Journal of Experimental Psychology: General
2008, Vol. 137, No. 3, 456–470

0096-3445/08/$12.00 DOI: 10.1037/a0012786

The Goal of Consistency as a Cause of Information Distortion

J. Edward Russo
Cornell University

Kurt A. Carlson
Duke University

Margaret G. Meloy
Pennsylvania State University

Kevyn Yong
Cornell University

Why, during a decision between new alternatives, do people bias their evaluations of information to support a tentatively preferred option? The authors test the following 3 decision process goals as the potential drivers of such distortion of information: (a) to reduce the effort of evaluating new information, (b) to increase the separation between alternatives, and (c) to achieve consistency between old and new units of information. Two methods, the nonconscious priming of each goal and assessing the ambient activation levels of multiple goals, reveal that the goal of consistency drives information distortion. Results suggest the potential value of combining these methods in studying the dynamics of multiple, simultaneously active goals.

Keywords: biases, choice, decision process, goals, information distortion

An unexpected finding of decision research has been the biased evaluation of new information to support a developing preference. When one alternative naturally emerges as the tentative leader during the choice process, decision makers typically interpret new information to favor that leader more than they should. This bias is unexpected because of the absence of any commitment to the emerging leader, the kind of commitment presumed to be essential to explanations of cognitive dissonance (Brehm, 1956; Festinger, 1964). Festinger (1964) himself claimed that information distortion (ID) could not occur until after a decision had been made: "While we have as yet cast no light on what cognitive processes do go on during the period of making a decision, we at least know that this one does not occur" (p. 31). In contrast to such predictions, ID before the decision has been found by Holyoak and Simon (1999) and Russo, Medvec, and Meloy (1996), among others. It has even been observed, though not by that name, in the Bayesian analysis of data (Boulding, Kalra, & Staelin, 1999).

The present work addresses one question: Why does ID occur during the choice process? That is, what drives decision makers, who presumably want to make a good decision, to bias their evaluation of new information to favor a tentatively preferred option? To begin to answer this question, we turn to a description of the bias itself and to a common method for measuring it.

J. Edward Russo, Johnson Graduate School of Management, Cornell University; Kurt A. Carlson, Fuqua School of Business, Duke University; Margaret G. Meloy, Smeal College of Business, Pennsylvania State University; Kevyn Yong, Johnson Graduate School of Management, Cornell University.

This research was supported in part by National Science Foundation Grant SES-00112039 to J. Edward Russo and Margaret G. Meloy and by a Faculty Research Grant from the Smeal College of Business to Margaret G. Meloy. We thank Hans Baumgartner and Jim Bettman for comments on a draft of the article.

Correspondence concerning this article should be addressed to J. Edward Russo, Johnson Graduate School of Management, Cornell University, Ithaca, NY 14850-6201. E-mail: jer9@cornell.edu

Information Distortion (ID)

Consider deciding between two unfamiliar restaurants described by several units of information, like menus, parking, service, and so on. Each unit of information, hereafter termed an attribute, describes both restaurants in a brief paragraph like the menu attribute below.

> **Restaurant K**'s daily menu contains a variety of entrees, including poultry, beef, vegetarian items and pasta dishes. In addition, they always have two fish specials, one vegetarian special, and one or two pasta specials. **Restaurant D**'s main entrée menu consists of a few poultry and beef dishes, several pasta dishes, and two different dinner salads. They have many dinner specials each day, however, including vegetarian items, one pasta special, several fish specials, and numerous game dishes such as venison, rabbit, duck and quail.

To track the emergence of preference and its impact on the interpretation of a new attribute, we employ a method known as the stepwise evolution of preference (SEP). The SEP method captures each attribute's diagnostic value for each decision maker (for an overview, see Meloy & Russo, 2004). By comparing these values with an unbiased evaluation of the same information (determined from either a pretest or a control group), it is possible to measure ID. For example, for decision makers who favor Restaurant K after having seen the first attribute, ID is manifest as an evaluation of the second attribute to favor Restaurant K more strongly than is the true diagnostic value of that attribute. Similarly, those decision makers who favor Restaurant D after the first attribute would exhibit distortion by evaluating the second attribute as overly favorable to Restaurant D.

The distortion of information during a decision seems to be nearly ubiquitous. It has been observed not only in numerous

decisions by college students (e.g., DeKay, Patino-Echeverri, & Fischbeck, 2006; Holyoak & Simon, 1999; Russo, Meloy, & Medvec, 1998) but also in decisions made by entrepreneurs (Boyle, Hanlon, & Russo, 2008), horserace bettors (Brownstein, Read, & Simon, 2004), nonstudent adults (DeKay et al., 2006), pharmaceutical sales representatives and public auditors (Russo, Meloy, & Wilks, 2000; Wilks, 2002), and prospective jurors (Carlson & Russo, 2001). Indeed, ID has proven to be remarkably robust. For instance, it occurs in real choices (Carlson & Pearo, 2004) and for choices in which decision makers are financially incented to evaluate the information objectively and accurately (Meloy, Russo, & Miller, 2006). That said, it is not entirely impervious to contextual influences. In particular, ID has been reduced by requiring participants to consider their attribute preferences in advance of the choice process (Carlson & Pearo, 2004) and by getting decision makers to switch the task from choosing the better of two alternatives to rejecting the less attractive of the two (Meloy & Russo, 2004).

The phenomenon of ID is systematic, increasing linearly with the confidence in the leading alternative (Russo et al., 1998, 2000). It is also substantial in magnitude, about twice as large as the postdecisional distortion of the same information (Russo et al., 1998). ID occurs not just in choice but also in judgment, that is, in the evaluation of single objects (Bond, Carlson, Meloy, Russo, & Tanner, 2007; Russo & Yong, 2007). In such situations, new information is distorted to support the current disposition toward an object, favorable/accept or unfavorable/reject, just as it is distorted to support the currently leading alternative in a choice. Finally, two recent articles showed that ID has material consequence for choice accuracy. Carlson, Meloy, and Russo (2006) showed that beginning with information that favors one alternative induces decision makers to initially favor that target. Then, ID biases the evaluation of subsequent information to support that manipulated leader, making it the final choice substantially more often than it should be. Russo, Carlson, and Meloy (2006) went a step further and showed that even a self-identified inferior alternative can be installed as the initial leader, then supported by ID, and chosen a majority of the time.

The question motivating the present work is why ID occurs. Specifically, what purpose could such a bias possibly serve? Does ID make the decision easier or the process less effortful because each new attribute is interpreted as favoring the current preference? That is, might at least one driving goal of ID be the reduction of cognitive effort? Alternatively, might the answer be as simple as recognizing that the fundamental task of choice is to distinguish the best alternative from all others? If so, then the essential task of a binary choice is to create separation between the two alternatives. Distorting new information to support the leading option serves this immediate goal of creating separation, appearing almost designed to do so. Finally, might a goal of consistency lead to greater ID? Specifically, does a desire for consistency between new and old information lead to distortion of the new information so that its interpretation fits better with the old information? These possibilities are represented by three different decision process goals (conservation of effort, separation, and information consistency). Each is a possible influence on the choice process in general, and a potential driver of ID in particular. In Decision Process Goals, we consider information processing goals broadly. Then, we marshal the evidence from prior research that has suggested why the aforementioned goals are the most logical drivers of ID and, therefore, why they are the focus of our investigation. Our experimental studies follow.

Decision Process Goals

There are several, quite different types of goals that have been studied in the context of decision making. Sometimes a goal is a desire for an attribute or a level of an attribute describing the options (Markman & Brendl, 2000; McElroy & Seta, 2007). When buying a minivan, such attribute goals might include high gas mileage, good crash test scores, or the presence of a DVD player. Other goals are longer term, general objectives, like building financial security or maintaining health (e.g., Beach, 1998; Grouzet et al., 2005; Schneider & Barnes, 2003). Social goals, like cooperation and competition, deal with our interactions with, self-presentation to, or evaluations of others (Bargh, 1997; Bargh, Gollwitzer, Lee-Chai, Barndollar, & Trötschel, 2001). In contrast to these types of goals, decision process goals are desired states of the process and its immediate consequences, not benefits delivered by a chosen action or object. (For a recent review of goals, including a listing of goal types, see Carlson et al., in press.)

The minimization of effort is perhaps the most common of such decision process goals (Chakravarti & Janiszewski, 2003; Payne, Bettman, & Johnson, 1990). Other recognized process goals are avoiding regret (Heitmann, Lehmann, & Herrmann, 2007) or any other negative feeling (Luce, Payne, & Bettman, 2001), finding reasons (Shafir, Simonson, & Tversky, 1993; Simonson, 1989), and separating the alternatives to support a clearer choice (Svenson, 1996; also Montgomery, 1983). Note that the label *process goals* includes goals whose achievement (or lack thereof) is perceived both during and at the conclusion of the process.[1] Thus, the search for reasons or the avoidance of regret are active during the decision process but may be felt most strongly at its conclusion. In the present work, we examine different decision process goals as potential answers to our focal question of why ID occurs. In particular, we propose three of the goals mentioned above as candidate causes of ID: the conservation of effort, the separation of alternatives, and the desire to be consistent in the interpretation of information.

Conserve Effort

Distortion of new information can save effort in two ways, by helping to bring the choice task to an earlier conclusion and by interpreting information in a confirmatory way. The former possibility, which derives from the idea that the choice process ends when the decision maker breaches a threshold of distinction between the options (Bockenholt, Dietrich, Aschenbrenner, &

[1] Some researchers make a further distinction within what we call process goals. The latter term is reserved for goals that are achieved more or less entirely within the process itself and that are independent of any alternative. Examples are conserving effort and minimizing internal conflict. In contrast, *criterion goals* are associated with an alternative but are distinct from its benefits, such as finding reasons that can justify its choice (Van Osselaer et al., 2005). Although we respect such further refinements of the taxonomy of goals, our purposes do not require any additional distinction among decision process goals.

