paradigm had the following basic parts—with modifications unique to each experiment. First, participants saw a pair of artworks and were told that a nationwide sample of college students had already rated each of the artworks. Participants were made aware that they would soon need to estimate the likelihood that one of the artworks was the more preferred artwork in the pair (as rated by the nationwide sample). Depending on the experiment, participants either did or did not need to make a prediction about the artwork, and a monetary manipulation was used to influence how desirable or undesirable it would be for one of the artworks to be the more preferred artwork. Next, participants were given an opportunity to select new information. Specifically, participants saw an information buffet that contained previews of eight comments about individual artworks. Each preview allowed people to discern the evaluative tone of the full comment (e.g., “Mountain Photo is a well detailed photo”). Of the eight comments, there were always two positive and two negative about each of the two artworks. Participants selected the comments they wanted to read, and they provided their first likelihood estimate. Finally, they read the full comments they selected from the buffet before providing another likelihood estimate.

The first study in this paper is important in its own right, but it also serves as a preliminary study that sets the stage for Studies 2 and 3, which directly test the two main research questions. Study 1 (and Follow-Up 1.1 and 1.2) did not involve a direct desirability manipulation. Instead, it measured people's information selections after they made a prediction and before they gave a confidence estimate—which has not been tested in previous published studies. In Study 2, we manipulated whether people stated a prediction and whether a particular outcome was made to be especially desirable. Study 3 (and Follow-Up 3.1) tested whether desirability biases detected in Study 2 were truly due to the increased desirability of an outcome or whether the bias was due to the fact that participants had high stakes in one outcome but not the other (Vosgerau, 2010).

Study 1

Study 1 tested the following hypothesis: After making a prediction between two neutral outcomes, people would tend to select new information that supports rather than contradicts their predicted outcome. The study also tested how this selection bias relates to subsequent confidence about the prediction. Presumably, a predicted outcome becomes somewhat more desirable after it has been predicted, because people like to be right in their predictions. This enhanced desirability is only one of several reasons people might tend to select new information that supported rather than contradicted their predicted outcome (additional reasons are discussed later; see also Scherer, Windschitl, & Smith, 2013).

Hence, Study 1 does not entirely isolate the role of desirability in fueling a selection bias, which we do in Studies 2 and 3.

However, Study 1 is a logical place to start our sequence of studies for three reasons. First, it nicely sets up our paradigm for readers, making the interpretation of later studies with more factors much easier. Second, although desirability was not directly manipulated in Study 1, desire for an outcome (precipitated by the act of explicitly predicting that outcome) is potentially quite relevant in the study and in everyday contexts where people make predictions and subsequently gain access to additional information. Third, we know of no published studies that have both tested for post-prediction selection biases and for a relation between such biases and subsequent confidence.

Readers familiar with research on overconfidence might be surprised by this last claim, so it requires further explication. Overconfidence can be conceptualized and measured in a number of ways (Klayman, Soll, González-Vallejo, & Barlas, 1999; Larrick, Burson, & Soll, 2007; Metcalfe, 1998; Moore & Healy, 2008; Ronis & Yates, 1987), and the classic paradigm involves soliciting confidence estimates from people about their predictions or answers to general knowledge questions (e.g., Dougherty, 2001; Dunning, Griffin, Milojkovic, & Ross, 1990; Lichtenstein, Fischhoff, & Phillips, 1982; McKenzie, 1997; Sieck, Merkle, & Van Zandt, 2007). A common explanation for overconfidence in this paradigm is that people are prone to confirmatory processes after they have given their prediction/answer, but before they state their confidence (Allwood & Johansson, 2004; Griffin & Brenner, 2004; Hoch, 1985; Klayman & Ha, 1987; Koriat, Lichtenstein, & Fischhoff, 1980; Metcalfe, 1998). Specifically, this view suggests that people are more likely to seek, attend to, or rely on evidence that supports rather than contradicts their prediction, and this bias leads to an inflation of confidence estimates. A study by Koriat et al. (1980) is often cited as support for this view. In their second study, participants became better calibrated if they were asked to list a reason why their prediction might be wrong before they made their probability estimate (see also, Hirt & Markman, 1995; Hirt & Sherman, 1985; Hoch, 1985; Lord, Lepper, & Preston, 1984). However, as noted by Griffin and Brenner (2004, page 186), the fact that considering-the-opposite interventions can reduce overconfidence does not necessarily mean that confirmatory search or evaluation processes were initially present and the main cause of overconfidence. Nor do such studies demonstrate that, when people are explicitly offered supporting and conflicting information about their prediction, they will select the former more than the latter. In Study 1 we remedied this gap in the literature by directly measuring potential bias in information selection after a prediction, and we assessed the relation of this bias to confidence estimates.

Method

Overview

Participants made a prediction about which artwork in a pair was more preferred in a nationwide sample. They then had an opportunity to select what information from a buffet they would like to read more about; the buffet contained previews/titles of pro and con comments (ostensibly written by other college students) about the artworks. Participants also made two confidence judgments: one following their selections from the buffet and one after they had read the full comments that they had selected. This cycle was repeated for four pairs of artwork.

Participants and design

The participants were 53 students from elementary psychology courses at the University of Iowa. Other than counterbalancing factors that were manipulated between participants (and did not

significantly influence the dependent variables), there were no manipulations in this study.

Procedure

Participants, at individual computers, began by reading onscreen instructions indicating that they were being tested for their ability to make accurate predictions about college students' aesthetic preferences. They were presented with a pair of artworks (see the below section entitled *Artwork Pairs* for details). Participants made a prediction about which artwork was more preferred by college students in a nationwide sample (by clicking on an artwork label). Next participants saw an *information buffet* containing eight titles of comments relevant to the artwork pair (see the below section entitled *Information Buffets* for details). Participants were told that the comments were written by University of Iowa students and that they should click on between three and seven titles to indicate which comments they would like to read later.

After selecting the titles, but before reading the full comments, participants were asked a confidence question. Specifically, they were reminded of their prediction for the artwork pair and were asked "What do you think is the probability that your prediction is correct?" They responded by placing a marker along a visual-analogue scale labeled from 0% to 100%. The marker also displayed its exact numeric location (e.g., it displayed "67%" if placed at that location). Instructions reminded participants that a response of 100% meant they were absolutely certain, a response of 50% meant they believed their chances of being right and wrong were equivalent, and a response below 50% meant they believed they should have selected the other artwork.

This entire sequence (prediction, information buffet selections, confidence estimation) was repeated for four artwork pairs.

Upon completion of this procedure for the four artwork pairs, participants were asked to indicate which artwork in each pair they personally preferred. Then they read the full comments of the subset of titles they had selected earlier. After reading each subset of comments, they were asked another confidence question. Specifically, they were reminded of their prediction for the artwork pair and were asked "Now that you've read the full comments, what do you think is the probability that your prediction is correct?" Their responses went on the same visual-analogue scale mentioned above.

Artwork pairs

There were four pairs of artworks in this study: a pair of abstract paintings, landscape photographs, abstract sculptures, and songs. The artworks were novel to the participants and selected from the internet. With the help of informal pilot testing, we tried to select artworks such that the two members of a pair were roughly equal in their appeal to college students. The left–right spatial ordering of art in each pair (for the visual art: paintings, photographs, and sculpture) was counterbalanced. For the pair of songs, a 100 s clip from each song was played in a counterbalanced order. The order in which the four pairs of artwork were presented to participants was also counterbalanced. The counterbalancing did not have a significant impact on any key results, and therefore is not discussed further.

Information buffets

Each buffet (one corresponding to each artwork pair) displayed titles of eight comments that were ostensibly written by local University students but were actually written by our lab group. Of the eight titles, two expressed positive evaluations of one of the artworks, two expressed negative evaluations of that artwork, two expressed positive evaluations of the other artwork, and two expressed negative evaluations of that artwork. The titles that were visible on the buffet were short sentences that clearly

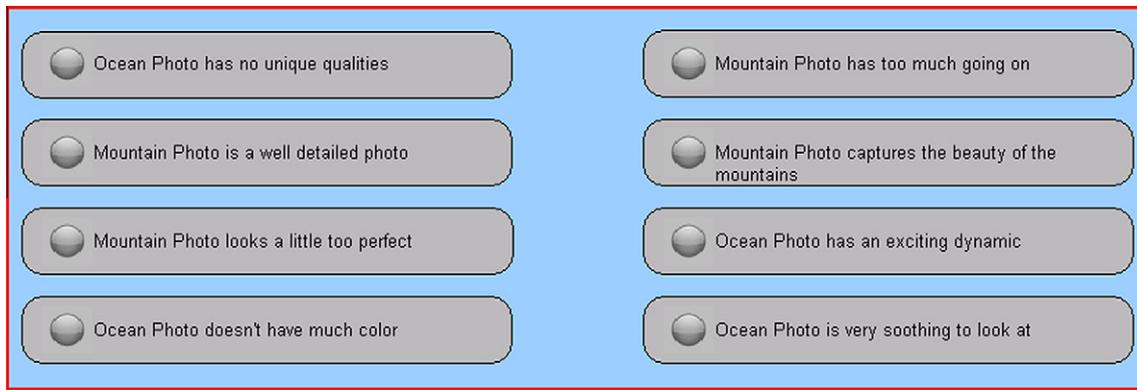


Fig. 1. This is a screenshot of the information buffet for the photography artwork. It contains titles of comments ostensibly written by local students. The titles appear in a random order for each participant. See Appendix for the full comments associated with these titles.

conveyed the valence of the full comment (e.g., “Mountain Photo is a well detailed photo”; “Mountain Photo looks a little too perfect”). The spatial ordering of the titles was randomized separately for each participant. The full-length comments were approximately 5–6 sentences in length. As an example, Fig. 1 displays a screenshot of the buffet for the photography artwork, and the Appendix displays the full comments associated with the titles on that buffet.

