Available online at www.sciencedirect.com
SCIENCE DIRECT®
Organizational Behavior and Human Decision Processes xxx (2006) xxx–xxx
ORGANIZATIONAL BEHAVIOR AND HUMAN DECISION PROCESSES
www.elsevier.com/locate/obhdp

Information distortion in the evaluation of a single option

Samuel D. Bond [a,*], Kurt A. Carlson [a], Margaret G. Meloy [b], J. Edward Russo [c], Robin J. Tanner [a]

[a] *Fuqua School of Business, Duke University, Durham, NC, USA*
[b] *Smeal College of Business, Penn State University, University Park, PA, USA*
[c] *Johnson Graduate School of Management, Cornell University, Ithaca, NY, USA*

Received 28 June 2005

Abstract

Extending previous work on biased predecisional processing, we investigate the distortion of information during the evaluation of a single option. A coherence-based account of the evaluation task suggests that individuals will form an initial assessment of favorability toward the option and then bias their evaluation of subsequent information to cohere with their initial disposition. Three experiments tested this hypothesis. Initial disposition was manipulated (Studies 1 and 3) or measured (Study 2), and attribute ratings were collected as indicators of information distortion. Results from all three experiments indicate that attribute evaluations were biased to favor initial dispositions. These findings provide evidence that information distortion is one cause of primacy effects in judgment and decision-making settings involving a single option.

Keywords: Predecisional bias; Information distortion; Primacy; Singular evaluation

Most of us believe that we are capable and impartial decision makers. We envision ourselves dispassionately collecting information about alternatives, analyzing this information in an unbiased fashion, and carefully drawing conclusions. Contradicting this flattering view, research has shown that our decisions are affected by a variety of factors that we are unable or unwilling to recognize. A prominent example is the phenomenon of primacy effects, whereby the starting point of a decision process has a disproportionate effect on its outcome. The present work focuses on a particular cause of this general primacy effect. Specifically, we test the idea that individuals, having not yet committed to a particular course of action, bias their interpretation of new information to support an initial disposition.

Expanding on recent interest in biased predecisional processing (Brownstein, 2003; Russo, Carlson, & Meloy, 2006; Simon, Krawczyk, & Holyoak, 2004), this article examines the evidence for primacy in the evaluation of a single, novel option. Further, we identify a process mechanism that generates the observed primacy effect. Central to this mechanism is the notion of cognitive coherence among the elements of a decision (Heider, 1946; Simon, Krawczyk, et al., 2004). We suggest that in the process of evaluating a single option, individuals maintain coherence in their emerging evaluation of the option by biasing their interpretation of new information. Specifically, we suggest that individuals will form an initial, tentative disposition toward the option, and subsequent information will be systematically distorted to cohere with that preliminary disposition.

The remainder of this paper is organized as follows: we first review work on judgmental primacy effects and the mechanisms that have been proposed to account for them. We focus on one of these mechanisms—subjective distortion of information—and review evidence for this bias in research on choice among options. Next, we

* Corresponding author. Fax: +1 919 660 7996.
E-mail address: sdb4@mail.duke.edu (S.D. Bond).

0749-5978/$ - see front matter © 2006 Elsevier Inc. All rights reserved.
doi:10.1016/j.obhdp.2006.04.009

suggest that the bias may extend to single-option evaluation, and we propose that a necessary first step is the emergence of an evaluative disposition. Three studies follow, each demonstrating how an initial evaluative disposition affects the interpretation of new information and influences downstream judgments and decisions.

Primacy effects in judgment and choice

For present purposes, we refer to primacy effects in a decision context as the disproportionate influence of information acquired early in a process on its final outcome (Lingle & Ostrom, 1981; Nickerson, 1998). The existence of primacy effects has been verified across a range of psychological disciplines, and the topic has long interested organizational researchers. For example, a number of scholars have investigated primacy in the personnel interview process, finding that recruiters' assessments of a candidate are powerfully affected by their initial impressions (Cable & Gilovich, 1998; Macan & Dipboye, 1990; cf. Sackett, 1982). Similarly, observer ratings of work performance have been found to be biased towards the information received earliest (Sinclair, 1988). In the study of negotiations, investigators have shown that initial offers substantially affect counteroffers, aspiration levels, and final settlement prices (Galinsky & Mussweiler, 2001; Kristenson & Garling, 1997; Ritov, 1996), and these effects have been demonstrated even when initial values are known to be unreliable (Whyte & Sebenius, 1997). Primacy in a courtroom setting has also received prominent attention, given that jurors are advised to avoid forming an opinion until all the evidence has been presented. Contrary to this procedural ideal, research indicates that jurors establish an initial representation of events based on prior beliefs and information presented first (Devine & Ostrom, 1985; Holstein, 1985), and that final verdicts are likely to align with these initial opinions (e.g., Lawson, 1968). In addition to areas just mentioned, the importance of early (and sometimes irrelevant) information in decision processing has been demonstrated in studies of probability estimation (Adelman, Tolcott, & Bresnick, 1993), asset valuation (Northcraft & Neale, 1987), goal setting (Hinsz, Kalnbach, & Lorentz, 1997), medical decision making (Barrows, Feightner, Neufeld, & Norman, 1978; Berwick, Fineberg, & Weinstein, 1981), and procedural justice (Lind, Kray, & Thompson, 2001; van de Bos, Lind, Vermunt, & Wilke, 1997). Primacy is among the most widely observed phenomena in basic and applied behavioral research.

Explanations of the primacy effect

Given the wide range of results that have been attributed to the operation of primacy, it is hardly surprising that a variety of theoretical explanations have been proposed. For ease of exposition, we group these accounts into four general categories: biased search from memory, biased search from the environment, biased weighting of information, and biased information evaluations. Although these categories are mutually exclusive, multiple mechanisms may share responsibility for any observed primacy effect.

The category of explanations most frequently invoked is that dealing with *biased search from memory*. Many of the researchers taking this approach combine models of information integration (e.g., Anderson, 1981) with models of memory storage and retrieval (e.g., Srull, 1981). For example, Hastie and Park (1986) distinguish between online judgments and judgments from memory, arguing that the former tend to exhibit a primacy effect of information order, while the latter tend to show a recency effect. By far the most prominent explanation within this category is the phenomenon of anchoring and adjustment, defined by Tversky and Kahneman (1979, p. 1128) as a process by which individuals "make estimates by starting from an initial value that is adjusted to yield a final answer…adjustments are typically insufficient." Although Tversky and Kahneman portrayed the process as one of incremental, incomplete adjustment, there continues to be debate about the extent to which adjustment occurs (Epley & Gilovich, 2004). This issue aside, there is now an emerging consensus that anchoring effects are driven by selective accessibility of anchor-consistent information in memory (Chapman & Johnson, 1999; Gilovich, 1991; Mussweiler & Strack, 1999; Strack & Mussweiler, 1997). Although much of the work on anchoring has explored arbitrary, irrelevant anchors, similar data patterns have been found in cases where the anchor is an initial opinion or belief. For example, work on the myside bias reveals that people recall reasons supporting the side of an issue they favor more easily than reasons supporting the opposition (Baron, 1995).

A second, related class of explanations focuses on *biased search from the environment*. Research on selective hypothesis testing has documented a tendency among problem solvers to investigate a hypothesis by examining only cases where it is likely to occur (Baron, Beattie, & Hershey, 1988; Klayman & Ha, 1987). This tendency, demonstrated most famously in Wason's (1960) card selection task, illustrates a more general failure to consider the diagnosticity of available information for confirming or disconfirming one's hypothesis (Doherty, Mynatt, Tweney, & Schiavo, 1979). Although cognitive limitations may play a role in the selective acquisition of information, motivational variables are also relevant (Kunda, 1990). For example, research on selective information exposure has shown that when disconfirmation of an existing belief is threatening, individuals become more discriminating in their search for information

(Frey, 1986). Consequently, individuals who form an opinion based on the earliest information received may actively avoid later information that conflicts with this opinion (Jonas, Schulz-Hardt, Frey, & Thelen, 2001).

A third class of explanations for primacy deals with *biased weighting of information* that is encountered during a judgment process. Biased weighting is central to the belief-updating model of Hogarth and Einhorn (1992), who describe an iterative mechanism for the effect of new information on overall evaluations: after being subjectively encoded, new information is compared to some reference point, weighted, and assimilated to one's prior evaluation. The extent to which early information or later information is overweighted in the final summary evaluation (leading to primacy or recency, respectively) depends on features of the judgment task, such as the number of items to be processed, their complexity, and the frequency of updating. If, for example, a decision task is short and relatively simple, the model predicts recency effects to predominate. Thus, when evaluating a project described by a few short attributes, one's belief that the project will succeed should be strongest when the most favorable information about the project is encountered last. Empirical observations have confirmed this effect of information order, even among judges trained in the task domain (Highhouse & Gallo, 1997; Tubbs, Gaeth, Levin, & Van Osdol, 1993).