Schmalhofer, 1991; Busemeyer & Townsend 1993), is not possible in our experiments because participants are given a fixed number of attributes to process. Since participants in our studies must examine all of these attributes, there is no way to conserve effort by examining only a subset of them. The second possible role of ID in conserving effort presumes that an attribute that favors the leading alternative is processed more easily than one that opposes it. Possibly the conflict between the current leader and the opposing information evokes more careful analysis, especially when that analysis might involve a switch of the leading option. Data in Meloy et al. (2006) enabled a test of whether those participants who exhibited more distortion also processed attributes more quickly. The correlation between participants' mean ID and their time processing an attribute (adjusted for differences in the time to read each attribute by subtracting the mean times of a control group) should have exhibited a negative value. This correlation was −.195 (Fisher's r to Z transformation, $Z = .1984$, $z = 8.87$, $p < .001$). Thus, there is some evidence suggesting that ID and effort conservation are related. However, this evidence is only an association between distortion and total time processing the attributes, with the direction of causality undetermined. It is possible that, instead of the goal of conserving effort causing an increase in distortion, something else increases distortion which then results in faster processing time. That is, more distortion could result in less effort, not the other way around.

Separation

Choosing between alternatives can be viewed as a process of separation. That is, decision making is a process of distinguishing one alternative from the others on the basis of each one's overall value (Fischer, Carmon, Ariely, & Zauberman, 1999). Each new piece of information can be used to aid this subjective separation or discrimination between options (Svenson, 1996). Indeed, achieving separation is such an obvious goal that it almost amounts to a restatement of the choice task itself with a different verb. Nonetheless, it also seems like a natural candidate for the driving force of ID. Distortion to support the currently leading alternative serves perfectly the goal of increasing the separation between the leading and trailing alternatives.

Consistency

There is a long history of theories based on a desire for consistency among beliefs (Heider, 1958; D. Simon & Holyoak, 2002; Thagard, 2000). There is empirical evidence to suggest that consistency operates as a goal (a) whose activation can be altered through manipulation (Bator, 1994) and (b) whose chronic level of activation differs systematically across people (Cialdini, Trost, & Newsom, 1995). For example, the goal of consistency applies to the formation of attitudes about others (Read & Miller, 1993) and to the assembly of information to support a decision (D. Simon, Snow, & Read, 2004). Specifically, might decision makers distort new information so that it agrees with their tentative preference for one alternative over the other? Indeed, Holyoak and Simon (1999) explicitly proposed that the predecisional distortion of information is driven by the goal of consistency.

Research Strategy

To test whether specific goals drive ID, we use two methods: priming to increase goal activation (Experiment 1) and assessment of the ambient levels of goal activation during the choice process (Experiments 2 and 3). In the former, goals are activated via subliminal stimulation to exceed ambient levels. If a particular goal drives ID, then the level of distortion should be higher when that goal is primed. Note that because the priming is subliminal, participants are not aware of it. Thus, any impact of goal priming on ID cannot be an artifact of awareness that the goal was activated. Experiment 1 uses this priming method to test for an effect on ID of the three goals above: consistency, effort, and separation.

We note that others have also primed processing goals. For example, in an initial task Gollwitzer, Heckhausen, and Stellar (1990) activated either a deliberation mindset by instructions to consider a personal change or an implementation mindset by instructions to plan the execution of a personal project. They found that the instructed mindset impacted an unrelated second task, either completing a story or choosing between test materials. Chartrand and Bargh (1996) nonconsciously activated (by using the Scrambled Sentences Test priming technique of Srull & Wyer, 1979) either the goal of memorization or impression formation and then examined impacts on recall. Note that although the primes (i.e., the words in the scrambled sentences) were supraliminal, the postexperimental debriefing revealed no awareness of the priming manipulation, that is, no awareness that the Scrambled Sentences Test was related to the subsequent task to memorize or to form an impression. Experiment 1 is the first study to prime decision process goals and to explore the linkage between those goals and ID.

In Experiments 2 and 3, we assess goal activation levels dynamically throughout the choice process. The goal assessment technique captures activation levels for multiple goals via a retrospective report. We then determine which goals are the best predictors of ID by regressing attribute-level ID on the goal activation levels for those same attributes. Note that our uses of nonconsciously priming goals (in Experiment 1) and consciously reporting them (in Experiments 2 and 3) are not contradictory. For the former, participants need only be unaware that a decision process goal was primed. The level of awareness of the source of goal activation is distinct from, and in no way conflicts with, awareness of the primed goal's level of activation. For the latter, we need only that goal activation levels can be reported when a sufficiently rich set of cues is provided. That is, the degree of conscious awareness of one's own goals need only be high enough that prompting raises it to a reportable level. A higher degree of awareness, possibly high enough to be reported without prompting, is neither necessary nor problematic.

Experiment 1

This experiment is a straightforward test of the effect of activating specific goals on ID. If a particular goal is a cause of ID, then increasing its activation should increase ID. The tested goals include the three noted above, be consistent, conserve effort, and create separation, as well as a fourth goal that is unrelated to ID,

remember the options. Increased activation of the remember goal is not expected to change the magnitude of distortion in any way. As such, it serves as a control goal. Because all three of the relevant goals are plausible candidates for driving distortion, only the absence of an effect of an irrelevant goal can assure that any impact of the first three is not an artifact of the goal activation method itself.

To activate each of the four goals individually, we use the method of tachistoscopic presentation of goal-related words (e.g., Bargh & Chartrand, 2000). That is, each goal is subliminally primed by the very brief presentation of words that are related to it. Such priming requires a manipulation check that each goal is successfully activated. The procedure we use to verify the effectiveness of the priming manipulation is a word naming task in which goal-related words, displayed visually, should be orally named more quickly than are words unrelated to the same goal (Abad, Noguera, & Ortells, 2003; de Groot, 1985; Neely, 1976, 1991; VanVoorhis & Dark, 1995). The use of this procedure requires a holdout sample devoted exclusively to the naming task. Therefore, Experiment 1 is divided into two parts, the focal test in which priming is followed by the choice task and the manipulation check in which priming is followed by the naming task.

Method

Design. A between-participants design used five experimental conditions for the focal choice task. There were three conditions in which the relevant goals were primed, one in which the irrelevant goal was primed, and a fifth unprimed control condition. In each of the four goal-primed conditions, a single goal was primed immediately before the decision task. Participants in the fifth condition, the unprimed control group, had no goal primed before they completed the decision task. The unprimed control condition provided a baseline level of distortion against which the levels observed in the four goal-primed conditions could be compared. For the manipulation check, there were also five priming–naming conditions that mirrored the design of the main part of the experiment.

Task. Participants were told that the purpose of the first part of the study was to examine the effects of distraction on decision making performance. Thus, participants in the four goal-primed experimental and manipulation check conditions were asked to complete in their heads a series of 15 single-digit additions. All digits were shown in the center of a computer screen. As participants gazed at the numbers, very brief flashes of other information, totally irrelevant to the addition task, appeared randomly in one of four positions around the center of the screen. As explained below, these flashes were the words and the mask that were intended to subliminally activate the four goals (consistency, effort, separation, and remember).

The decision task was a choice between two restaurants described by six attributes. The introduction to the choice explained that participants had to select one of two "white tablecloth" restaurants for a dinner for two and that information had been gathered about these restaurants to help them make the choice. The six attributes (ambience, amenities, dining guide description, hours of operation, location, and menu) were verbal descriptions that included all of the information about both restaurants, like the attribute comparing menus shown above. The six attributes were presented in a randomly determined order. That order was reversed for half of the participants. (There were no effects of attribute order on ID, so it is not discussed further.)

The SEP method required participants to provide after each attribute three responses that captured the progress of the choice process. The first was an evaluation of the diagnosticity of the attribute on a 9-point scale, anchored on opposite ends by the names of the two restaurants. Thus, on this evaluation scale 1 and 9 both meant *strongly favors* one restaurant or the other, and 5 meant *favors neither restaurant.* The second and third responses revealed which of the two restaurants was leading in overall preference and by how much. The latter was essentially a measure of confidence that the leading restaurant "will be your final choice after all the information has been seen," reported on a scale from 50 (*a complete toss-up*) to 100 (*absolutely certain*). Thus, the SEP method yielded an evaluation of each attribute, the identification of the currently leading alternative, and the confidence that this leader would be the eventual choice.

In the manipulation check conditions, participants completed the same additions (which exposed them to the same primes) and then read aloud, as quickly as possible, the word appearing in the middle of the computer screen. These words stayed on the computer screen until participants made an oral response. The initiation of this utterance was detected by a microphone. The response time (RT) was computed as the time from the onset of the presentation of the word to the initial sound sensed by the microphone. To familiarize them with the task, participants initially completed a practice phase in which they read aloud 10 nongoal-related words such as *halo, orchard,* and *zinc.* Participants then proceeded to the test phase where they were presented target words for the four goal-primed conditions and the unprimed control condition.

Participants. Participants were 395 undergraduates at a large northeastern university who were paid $10 for completing the study. Because of incomplete responses, 35 participants were dropped. All analyses were performed on the remaining 360 participants, who were roughly equally distributed across the five priming-choice conditions (range = 53–62), with a holdout sample of 75 that also was approximately equally distributed across the five priming–naming conditions of the manipulation check (range = 14–16).

