

Consumers' Beliefs about Product Benefits: The Effect of Obviously Irrelevant Product Information

TOM MEYVIS
CHRIS JANISZEWSKI*

When consumers try to assess the performance of a product on a key benefit, their information search often reveals both diagnostic information and irrelevant information. Although one would expect irrelevant information to have little impact on predictions of product performance, we present evidence that the irrelevant information systematically weakens consumers' beliefs that the product will provide the benefit. We show that this dilution effect persists after subjects have acknowledged the irrelevance of the additional information but that it does depend on whether the product information is processed with the desired benefit in mind. We conclude that consumers are selectively looking for information that suggests the product will deliver the desired benefit and that they categorize any additional evidence, be it irrelevant or disconfirming, as not confirming. As a consequence, irrelevant information weakens consumers' beliefs in the product's ability to deliver the benefit.

When consumers evaluate products or services, they often search for diagnostic information on specific product benefits. A natural consequence of this search for diagnostic information is exposure to obviously irrelevant information—information that consumers perceive as clearly uninformative about the desired benefit. Normatively, the presence of irrelevant information should not change consumers' assessment of the product's ability to deliver the desired benefit. Yet, studies on social judgment have demonstrated that adding obviously irrelevant information to diagnostic information may lead to less extreme judgments (e.g., De Dreu, Yzerbyt, and Leyens 1995; Fein and Hilton 1992; Nisbett, Zukier, and Lemley 1981). This so-called dilution effect suggests that obviously irrelevant product information may weaken consumers' belief in a product's ability to deliver a desired benefit.

Irrelevant product information may dilute consumers' beliefs through four different mechanisms. First, the averaging hypothesis proposes that adding irrelevant product infor-

mation reduces the weight consumers assign to the supportive information and, thus, weakens their belief in the product benefit (e.g., Anderson 1974). A second possibility is that consumers use the irrelevant product information because they are relying on conversational norms and assume that all information provided is relevant for assessing the product's ability to deliver the benefit (e.g., Grice 1975). Third, consumers may rely on the representativeness heuristic, using the similarity between the target product and typical products that provide the benefit to estimate the probability that the target product will also deliver the benefit (e.g., Kahneman and Tversky 1972, 1973). Thus, to the extent that irrelevant information makes a product less similar to typical products that provide the benefit, consumers should weaken their belief in the product benefit. Finally, consumers may also treat the idea that the product may deliver the benefit as a focal hypothesis, which they test in a biased fashion (e.g., Snyder and Swann 1978). Although most studies of biased hypothesis testing have demonstrated that adding nondiagnostic information increases belief strength (e.g., Hoch and Ha 1986), some researchers in the biased hypothesis testing literature have suggested that nondiagnostic information can also weaken beliefs (Fischhoff and Beyth-Marom 1983; Sanbonmatsu et al. 1998; Sanbonmatsu, Posavac, and Stasney 1997).

Although the first three mechanisms have received ample attention in the dilution literature, the biased hypothesis testing perspective has been largely ignored. Moreover, there

*Tom Meyvis is assistant professor of marketing at the Stern School of Business, New York University, New York, NY 10012-1126 (tmeyvis@stern.nyu.edu); Chris Janiszewski is the Jack Faricy Associate Professor of Marketing at the Warrington College of Business, University of Florida, Gainesville, FL 32611-7155 (chrisj@dale.cba.ufl.edu). The authors wish to thank Joe Alba and Alan Cooke for their helpful comments on an earlier draft of the article. This article has also benefited from insightful suggestions by Bob Wyer, John Lynch, the editor, associate editor, and the reviewers.

has yet to be a systematic test among the four mechanisms as competing explanations of the dilution effect. Therefore, we will formulate a detailed biased hypothesis testing account of the dilution effect and systematically test this new account against the three existing explanations. Together, the results of 10 experiments suggest that consumers first generate the hypothesis that the product will deliver the desired benefit and then selectively search for evidence that supports this hypothesis, classifying all information with respect to their search goal. Supportive information is classified as confirming, while any additional information is classified as not confirming. Consumers' beliefs in the hypothesis will be stronger (weaker) to the extent that their search for confirmatory evidence produces more confirming (not confirming) evidence.