Distortion of information

Although the theories above provide a broad range of explanations for primacy in certain decision settings, they do not exhaust the set of possibilities. The present research explores a fourth explanation, *distorted information evaluations*. The focus of this explanation is neither the method by which decision-relevant information is obtained nor the weight that it receives, but rather how this information is subjectively interpreted by the decision maker. For example, it is commonly accepted that individuals should not be affected by earlier opinions when determining the diagnostic value of new information (Chapman & Chapman, 1969; Troutman & Shanteau, 1977); in Bayesian terms, the interpretation of new information should remain independent of one's prior judgment (Evans & Over, 1996). However, numerous investigations have uncovered violations of this principle (in the context of Bayesian inference, see Boulding, Kalra, & Staelin, 1999). In general, information supporting a desired conclusion is absorbed noncritically, while opposing information is treated with "motivated skepticism" (Ditto & Lopez, 1992). In their work on belief persistence, Lord, Ross, and Lepper (1979) observed that opinions of individuals with pre-existing views became more polarized after viewing nondiagnostic information; that is, each side biased their evaluations of the same (neutral) evidence to fit their prior beliefs. Other research reveals that individuals who have made a decision tend to interpret subsequent information as supporting their choice (e.g., Elliot & Devine, 1994). Such postdecisional distortion is commonly attributed to the reduction of cognitive dissonance (Festinger, 1957), whereby individuals reduce the inconsistency between their attitudes or beliefs and their behavior ("I chose X, but this information favors Y") by adjusting their attitudes ("In fact, this information favors X.").

Distortion without prior commitment

More recently, a series of investigations has established that decision-relevant information can be distorted *in advance* of a decision, even in the absence of prior beliefs or preferences (Holyoak & Simon, 1999; Russo, Meloy, & Medvec, 1998). These studies reveal compelling violations of the cognitive algebra principle of *meaning invariance*, which states that unrelated pieces of information should be evaluated independently (Anderson, 1981). In typical demonstrations, individuals are informed of an impending choice and given descriptions of two or more unfamiliar alternatives. As part of the task, they are asked to evaluate each piece of information they encounter by indicating the degree to which it favors one option over the other. Results typically reveal that individuals form an early, tentative preference for one of the options, and later information is distorted to favor this alternative.

For example, Russo, Medvec, and Meloy (1996) asked participants to evaluate attribute information about two restaurants in the context of a decision between them. Comparing these attribute ratings to those of a control group, the authors showed that decision makers distorted the information to favor whichever restaurant was leading at the time. In related work, Svenson has described a variety of mechanisms by which attributes of the alternatives are restructured to support a preliminary choice (for a review, see Svenson, 1996). Most relevant to the present discussion is *attractiveness restructuring*, in which the subjective attractiveness of the alternatives is transformed, attribute by attribute, to favor the preliminary choice. These ideas are congruent with findings obtained in research involving medical decision making (Wallsten, 1981), advertising effects (Hoch & Ha, 1986), and impression formation (Seta & Hayes, 1994). Challenging the normative assumption, these studies reveal that individual pieces of uncorrelated information are not evaluated independently. Instead, evaluations are biased to support whichever alternative is currently preferred.

The operation of such distortion can have a powerful effect on choices. In a consumer setting, Carlson, Meloy, and Russo (2006) asked participants to evaluate a sequence of attributes describing two brands and then

select the brand that they preferred. By manipulating the order of the information sequence, it was possible to install one brand or the other as the initial leader. The manipulation resulted in a large leader-driven primacy effect, with the initially installed leader selected by roughly 70% of participants. Of particular relevance to the current article, the observed primacy effect was traced back to a systematic bias in the evaluation of attribute information.

Will information distortion extend to a single-option setting?

A variety of theories have been advanced to explain the occurence of biased predecisional processing in choices among options (Brownstein, 2003). In general, these theories share the assumption that predecisional processing is directed towards the specific goal of separating the various alternatives. For example, Svenson's (1992) differentiation–consolidation theory suggests that individuals first differentiate among the options to identify a superior target, then consolidate their preference by reconsidering and reevaluating the information. Similarly, Montgomery (1983, 1994) proposes that decision makers seek to create a "dominance structure," whereby the selected option dominates all others. In Janis and Mann's (1977) conflict theory, separation is motivated by anticipated regret, which leads individuals to engage in motivated processing.

In contrast to the examples reviewed above, many everyday evaluation tasks involve a single item. For example, when deciding whether to buy a product, accept an invitation to dinner, or play the state lottery, alternatives to the focal option may be so distant or poorly specified that they are not considered. Indeed, there has been growing interest in the special characteristics of singular evaluation. Investigators have uncovered a range of factors that are unique to this setting, including the importance of information that is easily evaluated (Hsee & Leclerc, 1998), a reduction in the need to justify (Okada, 2005), and a tendency to overweight negatively valenced attributes (Willemsen & Keren, 2004). Given that singular and multiple evaluation involve distinct modes of processing, it is not clear that information distortion should be expected in single-option settings. Moreover theoretical accounts relying on a separation goal are largely inapplicable to this domain: without multiple alternatives, neither the spreading component of differentiation–consolidation (Svenson, 1992) nor the relative superiority of a dominance relationship (Montgomery, 1983) seems possible.

Coherence and its processing implications

We propose a broader basis for information distortion that does not depend on a desire to separate discrete alternatives. Specifically, we suggest that the bias can be explained by the operation of coherence-directed processing (Holyoak & Simon, 1999; Simon, Snow, & Read, 2004). The concept of coherence derives from classic cognitive consistency theories (e.g., Heider, 1946), which propose, first, that mental processes are directed toward a state of equilibrium, and, second, that elements of knowledge are altered to achieve this state. Thagard (2000) discusses five types of coherence, which differ from each other based on the elements of knowledge involved. For example, explanatory coherence is characterized by a high degree of fit between hypotheses and evidence, while conceptual coherence is characterized by fit between stereotypes, attitudes, or impressions. Across all types of coherence, the fundamental assumption is that incoherent representations are inherently unstable, so coherent irrepresentations are constructed in their place. Therefore, the coherence approach to information distortion suggests that whenever a task requires the processing of new information, coherence is increased by interpreting this information to be consistent with existing knowledge.

In a series of recent investigations, Holyoak, Simon, and their colleagues have applied this coherence mechanism to decisions among alternatives. Borrowing from cognitive models of parallel constraint satisfaction (Read, Vanman, & Miller, 1997), their research has shown that decisions unfold in a series of *coherence shifts*, by which decision makers transform both new information and prior beliefs to align with each other and the emerging choice. For example, Simon, Krawczyk, et al. (2004) presented participants with a choice between two job offers. Over the course of the study, participants were asked to rate information about the jobs on three occasions: before, during, and after their actual decision. Results showed that attribute ratings became more correlated over time, and, as a result, more supportive of the chosen offer. However, like the theories cited above (Montgomery, 1994; Svenson, 1992), the model presented by Holyoak, Simon, and colleagues characterizes the decision-making process as one of separating discrete alternatives: "At each point of equilibrium, the subset of the task variables that support the emerging decision are strongly endorsed, whereas the variables that support the rejected decision are suppressed or rejected: in effect, the representations of the alternatives become spread apart" (Simon, Snow, et al., 2004, p. 816).

Expanding on the coherence-seeking account of decision making, we suggest that information distortion will occur even in the evaluation of a single option. Our fundamental assertion is that a tentative assessment of an option's appeal, an *evaluative disposition*, is a crucial element in the pursuit of cognitive coherence. We propose that information viewed early in the decision process

evokes an initial evaluative disposition,[1] and this disposition influences the interpretation of information viewed later. Just as new information in a binary-choice setting can be made to cohere with an emerging preference for one alternative, new information in a single-option setting can be made to cohere with an evaluative disposition toward that option. In the eyes of an individual with a highly unfavorable initial disposition, attributes which are objectively "neutral" towards the option may instead seem quite negative, and even "positive" attributes will appear less favorable. Accordingly, the pursuit of coherence causes systematic bias in evaluations of the attributes encountered and, as a consequence, in final evaluation of the option. In the language of older consistency models (e.g., Heider, 1946), the evaluative disposition represents a specific cognitive element, generated early in the decision process, to which other elements of the decision are aligned.