Materials. To activate the four goals, primes were needed. Selecting the words to use as primes involved a tradeoff between using a single prime that was central to the goal but that might have increased the risk of detection and using multiple primes that, because each was necessarily more distant from the goal, might not have as much power to activate it. For the consistency, effort, and remember goals, multiple primes were used. For the separate goal, for which it was difficult to find clearly associated candidate primes (e.g., *divide* and *distinguish*) in the specific context of a choice, we used *separate* itself as the single prime. The resulting primes were, for the consistency goal, *agreement, be coherent, be compatible, be congruent, be consistent,* and *well fitting.* For the effort goal, four primes were presented: *be efficient, be faster, be quick,* and

be speedy. Finally, for the remember goal, the primes were *memorize, recall, recognize, recollect, remind,* and *retain*.[2]

The naming task that checked for goal activation presented 6 associated words for each of the four goals. Six words irrelevant to all four goals were used for the unprimed control condition. The words used to test for the activation of the consistency goal were *agreement, coherence, compatible, congruence, consistent,* and *fitting*. For effort, the 6 words were *easy, efficient, effortless, faster, quickly,* and *speedy*. For separate, the corresponding words were *discriminate, distinguish, diverge, divide, separate,* and *split*. For remember, the words were *memorize, recall, recognize, recollect, remind,* and *retain*. Finally, for the unprimed control condition, the 6 words that were chosen to be irrelevant to all four goals were *above, collection, deepen, kitchen, overcome,* and *underline*. Efforts were made to keep word length approximately equal across the five conditions in the naming task. The 30 words were ordered so that each sequence of 6 consecutive words included 1 word from each of the five conditions, and no 2 words from the same condition were adjacent to each other. The control (no prime) group's RTs to this set of 30 words provided a baseline against which to compare RTs for those in the various goal-primed conditions.

Procedure. On arrival at the experimental laboratory, participants were randomly assigned to one of the five conditions in either the choice task or the manipulation check task and then were seated at individual computer stations. They were told that in order to study the effects that distractions can have on decision making, a "computer task" simulating a distraction would be used. Those in the choice conditions were then handed a packet containing the choice task and instructed to read the first two pages, which provided some general instructions and described a scenario in which they had to choose between one of two restaurants for dinner. Immediately after reading this cover story, participants were informed that they would be "distracted" by the ostensibly unrelated "computer task" involving the addition of numbers (which was used to accomplish goal priming). Those in the manipulation check conditions were treated identically, except that the priming was followed by the naming task. On completing the priming task, a message appeared on the computer screen instructing participants to alert the experimenter that they were finished with the distraction task. A 2-min delay was instituted as the experimenter ostensibly gathered the next set of materials, either for the continuation of the choice task or for the naming task. This delay was used to allow any purely semantic priming to dissipate, so that what remained was only the activation of a goal (e.g., Aarts, Custers, & Holland, 2007). To avoid suspicion, participants in the manipulation check conditions followed the naming task by completing the choice task, but these choice data were discarded as potentially tainted by the naming task.

In the cover story, participants were told that a sequence of numbers (between 1 and 9) would be briefly presented in the center of the computer screen, and their task was to add these numbers. After each sequence of 6 digits, participants were prompted to record the sum. Participants completed the first 3 sequences as practice trials followed by 12 more sequences, making a total of 90 digits seen. Participants were also told that during the addition task distractors in the form of random flashes would appear along with each number in one of four positions (above, below, right, and left). Participants were instructed to disregard these flashes as best they could while focusing their attention on the addition task. These apparently random flashes were, in fact, the subliminally presented primes. Each one appeared on the screen for 80 ms followed by a 50-ms mask (*MCXWQDK*); each was presented in one of four parafoveal locations at a visual angle between 2° and 6°. Note that because 1 prime occurred with every digit, 90 primes were shown in each condition. These were 15 replications of the six different phrases for the consistency and remember goals, 22–23 for the 4 primes of the effort goal, and, in the unique case of the separate goal, 90 presentations of the single word *separate*.

On completion of the addition task, participants in the choice conditions continued with the remainder of the restaurant choice while those in the manipulation check conditions proceeded to the naming task. The former read the six attributes in turn, responding after each to the three queries that formed the SEP method: (a) evaluation of the attribute's diagnosticity on the 1-to-9 scale, (b) identifying which of the two restaurants was leading, and (c) rating their confidence in that leading alternative. After responding to the last of the six attributes, participants made a final choice, accompanied by a rating of their confidence in that choice on a scale from 50 (*a complete toss-up*) to 100 (*absolutely certain*). On completing the decision task, participants responded to a questionnaire that included a funnel debriefing for suspicion and demographic questions.

The funnel debriefing involved a series of questions that were ordered from general to specific (e.g., Chartrand & Bargh, 1996). The questions revealed whether participants were suspicious of a relationship between the goal priming task and the choice process. First, participants answered an open-ended question about what they thought the purpose of the experiment was. Second, participants answered three yes–no questions asking them whether they were suspicious of the experimenter, of any tasks in the study, or of any information in the study. Third and last, participants were asked whether they noticed anything unusual as they progressed through the study and to indicate anything they may have noticed. Regarding the first question, none of the participants stated a theory about goals or about how goals influenced information processing. Moreover, no participant reported believing that the addition task was linked to the choice task in any way. Additionally, every participant answered "no" to the first two suspicion questions, and the very few who thought that the information might have been suspicious expressed reasons that were irrelevant to ID. Finally, there were no reports of "anything unusual" noticed during the study. We concluded that the priming/addition task was not perceived as related to the choice task and that the activation of the goals was genuinely subliminal.

Calculation of distortion. Each of the six attributes was written and pretested to be equivocal (i.e., to favor neither alternative),

[2] Such primes have usually been restricted to single words. Instead, we often used a two-word phrase for two reasons. First, phrases helped to clarify meaning. For instance, consistency has two common meanings: (a) coherence, correspondence, and lack of contradiction (what we intended); and (b) firmness, density, and viscosity (as in the consistency of pea soup). A phrase like *be congruent* resolved any uncertainty over the intended meaning of consistency. Second, the phrases often created a command, which made the prime more like a goal.

on average. That is, each attribute was designed so that it would achieve an average rating of 5 on the 9-point attribute evaluation scale. Although different participants could legitimately interpret an attribute as favoring either restaurant, pretests (conducted earlier with another student population) confirmed that average mean values fell in the range from 4 to 6 on the evaluation scale. As such, we concluded that the true unbiased diagnosticity of each attribute could be adequately approximated by the scale midpoint of 5. To compute the ID of a participant's rating of one attribute, we subtracted 5 from that participant's diagnosticity rating of the attribute. Then, the absolute difference was signed positively if the direction of the difference favored the restaurant that was leading before the attribute was evaluated, and it was signed negatively if the difference favored the restaurant that was trailing. For example, if the rating was 7 and the alternative that anchored the high end of the scale was leading prior to the attribute, then ID was computed to be +2 for that participant/attribute. If the restaurant that anchored the low end of the 1-to-9 scale was leading, then ID would have been –2. Note that distortion could not be calculated when there was no leader immediately prior to an attribute. Thus, there could not be any distortion of the first attribute, nor did we calculate distortion when the confidence rating of the previous attribute indicated no preference between the two restaurants (*50–50, a complete toss-up*).

Results

We computed participant-level distortions by averaging the attribute-level distortion for each participant for attributes two through six. As such, each participant provided one data point (a single distortion statistic) that captured his or her tendency to exhibit distortion during the choice process. To test the effect of each goal prime on ID required two steps. First, the benchmark level of ID was derived from the control group that received no goal priming before the restaurant choice. That value was .53, a level significantly above zero, $t(61) = 3.28$, $SE = .163$, $p < .01$. This magnitude of distortion meant that the average participant biased his or her evaluation of each attribute (second through sixth) .53 units on the 9-point diagnosticity scale to favor whichever option was leading.

The second step was the computation of distortion in the four goal-primed conditions. Those means are reported in Table 1, along with their differences from the control mean. A one-way analysis of variance revealed a significant main effect of condition, $F(4, 280) = 3.29$, $p < .05$. Planned comparisons using Dunnett's t test for the difference between the mean distortion of each primed group and that of the unprimed control group revealed a significant effect of priming only for one goal, consistency (difference = .64), Dunnett's $t(280) = 2.16$, $SE = .23$, $p < .05$ (two-tailed). Priming consistency led to a 119% increase in distortion. Finally, note that the unrelated goal, remember, did not yield a level of ID reliably different from that in the unprimed control condition (difference = .21), Dunnett's $t(280) = 0.87$, $SE = .24$, $p > .80$ (two-tailed). This null result removed the possibility that the priming technique itself, independent of which goal was primed, caused an increase in distortion.

Table 1
Distortion of Primed Goals in Study 1

Primed goal	Mean distortion	Difference from control
Consistency	1.17	0.64*
Effort	0.35	–0.17
Remember	0.74	0.21
Separate	0.79	0.26
Control (unprimed)	0.53	

* $p = .05$, two-tailed, Dunnett's t test of multiple means against a single (control) condition.

Manipulation check. The objective of the manipulation check was to verify that each prime had succeeded in activating the goal associated with it. Such activation should have appeared as faster mean RTs to naming goal-related words when primed than when not primed. The mean RTs in the primed and unprimed conditions, for each of the four goals, are presented in Table 2. One-tailed t tests of the differences in RT revealed reliably faster RTs as a result of goal priming for all four goals. For completeness, we also looked for reliable changes in RTs as a function of serial position of the 30 words and found no effect ($p > .50$). The priming was successful.

Discussion

The conclusions are straightforward. First, only the consistency goal seemed to drive ID. Priming the goals to conserve effort and create separation had no reliable effect on distortion. Second, the manipulation check showed that all four goals were successfully activated by their respective primes, eliminating the concern that failed priming may have accounted for the absence of increased ID for the effort goal and the separate goal.

Several considerations argue for caution in accepting the goal of consistency, and consistency alone, as the answer to our research question about what causes ID during a decision. First, the two candidate goals that were not empirically supported, effort and separate, seemed quite plausible a priori as drivers of ID. Rejecting them on the basis of one study seems premature. More important are several possible methodological reservations. Subliminal priming is a method new to the manipulation of decision process goals. It would seem desirable to use a second, different method and see whether it converges to the same result. An additional minor point is that the primed level of activation was presumably above what decision makers naturally experience. Maybe the consistency goal raised ID only because it was primed to a far-greater-than-normal level of activation. Such a concern would be allayed if the second method relied on the natural levels of goal activation.