Results

Preliminary analyses

Our intention was to use artwork such that participants would be roughly equally split in terms of predictions and personal preferences about the two pieces in each pair. We largely succeeded with three of the four pairs of artworks. For the photograph, sculpture, and song pairs, the percentage of participants predicting “Piece A” rather than “Piece B” (which are arbitrary labels) was near 50%—specifically 58.5%, 54.7%, and 47.2%, respectively. Personal-preference selections of Piece A were also near 50%—specifically 45.3%, 56.6%, and 35.8%, respectively for the three pairs. The exception was the paintings pair, for which only 7.5% predicted Piece A and only 18.9% preferred Piece A. As revealed shortly, the level of this imbalance does not seem to be related to the main findings. Also, not surprisingly, participants had somewhat of a tendency to predict the artwork that they personally preferred (this occurred for 72.6% of the cases overall).

Selection bias

Participants selected an average of 3.3 comments from each buffet. Table 1 displays the rates at which different types of information were selected overall and for each buffet. The main issue in

Study 1 was whether participants would tend select information that supported their prediction. Therefore, the metric for this selection bias is the percentage of selected comments that were positive about the predicted artwork or negative about the rejected artwork. If this percentage is greater than 50%, this means people were biased toward selecting supportive information (which could also be known as congenial or confirmatory information). The critical finding is that there was a robust overall bias; the average percent of supportive comments selected was 78.1%, which was significantly greater than 50%— $t(52) = 11.70, p < .001$. As can be seen in Table 1, this bias was significant for each individual buffet and did not differ much in magnitude across the four buffets. In short, for all buffets, people were biased toward selecting information that supported their prediction.

Confidence estimates and relations with selection bias

Unlike typical overconfidence studies, the present research was not designed to directly measure calibration and overconfidence—we required participants to make only four predictions, and we did not have the national-sample, art-preference data that would be required for determining the precise accuracy of those predictions. Instead our empirical focus was on the potential influence of a biased selection of information on confidence.

Recall that participants were asked about their confidence on two occasions for each prediction, once immediately after making selections from the buffet but before reading the subset of full comments (T1 confidence) and once after having read the subset of full comments that were selected (T2 confidence). Mean confidence estimates per artwork pair are displayed in Table 2. The overall mean for T1 confidence was 65.5% ($SD = 9.0$), which was significantly above 50%, $t(52) = 12.5, p < .001$. The overall mean

Table 1
Title selection percentages from buffets in Study 1.

Artwork/buffet type	A	B	C	D	Selection bias
Photograph	42.1	10.0	15.1	32.7	74.8***
Painting	44.8	5.7	13.1	36.4	81.2***
Sculpture	39.3	7.6	15.1	38.1	77.4***
Song	48.1	9.8	11.3	30.9	78.9***
Average	43.6	8.3	13.7	34.5	78.1***

Note: A = % of selected titles that were positive toward the predicted artwork, B = % that were negative toward the predicted artwork, C = % that were positive toward the non-predicted artwork, D = % that were negative toward the non-predicted artwork. The asterisks indicate that the values for the selection bias (A + D) were significantly different from 50%.
*** $p < .001$.

Table 2
Confidence/likelihood estimates for each artwork prediction in Study 1.

Artwork type	T1 estimate		T2 estimate		Change in estimates	
	M	SD	M	SD	M	SD
Photograph	65.5	11.0	70.1	15.5	4.6 ^a	17.4
Painting	68.2	14.1	70.9	16.7	2.76	17.3
Sculpture	62.0	14.0	73.3	18.1	11.3***	17.2
Song	66.4	14.2	71.8	15.4	5.4*	17.8
Average	65.5	9.0	71.5	11.0	6.0***	9.4

Note: Symbols in the change column refer to whether the changes were significantly different from 0.

^a $p < .10$.
* $p < .05$.
*** $p < .001$.

Table 3
Correlations between selection biases and confidence in Study 1.

Artwork type	<i>r</i> with T1 estimate	<i>r</i> with T2 estimate	<i>r</i> with Change in estimates
Photograph	.21	.03	–.10
Painting	.02	.33*	.30*
Sculpture	–.15	.26 ^a	.40***
Song	–.12	.30*	.35*

Note: All values are bivariate correlations. For example, the upper left value (.21) is the correlation between the selection bias for the photograph buffet and the T1 confidence estimate for the predicted photograph.

^a $p < .10$.

* $p < .05$.

** $p < .01$.

*** $p < .001$.

for T2 confidence was 71.5% ($SD = 11.0$), which constitutes a significant increase from T1, $t(52) = 4.66$, $p < .001$.

More important is the question of whether biases in information selection were related to T1 confidence, T2 confidence, and change in confidence from T1 to T2. It was theoretically possible that strong selection biases could be related to high T1 confidence because—even though participants had not yet read the full comments—they had already focused on a biased set of titles. The results, however, revealed no significant correlations between the degree of selection bias for a given art type and confidence about the relevant prediction at T1. See Table 3 for these correlations.

The results involving T2 confidence were different. For 3 of the 4 art types, there was a significant correlation between the degree of selection bias and confidence at T2 (see Table 3). Given those results, it is not surprising that the change in confidence from T1 to T2 (measured as a difference score) was significant for the same 3 of 4 art types. It appears that, to the extent that participants selected a biased set of titles off the buffet, reading the full-comment versions of the biased set led to increased confidence.

Discussion

Study 1 provides a novel demonstration of the following: After making a prediction between two neutral outcomes, people select new information that supports rather than contradicts their prediction, and this biased selection of information is significantly related to subsequent confidence about the prediction. Although we do not have accuracy data to draw the typical conclusions about overconfidence, this finding from Study 1 clearly has relevance to the overconfidence literature. Whereas various theories have posited or implied that people seek a disproportionate amount of prediction-consistent evidence and that this leads to overconfidence (see review by Griffin and Brenner (2004)), Study 1 was a far more direct in testing this notion than other studies have been—including the studies by Koriat et al. (1980), which are sometimes falsely presumed to have tested this notion (see earlier discussion).

Outcome desirability is one of the potential reasons why participants in Study 1 selected a biased set of information. Presumably, after having committed themselves to a prediction, people then desire that the outcome they predicted is the true outcome. Consequently, information supporting that outcome becomes appealing to read (e.g., Krizan & Windschitl, 2007; Scherer et al., 2013), and/or information conflicting with that outcome becomes dissonance provoking (Hart et al., 2009; Jonas et al., 2006; Kunda, 1990). However, alternatives to these motivated explanations are also tenable. For example, perhaps the

people tend to apply a nonmotivated positive-test strategy for testing any focal hypothesis—in this case, the one they predicted (Klayman & Ha, 1987; Metcalfe, 1998; Sieck et al., 2007; Snyder & Swann, 1978).² A third-variable explanation—involving pre-existing preferences—is also potentially relevant (see Scherer et al., 2013). For a given pair of artwork, some people would prefer Piece A rather than Piece B. This preference would impact their prediction and could also influence the information they select from the buffet. To isolate the role of desirability, we need a direct desirability manipulation, which is what we used in Studies 2 and 3. Before turning to Study 2, however, we very briefly present two follow-ups to Study 1. These follow-ups add modest but notable information about the boundary conditions for the findings (Follow-Up 1.1) and causal impacts (Follow-Up 1.2).

Follow-ups to Study 1

Follow-Up 1.1

All participants ($N = 15$) were informed, just prior to making buffet selections, that they would be able to revise their initial prediction based on what they learned from the student comments. Otherwise, all procedures were the same as for Study 1. The selection biases were again robust ($M = 66.1\%$, $SD = 18.9$, $t(14) = 3.3$, $p < .01$). This shows that even when participants have an opportunity to use buffet information to improve their predictions, they showed a tendency to select information that was supportive of their initial prediction.