Although distortion in single-option evaluations is predicted by our coherence account, this prediction does not obviously follow from a separation-based perspective. Consider, for example, an individual deciding whether to purchase a particular product, presented in isolation. What is the second alternative whose representation is modified to produce separation? A tempting answer is to consider the default (do nothing) as the second alternative. However, the default alternative lacks an attribute structure against which to compare the focal option, and it is this comparison process that prior research has implicated as the driver of distortion (e.g., Carlson et al., 2006). Alternatively, what if the individual is not requested to make a decision, but rather is asked to rate the product on a scale of attractiveness? Separation accounts would be even less relevant for such a task, which requires arriving at a particular point along a continuum. Therefore, if systematic distortion occurs in choice and rating tasks involving a single option, the coherence account provides a more parsimonious explanation than traditional, separation-based approaches.

Overview of the present studies

In the studies that follow, a variety of singular evaluation settings are employed to examine whether information is distorted in support of an initial disposition. Study 1 manipulates initial dispositions directly and explores their effects on the simple consumer decision of whether or not to purchase a product. Ratings of product-relevant information provide a measure of biased processing. Study 2 involves the evaluation of scholarship candidates from a pool of unknown applicants. Evaluative dispositions are allowed to develop naturally, and the evidence for distortion is examined in contexts of both binary choice (i.e., award/deny) and numerical judgment (i.e., rate the candidate). Finally, Study 3 explores willingness to pay for a risky prospect, using measures of perceived probability and value as indicators of bias. Importantly, all three studies involve settings in which participants (a) have no prior knowledge of the options and (b) cannot search for information selectively (from memory or from the environment). By focusing on situations where other mechanisms are unlikely to operate, we are able to determine whether distortion of information is sufficient for primacy to occur. We acknowledge, however, that distortion is just one of many channels by which primacy may be manifest.

To summarize, the empirical goals of this paper are threefold. First, we demonstrate primacy effects in single-option judgment and choice. Second, in accordance with our coherence explanation, we trace the source of these primacy effects to a biasing of information evaluations in support of an initial evaluative disposition. Finally, we utilize a variety of decision domains (product evaluation, impression formation, and risky prospect) to demonstrate the robustness of the information distortion paradigm.

Study 1: Single-option choice with an installed disposition

As a simple example of a single-option decision setting, consider the case of a consumer deciding whether to purchase a particular product, based on pieces of information acquired in sequential order. Study 1 tested the hypothesis that individuals faced with this decision will bias their information evaluations to support an initial disposition toward the product. The product examined was a personal digital assistant (PDA) described by six attributes, and the primary manipulation was the order in which this information was presented. We reasoned that seeing a positive (negative) attribute first would create a positive (negative) initial evaluative disposition toward the PDA. We expected that this leaning would influence the evaluation of subsequent attributes by the coherence process outlined above, and, as a result, more participants would opt to purchase the PDA when the first attribute was favorable.

Participants and design

One hundred eighteen undergraduates at a large university were paid $5 for participating in this and other,

[1] A growing body of social cognition research has established that evaluation is spontaneous and can happen below conscious awareness (Duckworth & Bargh, 2002; Ferguson & Bargh, 2003). Thus, even when individuals are intent on reserving judgment, they are likely to form an initial evaluative disposition of the focal option on the basis of the information they encounter first.

unrelated studies. A cover story informed participants that an aunt had recently given them a $150 birthday present, suggesting that the money be used to purchase a PDA. The stimulus packet contained a description of one PDA based on six attributes (*functionality*, *dimensions*, *data transfer capability*, *memory, price*, *and battery life*). Two target attributes, *functionality* and *price*, were used to install an initial disposition toward the PDA. The functionality attribute was written to provoke a moderately positive disposition:

The PDA comes with a complete suite of personal organization tools including: calendar, appointment book, phone number & address book, and To Do List. It has several useful applications preinstalled, such as a spreadsheet, word processor, email, and calculator, as well as a selection of classic games. It also has an on-screen keyboard (or can accommodate a portable keyboard) and has handwriting recognition software for easy use.

Conversely, the price attribute was written to provoke a moderately negative disposition (recall that only $150 was provided for the gift):

Until recently, this PDA sold for as much as $249.99, but it now sells for $199 at Mysimon.com and for $199.99 at CircuitCity.com. A 4-Pack of extra screen pens sells for $9.99 and an extended 24-Month Warranty can be purchased for $34.99.

The other four attributes were represented by similar narrative descriptions. However, these four attributes were designed and pretested to be neutral (on average) with respect to the purchase decision. We expected that whatever disposition was installed by the first attribute would gain support through biased evaluation of the remaining attributes, thereby influencing the eventual purchase decision in a manner consistent with primacy.

Procedure

All participants read initial instructions explaining the scenario of the cash gift and asking them to imagine that they had to decide whether to purchase the PDA. Attribute descriptions were presented sequentially on separate pages. For those in the favorable-first condition, the positive *functionality* attribute was presented first in the six-attribute sequence and the negative *price* attribute was presented fifth. In the unfavorable-first condition, the position of these two attributes was reversed. The four neutral attributes were randomly assigned to the other four positions, and this order was reversed for half of the participants for control purposes. The order of the four neutral attributes had no impact on any of the dependent variables and is not discussed further.

In order to track participants' preferences as the decision task progressed, we utilized the *stepwise evolution of preference* (SEP) method drawn from previous work on binary choice (Meloy & Russo, 2004; Russo et al., 1998). Applied to the single-option setting here, the critical features of the SEP method are: (1) presentation of information about the attributes one at a time, and (2) elicitation of three responses after each attribute. The first question, worded as follows, captured the extent to which participants felt the information was favorable to the product: "Consider only the information on *this* page and consider the extent to which you agree with the following statement: This information makes the PDA appeal to me." The nine-point response scale was anchored by 1 (strongly disagree) and 9 (strongly agree). The next two questions were designed to capture cumulative dispositions toward the PDA. First, participants estimated which way they were currently leaning: "Think about all the information you have received so far. If you were to think of your decision to buy the PDA as a horse race between buying and not buying the PDA, which would you consider to be leading right now?" (The horserace metaphor was used to convey that a lead was tentative and might change before the "race" was over.) Participants circled either "buy the PDA" or "do not buy the PDA." Next, participants indicated how confident they were in their leaning: "If you were given $10 to bet on which of these two positions would win the horse race, how much would you bet that your currently held position would win the race?" The scale was anchored by $5 (dead even) and $10 (clear winner).

Results and discussion

In order to verify that the target attributes created the desired initial dispositions, we examined participants' leanings immediately after the first attribute. The proportion of participants leaning toward purchasing the PDA at this point was significantly greater when the positive attribute (*functionality*) appeared first (.44) than when the negative attribute (*price*) appeared first (.22; $\chi^2 = 6.47$, $p < .01$). Thus, the two target attributes had their intended effect.

Next, we tested whether participants biased their attribute evaluations to support their initial dispositions. As predicted, the average rating of the four neutral attributes was significantly greater when the first attribute was favorable ($M = 6.06$) than when the first attribute was unfavorable ($M = 4.86$; $t(114) = 3.78$, $p < .001$), and this result also held for each attribute tested individually (all $p < .01$). Furthermore, the negative target attribute (*price*) was rated significantly more negatively when it appeared in the first position ($M = 4.06$) than when it appeared in the fifth position ($M = 5.02$; $t(114) = 62.56$, $p < .01$), and the positive attribute (*functionality*) was rated marginally more positively when it appeared first ($M = 6.43$) rather than fifth in the sequence ($M = 5.83$; $t(114) = 1.69$, $p < .09$). Together, these data indicate that

participants biased their evaluations of the attributes to cohere with their initial dispositions.

In addition, we examined participants' purchase decisions for evidence of primacy. As predicted, the proportion of participants indicating that they would purchase the PDA was significantly greater in the favorable-first condition (.48) than in the unfavorable-first condition (.22; $\chi^2 = 8.41$, $p < .01$). Thus, we have evidence not only that attribute evaluations can be biased in a single-option setting, but also that 'yes–no' choices are influenced by coherence-driven processing. It is noteworthy that choices aligned with initial dispositions even though information encountered after the first attribute was, on balance, objectively counter to initial leanings (four neutral attributes and one target attribute of opposite valence). The *subjective* balance of these subsequent attributes, however, was dramatically affected by the distortion process (as shown above). For example, the unfavorable *price* attribute was actually viewed positively by those with a favorable initial disposition.