A final reason for caution in accepting the above results is a concern commonly associated with the priming manipulation. Does priming have its effects on behavior through actual goal activation, as intended, or through the increased activation of some other nonmotivational construct, possibly no more than semantic associations (though the 2-min delay tried to eliminate that alternative) or something as powerful as a stereotype or a trait (Bargh, 1997)? We postpone discussion of whether a motivational versus a nonmotivational construct was activated by the primes in Experiment 1 until all the experimental evidence has been presented. We note here, however, that the issue

Table 2
Response Times in Naming Task (in Seconds) for Study 1

	Mean response time (ms)		
Primed goal	Primed	Control	Difference
Consistency	624	684	60†
Effort	574	640	76†
Remember	536	600	64†
Separate	587	693	106††

† $p < .05$, one-tailed. †† $p < .01$, one-tailed.

remains moot and that employing a method that does not rely on priming should prove useful.

For all of the above reasons, we sought a different method that could be used to confirm the above results. A complementary method would (a) permit direct measurement of the goal's activation rather than rely on an indirect approach (i.e., delay) and (b) would rely on the natural level, rather than an elevated level, of goal activation. Carlson (2001) has developed such a method, and we employed it in the next two studies.

Experiment 2

Experiments 2 and 3 retrospectively assess the activation of decision process goals in an attempt to further explore the relationship between the consistency goal and ID. The foundation of the method, called the retrospective assessment of goals (RAG), is a record of the choice process that serves as a memory aide to the decision maker. This record includes a concurrent verbal protocol, the original choice stimuli, and participants' responses to the three SEP questions after each attribute. With these as memory prompts, participants retrospectively recall and report (on a continuous scale) their activation levels of multiple goals at multiple points in the decision process. (Those multiple points correspond to the attributes of the alternatives.) The goal activation data obtained by this method enable a test of the relation between each goal and ID at the level of individual attributes.

Foundations for Retrospective Assessments of Goals

For the RAG method to work, it must be the case that multiple goals can be active simultaneously and that their activation levels fall along a continuum. Researchers often examine just one goal at a time. Especially when that one goal is primed, they naturally assume that active goals influence behavior and inactive (i.e., unprimed) goals do not (e.g., Bargh et al., 2001). However, there is increasing empirical evidence and theoretical acceptance that many goals can be active simultaneously (Carver & Scheier, 1990; Castelfranchi, Giardini, & Marzo, 2006; Louro, Pieters, & Zeelenberg, 2007; H. A. Simon, 1991, p. 362). Likewise, though goal activation is often treated as all or none, with pursued goals being active and nonpursued goals being inactive, a continuum of goal activation is also consistent with contemporary views of goal pursuit (e.g., Van Osselaer et al., 2005).

The claim that goal activation is continuous, with the lowest levels below conscious awareness, is compatible with metacognitive research, especially work on the feeling of knowing (Koriat, 2000; Nelson, 1996; Schwarz, 2004). Specifically, just as people can sense traces of knowledge that they cannot articulate, decision makers can retrospectively report the activation levels of goal traces when prompted but not when unaided. Thus, given that "the individual has partial (but not perfect) privileged access to his or her own idiosyncratic knowledge (Nelson, 1996, p. 107)," decision makers should also have privileged access to their own goal activation levels during a choice.

The success of the RAG method relies partly on the effectiveness of a replayed concurrent verbal protocol to cue the memory trace of the decision process. Note that though the collection of verbal protocols is familiar, as are techniques for training experimental participants to generate them (Ericsson & Simon, 1993), our use of a concurrent verbal protocol as a memory prompt is much less common. Nevertheless, there is evidence that verbal replay of a recorded episode is a valid way to help participants recall their thoughts and feelings from an episode (Ickes, Robertson, Tooke, & Teng, 1986). In fact, Carlson (2004) has shown that the protocol's replay is both a necessary and sufficient cue to memory. Without the participant's own verbal protocol as a prompt for recall, the goal activation levels cannot be validly retrieved. With the replay, however, goal activation levels are generally recoverable. Specifically, Carlson (2004) experimentally manipulated the activation of the consistency goal in a choice task. When participants were interrupted and asked for that goal's current activation level (without having generated a concurrent verbal protocol), they reported the expected difference between experimental conditions. Similarly, when asked retrospectively for the consistency goal's activation levels while listening to a recording of their verbal protocol, they reported the same expected difference between conditions. However, to eliminate demand accounts, observers were yoked to participants and listened to the corresponding participant's recorded verbal protocol. These observers were asked to assess the activation levels for different goals throughout each participant's choice process. The observers' estimated goal activation levels were not sensitive to the manipulation, presumably because the prompts did not activate valid memory traces.

The main purpose of Experiment 2 is to use the RAG method to test the relationship between goal activation and ID. For every attribute, the SEP method yields a measure of ID (calculated as in Experiment 1). For each of those same attributes, the RAG method yields activation levels for the several candidate goals. These data enable the regression of ID on goal activation levels. Whatever goals significantly predict ID are identified as the drivers of ID. That is, for each decision process goal, we examine changes in their level of activation across attributes to determine which are reliably correlated with the changes in ID across those same attributes.

Method

Design. Two conditions were needed in this experiment, a standard binary choice and a no-choice control. The latter enabled a precise measurement of the unbiased diagnostic value of each attribute (details provided below), which was deemed necessary because of the plan to examine attribute-level distortion.

Participants. Participants were 70 undergraduates at a large northeastern university who received course extra credit for participating in this study. Of these, 25 were randomly assigned to the no-choice control condition. One participant in the choice condition failed to report goal activation levels for all parts of the choice process. This left 44 choice participants on which all the analyses were performed. Choice participants were run individually, with a session lasting approximately 45 min.

Tasks and stimuli. Individuals in the binary choice condition selected between two pasta sauces for a meal with their roommates on the basis of three attributes (ingredients, nutrition, and taste). During the choice process, participants read all information aloud and spoke their thoughts aloud. These verbalizations were recorded for later replay (see below). As in Experiment 1, the attributes were narrative descriptions of the two pasta sauces. For example, the nutrition attribute was as follows:

> The nutritional label indicates that **Pasta Sauce M** has 80 calories in one 126 gram serving of the sauce. Each serving contains 20% of the Vitamin A and 27% of the Vitamin C recommended by the USDA as part of a nutritious daily diet. The nutritional label says that **Pasta Sauce Q** contains 22% of the Vitamin A and 25% of the Vitamin C recommended by the USDA as part of a nutritious daily diet. **Pasta Sauce Q** contains 85 calories in a 135 gram serving.

After a brief introductory scenario that explained the choice between the two pasta sauces, participants read the three attributes. After each one, they answered the three SEP questions used in Experiment 1. Finally, each participant selected their preferred pasta sauce.

Participants in the no-choice control condition also read the introduction and the three attributes and evaluated each attribute on a 1-to-9 scale. However, the cover story was altered slightly for this group so that every attribute described a different pair of pasta sauces by changing the attribute's identifying letters, like M and Q to R and C, and so on. Such prevention of any cumulative evaluation ensured that control participants did not transfer any distortion-inducing preference from the evaluation of one attribute to the evaluation of the next attribute. Note that these participants answered only the first of the three SEP questions, the attribute evaluation or diagnosticity, on the 1-to-9 scale. Changing the identifying letters of the two pasta sauces from attribute to attribute prohibited them from answering the last two questions (i.e., leader and leader confidence).

The goal assessment method. The RAG method consisted of three phases. First was a warm-up phase (see Ericsson & Simon, 1993, p. 240), during which participants read goal definitions and practiced speaking their thoughts out loud while describing a recent restaurant experience. The second phase involved the choice between the two pasta sauces, during which participants generated a concurrent verbal protocol that began with the cover story and continued through all three attributes. Third and finally, immediately after making their choices, participants listened to a recording of their individual verbal protocols while following along with the choice packet (which now contained their responses to the three SEP progress questions). While doing so, attribute by attribute, participants reported the activation level of each goal on a scale from 0 (*not at all active*) to 10 (*maximally active*). They responded to the instruction, "Think about how active each of the goals were in your mind as you considered the pasta sauces. Indicate the activation level for each goal by writing a number on the line before it, where 10 = Maximally Active, 0 = Not Active, and numbers between represent grades of activation."

In addition to the three candidate goals (consistency, effort, and separate) and the one unrelated goal from Experiment 1 (remember), nine other decision process goals were also measured throughout the choice process. These were the following: be certain, be complete (in the use of the available information), develop an impression (of each alternative), pick at random, enjoy the process, find reasons, learn my preferences, learn about the options, and pick the best. The use of a more complete set served two purposes. First, it provided a very general set of candidate goals that could be tested as potential drivers of ID (i.e., beyond the three tested in Experiment 1). They were derived from a review of the decision literature or generated as plausible decision process goals. The list was meant to include all reasonably common decision process goals. They were all tested by using the RAG method because, even though we could find no compelling reason to link any of the additional goals to ID, casting a wide net was not costly and might reveal a surprise. Note also that the priming method could have been used to test such a comprehensive set of goals, but only at a high cost in time and effort, driven by the need for a new group of participants for each goal to be tested. A second value from testing the extended set of goals was the diffusion of any focus on the consistency, effort, or separate goals. This reduced the chances that participants might infer which goals were the targets of the study.