Follow-Up 1.2

We wanted to verify that the selection biases of the type observed in Study 1 would indeed have causal impacts on confidence. The correlations between selection bias and confidence measures provided strong initial support for this idea (for at least 3 of the 4 artworks; see Table 3). To gather more direct support, we had 26 participants make 3 predictions (dropping the songs pair from Study 1), but instead of making buffet selections after a prediction, they were given researcher-determined buffet comments to read. Sometimes they received a set of 3 supportive comments and 1 conflicting comment (i.e., roughly equivalent to the selection biases observed in Study 1), sometimes they received 2 and 2 respectively (no bias), and sometimes they received 1 and 3, respectively (a reversed bias). After reading the buffet items for a given artwork pair, participants indicated their confidence in their prediction. The resulting average confidence estimates for the three conditions—selection bias, no bias, and reverse bias—were 78.5% ($SD = 10.3$), 67.0% ($SD = 12.5$), and 53.0% ($SD = 17.7$) respectively. Each mean was significantly different from the others in this within-subjects design ($ps < .001$). In short, this follow-up verifies that the type of selection biases seen in Study 1, where the average bias was 78.1%, would have robust causal consequences for confidence.

² A study by Radzevick and Moore (2008) shows the general potential for focalistic pursuit of information and its impact on likelihood judgment. When participants were allowed to selected statistical information about two unnamed football teams, participants tended to view more information about the team that they were arbitrarily told was their team. The selections of information were related to subsequent likelihood judgments regarding which team was the winning team. While the results show that selective (focalistic) information searches can have consequences for likelihood judgment, we note that participants did not make any initial prediction, and they had no idea, prior to picking a given piece of information, whether it would bode well or poorly for a team. Consequently, the bias observed is not necessarily tied to motivated pursuit of information supporting a predicted or desired outcome.

Study 2

The primary goal for Study 2 was to isolate and manipulate outcome desirability in order to test its effect on information selection and likelihood judgment/confidence. To isolate the effects of desirability, we used a straightforward monetary manipulation added to the same basic paradigm used for Study 1. This was crossed with a manipulation of whether or not people made their own predictions about the artworks. Hence, there were four cells for this design. When an artwork pair appeared, either the participant made a prediction or the computer randomly “selected” an artwork from the pair (we’ll refer to the predicted or selected artwork as the *target*). Then, depending on the monetary manipulation, the participant was told—just prior to the information buffet—that they would receive \$0 or \$10 if the target artwork was indeed the artwork preferred by most college students. The measures were essentially the same as those in Study 1 (i.e., buffet selections, T1 and T2 confidence questions).

We expected that having people make a prediction (vs. no prediction) and promising \$10 (vs. \$0) if the target was true would increase the selection bias. In other words, we expected two main effects and no interaction. Consistent with the earlier discussion, there are at least three potential reasons to expect the first main effect—i.e., the increased selection bias when making a prediction. These include the influence of desirability, a positive-test strategy, and a third-variable explanation. The more important expectation is about the main effect for the monetary manipulation. We expected that regardless of whether people were in a prediction or no-prediction condition, when they learned that they would earn \$10 if the target was true, this would cause them to desire that outcome. This in turn would influence their preference for information—namely increasing a preference for information suggesting that the target is true.

In addition to these expectations regarding the selection biases, we had compatible expectations regarding the confidence/likelihood judgments about the target—two main effects and no interaction. The first main effect of prediction-vs.-no-prediction is not particularly interesting; when people are allowed to make their own prediction, they will always be giving a likelihood estimate about the outcome that seemed more likely to them (because their prediction determines the target of the confidence question). The more critical expectation is that confidence would be higher in the \$10 condition than in the \$0 condition. This would constitute a rare example of a desirability bias detected on likelihood judgments within an experimental paradigm (for discussions see Bar-Hillel & Budescu, 1995; Krizan & Windschitl, 2007; Windschitl et al., 2010). Whereas previous studies have typically found no desirability effect on likelihood judgment, we believed the information-buffet paradigm is a particularly fertile context for desirability biases. The buffet provides people with a heterogeneous mix of new information, and people can choose to expose themselves to, or focus on, a subset of that information, which could thereby bias their reasoning about likelihood judgments.

Method

Participants and design

The participants were 102 students from elementary psychology courses at the University of Iowa. The design was a 2(prediction: yes or no) \times 2(desirability: \$0 or \$10) mixed factorial, with the second factor manipulated within participant.

Procedure

The procedures and materials were similar to those of Study 1, with the exception of modifications that were introduced to

accommodate the two manipulations. Another difference was that only two artwork pairs were used: landscape photographs and sculptures.

At the beginning of the study, all participants were informed that they might win money on some trials of the experiment. After other instructions, participants saw their first artwork pair. Participants in the prediction condition made a prediction about which artwork was preferred in a nationwide sample, just like Study 1. Participants in the no-prediction condition viewed the artwork pair for a minimum of 5 s, with knowledge that they would later make a confidence judgment about an as-yet-undetermined artwork from the pair. Immediately after viewing the artwork pair, the computer then appeared to randomly select an artwork from the pair. In short, some participants predicted an artwork and some participants witnessed the random selection of an artwork. Either way, we will refer to the predicted/selected artwork as the *target*.

Next, the potential monetary award was specified. The procedures for the prediction and no-prediction groups were the same. Namely, participants were told to click a button so that the computer could “randomly select a dollar value.” Unbeknownst to the participants, the value that was then displayed was always \$0 or \$10. If the value was \$10 (\$0), the computer then stated that if the target artwork was truly the one that was preferred in a nationwide sample, then the participant would receive \$10 (\$0)—otherwise nothing. In short, the monetary manipulation was intended to cause people in the \$10 condition to strongly desire that the target was, in fact, the correct artwork. This manipulation came after the prediction/selection of the target so that it could not influence the actual prediction people made (in the prediction condition).

Next, all participants made 3–7 selections from the buffet, just as in Study 1. The instructions provided some reminders “Again, if [target] was preferred by more students nationwide, you will receive [\$ amount]. Soon you will be asked to indicate your confidence that [target] was preferred nationwide...”

Then all participants provided a T1 confidence estimate. The wording of the confidence question was slightly different from that used in Study 1—e.g., “How likely do you think it is that, for this pair of photographs, [target] had the higher preference rating in the nationwide sample?” However, the 0–100% scale was the same.

After doing all of the above steps for one artwork pair and then the other, all participants read all the comments they had selected, provided T2 confidence estimates, and indicated their personal preference for each artwork pair. Counterbalancing manipulations ensured that the two artwork pairs served in the \$10 and \$0 conditions equally often, that the two artworks appeared in the first and second positions equally often, and that the \$10 and \$0 conditions were in the first and second positions equally often.

Results

Preliminary analyses

As in Study 1, our sample did not show overwhelming tendencies in their predictions or in their personal preferences for a particular photograph or sculpture. Within the group of participants making predictions, the percentage of participants predicting “Piece A” rather than “Piece B” (arbitrary labels) was 65.3% for the photographs and 51.0% for the sculptures. Across all participants, the percentage indicating that “Piece A” was personally more preferred was 65.7% for photographs and 42.2% for sculptures. Also, similar to Study 1, participants in the prediction condition tended to predict the artwork that they personally preferred (79.6% of the cases).

Table 4
Title selection percentages from buffets in Study 2.

Condition	A	B	C	D	Selection bias
Prediction/\$10	40.1	13.0	14.7	32.2	72.4***
Prediction/\$0	40.9	12.1	18.3	28.6	69.6***
No prediction/\$10	41.8	16.0	21.3	20.9	62.6***
No prediction/\$0	31.4	20.7	27.3	20.6	51.9

Note: A = % of selected titles that were positive toward the target artwork, B = % that were negative toward the target artwork, C = % that were positive toward the non-target artwork, D = % that were negative toward the non-target artwork. The asterisks indicate that the values for the selection bias (A + D) were significantly different from 50%.

*** $p < .001$.

Selection bias

Participants selected an average of 3.9 comments per buffet, and this count did not significantly differ as a function of any manipulations or interactions ($ps > .20$).

Table 4 displays the rates at which different types of information were selected, as a function of the four cells in our Prediction/No-Prediction \times \$0/\$10 design. The far-right column displays the magnitude of the selection bias for each cell—where selection bias is defined as the percentage of selected comments that were favorable about the target artwork or unfavorable about the non-target artwork. There was a significant selection bias (i.e., greater than 50%) for three of the four cells.

Our main concern was whether and how the magnitude of the selection bias differed as a function of the manipulations. We submitted selection bias scores to a 2(Prediction/No-Prediction) \times 2(\$0/\$10) mixed ANOVA. As expected, the interaction was not significant, $F(1,100) = 1.57$, $p > .10$, but the prediction and desirability main effects were significant, $F(1,100) = 13.4$, $p < .001$, and $F(1,100) = 4.5$, $p < .05$, respectively. We discuss these two main effects (with associated simple effects) in turn below.

The first main effect was of the following, expected pattern: Participants exhibited a greater selection bias in favor of the target when they themselves had predicted the target than when the target had been randomly selected by the computer. Importantly, a simple-effect test shows that the effect of the prediction/no-prediction factor was significant among participants in the \$0 condition—where outcome desirability was not an issue (again, see Table 4). In fact, in the \$0 condition when participants did not make their own prediction, there was no significant selection bias. This supports the position that the selection bias is not merely a tendency to select confirmatory information for whatever outcome had been pre-specified as a target.