THE EFFECT OF IRRELEVANT PRODUCT INFORMATION

Several studies in the consumer behavior literature have demonstrated that objectively irrelevant product information can influence consumer decisions. For instance, Hoch and Ha (1986) observed that a nondiagnostic, ambiguous product experience can increase the perceived quality of an advertised brand. Similarly, Carpenter, Glazer, and Nakamoto (1994) reported that a brand with a distinguishing but irrelevant attribute received a higher preference rating than the same brand without the attribute. In contrast, Simonson, Carmon, and O' Curry (1994) have shown that consumers are less likely to choose a brand that offers a promotion or feature that has no value to them. Finally, Brown and Carpenter (2000) replicated the findings of Carpenter et al. (1994) and Simonson et al. (1994) and demonstrated that the direction of the influence of the irrelevant information partially depends on the size of the choice set.

Although consumer behavior studies demonstrate that objectively irrelevant information influences consumer decisions, they may not tell us much about the influence of obviously irrelevant information on consumers' beliefs about product benefits. First, while the information presented in the preceding studies was objectively irrelevant, subjects did not perceive the information as irrelevant and they knowingly relied on it. Subjects in Hoch and Ha (1986) perceived the ambiguous product experience as diagnostic information and interpreted their experiences as confirming advertising claims. Similarly, although subjects in Carpenter et al. (1994) could infer that the differentiating attribute was irrelevant, they still used it to make positive quality inferences (see also Broniarczyk and Gershoff 1997). Likewise, subjects in Simonson et al. (1994) identified the irrelevant, unwanted feature as a reason for not choosing the target brand, and subjects in Brown and Carpenter (2000) referred to the trivial information to justify their choices.

Second, previous studies have failed to examine product-specific judgments, such as beliefs in a product's ability to deliver a particular benefit. With the exception of Hoch and Ha (1986), all previous studies examined the influence of

irrelevant information in a choice or preference context. This distinction is important because the choice context is essential for explaining the effect of the irrelevant information. The irrelevant information influences consumer decisions because it differentiates the target brand from the other brands (Carpenter et al. 1994) and, thus, it provides consumers with a justification for their decision and a way to resolve the choice conflict (Brown and Carpenter 2000; Simonson et al. 1994). Subjects' choice protocols (Brown and Carpenter 2000; Simonson et al. 1994) and the moderating effect of the choice context (Brown and Carpenter 2000) provide strong support for these interpretations.

In sum, although previous findings clearly demonstrate that objectively irrelevant product information can influence consumer choices and preferences, they do not inform us whether specific benefit judgments can be influenced by information that consumers perceive as obviously irrelevant with respect to these benefits. A more promising perspective on this research question may be offered by the psychological literature on the dilution effect.

THE DILUTION EFFECT

Dilution occurs when a person's consideration of irrelevant information leads to a less extreme judgment. The dilution effect has been found using both nonsocial stimuli with probability judgments and social stimuli with natural judgments. For example, Troutman and Shanteau (1977) showed subjects draws of beads from one of two boxes with different proportions of red and white beads but equal proportions of blue beads. Subjects moderated their estimate of the probability that the red box was being sampled when a draw of red beads (i.e., diagnostic information) was followed by a draw of blue beads (i.e., irrelevant information). Similarly, Zukier and Jennings (1983) found that jurors were less likely to find a man guilty of murdering his aunt when diagnostic information (e.g., "He was known to have argued with his aunt") was supplemented with irrelevant information (e.g., "The defendant is of average height and vision").

Explanations of the dilution effect include averaging models, conversational norms, and the representativeness heuristic. We will first review these existing explanations and then propose a new explanation based on the biased hypothesis testing literature.