In describing the features of any preference formation task, a distinction should be made between the number of items being evaluated and the number of available responses. Even when there is a single target of evaluation, *choice* tasks (like that used in Study 1) present decision makers with at least two response options, (e.g., to buy or not buy a PDA). Accordingly, one might argue that a separation goal could be active in these tasks, because decision makers are motivated to create a distinction between the available responses. Equally common, however, are cases of preferential *judgment*, where potential responses lie along a continuum. For example, the assessment of a product's dollar value involves an unlimited number of possible responses. It is for decisions of the latter type that the notion of separation is least applicable. However, coherence should operate in either situation, biasing evaluations in favor of an initial disposition. Evidence supporting this contention can be found in a series of studies described by Montgomery (1999); participants were given multiple options and instructed either to choose one option or assign preference ratings to each of them. Attribute evaluations were taken prior to the final decision, and these evaluations revealed surprisingly similar levels of bias across the choice and rating groups. In order to address this issue in a single-option setting, Study 2 includes both choice and judgment tasks.

Study 2: Choice, judgment, and spontaneous preferences

Study 1 demonstrated that individuals deciding whether to accept or reject a focal option will bias their attribute evaluations, creating coherence between new information and an initial evaluative disposition. Study 2 attempts to extend this finding in two respects. First and more importantly, whereas the decision task of Study 1 was a forced choice to accept or reject the option, Study 2 uses both choice and judgment instructions. The judgment task instructions allow us to explore whether distortion occurs when the processing objective is a single judgment from a continuous distribution of potential responses. Second, the procedure used in Study 2 does not manipulate initial dispositions directly, but instead allows them to develop spontaneously as a result of idiosyncratic preferences.

Participants and design

Fifty-four undergraduates participated in the study for extra course credit. Their task was to evaluate four scholarship candidates, one candidate at a time, by examining information on six relevant attributes. Participants were divided evenly into two groups: *choice* and *judgment*.

As in Study 1, all participants read and evaluated descriptions of all six attributes. In the choice condition, participants indicated whether they were leaning towards awarding or denying the scholarship to the candidate after each attribute; at the final stage, these participants provided a binary decision (award/deny). In the judgment condition, participants read and evaluated the same information but were instructed at each step to provide a rating of the applicant's qualifications for a scholarship; at the final stage, these participants provided a summary judgment of overall merit.

Procedure

Participants read a role-playing scenario explaining the task and the criteria used to assess accuracy. The scenario involved "auditioning" to become the undergraduate representative on a scholarship awards committee. Participants were asked to evaluate the profiles of four scholarship applicants, one applicant at a time, to determine the applicant's worthiness of receiving a scholarship. Six attributes of information were provided for each applicant: *performance on exams* (e.g., SAT, high school GPA), *work experience*, *extra-curricular activities*, *letters of recommendation*, *results of a personal interview*, and *quality of application essay*. Half of the subjects reviewed these six dimensions in a randomly chosen order, and half reviewed the same information in reverse order. Because no order effects were observed, this counterbalancing will not be discussed further.

As in Study 1, the procedure utilized a modified version of the SEP method to track evolving preferences for both the choice and the judgment groups. After reading each attribute, participants in both conditions rated its favorability: "Please consider the information that you just received. Rate it on the 1-to-9 scale below according to your best judgment." The endpoints of the scale were

labeled 'deny a scholarship' and 'award a scholarship.' Following this question, participants responded to items (described below) measuring their current evaluative disposition towards the candidate.

In the choice condition, participants were given the task of determining whether the candidate should be awarded or denied a scholarship. Following the favorability rating of each attribute, evaluative disposition was measured by the following: "Based on all of the information you've seen so far, are you leaning toward or away from awarding this candidate a scholarship?" Next, confidence in the leaning was assessed by a betting analogy similar to that of Study 1: "If you were given $10 to place a bet on whether you would award the candidate a scholarship, how would you split the $10 between 'awarding' and 'denying'?" The scale ranged from $5 on each ("dead even") to $10 ("clear choice"). After reviewing all six attributes, participants indicated whether the applicant should be awarded or denied a scholarship—a binary decision.

In the judgment condition, participants evaluated the same information but were instructed to rate the applicant's overall qualifications for a scholarship on a scale from 1 to 9. Participants received the following instructions regarding the judgment task:

Please really think about what it means to score one candidate as a 6, one candidate as a 5, and another candidate as a 7. This difference is important because the scholarship awards committee has a wide range of scholarships available. In some cases, restrictions have been placed on eligibility by the donors. The exact rating helps identify the more qualified candidate if there are multiple candidates who are eligible for scholarship funds and who qualify for multiple scholarships.

After evaluating each attribute on the scale described above, these participants indicated their current disposition by responding to the following: "Based on all of the information you have seen so far, what rating are you most leaning toward for how likely this candidate is to receive a scholarship?" Responses were given on a nine-point scale (1 = will certainly not receive a scholarship, 9 = will certainly receive a scholarship). These ratings indicated both the direction and the magnitude of a participant's leaning. After all six attributes were reviewed, participants provided a final summary rating of the candidate's merit on a nine-point scale (1 = the candidate should be denied a scholarship; 9 = the candidate should be awarded a scholarship).

Results

If an initial disposition caused participants to bias their interpretation of later information, this bias would be reflected in attribute evaluations that were skewed toward that initial disposition. Participants in each condition were divided into two groups based on their evaluative disposition after the first attribute. As predicted, evaluations of subsequent attributes reflected systematic distortion in the direction of the initial disposition. Table 1 shows the mean evaluations of both the choice group and the judgment group for each attribute.

In the choice condition, when the initial leaning favored denying the scholarship, subsequent attributes were rated systematically lower ($M = 5.10$) than when the initial learning favored awarding a scholarship ($M = 6.42$; $t(104) = 6.69$, $p < .001$). The judgment condition yielded similar results. When the initial leaning was in the lower half of the rating scale, signifying an unfavorable disposition towards the candidate, the average rating for the remaining attributes was 4.77. When the initial leaning was in the upper half of the scale, signifying a favorable disposition, the average rating of the remaining attributes was significantly higher at 6.61 ($t(108) = 7.27$, $p < 0.01$). These results suggest that an initial disposition, based on a single piece of information, caused individuals to distort subsequent information to be coherent with that disposition. Furthermore, as

Table 1
Attribute evaluations for choice and judgment groups in Study 2

	Choice group			Judgment group		
	Initial disposition against	Initial disposition towards	Difference[b]	Initial disposition against	Initial disposition towards	Difference[b]
Activities	5.17	6.48	1.31	4.71	6.61	1.90
Interview	4.06	5.84	1.78	4.14	6.14	2.00
Letters	4.83	6.76	1.93	5.07	6.82	1.75
Exams	5.61	6.90	1.29	5.14	6.88	1.74
Statement	4.78	6.04	1.26	4.36	6.23	1.87
Work Experience	5.22	6.57	1.35	4.79	6.88	2.09
Decision to Award[a]	16.7%	81.8%	65.1%	42.9%	92.7%	49.8%

[a] For the choice group, decision to award is the % of participants awarding the scholarship. For the judgment group, decision to award is the % of participants whose final summary rating was greater than the midpoint of five. Note that because attribute order was not manipulated in Study 2, it is not surprising that final decisions are correlated with initial dispositions.

[b] All differences significant at $p < .01$.

shown in Table 1, an examination of final decisions reveals substantial primacy effects on both the decision to award or deny (choice condition) and final summary evaluations (judgment condition).

Finally, was there a difference in the amount of distortion that occurred in the choice and judgment tasks? For every attribute, we calculated the deviation between each participant's rating on the nine-point scale and the neutral value of 5 (i.e., indifference). This deviation was signed positively if the evaluation favored the initial disposition and negatively if it did not. Averaging these deviations for each participant allows us to compare the magnitude of the bias in the two conditions. Across subjects in the choice condition, the average distortion was 1.13 units in favor of the initial disposition ($t(26) = 8.13$, $p < .001$), and the analogous value in the judgment condition was 1.42 units ($t(27) = 11.09$, $p < .001$). The difference in these magnitudes was not significant ($t(53) = 1.51$, ns). Although null effects must be interpreted with caution, the fact that the bias was not more prominent in the choice condition is difficult to reconcile with separation-based accounts.

Discussion

The results from the choice condition of Study 2 mirrored the results of Study 1, suggesting that spontaneously formed evaluative dispositions lead to the same information distortion as those manipulated experimentally. Furthermore, the judgment condition provided initial evidence that distortion occurs in evaluative tasks requiring a summary judgment.