The second main effect was of the following, expected pattern: Participants exhibited a greater selection bias in the \$10 condition than the \$0 condition. Although the Prediction \times Desire interaction was not significant, paired t -tests suggest that this desirability effect (\$0 vs. \$10) was primarily driven by participants in the no-prediction condition, $t(52) = 2.17$, $p < .05$, rather than the prediction condition, $t(48) = 0.72$, $p > .10$. The former simple effect is important and shows that even when people have not made a prediction, simply desiring an outcome can cause people to select a biased set of information to inform their later likelihood judgments.

Confidence estimates and relations with selection bias

Table 5 displays mean T1 and T2 confidence estimates. Recall that T1 estimates were made after the buffet (where they read and selected from evaluative comment titles) but before participants read the selected subset of full-length comments. We submitted confidence estimates to a 2(Prediction/No-Prediction) \times 2(\$0/\$10) \times 2(T1/T2) mixed ANOVA. As expected, the main effects for the prediction and desirability factors were significant,

Table 5
Confidence/likelihood estimates about targets in Study 2.

Condition	T1 estimate		T2 estimate		Change in estimates	
	M	SD	M	SD	M	SD
Prediction/\$10	76.4	12.8	74.7	15.0	-1.7	11.5
Prediction /\$0	64.9	15.6	67.9	15.5	3.0	14.2
No prediction/\$10	61.9	21.8	62.7	21.2	0.8	16.8
No prediction/\$0	53.3	22.6	55.6	22.6	2.3	17.4

$F(1,100) = 23.09$, $p < .001$ and $F(1,100) = 14.71$, $p < .001$, respectively. The T1/T2 effect was not significant, nor were any interactions ($ps > .10$).

The main effect for the prediction factor is not surprising; it reflects that participants gave higher likelihood estimates about targets they had predicted than about targets that the computer randomly selected. More importantly, the main effect for desirability reveals more optimism about the target when it was highly desired (\$10 condition rather than \$0 condition). Simple effect tests reveal that this desirability effect was significant at both T1 ($p < .001$) and T2 ($p < .01$), and both in the prediction condition ($p < .001$) and the no-prediction condition ($p < .05$).

Also important is how the biases in information selection were related to T1 confidence, T2 confidence, and change in confidence. Table 6 displays the relevant correlations. Consistent with Study 1, the selection bias was not significantly related to initial T1 confidence among participants in the prediction condition (for both the \$0 and \$10), but it was related to T2 confidence and change in confidence. This again supports the idea that reading a self-selected and biased subset of comments can influence confidence about one's prediction.

A different pattern was evident in the no-prediction condition. The magnitude of the selection bias was significantly related to both T1 and T2 confidence, but it was not related to change in confidence. The correlation with T1 confidence suggests that even though participants had not yet read the full version of the comments they had selected, merely perusing and selecting a biased set of titles influenced their level of T1 confidence. An alternative (or co-contributing) explanation is that people who were highly confident about the target being correct might have found comments suggesting a contradictory view unworthy of much attention.

The fact that the magnitude of the selection bias was related to T1 confidence among one group of participants but not the other was not anticipated and was initially puzzling. We did, however, develop a plausible explanation in hindsight. Participants in the prediction group had worked through the process of making a prediction before seeing the buffet and T1 confidence measures. Therefore, when making a T1 confidence estimate, the influence of the short buffet titles might have been drowned out by

Table 6
Correlations between selection biases and confidence in Study 2.

Condition	r with T1 estimate	r with T2 estimate	r with change in estimates
Prediction/\$10	.12	.35*	.32*
Prediction/\$0	.17	.43**	.28 ^a
No prediction/\$10	.37**	.42**	.04
No prediction/\$0	.51***	.37**	-.17

Note: All values are bivariate correlations. For example, the upper left value (.12) is the correlation between the selection bias and T1 confidence among participants in the prediction/\$10 condition.

^a $p < .10$.

* $p < .05$.

** $p < .01$.

*** $p < .001$.

characteristics of the prediction-deliberation process that the participants had completed. By the time they were asked to give a T2 estimate, the prediction-deliberation process was relatively distant, allowing the biased set of full length comments to have a strong influence on T2 confidence. In the no-prediction condition, however, participants did not have to go through a prediction-deliberation process prior to the buffet and T1 confidence measures; they were simply informed which artwork was the target. Participants might have viewed the artwork less actively, and the buffet titles suggested new insight on how to think about the artworks. That is, because the buffet titles were evaluative, a bias to attend to a subset of titles had an immediate impact on T1 confidence. For example, if a participant learned that he or she would gain money if the ocean photograph is the more preferred photograph, then a bias towards reading and selecting supportive buffet titles (e.g., “The ocean photo has an exciting dynamic.” “The mountain photo has too much going on.” See Fig. 1) would tend to shape confidence about the ocean photo even at T1. The bias then carried through to T2 confidence—perhaps reified but not further enhanced by exposure to the full length comments.

For the no-prediction group, we also tested for evidence that the selection bias mediated the relationship between desirability and T1 confidence using procedures relevant to within-subject designs (see Judd, Kenny, & McClelland, 2001). As already reported, the desirability manipulation had significant effects on both the selection bias and T1 confidence estimates, which is considered necessary for evidence of mediation. The centered sums of the selection biases in the \$0 and \$10 conditions ($X_{\$10-\$0}$) and the differences in the respective selection biases ($X_{\$10-\$0}$) were regressed on the differences in T1 confidence ($Y_{\$10-\$0}$). This analysis produced evidence consistent with mediation—the differences in the selection biases were significant predictors of differences in T1 confidence (standardized coefficient = .35, $t(50) = 2.61$, $p < .05$). The intercept from the resulting model, which represents the effect of desirability beyond that carried by the selection bias, was not significant, $t(50) = 1.41$, $p = .17$. These results are consistent with our mediational account.

Discussion

Study 2 produced three important findings. First, the selection bias for evidence supporting a target outcome was greater when the target was a participant’s own prediction than when it was randomly selected by the computer. This rules out the notion that participants would tend to select confirmatory information for any specific target outcome. This does not bode well for attributing post-prediction selection biases to a generic cognitive strategy (something akin to a positive-test strategy; see Klayman & Ha, 1987; Snyder & Swann, 1978).

Second, people exhibited a greater target-supportive selection bias when the target outcome was highly desirable (i.e., when told they would receive \$10 if the target outcome was the true outcome). This desirability bias suggests that motivational factors can play a significant role in information searches regarding uncertain outcomes. Although this process has been discussed in previous papers on the desirability bias (see review by Krizan and Windschitl (2007)), this is the first time it has been directly tested and demonstrated. Notably, this effect was significant only in the no-prediction condition. One reason why this effect was non-significant in the prediction condition might be that the desirability associated with the prediction itself (i.e., wanting to be right in one’s prediction) caused some degree of selection bias, and learning that one could get an extra \$10 if the prediction was right didn’t significantly add to the impact of the existing desirability (i.e., the two doses of desirability were not fully additive).

Third, the desirability manipulation also had a significant effect on confidence/likelihood judgments. This constitutes a rare find, as

many studies have tested and failed to find a desirability bias on likelihood judgments within an experimental paradigm (for discussions see Bar-Hillel & Budescu, 1995; Krizan & Windschitl, 2007; Vosgerau, 2010; Windschitl et al., 2010). Whereas previous studies have typically found no desirability effect on likelihood judgment, we believed the information-buffet paradigm would be a particularly fertile context for desirability biases. The buffet provides people with a heterogeneous mix of new information, and people could choose to focus on a subset of that information, which could thereby bias their reasoning about likelihood judgments. Consistent with this notion, the extent to which participants exhibited a selection bias was significantly related to likelihood judgments about the target outcome.

The results of the desirability manipulation in Study 2 take on additional importance in light of recent arguments and evidence provided by Vosgerau (2010). Vosgerau argued that there is a stake-likelihood effect that is perhaps more influential than a desirability bias. According to the stake-likelihood hypothesis, when people have a positive or negative stake in the outcome of an event, they might misattribute arousal regarding the stake itself to the likelihood of the outcome. Therefore, relative to a case in which nothing is at stake, people might inflate their likelihood judgments about outcomes that would be highly desirable or undesirable. According to this view, empirical studies that compare likelihood estimates for undesirable vs. desirable outcomes are bound to produce null effects, since the stake-likelihood effect would work the same under undesirable and desirable stakes. This view would also suggest that any significant effects in studies comparing likelihood estimates for neutral vs. desirable outcomes might be explained by a stake-likelihood process rather than a desirability bias. This leads to why the findings from Study 2 are important. In Study 2, we found that not only did our desirability manipulation influence likelihood estimates, it also influenced information search. The stake-likelihood explanation does predict that likelihood estimates would be influenced (via arousal misattribution), but it does not predict the influence on information search. Therefore, the desirability-bias account appears to be a better explanation for the set of results for Study 2.