Results

Unbiased diagnosticity and distortion. The no-choice control group's mean evaluations of the second and third attributes provided unbiased estimates of these attributes' true diagnostic values. (Recall that the first attribute could not be distorted because participants could have no preferred option when they encountered it, making its true diagnostic value immaterial.) The mean unbiased evaluations of the second and third attributes (nutrition and taste) were 4.59 and 5.76, respectively. With these means, we computed distortion as follows. For each attribute, we calculated the absolute difference between that attribute's evaluation (individually for each participant) and the unbiased evaluation of that attribute from the no-choice control group's mean rating. As in Experiment 1, we signed each deviation positively (or negatively) if it favored the leading (or trailing) pasta sauce. The average participant-level ID was .67 (on the 9-point evaluation scale). This value was reliably greater than zero, $t(44) = 3.83$, $SE = .174$, $p < .001$.

Predicting distortion by goal activation levels. The main objective of this experiment was to use the RAG method to identify which goals predicted ID. Our approach was to regress attribute-level distortion on goal activation levels of all 13 goals in a multivariate regression. The necessary regression analysis could have been executed in two ways, *lagged*, in which a goal's activation during the previous attribute predicted the ID of the current attribute, or *concurrent*, in which the goal's activation and ID were drawn from the same attribute. Both analyses may have been valid. Favoring the lagged approach was the support for causality provided by a time sequence: Because goals direct future behavior, goal activation at time $t - 1$ during the choice process should

influence behavior at time t. Because we had no evidence about when ID occurred during the processing of an attribute, the timing of causality could not be specified with precision. However, allowing some lag between goal activation and ID seemed appropriate. Favoring a concurrent analysis was the segment-based structure of the choice process and the RAG method's measurement of both goal activation and distortion within the same segment. The processing of these segments, which were based on the separate attributes, lasted for roughly 1 min each. Thus, it is possible that goal activation levels during an attribute could have influenced processing (and, thus, distortion) within that segment of the choice. Because we could find no a priori argument for rejecting either rationale, we computed both the one-period lagged regression model and the concurrent model.[3]

Lagged regression analysis. Attribute-level distortion was regressed on the one-interval lagged activation levels of the 13 goals across participants and attributes. This model revealed that activation of the consistency goal was the sole predictor of distortion of the next attribute (est. $\beta_{consistency} = .152$, $SE = .072$), $t(70) = 2.10$, $p < .05$. None of the remaining 12 goals was a significant predictor of distortion (all $ps > .05$). Moreover, neither of the other two candidate goals even had the correct sign (est. $\beta_{effort} = -.006$; est. $\beta_{separate} = -.208$). That is, the only goal whose activation level significantly influenced the evaluation (and distortion) of the immediately following attribute was the consistency goal.

The large number of goals used as predictors in this equation might have led to relatively poorly estimated coefficients. Specifically, if the goal activation levels were correlated, then the possibility existed that the standard errors on the regression coefficients for the separate and conserve effort goals were inflated. (Note that multicollinearity did not yield coefficients that were biased, but it could have undermined the ability to assess the exact relationship between these goals and ID.) As a test for correlated predictors, we computed the variance inflation factors for the coefficients in the model above. All factors were between 1.3 and 2.5 ($VIF_{effort} = 1.6$, $VIF_{consistency} = 1.3$, $VIF_{separate} = 2.2$), indicating that relatively little of the standard error in the estimates came from multicollinearity (Kennedy, 1994; Velleman & Welsch, 1981). Nevertheless, to examine the possibility that the null results for conserve effort and separate stemmed from correlated predictors, we estimated a lagged multiple regression analysis of ID on only the 4 goals examined in Experiment 1 (i.e., consistency, effort, separate, and remember). This yielded the same result as the full 13-goal regression equation. The consistency goal was a significant predictor of ID (est. $\beta = .123$), $t(79) = 1.97$, $SE = .062$, $p < .05$, while none of the other 3 goals reliably predicted ID (all $ps > .20$).

Concurrent regression analysis. We then estimated the concurrent regression model (i.e., the regression of ID on all 13 goals' concurrent activation levels). This model revealed results similar to the lagged model above. The consistency goal was positively associated (though without reaching a conventional level of statistical significance) with distortion of the current attribute in the 13-goal multivariate regression (est. $\beta = .115$), $t(70) = 1.30$, $SE = .088$, $p < .10$ (one-tailed). Again, none of the other goals was a significant predictor of ID (all $ps > .10$).

In all of the models above, attribute-level distortions consisted of two observations (i.e., Attributes 2 and 3) for each participant. A drawback of this approach was that it assumed that the two observations from each participant were independent. Unfortunately, the small number of observations per participant (2) and the large number of regressors (13) prevented us from testing for participant-level effects by estimating a full model with participant-level main effects and interactions (Lorch & Myers, 1990). Nevertheless, we could explore the possible failure of independence in an alternate manner. Specifically, for dependence among the observations to be a serious concern, it would have necessarily been the case that the attribute-level distortions were positively correlated within participants. However, when we tested for this correlation, it was not even positive ($r = -.008$, $p = .96$). That is, a participant's distortion of the second attribute did not predict his or her distortion of the third attribute. As such, we can conclude that the results above were not driven by participant-level effects.

Discussion

The most important conclusion from this experiment is that the findings of Experiment 1 were corroborated by using a different method. Thus, the claim that the consistency goal is a key driver of ID is bolstered by the convergent results from two different methods. Nonetheless, given the lack of experience with the RAG method, a replication seems prudent.

A follow-up study also offers the opportunity to delve deeper into participants' interpretation of the goal activation question. Reported activation levels are supposed to reflect how much participants intended to achieve each goal. However, some participants may have reported instead what they thought they had actually achieved. To more clearly distinguish between the reported goal activation levels as measures of intention, which is what they are supposed to capture, and measures of achievement, which may have been confused with intention, participants in Experiment 3 were asked to provide both measures.

Experiment 3

This experiment replicates Experiment 2 with three exceptions. First, we address the possibility that participants in Experiment 2 may have confounded two aspects of a goal's activation. As noted above, to address this possible confusion, Experiment 3 measures both intended and achieved activation for each goal throughout the process. We reason that a norm of nonredundancy will compel participants to distinguish the level of intended goal activation from the level of achieved goal activation.

Second, goal activation levels are collected for just three goals. None of the nine goals added in Experiment 2 to the four goals primed in Experiment 1 showed any systematic relation to ID. Further, trimming the set to three goals and replicating the results of Experiment 2 should eliminate the possibility that having so many goals had somehow undermined the participants' accurate

[3] Before estimating these regression models, we ran two analyses of variance (one on lagged distortion and one on concurrent distortion), with participant as a main effect and goal activation levels (concurrent and lagged) as covariates. The results revealed no effect of participant for either the lagged or concurrent goal activation levels (both $ps > .45$). As a consequence, the multivariate regressions were estimated without participant-level effects.

reporting of their activation levels. In addition to the consistency goal, we gave the separation goal a last chance to be revealed as a cause of ID. We also included an unrelated goal, using remember once again.

The final change from Experiment 2 is a switch of choice domain to two resort hotels that were described by six attributes, not three. The lengthened choice process enables more observations of ID, five instead of two.

Method

Participants. Participants were 60 undergraduates at a large northeastern university who were compensated with course extra credit. Each participant was randomly assigned to either the choice condition ($n = 35$) or to the no-choice control condition ($n = 25$). Participants in the choice condition were run individually, with each session lasting about 1 hr.

Task. As in Experiment 2, those in the no-choice control condition read and evaluated the six attributes (atmosphere, beach, location, pool, rooms, and service) under the scenario that each attribute described a different pair of hotels. Those in the choice condition read the stimuli aloud and spoke their thoughts into a voice recorder as they considered the six attributes in sequence en route to making a choice between the hotels (for a hypothetical spring break vacation with their friends). After reading each attribute, participants in this choice condition provided the three responses required by the SEP method.

On completing the choice process, participants listened to the recording of their verbal protocol and reported both their level of goal activation (meaning the intention to achieve the goal) and their level of goal achievement. The goal achievement question was worded, "How successful were you in achieving each of the following goals after you had evaluated this attribute? Use the 0 to 10 scale below, where 0 = not at all successful and 10 = extremely successful, to indicate how much success you had in achieving each of the three goals." The goal activation question was worded exactly as in Experiment 2 with the identical response scale. The essential phrase was simply, "Indicate the activation level for each goal by writing a number . . . ." Participants reported both measures for each of the three goals and for each of the seven stages of the choice process (viz., the cover story followed by the six attributes).

Results

The key evaluations needed to calculate distortion are those of the second through sixth attributes. The mean evaluations of the no-choice control condition for these attributes were atmosphere = 3.83; beach = 5.03; pool = 5.43; rooms = 5.77; services = 3.80. Using these means to calculate attribute-level distortion of the choice participants (and averaging over the five distortion measures for each participant and across participants), the mean participant-level distortion was 0.74. As in the two previous studies, this value was reliably greater than zero, $t(34) = 2.97$, $SE = .251$, $p < .01$.

Lagged and concurrent regression analyses. Next, we performed the lagged multiple regression of attribute-level distortion for Attribute$_t$ on the activation of the three goals on Attribute$_{t-1}$. As in Experiment 2, the only goal that significantly predicted distortion was the consistency goal (est. $\beta = .259$, $SE = .102$), $t(133) = 2.55$, $p < .01$. Activation levels of the other two goals showed no reliable relation to ID (both $ps > .75$), and the sign on the coefficient for the separation goal was again negative (the wrong direction). We also performed the concurrent regression of ID on the reported activation levels of the three goals. Again, only the consistency goal exhibited a significant relation to ID (est. $\beta = .291$), $t(153) = 2.87$, $SE = .101$, $p < .01$ (for both the separation and remember goals, $ps > .35$).