Study 3 was designed to extend the investigation in two ways. First, even though the judgment task of Study 2 presented a distribution of possible responses, these responses were framed in terms of awarding or denying a scholarship. Therefore, participants may still have interpreted the task to some extent as a choice between two courses of action. For this reason, Study 3 involves a task that would be difficult to reframe in binary terms. Second, the attribute descriptions provided in Studies 1 and 2 were verbal and multifaceted, allowing substantial flexibility in their interpretation. Study 3 explores the effects of an initial disposition in a setting where attribute information is numeric and considerably more concrete.

Study 3: Valuing a lottery

The demonstrations in Studies 1 and 2 relied on tasks that involve multiple kinds of uncertainty. During each study, information was provided about relatively complex attributes, and a personal assessment of the information was required. Therefore, participants faced uncertainty both in the attribute information itself and in knowledge of their own preferences. Presumably, the pursuit of cognitive coherence is facilitated in such an environment, because there is sufficient ambiguity for evaluations to be easily distorted.

A different situation arises for judgment tasks involving the evaluation of an uncertain prospect. One of the simplest examples is a gamble with a specified monetary payoff and probability of winning. Given that the probability and payoff are the central "attributes" of such a prospect, should we expect these attributes to be distorted in support of an initial disposition? On the one hand, the numeric unidimensionality of the payoff and probability may inhibit distortion of this information. On the other hand, even probabilities and payoffs must be subjectively interpreted to determine overall preference. Tversky and Kahneman's *prospect theory* (1979) proposes that payoffs or losses are evaluated according to a reference point, which is itself a type of subjective frame. Actual amounts are converted to subjective values by a value function, so that losses are typically felt more strongly than equivalent gains. Similarly, actual probabilities are converted by a weighting function so that, for example, low probabilities are overweighted and high probabilities are underweighted.

In the present context, the value and weighting functions provide a means by which information can be distorted to achieve coherence. For example, upon discovering that the chance of winning is small, individuals might form an unfavorable initial assessment of a gamble. When presented with objectively positive information on the next attribute (a high payoff value), these individuals may distort their evaluation of the attribute downward in order to maintain coherence with their negative disposition.

Participants in the third study read about a simple lottery and evaluated its two attributes, *probability of winning* and *prize value*. As in Study 1, the order of information was manipulated between groups; we expected that initial exposure to positive (negative) information would create a favorable (unfavorable) disposition towards the gamble. By obtaining direct assessments of the probability and value information, we were able to determine whether the attributes were evaluated differently by participants in the two conditions. Such a result would demonstrate that initial dispositions can affect judgments involving precise, unambiguous information.

Participants and design

Participants were 102 MBA students at a large university. The study was contained within a set of unrelated experiments. Stimuli were designed so that either the unfavorable *probability of winning* (.025) or the favorable *prize value* ($200) came first, with the other attribute following. Participants were randomly assigned to one of the two versions.

Procedure

Participants were given a one-page, pen-and-paper questionnaire describing a hypothetical lottery to be held at the university. Following the introduction, information was provided about the first attribute (either *probability of winning* or *prize value*, depending on condition). After evaluating this information (see below), all participants viewed a paragraph of information describing the lottery; e.g., the paragraph stated that winners would be chosen randomly by a representative of the university. Next, participants viewed information regarding the second attribute (*prize value* or *probability of winning*, depending on condition).

Immediately following the *probability* and *prize* attributes, evaluation questions were presented (participants evaluated the first attribute before seeing the second). Specifically, after reading about the prize ($200), participants were asked, "How valuable would the $200 be to you?" Responses were collected on a nine-point scale (1 = not at all valuable; 9 = enormously valuable). Likewise, after reading about the probability of winning (.025), participants were asked "To what extent do you believe this is a good chance of winning?" and gave their response on a nine-point scale (1 = extremely poor chance of winning; 9 = extremely good chance of winning).

Results and discussion

In order to examine the hypothesis that perceptions of probability and payoff information would be affected by the order in which they were presented, we compared the attribute evaluations across the two conditions. Fig. 1 shows the mean evaluations of *probability* and *prize value* for both attribute orders. As expected, the perceived value of the relatively attractive $200 prize was significantly impacted by the order in which attributes were presented. When the prize amount was presented first, participants rated it significantly more valuable than when they first viewed the small probability of winning ($M = 5.14$ vs 4.19, $t(100) = 2.80$, $p < .01$).

Similarly, evaluations of probability depended on attribute order. The probability of winning was rated higher by participants who had first observed the favorable prize information than by participants who had not yet done so ($M = 2.54$ vs 1.96, $t(100) = 2.32$, $p < .05$). In keeping with Studies 1 and 2, these results suggest that participants' attribute evaluations were biased by their initial disposition toward the gamble, which was itself determined by whether the prize or probability information was seen first. This effect is even more striking in light of the quantitative, concrete nature of the attributes underlying the lottery.

In all three studies, attribute evaluations were measured directly during the information evaluation process. A possible concern is that the progress questions themselves might have increased participants' commitment to their initial evaluative disposition. In order to explore this possibility, we replicated Study 3 without the attribute evaluation questions. As an indirect measure of biased attribute evaluations, we examined the willingness of participants to pay for a chance to play the lottery. If, as we suggest, initial dispositions bias the evaluation of subsequent information, then willingness to pay should be higher when the cash prize attribute appears first than when the low probability of winning appears first.

Participants in the replication were 51 full-time MBA students at a large university. As before, participants were provided information about a hypothetical lottery with a favorable prize value ($200) and unfavorable likelihood of winning (.025), and the order of this attribute information was varied. Unlike Study 3, participants were not asked to evaluate the individual attributes. Instead, they simply reported their willingness-to-pay (WTP) for the lottery after reading all the information. In order to avoid scale compatibility effects (Tversky, Sattath, & Slovic, 1988), a nonmonetary, effort-based measure of WTP was created. Prior to the presentation of the lottery, participants were instructed to sum six columns of single-digit numbers (e.g., 1 + 0 + 1 + 3 + 2 + 3 + 0). Later, having reviewed the lottery information, participants provided their WTP by indicat-

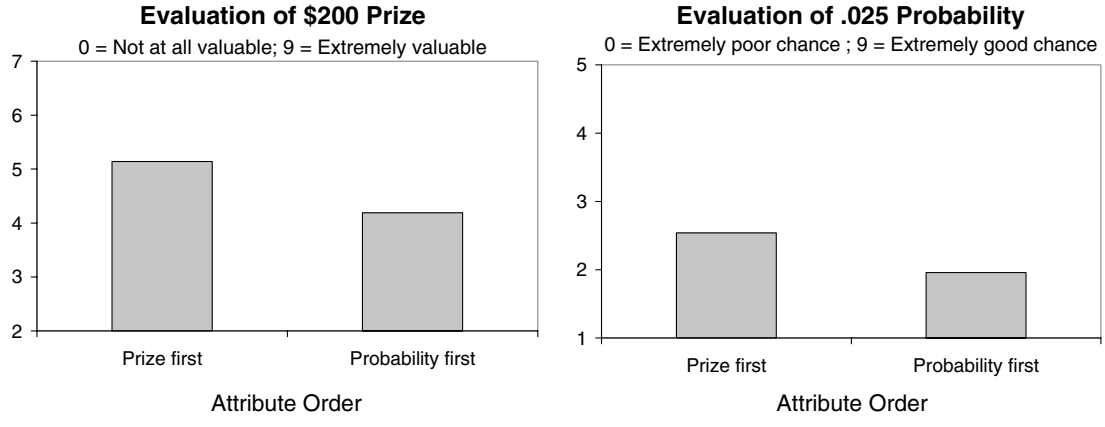

Fig. 1. Attribute evaluations for Study 3.

ing the number of such columns they would be willing to sum for a ticket.

Participants did not find the addition task difficult, and 48 out of 51 summed every column correctly (the analysis includes all participants). As expected, there was a significant effect of attribute order on WTP for the lottery ($t(49) = 2.20$, $p < .05$). Participants reported greater WTP when the favorable *prize* attribute appeared first ($M = 45.5$ columns) than when the less favorable *probability* attribute appeared first ($M = 26.2$). These data support our contention that even in the absence of attribute evaluation questions, participants established an initial disposition toward the lottery, leading to distortion of subsequent information and eventually impacting their assessment of the prospect's worth.