Averaging

The averaging explanation proposes that the irrelevant product information reduces the impact of the supportive information. Averaging is the most popular account of the dilution effect within the nonsocial judgment literature and can be classified into two groups of models (Birnbaum and Mellers 1983; Lichtenstein, Earle, and Slovic 1975; Shanteau 1975; Troutman and Shanteau 1977). A first group of algebraic averaging models assumes that each attribute's weight is adjusted according to the weights of the other attributes being considered (e.g., Anderson 1971, 1974; Birnbaum and Mellers 1983). Therefore, if the irrelevant

information receives a nonzero weight, adding irrelevant information can weaken the impact of the diagnostic information, thus diluting people's responses. A second set of averaging models assumes that people make separate predictions based on each piece of information and average these predicted outcomes (e.g., Lichtenstein et al. 1975; Lopes 1987; Shanteau 1975). For example, Shanteau (1975) told subjects there were two boxes with different proportions of red and white beads. A draw of a single red bead indicated that there was a 60% chance that the predominantly red box was being sampled, while a draw of equal numbers of red and white beads implied that there was a 50% chance that either box was being sampled. Subjects who saw a red bead sample followed by a neutral sample adjusted their predictions downward to a probability between 50% and 60%.

Conversational Norms

An alternative account from the social judgment literature argues that dilution is an experimental artifact resulting from subjects' mistaken reliance on conversational norms (Schwarz et al. 1991; Slugoski and Wilson 1998; Tetlock, Lerner, and Boettger 1996). In intentional communication, a number of conversational norms are assumed to be respected, one being that all information is relevant for the goal of the conversation. This has been referred to as the maxim of relation (Grice 1975) or the principle of relevance (Sperber and Wilson 1986). However, this norm is violated in the experimental context of the dilution studies, thereby leading subjects to erroneous inferences (e.g., "The experimenter provides this information, so it must be relevant"). For example, Tetlock et al. (1996) asked subjects to predict the GPA of a hypothetical student and manipulated subjects' accountability as well as the activation of conversational norms. They observed that, for accountable subjects, the dilution effect disappeared when conversational norms had been deactivated by mentioning that the computer had randomly selected the presented information. However, for non-accountable subjects, the dilution effect persisted when conversational norms had been deactivated. This suggests that conversational norms may contribute to the dilution effect but that they are not necessary for the effect to occur.

Representativeness

The most popular account of the dilution effect within the social judgment literature relies on the representativeness heuristic (Fein and Hilton 1992; Hilton and Fein 1989; Locksley, Hepburn, and Ortiz 1982; Nisbett et al. 1981; Tetlock and Boettger 1989; Zukier 1982). The representativeness heuristic is a strategy by which subjects use the similarity between the available information about the individual and the to-be-predicted behavior to estimate the probability that the individual will display the behavior (Kahneman and Tversky 1972, 1973). It is assumed that diagnostic information is highly representative of the behavior, while irrelevant information is not. Therefore, adding irrelevant information to diagnostic information makes the

individual less representative of the behavior and attenuates the judgment. For example, a man who has a drinking problem is more representative of the stereotypical child abuser than a man who has a drinking problem and has two fingers missing on his left hand (Nisbett et al. 1981).

BIASED HYPOTHESIS TESTING

Although the dilution literature has focused on the three preceding explanations, the dilution effect may also be the result of biased hypothesis testing. Many studies have demonstrated that people test hypotheses in a biased fashion. For instance, people often consider the implications of evidence for the focal hypothesis but ignore the implications of the evidence for the alternative hypothesis (e.g., Trope and Liberman 1996). Although such selective hypothesis testing can unduly increase confidence in the focal hypothesis (e.g., Fischhoff and Beyth-Marom 1983), the same strategy may unduly decrease confidence in the hypothesis when the evidence is unlikely given either hypothesis (Fischhoff and Beyth-Marom 1983; Sanbonmatsu et al. 1997; Sanbonmatsu et al. 1998). Fischhoff and Beyth-Marom (1983) suggest that, by selectively focusing on a single hypothesis, people may decrease their confidence in the hypothesis when encountering unlikely irrelevant information, such as the information presented by Nisbett and colleagues (1981). Yet, although selective hypothesis testing can explain the dilution effect when the irrelevant information is unlikely given either hypothesis, it cannot account for demonstrations of the dilution effect that use typical irrelevant information. For example, Zukier and Jennings (1983) observed dilution effects with irrelevant information that is likely given either hypothesis (e.g., the defendant is of average height and vision). Thus, we propose that while selectivity is a necessary assumption for a biased hypothesis testing account of the dilution effect, it is not sufficient. Accordingly, we include additional assumptions in the proposed biased hypothesis testing process.