Intended versus achieved goal activation. Participants reported both intended and achieved goal activation levels to highlight the distinction between them and enable the assessment of intention unconflated with actual achievement. One source of evidence of conceptual independence between goal intention and achievement was the correlation between these measures. Although moderate-to-high values might have been expected (as the level of success in implementing intention produced equivalent achievement), low correlations for all three goals would have supported independence. Low values would also have eliminated any concern about a demand effect in which participants felt that they should report similar values for both measures. The correlations between normalized goal activation and goal achievement were remarkably low ($r_{separate} = .103$, $p = .218$; $r_{consistency} = .200$, $p = .001$; and $r_{remember} = .214$, $p = .009$). Thus, less than 10% of the variance in the goal activation (i.e., intention) was accounted for by goal achievement. In other words, participants successfully distinguished between their intention to achieve a goal and how much of it they actually achieved.

Although activation and achievement were weakly correlated, a priori they seemed to us so closely related that participants might have confused them. This discrepancy between expectation and evidence prompted us to ask: What was their relation to each other and the relation of each to ID? One possibility was that the effect of the consistency goal's activation on ID was mediated by the level of achievement of that goal. That is, the (intended) goal activation determined the degree of success or achievement; that level of achievement, in turn, determined the amount of ID that was observed. The design of the current study allowed us to test this activation–achievement–distortion model in which achievement mediated the effect of activation on ID. The necessary conditions to support such a model were (a) goal activation predicted ID, (b) goal activation predicted goal achievement, (c) goal achievement had a unique influence on ID, and (d) the effect of goal activation on ID weakened when goal achievement was added to the model. The first of these results was established in the multiple regression above, and it remained in the zero-order (i.e., concurrent) model in which ID is predicted by activation of the consistency goal ($\beta = .284$, $SE = .091$, $p < .002$). The second relationship was also supported via univariate regression ($\beta = .501$, $SE = .070$, $p < .001$). The third and fourth relationships were supported by the multiple regression of ID on activation and achievement of the consistency goal ($\beta_{activation} = -.016$, $SE = .097$, $p = .87$; $\beta_{achievement} = .556$, $SE = .095$, $p < .001$). Finally, a Sobel test provided additional support for the mediation model (Sobel = 4.545, $p < .001$). In summary, the influence of the consistency goal's (intended) activation on ID appeared to be fully mediated by degree to which this goal was achieved.

Discussion

Once again, only the consistency goal emerged as a driver of ID. Reducing the long list of goals from 13 goals in Experiment 2 to 3 goals in the current study had no substantive impact. Distinguishing the two interpretations of goal activation, intention (which we continue to call activation) and achievement, also had no effect on the role of consistency as the sole driver of ID. Indeed, goal achievement was not only a distinct construct for participants, but its role as a mediator between goal activation and the resulting ID was revealed. Thus, Experiment 2's main result was replicated.

The RAG method contrasts with the subliminal priming technique of Experiment 1. This new method measures rather than manipulates goal activation, and it measures these goal activations at ambient levels during distinct time segments of the choice process. The clear replication of the result obtained from the priming method supports the validity not only of that substantive result but also of the RAG method itself.

General Discussion

The main result of all three experiments is that the desire for consistency between units of information leads to the distortion of that information to support the leading alternative. That is, the tendency for individuals to bias their evaluation of new information to support an emerging preference is driven (at least in part) by the desire to see the separate units of information as consistent with each other. This role of consistency echoes work during the 1940s through the 1960s (and beyond) on the consistency between two beliefs and, especially, between an attitude and a related behavior (D. Simon & Holyoak, 2002).

Consistency as a Goal

Throughout this article, we have referred to this desire for consistency as a goal. Although we did not design our studies specifically to answer the question of whether the consistency construct is a goal, the combination of our studies and previously published work offer substantial relevant evidence. Three classes of alternative interpretations seem most pertinent, purely semantic priming, direct behavioral priming, and some form of construal. The first two of these are alternative accounts for Experiment 1, so we consider them first.

Semantic and behavioral priming. The idea behind the semantic priming account is that priming with a set of related words makes both the primed words and any words associated with them more accessible in memory. Thus, though the naming data from Experiment 1 are consistent with the interpretation that all four primed constructs are goals, they are equally compatible with a purely semantic priming account. Thus, the naming data are by themselves unable to differentiate between whether goal priming or semantic priming occurred in Experiment 1.

The evidence that impairs the semantic priming explanation is the increase in ID observed. The change in behavior suggests that the prime activated a goal (with its power to drive behavior) and not merely the meaning of consistency-related words. We do not intend to assert a general claim that all changes in behavior derive from goals. To the contrary, a goal's activation can be constant but something else, such as new information or a change of context, can drive the observed shift in behavior. However, given that the only difference between the consistency-primed and unprimed control conditions in Experiment 1 was priming, the presence of a change in an actual behavior (ID) infirms a semantic association explanation.

Another interpretation, behavioral priming, however, cannot be eliminated as readily. Bargh et al. (2001) suggested that word primes can activate a mediating construct that causes the behavioral change and that this construct need not be a goal. In their example, the mediator is a stereotype. That is, primed words activate a stereotype, such as that of a rude person or a polite person (Bargh, Chen, & Burrows, 1996), which then automatically elicits stereotype-compatible behavior. A major problem with applying a behavioral priming account when ID is the focal behavior is the difficulty of identifying likely candidates for that mediating construct. That is, in the context of ID, what is the parallel to stereotypes as the mediating construct, particularly when decision makers are unaware of ID when it occurs (e.g., Russo et al. 2006, 2000)? We find it implausible that the consistency primes, phrases like *agreement, be coherent,* and *well fitting,* directly altered ID without the intervening activation of the consistency goal. Nonetheless, our failure to find such a nongoal mediator is hardly proof that one does not exist. Thus, the data from Experiment 1 cannot definitively rule out the behavioral priming account. Doing so requires comparing the expected time course of the effect of such priming with the expected time course of goal activation.

Temporal pattern of activation. The long accepted temporal pattern of goal activation is persistence until the goal is satisfied, followed by a rapid decline. Lewin famously introduced the anecdote of waiters who remember many orders flawlessly, but only until the customers leave. One of his students, B. Zeigarnik, experimentally demonstrated this persistence–decay pattern in the effect that bears her name (Zeigarnik, 1967). Russo et al. (1998) traced the pattern of ID both before and after choice. Of course, ID is not the goal of consistency, but given the relation between them observed in the three present studies, the two should track each other temporally. Prior to choice, ID exhibits an increasing linear relation to the confidence in the leading alternative. The greater this confidence, the greater the ID of the next attribute. Because, on average, confidence increases with proximity to the choice, ID increases as well. After the choice has been made, ID drops precipitously and monotonically. Thus, the pattern of ID over time is compatible with the expected pattern of goal activation over time, namely a nearly immediate cessation of its activation on achievement. We note that the pattern reported in Russo et al. (1998) is not compatible with either semantic or behavioral priming, both of which are predicted to decay immediately after priming ceases, not after the choice is made.

One other piece of evidence informs the competition between the goal-based explanation of our results and the nonmotivational explanations based on semantic or behavioral priming. Recall that there was a 2-min interval between priming and the naming and choice tasks. This delay exceeds the few seconds typically assumed for semantic activation to decay (K. I. Forster, Booker, Schacter, & Davis, 1990; Higgins, Bargh, & Lombardi, 1985; Mckone, 1995). Because activation that remained after 2 min is unlikely to have been purely semantic, the observed effect on ID seems more compatible with the activation of the consistency

goal.[4] Again, however, we note that the argument is not conclusive. Semantic decay does not always occur within a few seconds. For instance, Srull and Wyer (1979) showed effects lasting 24 hr when a trait like hostile or kind was primed in a person perception task. Although the bulk of the evidence seems to favor the goal interpretation, data from priming studies, including ours, do not easily rule out alternative explanations.

Participant construal. Another account of our results that is not goal-based rests on participants' construals or understanding of the requirements of the experiment. This account holds that participants' desires to seek consistency are not direct motivation but rather are a means to achieving the higher order objective of behaving in accord with one's beliefs about what is normal, desirable, or appropriate in the eyes of others. If so, then the observed correspondence between distortion and both the consistency prime (in Experiment 1) and the measured consistency goal (in Experiments 2 and 3) would be, simply put, artifacts.

Although there is no way to ever know for certain what participants believed they should be doing in a particular study, the standard test for this form of construal is a funnel debriefing (Chartrand & Bargh, 1996). If participants were interested in behaving appropriately during the choice process of Experiment 1, then they should be willing to reveal what they believed was appropriate to the study administrator. However, as we noted, a comprehensive funnel debriefing revealed no evidence of this sort.

What do the data from Experiments 2 and 3 have to say about construal? Recall the reported levels of the consistency goal correlated significantly with ID in both experiments, but other goals did not correlate with ID (12 others in Experiment 2 and 2 others in Experiment 3). Why would the consistency goal alone have exhibited a systematic relation with ID? That is, how would participants know which of the goals "should" be linked to ID when evidence from past studies has found that participants are nearly completely unaware of ID (Russo et al. 2006, 2000)? There were numerous other plausible goals (e.g., separation and build confidence). Further, even if participants somehow knew the consistency goal was the "appropriate" goal, they would also need insight into how much they had distorted each piece of information. Only with such insight could participants alter the reported activation of the consistency goal accordingly and be aware of the dynamics of their distortion during the choice process to produce the results obtained by the lagged models. Specifically, when reporting the goal activation level on a particular attribute, they would have to be able to recall the extent to which they distorted the next attribute during the choice process, so that they could report an appropriate attribute-specific goal activation level for the consistency goal. We find this highly implausible.