General discussion

Although the three studies spanned a variety of contexts, they present a similar picture of single-option evaluation. Central to each experiment was the emergence of a valenced disposition toward the option being considered, whether directly manipulated (in Studies 1 and 3) or appearing naturally during the evaluation process (in Study 2). In each case, this disposition led to a predictable bias in the evaluation of subsequent information. Study 1 revealed the operation of this bias in the context of a product purchase decision. The perceived favorability of the first information presented was predictive of both future information evaluations and final choice. Study 2 expanded the question by examining both choice and judgment tasks in a setting of interpersonal evaluation. Results showed a tendency for attribute perceptions to be distorted in favor of a naturally formed initial disposition, and this tendency was equally present in choice and judgment conditions. Study 3 demonstrated the impact of an evaluative disposition on the evaluation of a risky prospect, showing that even such unambiguous attributes as cash value and numeric probability are subject to biased interpretation. Together, these findings provide compelling evidence that the pursuit of coherence influences information processing during evaluation of a single option.[2]

[2] A potential concern in the present studies is that participants may have translated their decision tasks into binary form. Recall, however, that Study 2 examined this issue directly by utilizing judgment and choice conditions, and no difference was observed in the magnitude of bias across these groups. Also, the outcome measure of Study 3 (willingness-to-pay) was clearly treated as a continuous scale—only 3 of 51 participants reported a WTP of 0. Finally, even if one were to dichotomize the evaluation task of Study 1 into one of "buying" vs "not buying" a PDA, the latter is an imprecise option that aligns poorly with the attributes presented. To our knowledge, information distortion in such decisions has not been examined in prior work.

Relation to other primacy effects

The present studies measured the downstream effect of an evaluative disposition on choices and judgments involving a single option. The results in each case resembled traditional primacy effects in that the earliest information acquired about an option influenced final evaluations. Although we believe our studies provide interesting demonstrations of primacy in a singular evaluation setting, our principal concern has been to identify a process mechanism driving this phenomenon.

Earlier, we noted four general categories that have been used to explain the occurrence of primacy effects: biased search from the environment, biased search from memory, biased weighting, and distorted information evaluations. Because the last of these explanations has been the focus of this paper, it is important to consider whether the other three categories were relevant to the studies presented here. The first of these, *biased search from the environment*, is clearly inapplicable to the settings utilized in our studies. Participants were given no control over the attribute information they encountered, so explanations based on positive hypothesis testing (Klayman & Ha, 1987) or selective information exposure (Frey, 1986) are difficult to maintain.

Similarly, our studies are not easily reconciled with primacy accounts depending on *biased search from memory*. Because the options in all three studies were designed to be novel, even participants with experience in the task domains could not have possessed relevant knowledge about the particular options presented. Therefore, a biased retrieval explanation would need to focus on information acquired during the process itself. We focus here on anchoring and adjustment accounts, and to do so, we consider three different construals of the anchoring phenomenon.

In its classic interpretation, anchoring refers to the influence of random or uninformative starting points on judgments made in close temporal proximity (Chapman & Johnson, 1999; Tversky & Kahneman, 1979). Reconciling this idea with the present data would require a very complicated framework. In Study 2, for example, participants saw the same information (and presumably the same arbitrary anchors), in the same order, but their attribute evaluations diverged in a systematic manner. A second, somewhat broader construal of anchoring might suggest that relevant anchors were produced by participants themselves via their responses to the initial evaluation questions. If so, participants may have insufficiently adjusted from these anchors when determining their later evaluations. However, this account is also unsatisfactory, most notably in the follow-up to Study 3, where evaluation questions were not asked but the primacy effect was still observed. Finally, a very general view of anchoring might suggest that the evaluative disposition itself served as an implicit, internally generated starting

point for subsequent evaluations of the option. In fact, this broad construal is reasonably compatible with our framework. Note, however, that our focus is not simply on global evaluations of the option, but rather on the specific mechanism by which coherence is maintained between initial, tentative evaluations and newly encountered information. Therefore, in fitting the present studies to an anchoring framework, biased interpretation of attribute information should be recognized as a specific cause of underadjustment.

Finally, the primacy effects we demonstrate may seem suited to explanations involving *biased weighting of information*. Thus, the present studies can be considered in light of existing information integration models, particularly the belief-updating model of Hogarth and Einhorn (1992). A specific prediction of the belief-updating model is that for short tasks processed in a step-by-step manner, *recency* effects will be observed in final evaluations. Studies 1 and 2 fit this description: they consisted of only 5–6 attributes, and participants gave an updated summary evaluation after each attribute. Therefore, the belief-updating model may seem difficult to reconcile with our findings; in Study 1, for example, when the unfavorable *price* attribute occurred late in the sequence, participants were more (not less) likely to purchase the PDA. However, the discrepancy can be reconciled by looking at the attribute evaluations themselves. The *price* attribute was not evaluated negatively (on average) by those who saw it late in the sequence; instead, the distortion process was potent enough to make the attribute actually appear favorable. In keeping with the Hogarth and Einhorn (1992) model, it may well be that the *price* attribute was overweighted by these participants because of its serial position. Nonetheless, a weighting model alone would not predict our result, which relies on the fact that attribute evaluations are distorted to cohere with an existing disposition. This example reinforces our view that existing frameworks would benefit from attempts to incorporate the subjective encoding of information.[3]

Coherence vs other drivers of information distortion

In contrast to the explanations discussed above, we believe that our studies represent situations in which the *distortion of new information* is a fundamental mechanism underlying primacy effects. In doing so, this research adds to the growing body of work revealing. Further more, by extending past research on predecisional processing to include the evaluation of single options, our work bears specific implications for existing theories of information distortion. Even if a separation goal is critical in the forced-choice tasks utilized by most prior research, such a goal is largely inapplicable to the singular evaluation tasks presented here. Thus, the distortion observed in our studies cannot be explained as an attempt to create separation among alternatives.

An alternative position, supported by the present studies, is that the pursuit of coherence is a fundamental driver of predecisional information distortion. According to this perspective, decision makers form an initial evaluative disposition and bias their evaluation of subsequent information to cohere with that disposition. By utilizing singular evaluation settings, our studies provided a convenient means of distinguishing coherence- and separation-based accounts. However, it is important to note that explanations based on coherence are easily extended to multi-option settings; in fact, scholars have already proposed mechanisms by which coherence may contribute to predecisional "spreading" of alternatives (e.g., Holyoak & Simon, 1999). Rather than argue that separation-based approaches are inaccurate, we suggest that they may be better understood as contextually specific motivations (i.e., in multi-option tasks), serving a broader desire for coherence.

Testing the relative impact of separation and coherence motivations will require novel experimental designs and measurement methods, so we leave this issue to future investigation. Similarly, we have no way of knowing whether coherence operated in the present studies as an end in itself or as a conduit to other goals. For example, the pursuit of coherence might support the desire of individuals to feel certain about their actions (Mills, 1965), to reduce effort expenditure (Payne, Bettman, & Luce, 1998), or to bring about a state of cognitive closure (Kruglanski & Webster, 1996). The extent to which information distortion facilitates these and other decision goals requires further exploration.

Opportunities for future research

Given the robustness of information distortion observed in the present studies, it is tempting to ponder how the effect of an initial disposition might be mitigated. In binary choice settings, research has tested a number of potential moderators of information distortion, but these attempts have yielded little success. Expertise in the decision domain has not been found to eliminate distortion (Russo, Meloy, & Wilks, 2000), nor have specific instructions asking decision makers to reserve judgment (Carlson & Russo, 2001; Simon, Snow, et al., 2004). The effect has been demonstrated in important decision domains of personal relevance as well as low-involvement tasks with nothing at stake. Explora-

[3] In Study 2, the recency prediction of the belief-updating model can be reconciled with our findings by a similar argument. In Study 3, where summary evaluations were only assessed at the end of the process, the model predicts primacy to predominate, and this was in fact observed. We reiterate, however, that our focus is more on the information evaluations themselves than their downstream consequences.

tion of potential debiasing techniques would provide a valuable addition to this area.

Pertinent to this issue are a number of organizational studies which demonstrate recency effects in short, simple evaluation tasks (Farr, 1973; Highhouse & Gallo, 1997; London & Hakel, 1974). For example, Highhouse and Gallo (1997) asked trained raters to evaluate a candidate on the basis of two work samples. Ratings were requested on a variety of dimensions (oral communication, tolerance for stress, etc.), and the samples were constructed so that performance along these dimensions was positive in one sample and negative in the other. By manipulating the order in which the two samples were presented, the authors demonstrated a recency effect on final evaluations of the candidate. Note, however, an important contrast with our experiments: when observing the second work sample, participants in the Highhouse and Gallo (1997) study saw new information regarding all of the target dimensions, and this information directly contradicted that of the first work sample. In contrast, new information in our experiments always took the form of an entirely new attribute. Taken together, these results suggest that distortion may be inhibited for new information that specifically opposes earlier information on the same attribute. This possibility is worthy of further investigation.