Study 3

Despite the conclusion we just discussed regarding the stake-likelihood hypothesis, one could reasonably imagine a more general version of this hypothesis. Namely, one could suppose that whenever stakes are high (whether in a positive or negative direction), people will become inclined to search for evidence supporting the outcome. Another way of framing this is to say that people’s tendency to select confirming information about an outcome will increase whenever the outcome would be especially good or bad. This possibility—while intuitively plausible—is critically different from our characterization of how (and why) desirability might influence information search and likelihood judgments.

To address this possibility, we conducted Study 3. In Study 3, we used the same general paradigm and we focused exclusively on a no-prediction context. We had two goals in mind: (1) to replicate the effects of desirability that were detected in Study 2, and (2) to test the impact of learning that a target outcome would be undesirable (result in a loss of money). One might expect that when people learn that an outcome is undesirable, they would become keenly interested in information suggesting the outcome might happen. However, consistent with the idea that people prefer a more positive orientation, we expected that participants’ interest in information supporting a target would be reduced when the target is undesirable. In other words, we expected participants to select more disconfirming information when the target outcome was undesirable rather than neutral (or desirable).

Method

Participants

The participants were 96 students from elementary psychology courses at the University of Iowa.

Procedure and design

The procedures were similar to those in the no-prediction condition of Study 2 with some key exceptions. First, participants were warned that they could win or lose money during the study. If they won more than they lost, we would pay them the amount, and if they lost more than they won, they would not have to pay anything. Second, all four artwork pairs from Study 1 were used, with a minor alteration to one of the paintings.³ Third, the photographs and sculpture pairs always came first, followed by the paintings and songs. Finally, like in Study 2, each participant experienced a trial in which \$10 would be received if the target artwork was the artwork preferred nationwide, and there was a comparable trial in which \$0 would be received (called the +\$10 and +\$0 trials, respectively). Unlike Study 2, each participant also experienced a trial in which \$8 would be lost if the target artwork was the artwork preferred nationwide, and there was a comparable trial in which \$0 would be “lost” (called the -\$8 and -\$0 trials, respectively). We chose \$8 rather than \$10 for the loss value so that more participants would finish their sessions with at least some money in hand. For half the participants, the +\$10 and +\$0 trials came first (in counterbalanced order), and for half the participants, the -\$8 and -\$0 trials came first (again in counterbalanced order).

In summary, each participant experienced the four types of trials (+10, +0, -\$8, and -\$0). The +0 and -0 trials were logically very similar, but not fully collapsible. Therefore, the most effective way of characterizing the design is that it allowed for two main comparisons: One between the +\$10 and +\$0 trials, and one between the -\$8 and -\$0 trials.

Results

Selection bias

Participants selected an average of 4.1 comments per buffet, and this count did not differ significantly as a function of the desirability manipulation. Table 7 displays the rates at which different types of information were selected. The far-right column displays the magnitude of the selection bias for each cell. There were two critical findings. First, replicating Study 2, the index for the selection bias was significantly greater in the +\$10 cell than in the +\$0 cell, $t(95) = 3.01$, $p < .01$. Second, the index in the -\$8 condition (which was 45.2% and not significantly different from 50%) was significantly smaller than the index in the -\$0 condition, $t(95) = 2.25$, $p < .05$. In other words, participants were more prone to select supportive information about an outcome when it was desirable rather than neutral, and they were less prone to select supportive information about an outcome when it was undesirable rather than neutral.

Confidence/likelihood estimates and relations with selection bias

Table 8 displays mean T1 and T2 confidence estimates. Replicating Study 2, a $2(+\$0/+ \$10) \times 2(T1/T2)$ mixed ANOVA produced a significant effect for desire, $F(1,95) = 7.37$, $p < .01$, and the T1/T2 main effect and the interaction terms were not significant, both

Table 7
Title selection percentages from buffets in Study 3.

Desirability condition	A	B	C	D	Selection bias
+\$10	37.4	17.3	22.7	22.6	60.0***
+\$0	30.3	22.3	29.7	17.7	48.0
-\$8	29.2	23.9	30.8	16.0	45.2 ^a
-\$0	32.6	23.0	22.6	21.7	54.3

Note: A = % of selected titles that were positive toward the target artwork, B = % that were negative toward the target artwork, C = % that were positive toward the non-target artwork, D = % that were negative toward the non-target artwork. The asterisks and superscript indicate results of *t*-tests comparing the values for the selection bias (A + D) to 50%.

^a $p < .10$.
*** $p < .001$.

Table 8
Confidence/likelihood estimates about targets in Study 3.

Desirability condition	T1 estimate		T2 estimate		Change in estimates	
	M	SD	M	SD	M	SD
+\$10	62.5	22.3	62.2	24.2	-0.3	16.4
+\$0	53.9	23.1	54.7	22.0	0.9	20.2
-\$8	48.6	23.1	49.2	22.3	0.6	18.3
-\$0	57.4	21.7	58.2	21.9	0.8	19.3

$F_s < 1$. In short, people tended to give higher likelihood estimates in the +\$10 condition than in the +\$0 condition.

The novel question regarding likelihood judgments was what would happen in the -\$8 and -\$0 conditions. A $2(-\$0/-\$8) \times 2(T1/T2)$ mixed ANOVA produced a significant effect for desire, $F(1,95) = 9.37$, $p < .01$, and the T1/T2 main effect and the interaction terms were not significant, both $F_s < 1$. Critically, the direction of the desirability effect was consistent with a desirability-bias interpretation, rather than a stake-likelihood interpretation. Participants gave lower likelihood estimates in the -\$8 condition than in the -\$0 condition (see Table 8).

Table 9 displays the correlations for determining how information selection was related to confidence. Consistent with the no-prediction condition in Study 2, the magnitude of the selection bias was significantly related to T1 and T2 confidence, and this was true within all four desirability cells. The relationship between selection bias and change in confidence was significant within one of the four cells (and nearing significance for another).

Given the strong correlations between the selection bias and T1 confidence, we conducted the same mediation tests as reported for Study 2 (see Judd et al., 2001). First we tested for evidence that the selection bias mediated the +\$10 vs. +\$0 effect on confidence. Consistent with mediation, the differences in the selection biases were significant predictors of differences in T1 confidence (standardized coefficient = .51, $t(93) = 5.74$, $p < .001$). The intercept from the

Table 9
Correlations between selection biases and confidence in Study 3.

Desirability condition	<i>r</i> with T1 estimate	<i>r</i> with T2 estimate	<i>r</i> with change in estimates
+\$10	.45***	.62***	.32**
+\$0	.45***	.56***	.10
-\$8	.42***	.54***	.13
-\$0	.43***	.58***	.17 ^a

Note: All values are bivariate correlations. For example, the upper left value (.45) is the correlation between the selection bias and T1 confidence in the +\$10 condition.

^a $p < .10$.
* $p < .05$.
** $p < .01$.
*** $p < .001$.

³ We cropped one of the paintings to better equate the appeal of the two paintings. This seems to have been somewhat successful; when participants were asked about their personal preferences among the two paintings, the selection rates were relatively balanced (55.2% vs. 44.8%)—unlike in the imbalance reported in the Preliminary analyses section of Study 1.

resulting model, which represents the effect of desirability beyond that carried by the selection bias, was not significant, $t(93) = 1.37$, $p > .10$. Second, we tested for evidence that the selection bias mediated the $-\$0$ vs. $-\$8$ effect on confidence. Again consistent with mediation, the differences in the selection biases were significant predictors of differences in T1 confidence (standardized coefficient = .53, $t(93) = 6.10$, $p < .001$). The intercept from the resulting model was borderline significant, $t(93) = 1.8$, $p = .07$.

Discussion

Study 3, like Study 2, showed that when the target outcome is made desirable—by virtue of a promise to win money—this fueled a tendency to select evidence supportive of the target, and it inflated likelihood judgments about the target. Study 3 also showed that when the target outcome is made undesirable—by virtue of a threat to lose money—the tendency to select evidence supporting the target decreases, as do the likelihood judgments about the target. Both the desirability-bias hypothesis and the stake-likelihood hypothesis (modified from Vosgerau (2010) as discussed above) would anticipate the first set of results comparing desirable and neutral conditions. However, the second set of results involving the undesirable condition is inconsistent with the stake-likelihood hypothesis and instead supports the desirability-bias hypothesis. Before turning to the General Discussion, we will briefly discuss results from a follow-up to Study 3.