At first glance, it might seem that the results of Experiments 2 and 3 are incompatible with the nonconscious goal priming of Experiment 1. However, there is an important distinction between activating a goal outside of awareness and its having an influence on behavior with the individual knowing the goal is active. The former simply means that the goal was activated by a technique that left the participant unable to link the goal's activation to the technique (as confirmed by the funnel debriefing). That is, in much the same way that moods can have unknown or "mysterious" origins (Chartrand, Van Baaren, & Bargh, 2006), we believe that active goals can also have unrecognized origins. Thus, the key requirement of the nonconscious priming used in Experiment 1 is that the participant be unaware of how the goal became active. We do not need the stronger requirement of nonconscious goal pursuit, namely that the goal influences behavior without individuals being aware that it is active. Indeed, the results of our last two studies indicate that participants can become aware of their goal activation levels with adequate prompting of memory, and these levels can be accessed and reported. With respect to such access, Nisbett and Wilson (1977), who are often credited with establishing the mantra that people cannot know their thoughts, were open to the idea that individuals might know their goals even though they did not know what caused their goals to become active: "An individual may know . . . that he was or was not pursuing a particular intention. An observer, lacking such private knowledge of content, might often be more prone to error in his assumptions about the causes of an individual's behavior than the individual himself" (Nisbett & Wilson, 1977, p. 256). It is this idea that goals can be nonconsciously activated and yet be somewhat accessible that reconciles the methods and results of Experiment 1 with those of Experiments 2 and 3 (i.e., that the consistency goal can be nonconsciously activated via priming yet its level of activation can be retrospectively recalled).

Consistency as a goal or a property of the cognitive system. If we have infirmed, though not excluded, the competing explanations for our results, might there still be another interpretation that leaves in doubt the status of consistency as a goal? We see one especially plausible alternative.

The unique role of consistency echoes work from the 1940s through present day on the consistency between two beliefs and, especially, between an attitude and a related behavior (D. Simon & Holyoak, 2002). Consistency as a concept seems so widespread and fundamental that it might be viewed as a property of the cognitive system, possibly one whose chronic strength differs stably across people (Cialdini et al., 1995). The view of consistency as a goal, though conceptually different from this alternative view, cannot be distinguished from it based on our experimental evidence. If consistency is a property of the cognitive system, there is no reason why its activation cannot be enhanced through priming and no reason why its level of activity cannot be detected and reported.

[4] Unfortunately one often used strategy for distinguishing genuine goal activation, a delay between the goal priming and the focal task (Bargh et al., 2001), was deemed inappropriate for our choice task. The delay is typically tied to an unrelated secondary task. This facilitates the decay of nonmotivational constructs (manifest as a decline in the effect of the prime on behavior), while increasing the activation of motivational constructs like goals (yielding the same or greater effect of the prime on behavior; Marsh et al., 1998). Unfortunately, this method presents a problem because many intervening tasks may alter the activation of the processing goals. An intervening task related to the primed goal might cause that goal's activation to decline if the individual feels sufficient progress was made toward the goal during the task (J. Forster, Liberman, & Higgins, 2005; Marsh et al., 1998). The activation of a primed goal might also decline for the opposite reason, namely that individuals disengage the goal because of failure to make progress on it during the intervening task (Carver & Scheier, 1990; Dweck & Leggett, 1988). Consider just the effort goal. It is essentially impossible to find an intervening task that is guaranteed not to affect its level of activation. As such, the delay approach was deemed inappropriate for verifying that our primes activated motivational constructs (i.e., goals).

Although our studies were not designed to test alternative explanations for the observed results, the evidence for interpreting those findings as driven by the consistency goal and not as semantic priming, behavioral priming, or a demand artifact (i.e., the construal account) seems persuasive without being conclusive. Nonetheless, only further experimentation can eliminate these competing accounts and explore the intriguing possibility that consistency is a property of the cognitive system.

Goal Dynamics

Our main finding, that the decision process goal of consistency drives ID, may be viewed as an endorsement of the value of a goal-based analysis of decision (and other) processes. As goals play a more central role in these analyses, two aspects of goal dynamics may become prominent. The first is the pattern of activation levels over the course of the decision process. For instance, when a goal has been achieved, it is generally agreed that its activation drops (e.g., Marsh, Hicks, & Bink, 1998). Alternatively, if progress in achieving a focal goal breaches a criterion threshold of "failure," individuals will direct resources from this goal to other goals, resulting in a decline in reported activation of the focal goal (Carver & Scheier, 1990). Carver and Scheier (1990) have been silent as to whether such redirection of resources can occur within a choice process. Perhaps this silence derives partly from the inability (to date) to measure goal activation levels throughout a choice process.

A second aspect of goal dynamics is the impact of one goal on another (Louro et al., 2007). Carlson (2006), in an unpublished study, calculated an intercorrelation matrix for the activation levels of multiple decision process goals. Of the observed significant correlations, most were unsurprising. For example, the goals of completeness (in the use of information) and conserve effort were negatively correlated, while the same goal of completeness was positively correlated with the consistency goal. More surprising, however, was the positive correlation between conserve effort and the goal to find reasons. Is the desire to minimize effort facilitated by finding clear reasons for choosing one alternative, suggesting that the effort goal drives the reasons goal? Alternatively, do the two merely co-occur, with no causal relation between them? Further study is needed to examine the correlation between the activation trace of Goal 1 with the lagged (by one time unit) trace of Goal 2.

Many questions about the time course of a goal's activation or about the interactions among two or more goals might be answered by the RAG method. Thus, the method's ability to track the time course of a goal's activation (at several intervals, not continuously) enables a test of a theoretically predicted decline of a goal's activation (e.g., Marsh et al., 1998) and, possibly, the measurement of that decline as well. Similarly, the method may help to answer such questions about the causal relationships among the simultaneously active process goals.

In summary, the theoretical view that multiple goals are active simultaneously and that their activation falls on a continuum (rather than being either fully active or not active at all) may be complemented by the RAG method. What seems especially promising is the tracing of goal activation dynamics, which may spur the generation of richer theories that are not only goal-based but also include hypotheses about the change over time in the activation of multiple goals.

Conclusion

The main objective of this work was to explain, at the level of decision process goals, why people who want to make good decisions bias their interpretation of information to support whichever option is currently leading. The result was surprisingly clear, namely the pursuit of the goal of consistency. To identify this cause of ID, two methods were used. Subliminal priming is a familiar method that has been underutilized in the domain of decision process goals. The RAG method is new, with some intriguing prospects for useful application. Thus, the conclusions from this work are both substantive and methodological, with both focusing on the role of goals.

Goals are fundamental to many explanations of decision and other cognitive phenomena (e.g., scientific discovery, Klahr & Simon, 1999). Yet the role of goals is typically confined to a post hoc explanation of the observed results. The equivalent goal-based hypotheses are not tested. Might one reason for the preponderance of theory over testing be the lack of methods for measuring the level of goal activation? The RAG method seems to offer one approach for measuring the ambient level of goal activation, and it can do so for the multiple goals that might be simultaneously active during processing. For many tasks, such measurement may enable testing of goal-based hypotheses that have never been verified before.

References

Aarts, H., Custers, R., & Holland, R. W. (2007). The nonconscious cessation of goal pursuit: When goals and negative affect are coactivated. *Journal of Personality and Social Psychology, 92,* 165–178.

Abad, M. J. F., Noguera, C., & Ortells, J. J. (2003). Influence of prime–target relationship on semantic priming effects from words in a lexical-decision task. *Acta Psychologica, 113,* 283–295.

Bargh, J. A. (1997). The automaticity of everyday life. In R. S. Wyer, Jr. (Ed.), *The automaticity of everyday life: Advances in social cognition* (Vol. 10, pp. 1–61). Mahwah, NJ: Erlbaum.

Bargh, J. A., & Chartrand, T. L. (2000). The mind in the middle: A practical guide to priming and automaticity research. In H. T. Reis & C. M. Judd (Eds.), *Handbook of research methods in social and personality psychology* (pp. 253–285). New York: Cambridge University Press.

Bargh, J. A., Chen, M., & Burrows, L. (1996). The automaticity of social behavior: Direct effects of trait concept and stereotype activation on action. *Journal of Personality and Social Psychology, 71,* 230–244.

Bargh, J. A., Gollwitzer, P. M., Lee-Chai, A., Barndollar, K., & Trötschel, R. (2001). The automated will: Nonconscious activation and pursuit of behavioral goals. *Journal of Personality and Social Psychology, 81,* 1014–1027.

Bator, R. J. (1994). *Priming a consistency motivation enhances cognitive dissonance effects.* Unpublished master's thesis, Arizona State University.

Beach, L. R. (1998). *Image theory: Theoretical and empirical foundations.* Mahwah, NJ: Erlbaum.

Bockenholt, U., Dietrich, A., Aschenbrenner, M., & Schmalhofer, F. (1991). The effects of attractiveness, dominance, and attribute differences on information acquisition in multi-attribute binary choice. *Organizational Behavior and Human Decision Processes, 49,* 258–281.

Bond, S. D., Carlson, K. A., Meloy, M. G., Russo, J. E., & Tanner, R. J.

(2007). Precommitment bias in the evaluation of single options. *Organizational Behavior and Human Decision Processes, 102,* 240–254.

Boulding, W., Kalra, A., & Staelin, R. (1999). The quality double whammy. *Marketing Science, 18,* 463–484.

Boyle, P. J., Hanlon, D., & Russo, J. E. (2008). *Distortion of information to support an entrepreneurial decision.* Unpublished manuscript, Central Washington University.

Brehm, J. W. (1956). Postdecision changes in the desirability of alternatives. *Journal of Abnormal and Social Psychology, 52*(3), 384–389.

Brownstein, A. L., Read, S. J., & Simon, D. (2004). Effects of individual expertise and task importance on pre-decision reevaluation of alternatives. *Personality and Social Psychology Bulletin, 30,* 819–904.

Busemeyer, J. R., & Townsend, J. T. (1993). Decision field theory: A dynamic-cognitive approach to decision making in an uncertain environment. *Psychological Review, 100,* 432–459.

Carlson, K. A. (2001). *In search of consumer goals.* Unpublished doctoral dissertation, Cornell University.

Carlson, K. A. (2004, February). *Measuring the activation of consumers' decision process goals.* Paper presented at the meeting of the Society for Consumer Psychology, San Francisco, CA.