Finally, the findings of Study 3 may be especially relevant to researchers in the field of risky decision making. In a straightforward lottery task, willingness-to-pay was found to depend on the order in which probability and value information were presented. The generalizability of these effects to other gambling scenarios is worth investigating, as their presence would have implications for the assessment of risk preferences.

Practical implications and precautions

The primacy mechanism that we describe presents a powerful tool for persuasion, and some practitioners may be aware of its potential. For example, a clear implication of this research is that communicators should try to create a favorable evaluative disposition as early as possible. As demonstrated in Study 1, positive initial information can lead to a "snowballing" effect, whereby subsequent attributes are distorted to align with an initially positive view. Such a technique is commonly observed in sales pitches, political speeches, etc., implying that some communicators understand the importance of initial disposition (if not the mechanism by which it operates). In addition, attempts may be made to "endow" a positive disposition outside the persuasion setting itself. For example, businesses may utilize a portfolio of well-liked products, successful advertising, or a positive corporate image to manipulate the disposition of customers even before specific information about a product is presented.

An important concern for communicators wishing to apply these principles is the extent to which the targets of influence sense that they are being manipulated. The "persuasion knowledge model" of Friestad and Wright (1994) suggests that individuals will respond defensively to a perceived persuasion attempt, acting to increase their control of the situation. Applied here, the necessary condition for such a defense is that individuals be aware of their tendency to distort information to support an early disposition. Recent evidence suggests that individuals lack this awareness during binary choices, even when the bias operates to favor an inferior alternative over a superior one (Russo et al., 2006). However, if targets could be made aware of their tendency to distort, then it might be possible to inhibit the effects we have observed. Moreover, if individuals become suspicious of a persuasion attempt, they may develop an unfavorable evaluative disposition, leading to negative distortion in the processing of new information. The interaction of persuasion knowledge and predecisional processing remains an open question.

Conclusion

Although this work focuses on the negative consequences of coherence-motivated processing, there may be some adaptive value to the bias. For example, if attributes are positively correlated in everyday decision tasks, then maintaining one's initial disposition may offer the benefit of efficiency with little downside. Alternatively, perceived coherence may provide benefits to mental well-being, either directly or through a feeling of confidence in one's judgment. However, in terms of accuracy alone, the systematic misinterpretation of attribute information is clearly nonnormative. In the studies presented here, individuals made predictable errors in person perception and product valuation. In more significant decisions, we would expect the same bias to operate, distorting evaluations of a job opportunity, a surgical procedure, or a prospective partner. By enhancing descriptive and prescriptive understanding of the distortion mechanism, researchers may be able to improve the process underlying such consequential decisions.

References

Adelman, L., Tolcott, M. A., & Bresnick, T. A. (1993). Examining the effect of information order on expert judgment. *Organizational Behavior and Human Decision Processes, 56*(3), 348–369.

Anderson, N. H. (1981). *Foundations of information integration theory*. New York: Academic Press.

Baron, J., Beattie, J., & Hershey, J. C. (1988). Heuristics and biases in diagnostic reasoning II: congruence, information, and certainty. *Organizational Behavior and Human Decision Processes, 42*, 88–110.

Baron, J. (1995). Myside bias in thinking about abortion. *Thinking and Reasoning, 7*, 221–235.

Barrows, H. S., Feightner, J. W., Neufeld, V. R., & Norman, G. R. (1978). *Analysis of the clinical methods of medical students and physicians*. Hamilton, Ontario: McMaster University School of Medicine.

Berwick, D. M., Fineberg, H. V., & Weinstein, M. C. (1981). When doctors meet numbers. *American Journal of Medicine, 71*, 991.

Boulding, W., Kalra, A., & Staelin, R. (1999). The Quality Double Whammy. *Marketing Science, 18*(4), 463–484.

Brownstein, A. L. (2003). Biased predecisional processing. *Psychological Bulletin, 129*(4), 545–568.

Cable, D. M., & Gilovich, T. (1998). Looked or overlooked? Prescreening decisions and postinterview evaluations. *Journal of Applied Psychology, 83*(3), 501–508.

Carlson, K. A., Meloy, M. G., & Russo, J. E. (2006). Leader-driven primacy: Using information order to affect choice. *Journal of Consumer Research, 32*(4), 513–518.

Carlson, K. A., & Russo, J. E. (2001). Biased interpretation of evidence by mock jurors. *Journal of Experimental Psychology: Applied, 7*(2), 91–103.

Chapman, L. J., & Chapman, J. P. (1969). Illusory correlation as an obstacle to the use of valid psychodiagnostic signs. *Journal of Abnormal Psychology, 74*, 271–280.

Chapman, G. B., & Johnson, E. J. (1999). Anchoring, activation, and the construction of values. *Organizational Behavior and Human Decision Processes, 79*(2), 115–153.

Ditto, P. H., & Lopez, D. F. (1992). Motivated skepticism: use of differential decision criteria for preferred and nonpreferred conclusions. *Journal of Personality and Social Psychology, 63*(4), 568–584.

Devine, P. G., & Ostrom, T. M. (1985). Cognitive mediation of inconsistency discounting. *Journal of Personality and Social Psychology, 49*(1), 5–21.

Doherty, M. E., Mynatt, C. R., Tweney, R. D., & Schiavo, M. D. (1979). Pseudodiagnosticity. *Acta Psychologica, 43*(2), 111–121.

Duckworth, K. L., & Bargh, J. A. (2002). The automatic evaluation of novel stimuli. *Psychological Science, 13*(6), 513–519.

Elliot, A. J., & Devine, P. G. (1994). On the motivational nature of cognitive-dissonance: Dissonance as psychological discomfort. *Journal of Personality and Social Psychology, 67*(3), 382–394.

Epley, N., & Gilovich, T. (2004). Are adjustments insufficient? *Personality and Social Psychology Bulletin, 30*(4), 447–460.

Evans, J. B. T., & Over, D. E. (1996). *Rationality and reasoning*. Hove: Psychology Press.

Farr, J. L. (1973). Response requirements and primacy-recency effects in a simulated selection interview. *Journal of Applied Psychology, 57*(3), 228–233.

Ferguson, M. J., & Bargh, J. A. (2003). The constructive nature of automatic evaluation. In J. Musch & K. C. Klauer (Eds.), *The psychology of evaluation: Affective processes in cognition and emotion* (pp. 169–188). Mahwah, NJ: Erlbaum.

Festinger, L. (1957). *A theory of cognitive dissonance*. Oxford: Row, Peterson.

Frey, D. (1986). Recent research on selective exposure to information. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 19, pp. 41–80). New York: Academic Press.

Friestad, M., & Wright, P. (1994). The persuasion knowledge model: how people cope with persuasion attempts. *Journal of Consumer Research, 21*(1), 1–31.

Galinsky, A. D., & Mussweiler, T. (2001). First offers as anchors: the role of perspective-taking and negotiator focus. *Journal of Personality and Social Psychology, 81*(4), 657–669.

Gilovich, T. D. (1991). *How we know what isn't so: The fallibility of human reason in everyday life*. New York: Free Press.

Hastie, R., & Park, B. (1986). The relationship between memory and judgment depends on whether the judgment task is memory-based or on-line. *Psychological Review, 93*(3), 258–268.

Heider, F. (1946). Attitudes and cognitive organization. *The Journal of Psychology, 21*(1), 107–112.

Highhouse, S., & Gallo, A. (1997). Order effects in personnel decision making. *Human Performance, 10*(1), 31–46.

Hinsz, V. B., Kalnbach, L. R., & Lorentz, N. R. (1997). Using judgmental anchors to establish challenging self-set goals without jeopardizing commitment. *Organizational Behavior and Human Decision Processes, 71*, 287–308.

Hoch, S. J., & Ha, Y. W. (1986). Consumer learning—advertising and the ambiguity of product experience. *Journal of Consumer Research, 13*(2), 221–233.

Hogarth, R. M., & Einhorn, H. J. (1992). Order effects in belief updating: the belief-adjustment model. *Cognitive Psychology, 24*(1), 1–55.

Holstein, J. A. (1985). Jurors' interpretations and jury decision making. *Law and Human Behavior, 9*(1), 83–99.

Holyoak, K. J., & Simon, D. (1999). Bidirectional reasoning in decision making by constraint satisfaction. *Journal of Experimental Psychology: General, 128*(1), 3–31.

Hsee, C. K., & Leclerc, F. (1998). Will products look more attractive when presented separately or together? *Journal of Consumer Research, 25*(2), 175–186.

Janis, I. L., & Mann, L. (1977). *Decision making: A psychological analysis of conflict, choice, and commitment*. New York: Free Press.