Follow-Up Study 3.1

We wanted to replicate the key finding that people were less prone to seek supportive information and give high likelihood judgments about an undesirable outcome (vs. a desirable outcome). We also wanted to introduce a slight modification to the paradigm to see if it would offer any support for the stake-likelihood hypothesis. In the undesirable (or $-\$8$) condition of Study 3, participants saw the target artwork listed with $-\$8$ and the other with $\$0$. If a participant's sense of what is desirable and undesirable is primarily relative, then the $\$0$ outcome could be viewed as desirable because it is better than the $-\$8$ outcome. The $-\$8$ outcome might essentially lose its salience as a *negative* outcome. With this in mind, we wondered whether results would be more supportive of stake-likelihood hypothesis in a follow-up study that presented artworks in sets of 4, rather than in pairs. Hence, in the $-\$8$ condition of the follow-up, the target outcome was listed with $-\$8$, whereas the other three artworks in the set were each listed with $\$0$. With this situation, perhaps a $-\$8$ outcome would more readily stand out as negative, which is presumably conducive for a stake-likelihood effect. We also had a $\$10$ condition and a $\$0$ condition. In moving from 2 to 4 artworks per set, we also introduced a slight change to the buffets (to contain one positive and one negative comment about each artwork). There were 50 participants, and the design was within-subject.

In large measure, the results led to same overall conclusions as did those from Study 3. Consistent with the desirability-bias hypothesis (but not stake-likelihood), there was a larger selection bias in the $+\$10$ condition ($M = 55.6\%$, $SD = 27.1\%$) than the $-\$8$ condition ($M = 42.8\%$, $SD = 23.6\%$), $t(49) = 2.44$, $p < .05$. The mean in the $\$0$ condition ($M = 44.1\%$, $SD = 23.0\%$) was between those for the $+\$10$ and $-\$8$ conditions, and significantly different from the $+\$10$ condition, $t(49) = 2.61$, $p < .05$. Regarding likelihood judgments, the mean (T1 and T2 combined) was higher in the $+\$10$ condition ($M = 53.7\%$, $SD = 24.1\%$) than the $-\$8$ condition ($M = 38.4\%$, $SD = 23.9\%$), $t(49) = 3.44$, $p < .001$. The mean in the $\$0$ condition ($M = 49.3\%$, $SD = 22.9\%$) was between those for $+\$10$ and $-\$8$ conditions, and significantly different from the $-\$8$ condition,

$t(49) = 2.18$, $p < .05$, but not from the $\$10$ condition, $t(49) = 0.91$, $p = .37$.

In short, we made a modest change to make a $-\$8$ outcome stand out as undesirable rather than easily reframed as merely less desirable. Even under these conditions, the results supported the desirability-bias hypothesis and not the stake-likelihood hypothesis.

General discussion

At the start of this paper, we introduced an example of Alex, who just learned from her financial advisor that she would make more money if two companies merged. In other words, Alex has a newly established desire for a particular outcome. We suggested that a bird's-eye view of the literature on motivated reasoning and desirability bias might lead to quick assumptions about how this desire for an outcome would influence Alex's processing of new information and her expectations. However, with a closer look at the literature, it became clear that the specific possibility that desire itself—unconfounded with other factors such as prior expectations—would affect information search and optimism has not been adequately tested. We noted that there are plausible theoretical challenges to the idea that an optimistic bias in processing would dominate (e.g., with a moment of truth at hand, accountability and bracing concerns might be prominent; Gilovich et al., 1993; Shepperd, Findley-Klein, Kwavnick, Walker, & Perez, 2000; Sweeny, Carroll, & Shepperd, 2006; van Dijk, Zeelenberg, & van der Pligt, 2003). Moreover, most direct tests of the idea that desire would impact likelihood judgments had suggested that outcome desirability would have no effect (e.g., Price & Marquez, 2005; Windschitl et al., 2010; see meta-analysis by Krizan and Windschitl (2007)) or perhaps even a pessimistic effect (Vosgerau, 2010).

To respond to these unresolved issues in the literature, we developed a paradigm with three key features: (1) direct experimental manipulations of outcome desirability, (2) presentations of new heterogeneous information after participants learned about outcome desirability, and (3) measures of both information selection and likelihood judgment. We also developed a variation on this paradigm—in which participants made their own predictions—to test some unresolved issues related to the overconfidence literature (Studies 1 and 2). In the end, our empirical work produced clear evidence of motivated searches and inflated optimism caused by outcome desirability.

More specifically, Study 1 demonstrated that people tend to select information that favors rather than opposes their predictions—even though the predictions concerned entirely novel outcomes. Critically, the magnitude of this selection bias was predictive of the change in participants' confidence from before to after reading the information they selected from the buffet. The causal connection was verified in the second follow-up to Study 1 (Follow-Up 1.2). The first follow-up to Study 1 (Follow-Up 1.1) demonstrated that the selection biases persisted even when people were made aware that they would be able to update their prediction given new information that they read from the buffet.

The findings from Study 1 and the follow-ups are compatible with, but distinct from, existing theory and research on overconfidence. Researchers have previously discussed and theorized about the possibility that people's overconfidence stems from their tendency to focus primarily on reasons why they might be right rather than why they might be wrong in their prediction. However, tests of this notion have been indirect ones. For example, perhaps the most widely cited example comes from Koriat et al. (1980) who tested and found that a consider-the-opposite intervention led to slightly better calibration in confidence estimates, presumably because people were prompted to think of evidence (against their

prediction) that they otherwise would neglect. Our hypothesis was about selections when a heterogeneous set of new information is made available. Study 1 and its follow-ups showed that when people are provided with a heterogeneous set of new information, they will engage in biased selection consistent with their initial prediction, which ultimately will feed into their subsequent confidence.

Whereas Study 1 did not fully isolate the role of desirability, Studies 2 and 3 did by means of a straightforward monetary manipulation. When a target outcome that was randomly specified by the computer was made desirable (because \$10 was promised if it was true), this enhanced the tendency of people to select information that supported it and also increased the judged likelihood of it being true. Study 3 also found that when a target outcome was made undesirable rather than desirable, participants were more likely to select information against the outcome, and their likelihood judgments tended to be lower (see also Follow-Up 3.1). The effects of these manipulations on likelihood judgments were seen immediately after participants made biased selections from the information buffet—even before they read the full-length versions of the comments they had selected. This finding suggests that even biased perusing and selection of short evaluative titles was enough to fuel biased optimism. Mediation analyses provided evidence consistent with the notion that the effect of desire on likelihood judgment was mediated by selection biases.

Results from Studies 2 and 3 also suggest that the post-prediction selection bias is not attributable to a generic cognitive tendency or positive test strategy for any hypothesis. Namely, the participants who were assigned a target hypothesis by the computer and were promised nothing (\$0) did not show any selection bias. Furthermore, the fact that the monetary desirability manipulation produced reliable effects on selection biases in Studies 2 and 3 bodes well for the notion that a related desire—namely a desire to be right about one's own prediction—could be a key contributor to selection biases that occur after predictions (see Blanton, Pelham, DeHart, & Carvalho, 2001; Scherer et al., 2013). It is interesting to note (but open to multiple interpretations) that the effect of the prediction-vs.-no-prediction manipulation on selection bias in Study 2 was much stronger than the effect of the \$10-vs.-\$0 manipulation.

Wishful thinking and stakes likelihood

Our findings are critical for the literature on the desirability bias because there are many studies that attempted, but failed, to find evidence of enhanced likelihood judgments as a function of outcome-desirability manipulations (for review see Krizan & Windschitl, 2007). One set of authors let their titles express the state of affairs: “The Elusive Wishful Thinking Effect” (Bar-Hillel & Budescu, 1995) and “Wishful Thinking in Predicting World Cup Results: Still Elusive” (Bar-Hillel et al., 2008). However, Studies 2, 3, and 3.1 change matters, and wishful thinking now seems less elusive than it did before.

Moreover, the present findings suggest some potential boundary conditions on the recently developed stake-likelihood hypothesis (Vosgerau, 2010). Again, as originally described, Vosgerau's hypothesis posited that arousal from a negative (or positive) stake in an outcome can be misattributed and inflate likelihood judgment. He showed, for example, that judgments regarding probability questions were higher when those questions were printed on vibrant-pink paper rather than grey paper. In one study (Study 3), Vosgerau also directly manipulated desirability of outcomes. In key conditions, participants were told they would either win or lose a valued prize (a shot glass with a university emblem) if four simulated dice rolls yielded exactly two sixes. Likelihood judgments for these conditions (combined) were elevated relative to those from a neutral/control condition, which supports the stake-likelihood hypothesis. As noted earlier, the overall pattern of comparisons

between the desirable and undesirable conditions suggested some pessimism; see Vosgerau's discussion of his Study 3.

All this begs the following question: Why did Vosgerau (2010) third study fail to produce evidence of a desirability bias, whereas our studies produced robust desirability biases? Perhaps part of the difference is that Vosgerau's third study concerned stochastic outcomes (dice rolls). As discussed in their review, Krizan and Windschitl (2007) revealed that results of studies on desirability bias tend to be different for studies involving stochastic outcomes and nonstochastic outcomes (see also Windschitl et al., 2010). Perhaps even more relevant, is the fact that participants in Vosgerau's study, unlike those in our studies, received no new information to select from. That is, biased information selection was not in play. Perhaps if people can look for optimistic information, as they could in the present studies, they do—even if cautiously. Having found new evidence to be optimistic about, likelihood judgments were inflated. These ideas remain speculation at this time. We think it is important to note that the stake-likelihood hypothesis and the desirability-bias hypothesis can both be valid—depending on the context in which they are applied. In the present context, however, the effects of the desirability bias were clearly stronger than any type of stake-likelihood effect.