Carlson, K. A. (2006). [Actor–observer differences in the reporting of decision process goals]. Unpublished raw data.

Carlson, K. A., Janiszewski, C., Keeney, R. L., Krantz, D. H., Kunreuther, H. C., Luce, M. F., Russo, J. E., et al. (in press). A theoretical framework for goal-based choice and for prescriptive analysis. *Marketing Letters.*

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

Carlson, K. A., & Pearo, L. K. (2004). Limiting predecisional distortion by prior valuation of attribute components. *Organizational Behavior and Human Decision Processes, 94,* 48–59.

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

Carver, C. S., & Scheier, M. F. (1990). Principles of self-regulation: Action and emotion. In R. Sorrentino & E. T. Higgins (Eds.), *Handbook of motivation and cognition: Foundations of social behavior* (Vol. 4, pp. 3–52). New York: Guilford Press.

Castelfranchi, C., Giardini, F., & Marzo, F. (2006). Relationships between rational decisions, human motives, and emotions. *Mind & Society, 5,* 173–197.

Chakravarti, A., & Janiszewski, C. (2003). The influence of macro-level motives on consideration set composition in novel purchase situations. *Journal of Consumer Research, 30,* 244–258.

Chartrand, T. L., & Bargh, J. A. (1996). Automatic activation of impression formation and memorization goals: Nonconscious goal priming reproduces effects of explicit task instructions. *Journal of Personality and Social Psychology, 71,* 464–478.

Chartrand, T. L., Van Baaren, R. B., & Bargh, J. A. (2006). Linking automatic evaluation to mood and information processing style: Consequences for experienced affect, impression formation, and stereotyping. *Journal of Experimental Psychology: General, 135,* 70–77.

Cialdini, R. B., Trost, M. R., & Newsom, J. T. (1995). Preference for consistency: The development of a valid measure and the discovery of surprising behavioral implications. *Journal of Personality and Social Psychology, 69,* 318–328.

de Groot, A. M. (1985). Word-context effects in word naming and lexical decision. *Quarterly Journal of Experimental Psychology: Human Experimental Psychology, 37A,* 281–297.

DeKay, M. L., Patino-Echeverri, D., & Fischbeck, P. S. (2006). *Distortion of outcome probability information in risky decisions.* Unpublished manuscript, Carnegie Mellon University.

Dweck, C. S., & Leggett, E. L. (1988). A social–cognitive approach to motivation and personality. *Psychological Review, 95,* 256–273.

Ericsson, K. A., & Simon, H. A. (1993). *Protocol analysis: Verbal reports as data* (rev. ed.). Cambridge, MA: MIT Press.

Festinger, L. (1964). *Conflict, decision, and dissonance.* Stanford, CA: Stanford University Press.

Fischer, G. W., Carmon, Z., Ariely, D., & Zauberman, G. (1999). Goal-based construction of preferences: Task goals and the prominence effect. *Management Science, 45,* 1057–1075.

Forster, J., Lieberman, N. N., & Higgins, E. T. (2005). Accessibility from active and fulfilled goals. *Journal of Experimental Social Psychology, 41,* 220–239.

Forster, K. I., Booker, J., Schacter, D. L., & Davis, C. (1990). Masked repetition priming: Lexical activation or novel memory trace? *Bulletin of the Psychonomic Society, 28,* 341–345.

Gollwitzer, P. M., Heckhausen, H., & Steller, B. (1990). Deliberative and implemental mind-sets: Cognitive tuning toward congruous thoughts and information. *Journal of Personality and Social Psychology, 59,* 1119–1127.

Grouzet, F. M. E., Kasser, T., Ahuvia, A., Fernandez Dols, J. M., Kim, Y., Lau, S., et al. (2005). The structure of goal contents across cultures. *Journal of Personality and Social Psychology, 89,* 800–816.

Heider, F. (1958). *The psychology of interpersonal relations.* New York: Wiley.

Heitmann, M., Lehmann, D. R., & Herrmann, A. (2007). Choice goal attainment and decision and consumption satisfaction. *Journal of Marketing Research, 46,* 234–250.

Higgins, E. T., Bargh, J. A., & Lombardi, W. (1985). The nature of priming effects on categorization. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 11,* 59–69.

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

Ickes, W., Robertson, E., Tooke, W., & Teng, G. (1986). Naturalistic social cognition: Methodology, assessment, and validation. *Journal of Personality and Social Psychology, 51,* 66–82.

Kennedy, P. (1994). *A guide to econometrics.* Cambridge, MA: MIT Press.

Klahr, D., & Simon, H. A. (1999). Studies of scientific discovery: Complementary approaches and convergent findings. *Psychological Bulletin, 125,* 524–543.

Koriat, A. (2000). The feeling of knowing: Some metatheoretical implications for consciousness and control. *Consciousness and Cognition, 9,* 149–171.

Lorch, R. F., & Myers, J. L. (1990). Regression analyses of repeated measures data in cognitive research. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 16,* 149–157.

Louro, M. J., Pieters, R., & Zeelenberg, M. (2007). Dynamics of multiple-goal pursuit. *Journal of Personality and Social Psychology, 93,* 174–193.

Luce, M. F., Payne, J. W., & Bettman, J. R. (2001). The emotional nature of trade-offs. In S. J. Hoch & H. C. Kunreuther (Eds.), *Wharton on making decisions* (pp. 17–35). New York: Wiley.

Markman, A. B., & Brendl, C. M. (2000). The influence of goals on value and choice. In D. L. Medin (Ed.), *The psychology of learning and motivation: Advances in research and theory* (Vol. 39, pp. 97–128). San Diego, CA: Academic Press.

Marsh, R. L., Hicks, J. L., & Bink, M. L. (1998). Activation of completed, uncompleted and partially completed intentions. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 24,* 350–361.

McElroy, T., & Seta, J. J. (2007). Framing the frame: How task goals determine the likelihood and direction of framing effects. *Judgment and Decision Making, 2,* 251–256.

Mckone, E. (1995). Short-term implicit memory for words and non-words. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 21,* 1108–1126.

Meloy, M. G., & Russo, J. E. (2004). Binary choice under instructions to

select versus reject. *Organizational Behavior and Human Decision Processes, 93,* 114–128.

Meloy, M. G., Russo, J. E., & Miller, E. (2006). Monetary incentives and mood. *Journal of Marketing Research, 43,* 267–275.

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

Neely, J. H. (1976). Semantic priming and retrieval from lexical memory: Evidence for facilitatory and inhibitory processes. *Memory & Cognition, 4,* 648–654.

Neely, J. H. (1991). Semantic priming effects in visual word recognition: A selective review of current findings and theories. In D. Besner & G. W. Humphreys (Eds.), *Basic processes in reading: Visual word recognition* (pp. 264–336). Hillsdale, NJ: Erlbaum.

Nelson, T. O. (1996). Consciousness and metacognition. *American Psychologist, 51,* 102–116.

Nisbett, R. E., & Wilson, T. D. (1977). Telling more than we can know: Verbal reports on mental processes. *Psychological Review, 84,* 231–259.

Payne, J. W., Bettman, J. R., & Johnson, E. J. (1990). The adaptive decision maker: Effort and accuracy in choice. In R. M. Hogarth (Ed.), *Insights in decision making: A tribute to Hillel J. Einhorn* (pp. 129–152). Chicago, IL: The University of Chicago Press.

Read, S. J., & Miller, L. C. (1993). Rapist or "regular guy": Explanatory coherence in the construction of mental models of others. *Personality and Social Psychology Bulletin, 19,* 526–540.

Russo, J. E., Carlson, K. A., & Meloy, M. G. (2006). Choosing an inferior option. *Psychological Science, 17,* 899–904.

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.

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

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., & Yong, K. (2007). *The distortion of information to support an emerging evaluation of risk.* Manuscript submitted for publication.

Schneider, S. L., & Barnes, M. D. (2003). What do people really want? Goals and context in decision making. In S. L. Schneider & J. Shanteau (Eds.), *Emerging perspectives on judgment and decision making research. Cambridge series on judgment and decision making* (pp. 394–427). New York: Cambridge University Press.

Schwarz, N. (2004). Metacognitive experiences in consumer judgment and decision making. *Journal of Consumer Psychology, 14,* 332–348.

Shafir, E., Simonson, I., & Tversky, A. (1993). Reason-based choice. *Cognition, 49,* 11–36.

Simon, D., & Holyoak, K. J. (2002). Structural dynamics of cognition: From consistency theories to constraint satisfaction. *Personality and Social Psychology Review, 6,* 283–294.

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,* 814–837.

Simon, H. A. (1991). *Models of my life.* New York: Harper Collins.

Simonson, I. (1989). Choice based on reasons: The case of attraction and compromise effects. *Journal of Consumer Research, 16,* 158–174.

Srull, T. K., & Wyer, R. S. (1979). The role of category accessibility in the interpretation of information about persons: Some determinants and implications. *Journal of Personality and Social Psychology, 37,* 1660–1672.

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,* 252–267.

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

Van Osselaer, S., Ramanathan, S., Campbell, M. C., Cohen, J., Dale, J. K., Herr, P. M., et al. (2005). Choice based on goals. *Marketing Letters, 16,* 335–346.

VanVoorhis, B. A., & Dark, V. J. (1995). Semantic matching, response mode, and response mapping as contributors to retroactive and proactive priming. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 21,* 913–932.

Velleman, P. F., & Welsch, R. E. (1981). Efficient computing of regression diagnostics. *The American Statistician, 35,* 234–242.

Wilks, J. (2002). Predecisional distortion of evidence as a consequence of real-time audit review. *The Accounting Review, 77,* 51–71.

Zeigarnik, B. V. (1967). On finished and unfinished tasks. In W. D. Ellis (Ed.), *A sourcebook of Gestalt psychology* (pp. 300–314). New York: Humanities Press.

Received September 4, 2007
Revision received May 6, 2008
Accepted May 6, 2008 ■