Jonas, E., Schulz-Hardt, S., Frey, D., & Thelen, N. (2001). Confirmation bias in sequential information search after preliminary decisions: an expansion of dissonance theoretical research on selective exposure to information. *Journal of Personality and Social Psychology, 80*(4), 557–571.

Klayman, J., & Ha, Y. W. (1987). Confirmation, disconfirmation, and information in hypothesis testing. *Psychological Review, 94*(2), 211–228.

Kristenson, H., & Garling, T. (1997). The effects of anchor points and reference points on negotiation process and outcome. *Organizational Behavior and Human Decision Processes, 71*(1), 85–94.

Kunda, Z. (1990). The case for motivated reasoning. *Psychological Bulletin, 108*(3), 480–498.

Kruglanski, A. W., & Webster, D. M. (1996). Motivated closing of the mind: "Seizing" and "freezing". *Psychological Review, 103*(2), 263–283.

Lawson, R. (1968). Order of presentation as a factor in jury persuasion. *Kentucky Law Journal, 56*(1), 523–555.

Lind, E. A., Kray, L., & Thompson, L. (2001). Primacy effects in justice judgments: testing predictions from fairness heuristic theory. *Organizational Behavior and Human Decision Processes, 85*(2), 189–210.

Lingle, J. H., & Ostrom, T. M. (1981). Principles of memory and cognition in attitude formation. In R. E. Petty, T. M. Ostrom, & T. C. Brock (Eds.), *Cognitive responses in persuasive communications: A text in attitude change* (pp. 399–420). Hillsdale, NJ: Erlbaum.

London, M., & Hakel, M. D. (1974). Effects of applicant stereotypes, order, and information order on interview impressions. *Journal of Applied Psychology, 59*(2), 157–162.

Lord, C. G., Ross, L., & Lepper, M. R. (1979). Biased assimilation and attitude polarization: the effects of prior theories on subsequently considered evidence. *Journal of Personality and Social Psychology, 37*(11), 2098–2109.

Macan, T. H., & Dipboye, R. L. (1990). The relationship of interviewers' preinterview impressions to selection and recruitment outcomes. *Personnel Psychology, 43*(4), 745–768.

Meloy, M. G., & Russo, J. E. (2004). Binary choice under instructions to select versus reject. *Organizational Behavior and Human Decision Processes, 93*(2), 114–128.

Mills, J. (1965). The effect of certainty on exposure to information prior to commitment. *Journal of Experimental Social Psychology, 1*(4), 348–355.

Montgomery, H. (1983). Decision rules and the search for a dominance structure: towards a process model of decision making. In P. C. Humphreys, O. Svenson, & A. Vari (Eds.), *Analysing and aiding decision processes* (pp. 343–369). Amsterdam: North-Holland.

Montgomery, H. (1994). Towards a perspective theory of decision making and judgment. *Acta Psychologica, 87*, 155–178.

Montgomery, H. (1999). Decision making and action: The search for a good structure. In H. Montgomery & P. Juslin (Eds.), *Judgment and decision making: Neo-Brunswickian and process-tracing approaches* (pp. 147–173). Mahwah, NJ: Erlbaum.

Mussweiler, T., & Strack, F. (1999). Hypothesis-consistent testing and semantic priming in the anchoring paradigm: a selective accessibility model. *Journal of Experimental Social Psychology, 35*(2), 136–164.

Nickerson, R. S. (1998). Confirmation bias: a ubiquitous phenomenon in many guises. *Review of General Psychology, 2*(2), 175–220.

Northcraft, G. B., & Neale, M. A. (1987). Experts, amateurs, and real estate: an anchoring-and-adjustment perspective on property pricing decisions. *Organizational Behavior and Human Decision Processes, 39*(1), 84–97.

Okada, E. M. (2005). Justification effects on consumer choice of hedonic and utilitarian goods. *Journal of Marketing Research, 42*(1), 43–53.

Payne, J. W., Bettman, J. R., & Luce, M. F. (1998). Behavioral decision research: An overview. In M. H. Birnbaum (Ed.), *Measurement, judgment, and decision making* (pp. 303–359). San Diego, CA: Academic Press.

Read, S. J., Vanman, E. J., & Miller, L. C. (1997). Connectionism, parallel constraint satisfaction processes, and gestalt principles: (Re)Introducing cognitive dynamics to social psychology. *Personality and Social Psychology Review, 1*(1), 26–53.

Ritov, I. (1996). Anchoring in simulated competitive market negotiation. *Organizational Behavior and Human Decision Processes, 67*(1), 16–25.

Russo, J. E., Carlson, K. A., & Meloy, M. G. (2006). Choosing an inferior alternative. *Psychological Science* (forthcoming).

Russo, J. E., Medvec, V. H., & Meloy, M. G. (1996). The distortion of information during decisions. *Organizational Behavior and Human Decision Processes, 66*(1), 102–110.

Russo, J. E., Meloy, M. G., & Wilks, T. J. (2000). Predecisional distortion of information by auditors and salespersons. *Management Science, 46*, 13–27.

Russo, J. E., Meloy, M. G., & Medvec, V. H. (1998). Predecisional distortion of product information. *Journal of Marketing Research, 35*(4), 438–452.

Sackett, P. R. (1982). The interviewer as hypothesis tester: the effects of impressions of an applicant on interviewer questioning strategy. *Personnel Psychology, 35*(4), 789–804.

Seta, C. E., & Hayes, N. (1994). The influence of impression formation goals on the accuracy of social memory. *Personality and Social Psychology Bulletin, 20*(1), 93–101.

Simon, D., Krawczyk, D. C., & Holyoak, K. J. (2004). Construction of preferences by constraint satisfaction. *Psychological Science, 15*(5), 331–336.

Simon, D., Snow, C. J., & Read, S. J. (2004). The redux of cognitive consistency theories: evidence judgments by constraint satisfaction. *Journal of Personality and Social Psychology, 86*(6), 814–837.

Sinclair, R. C. (1988). Mood, categorization breadth, and performance appraisal: the effects of order of information acquisition and affective state on halo, accuracy, information retrieval, and evaluations. *Organizational Behavior and Human Decision Processes, 42*(1), 22–46.

Srull, T. K. (1981). Person memory: some tests of associative storage and retrieval models. *Journal of Experimental Psychology: Human Learning and Memory, 7*, 440–463.

Strack, F., & Mussweiler, T. (1997). Explaining the enigmatic anchoring effect: Mechanisms of selective accessibility. *Journal of Personality and Social Psychology, 73*(3), 437–446.

Svenson, O. (1992). Differentiation and consolidation theory of human decision making: a frame of reference for the study of pre- and post-decision processes. *Acta Psychologica, 80*(1–3), 143–168.

Svenson, O. (1996). Decision making and the search for fundamental psychological regularities: what can be learned from a process perspective? *Organizational Behavior and Human Decision Processes, 65*(3), 252–267.

Thagard, P. (2000). *Coherence in thought and action*. Cambridge, MA: MIT Press.

Troutman, C. M., & Shanteau, J. (1977). Inferences based on nondiagnostic information. *Organizational Behavior and Human Performance, 19*, 43–55.

Tubbs, R. M., Gaeth, G. J., Levin, I. P., & Van Osdol, L. A. (1993). Order effects in belief updating with consistent and inconsistent evidence. *Journal of Behavioral Decision Making, 6*(4), 257–269.

Tversky, A., & Kahneman, D. (1979). Prospect theory: analysis of decision under risk. *Econometrica, 47*(2), 263–291.

Tverksy, A., Sattath, S., & Slovic, P. (1988). Contingent weighting in judgment and choice. *Psychological Review, 95*, 371–384.

van de Bos, K., Lind, E. A., Vermunt, R., & Wilke, H. A. M. (1997). How do I judge my outcome when I do not know the outcomes of others? The psychology of the fair process effect. *Journal of Personality and Social Psychology, 72*(5), 1034–1046.

Wallsten, T. S. (1981). Physician and medical student bias in evaluating diagnostic information. *Medical Decision Making, 1*, 145–164.

Wason, P. C. (1960). On the failure to eliminate hypotheses in a conceptual task. *Quarterly Journal of Experimental Psychology, 12*, 129–140.

Whyte, G., & Sebenius, J. K. (1997). The effect of multiple anchors on anchoring in individual and group judgment. *Organizational Behavior and Human Decision Processes, 69*(1), 75–85.

Willemsen, M. C., & Keren, G. (2004). The role of negative features in joint and separate evaluation. *Journal of Behavioral Decision Making, 17*(4), 313–329.