Conclusion

When people have an opportunity to select new information after having made a prediction or after having discovered they would benefit from a particular outcome, they have an important opportunity. If they have made a prediction, they can use the new information as a basis for a revision of the prediction or at least a dampening of initial optimism about it (when warranted). If the context did not require an initial prediction, they have a blank slate on which to develop a realistic sense of optimism. In our studies, however, participants tended toward checking for optimistic information, and this bias fueled optimism. Decision makers in business, government, military, health, legal, and other fields face similar situations: They make predictions and need to assess the likelihood of outcomes that they might see as desirable or undesirable. They also often have easy access to new information—sometimes in abundant amounts. Even if the buffet of available information is unbiased, the selection of information appears to be a process that is ripe for initial bias.

Appendix A

This appendix shows the titles and full comments listed for the photography artwork. The titles appeared in the buffet (see Fig. 1) and were randomized separately for each participant. Participants read the full comments of only those titles they selected from the buffet.

Title: Ocean Photo is very soothing to look at.

Full Comment: Looking at this picture, you can almost hear the sound of the ocean and the crashing of the waves against the rocks. It looks so peaceful, and it leaves you wishing that you could be standing there watching such a marvelous example of the power of nature. I have never seen the ocean look so powerful and soothing at the same time.

Title: Ocean Photo has an exciting dynamic.

Full Comment: This photo is exciting for me to look at because it seems to be set to motion. It has captured the beauty of a sunset on the water as genuinely as photograph could without actually being there to see it for yourself. The reflection that the sun has on the water must have taken a lot of patience from the photographer as it seems to be done almost at the perfect moment. The

wave crashing on the rock was the focal point for me as it really gives you a sense of what the setting is really like; any closer and I could hear the wave break.

Title: Ocean Photo doesn't have much color.

Full Comment: You would think that for an amazing outdoor setting like this that the picture would naturally just take your breath away. Maybe it's the bland and hazy background, but this picture is actually a little boring to look at. Maybe if the light had been better or the photographer has chosen a different angle, the picture would have turned out better.

Title: Ocean Photo has no unique qualities.

Full Comment: I am not drawn to this picture at all. It is very dark and the lighting to me is not quite right. I also do not think the scene itself is very impressive, sharp pointy rocks are not that appealing to me. I get bored when I look at the picture and I am not sure what I am supposed to be focusing on. For instance the ocean does not really stand out, maybe if the brightness of the ocean was more apparent it would be more attractive.

Title: Mountain Photo captures the beauty of the mountains.

Full Comment: This photo successfully captures the beauty of the mountain range. The coloration of the different elements of land complements themselves well. It looks peaceful and pure. It makes me want to be outside and enjoying all the wonderful sites to see. This photo is without a doubt a magnificent masterpiece.

Title: Mountain Photo is a well detailed photo.

Full Comment: This photo has a lot of detail because it portrays several elements like the water, mountains, and forests. The camera angle is also very clever because it captures everything and shows the reflection of the mountain on the water with the view of some rocks submerged in water. It also shows half of the mountain in shadow and the other half very bright.

Title: Mountain Photo looks a little too perfect.

Full Comment: The picture looks digitally enhanced which makes it just look fake. It gives you a sense of this ideal image of nature. Also in the photo an area of trees are really dark bringing the beautiful colors in this piece down. The trees definitely stand out as a negative spot, so maybe taking a different angle of the scene it would eliminate the dark spot of the trees.

Title: Mountain Photo has too much going on.

Full Comment: Though the picture is no doubt aesthetically appealing, it just does not do it for me. I wish the photographer would have just focused on one of the natural surroundings because it is a little too much to take in at once. The trees, the lake, the mountains, the sky, the reflection. . .any one of these would have made a great photograph. With this photo I have way too much to look at and need a central focal point because right now I am clueless to what that could be. I would also be interested to learn if the photographer enhanced any of the photo's images; it looks a little fake to me.

References

- Allwood, C. M., & Johansson, M. (2004). Actor-observer differences in realism in confidence and frequency judgments. *Acta Psychologica*, *117*, 251–274.
- Armor, D. A., & Sackett, A. M. (2006). Accuracy, error, and bias in predictions for real versus hypothetical events. *Journal of Personality and Social Psychology*, *91*, 583–600.
- Babad, E. (1997). Wishful thinking among voters: Motivational and cognitive influences. *International Journal of Public Opinion Research*, *9*, 105–125.
- Balcetis, E. (2008). Where the motivation resides and self-deception hides: How motivated cognition accomplishes self-deception. *Social and Personality Psychology Compass*, *2*(1), 361–381.
- Balcetis, E., & Dunning, D. (2006). See what you want to see: Motivational influences on visual perception. *Journal of Personality and Social Psychology*, *91*, 612–625.
- Bar-Hillel, M., & Budescu, D. V. (1995). The elusive wishful thinking effect. *Thinking and Reasoning*, *1*, 71–104.
- Bar-Hillel, M., Budescu, D. V., & Amar, M. (2008). Wishful thinking in predicting World Cup results: Still elusive. In J. I. Krueger (Ed.), *Rationality and social responsibility*. Psychology Press.
- Blanton, H., Pelham, B. W., DeHart, T., & Carvallo, M. (2001). Overconfidence as dissonance reduction. *Journal of Experimental Social Psychology*, *37*, 373–385.
- Carlson, K. A., & Russo, J. E. (2001). Biased interpretation of evidence by mock jurors. *Journal of Experimental Psychology: Applied*, *7*, 91–103.
- Clark, J. K., & Wegener, D. T. (2008). Unpacking outcome dependency: Differentiating effects of dependency and outcome desirability on the processing of goal-relevant information. *Journal of Experimental Social Psychology*, *44*, 586–599.
- DeKay, M. L., Patiño-Echeverri, D., & Fischbeck, P. S. (2009). Distortion of probability and outcome information in risky decisions. *Organizational Behavior and Human Decision Processes*, *109*, 79–92.
- Dawson, E., Gilovich, T., & Regan, D. T. (2002). Motivated reasoning and performance on the Wason selection task. *Personality and Social Psychology Bulletin*, *28*, 1379–1387.
- Ditto, P. H., Munro, G. D., Apanovitch, A. M., Scepansky, J. A., & Lockhart, L. K. (2003). Spontaneous skepticism: The interplay of motivation and expectation in responses to favorable and unfavorable medical diagnoses. *Personality and Social Psychology Bulletin*, *29*, 1120–1132.
- Dougherty, M. R. P. (2001). Integration of the ecological and error models of overconfidence using a multiple-trace memory model. *Journal of Experimental Psychology: General*, *130*, 579–599.
- Dunning, D. (2003). The zealous self-affirmer: How and why the self lurks so pervasively behind social judgment. In S. J. Spencer, S. Fein, M. P. Zanna, & J. M. Olson (Eds.), *Motivated social perception: The Ontario symposium* (Vol. 9, pp. 45–72). Mahwah, NJ: Lawrence Erlbaum.
- Dunning, D., Griffin, D. W., Milojkovic, J. D., & Ross, L. (1990). The overconfidence effect in social prediction. *Journal of Personality and Social Psychology*, *58*, 568–581.
- Festinger, L., & Carlsmith, J. M. (1959). Cognitive consequences of forced compliance. *The Journal of Abnormal and Social Psychology*, *58*, 203–210.
- Gilovich, T., Kerr, M., & Medvec, V. H. (1993). Effect of temporal perspective on subjective confidence. *Journal of Personality and Social Psychology*, *64*, 552–560.
- Granberg, D., & Brent, E. (1983). When prophecy bends: The preference – Expectation link in U.S. presidential elections, 1952–1980. *Journal of Personality and Social Psychology*, *45*, 477–491.
- Griffin, D., & Brenner, L. (2004). Perspectives on probability judgment calibration. In D. J. Koehler & N. Harvey (Eds.), *Blackwell handbook of judgment and decision making* (pp. 177–198). Malden, MA, US: Blackwell Publishing.
- Hart, W., Albarracín, D., Eagly, A. H., Brechan, I., Lindberg, M. J., & Merrill, L. (2009). Feeling validated versus being correct: A meta-analysis of selective exposure to information. *Psychological Bulletin*, *135*, 555–588.
- Hirt, E. R., & Markman, K. D. (1995). Multiple explanation: A consider-an-alternative strategy for debiasing judgments. *Journal of Personality and Social Psychology*, *69*, 1069–1086.
- Hirt, E. R., & Sherman, S. J. (1985). The role of prior knowledge in explaining hypothetical events. *Journal of Experimental Social Psychology*, *21*, 519–543.
- Hoch, S. J. (1985). Counterfactual reasoning and accuracy in predicting personal events. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *11*, 719–731.
- Holton, B., & Pyszczynski, T. (1989). Biased information search in the interpersonal domain. *Personality and Social Psychology Bulletin*, *15*, 42–51.
- Irwin, F. W. (1953). Stated expectations as functions of probability and desirability of outcomes. *Journal of Personality*, *21*, 329–335.
- Jonas, E., Schulz-Hardt, S., Fischer, P., & Frey, D. (2006). Biased information seeking after decisions: Cognitive and motivational biases. In D. Frey, H. Mandl, & L. von Rosenstiel (Eds.), *Knowledge and action* (pp. 23–51). Ashland, OH, US: Hogrefe & Huber Publishers.
- Judd, C. M., Kenny, D. A., & McClelland, G. H. (2001). Estimating and testing mediation and moderation in within-subject designs. *Psychological Methods*, *6*, 115–134.
- Kay, A. C., Jimenez, M. T., & Jost, J. T. (2002). Sour grapes, sweet lemons, and the anticipatory rationalization of the status quo. *Personality and Social Psychology Bulletin*, *28*, 1300–1312.
- Klayman, J., & Ha, Y. W. (1987). Confirmation, disconfirmation, and information in hypothesis testing. *Psychological Review*, *94*, 211–228.
- Klayman, J., Soll, J. B., González-Vallejo, C., & Barlas, S. (1999). Overconfidence: It depends on how, what, and whom you ask. *Organizational Behavior and Human Decision Processes*, *79*, 216–247.
- Koriat, A., Lichtenstein, S., & Fischhoff, B. (1980). Reasons for confidence. *Journal of Experimental Psychology: Human Learning and Memory*, *6*, 107–118.
- Krizan, Z., Miller, J. C., & Johar, O. (2010). Wishful thinking in the 2008 U.S. presidential election. *Psychological Science*, *21*, 140–146.
- Krizan, Z., & Windschitl, P. D. (2007). The influence of outcome desirability on optimism. *Psychological Bulletin*, *133*, 95–121.
- Krizan, Z., & Windschitl, P. D. (2009). Wishful thinking about the future: Does desire impact optimism? *Social and Personality Psychology Compass*, *3*, 227–243.
- Kunda, Z. (1990). The case for motivated reasoning. *Psychological Bulletin*, *108*, 480–498.
- Kunda, Z., & Sanitioso, R. (1989). Motivated changes in the self-concept. *Journal of Experimental Social Psychology*, *25*, 272–285.
- Larrick, R. P., Burson, K. A., & Soll, J. B. (2007). Social comparison and confidence: When thinking you're better than average predicts overconfidence (and when it does not). *Organizational Behavior and Human Decision Processes*, *102*, 76–94.

- Lench, H. C., & Ditto, P. H. (2008). Automatic optimism: Biased use of base rate information for positive and negative events. *Journal of Experimental Social Psychology, 44*, 631–639.
- Lichtenstein, S., Fischhoff, B., & Phillips, L. D. (1982). Calibration of probabilities: State of the art to 1980. In D. Kahneman, P. Slovic, & A. Tversky (Eds.), *Judgment under uncertainty: Heuristics and biases* (pp. 306–334). New York: Cambridge University Press.
- Lord, C. G., Lepper, M. R., & Preston, E. (1984). Considering the opposite: A corrective strategy for social judgment. *Journal of Personality and Social Psychology, 47*, 1231–1243.
- Marks, R. W. (1951). The effect of probability, desirability, and “privilege” on the stated expectations of children. *Journal of Personality, 19*, 332–351.
- Massey, C., Simmons, J. P., & Armor, D. A. (2011). Hope over experience: Desirability and the persistence of optimism. *Psychological Science, 22*, 274–281.
- McDonald, H. E., & Hirt, E. R. (1997). When expectancy meets desire: Motivational effects in reconstructive memory. *Journal of Personality and Social Psychology, 72*, 5–21.
- McKenzie, C. R. M. (1997). Underweighting alternatives and overconfidence. *Organizational Behavior and Human Decision Processes, 71*, 141–160.
- Metcalfe, J. (1998). Cognitive optimism: Self-deception or memory-based processing heuristics? *Personality and Social Psychology Review, 2*, 100–110.
- Moore, D. A., & Healy, P. J. (2008). The trouble with overconfidence. *Psychological Review, 115*, 502–517.
- Price, P. C., & Marquez, C. (2005). *Wishful thinking in the prediction of a simple repeatable event: Effects of deterministic versus probabilistic predictions*. Unpublished manuscript. California State University, Fresno.
- Radzevick, J. R., & Moore, D. A. (2008). Myopic biases in competitions. *Organizational Behavior and Human Decision Processes, 107*, 206–218.
- Roese, N. J., & Olson, J. M. (2007). Better, stronger, faster: Self-serving judgment, affect regulation, and the optimal vigilance hypothesis. *Perspectives on Psychological Science, 2*, 124–141.
- Ronis, D. L., & Yates, J. F. (1987). Components of probability judgment accuracy: Individual consistency and effects of subject matter and assessment method. *Organizational Behavior and Human Decision Processes, 40*, 193–218.
- Russo, J. E., Carlson, K. A., Meloy, M. G., & Yong, K. (2008). The goal of consistency as a cause of distortion. *Journal of Experimental Psychology: General, 137*, 456–470.
- Russo, J. E., Medvec, V. H., & Meloy, M. G. (1996). The distortion of information during decisions. *Organizational Behavior and Human Decision Processes, 66*, 102–110.
- Pyszczynski, T., & Greenberg, J. (1987). Toward an integration of cognitive and motivational perspectives on social inference: A biased hypothesis-testing model. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 20, pp. 297–333). San Diego, CA: Academic Press.
- Scherer, A. M., Windschitl, P. D., & Smith, A. R. (2013). Hope to be right: Biased information seeking following arbitrary and informed predictions. *Journal of Experimental Social Psychology, 49*, 106–112.
- Shepperd, J. A., & McNulty, J. K. (2002). The affective consequences of expected and unexpected outcomes. *Psychological Science, 13*, 85–88.
- Shepperd, J. A., Findley-Klein, C., Kwavnick, K. D., Walker, D., & Perez, S. (2000). Bracing for loss. *Journal of Personality and Social Psychology, 78*, 620–634.
- Shepperd, J. A., Ouellette, J. A., & Fernandez, J. K. (1996). Abandoning unrealistic optimism: Performance estimates and the temporal proximity of self-relevant feedback. *Journal of Personality and Social Psychology, 70*, 844–855.
- Sieck, W. R., Merkle, E. C., & Van Zandt, T. (2007). Option fixation: A cognitive contributor to overconfidence. *Organizational Behavior and Human Decision Processes, 103*, 68–83.
- Snyder, M., & Swann, W. B. (1978). Hypothesis-testing processes in social interaction. *Journal of Personality and Social Psychology, 36*, 1202–1212.
- Sweeny, K., Carroll, P. J., & Shepperd, J. A. (2006). Thinking about the future: Is optimism always best? *Current Directions in Psychological Science, 15*, 302–306.
- Sweeny, K., & Krizan, Z. (in press). Sobering up: A quantitative review of temporal declines in optimism. *Psychological Bulletin*.
- Taylor, S. E., & Brown, J. D. (1988). Illusion and well-being: A social psychological perspective on mental health. *Psychological Bulletin, 103*, 193–210.
- Trope, Y., & Liberman, A. (1996). Social hypothesis-testing: Cognitive and motivational mechanisms. In E. T. Higgins & A. W. Kruglanski (Eds.), *Social psychology: Handbook of basic principles*. New York: Guilford Press.
- Tyler, J. M., & Rosier, J. G. (2009). Examining self-presentation as a motivational explanation for comparative optimism. *Journal of Personality and Social Psychology, 97*, 716–727.
- van Dijk, W. W., Zeelenberg, M., & van der Pligt, J. (2003). Blessed are those who expect nothing: Lowering expectations as a way of avoiding disappointment. *Journal of Economic Psychology, 24*, 505–516.
- Vosgerau, J. (2010). How prevalent is wishful thinking? Misattribution of arousal causes optimism and pessimism in subjective probabilities. *Journal of Experimental Psychology: General, 139*, 32–48.
- Weber, E. U. (1994). From subjective probabilities to decision weights: The effect of asymmetric loss functions on the evaluation of uncertain outcomes and events. *Psychological Bulletin, 115*, 228–242.
- Windschitl, P. D., Smith, A. R., Rose, J. P., & Krizan, Z. (2010). The desirability bias in predictions: Going optimistic without leaving realism. *Organizational Behavior and Human Decision Processes, 111*, 33–47.
- Wyer, R. S., & Frey, D. (1983). The effects of feedback about self and others on the recall and judgments of feedback-relevant information. *Journal of Experimental Social Psychology, 19*, 540–559.