First, we assume that consumers are more likely to test the hypothesis that the product will deliver the benefit rather than the hypothesis that the product will not deliver the benefit. A bias toward considering the positive hypothesis is consistent with evidence that the acceptance of an idea is part of the automatic comprehension of that idea, that acceptance occurs before rejection, and that acceptance is less effortful than rejection (Gilbert 1991; Gilbert, Tafarodi, and Malone 1993). Moreover, advertising claims can explicitly create the hypothesis that the product will deliver the benefit (Ha and Hoch 1989; Hoch and Ha 1986). Finally, consumers are more likely to benefit from identifying products that will deliver the desired benefit than by identifying products that will not deliver the benefit.

Second, we assume that consumers will search for evidence in a biased fashion. Consumers will selectively search for confirming evidence (i.e., evidence that suggests the product will deliver the benefit). This assumption is supported by many demonstrations of biased search for confirming evidence (e.g., Shaklee and Fischhoff 1982; Snyder

and Cantor 1979; Snyder and Swann 1978) or positive test cases (e.g., Klayman and Ha 1987).

Third, we assume that consumers will classify the product information with respect to their search goal. The extensive literature on goal-derived categories (e.g., Barsalou 1983; Ratneshwar, Pechmann, and Shocker 1996) demonstrates that categorization is often based on temporarily salient goals. Similarly, consumers who are looking for information that suggests the product will deliver the benefit may classify all information with respect to this search goal. Thus, we assume that consumers classify information as either confirming (i.e., the type of information they were searching for) or not confirming (i.e., not the type of information they were searching for). Finally, we assume that consumers will selectively use the product information to test the focal hypothesis, while ignoring the implications for the alternative hypothesis. Several studies on selective or pseudodiagnostic hypothesis testing confirm that people only consider the focal hypothesis when interpreting evidence (Beyth-Marom and Fischhoff 1983; Fischhoff and Beyth-Marom 1983; Sanbonmatsu et al. 1997; Sanbonmatsu et al. 1998; Trope and Liberman 1996). Thus, when consumers classify ambiguous information as confirming that the product will deliver the benefit, they will strengthen their belief in the product benefit, even though the information is equally supportive of the hypothesis that the product will not deliver the benefit (e.g., Ha and Hoch 1989; Hoch and Ha 1986). Conversely, when consumers classify obviously irrelevant information as not confirming that the product will deliver the benefit, they will weaken their belief in the product benefit, even though the information also does not confirm that the product will not deliver the benefit.

Summary

In summary, we propose that consumers test the hypothesis that the product will deliver the benefit. Consumers selectively search for information that suggests that the product will deliver the benefit, and they classify product information with respect to this search goal. Obviously supportive and ambiguous information is classified as confirming, while counterdiagnostic and obviously irrelevant information is classified as not confirming. When information is classified as confirming, it strengthens consumers' beliefs that the product will deliver the benefit. When information is classified as not confirming, it weakens consumers' beliefs that the product will deliver the benefit.

We present 10 experiments that systematically examine the different explanations of the dilution effect. In experiments 1 and 1A, we demonstrate the dilution effect, show its robustness across product categories, presentation orders, and belief measures, and provide evidence against a distraction-of-attention account. In experiment 2, we demonstrate that the observed dilution effect cannot be accounted for by an averaging model. Experiments 3 and 3A provide evidence that is inconsistent with a conversational norms explanation of the dilution effect, whereas experiment 4 shows evidence that is inconsistent with the representative-

ness account. The remaining experiments systematically manipulate the assumptions of the proposed biased hypothesis testing mechanism. Experiments 5 and 5A demonstrate that the dilution effect only occurs when the irrelevant information is processed with the desired benefit in mind, whereas experiment 6 shows that the effect disappears when consumers consider the implications of the evidence for both the focal and alternative hypotheses. Finally, experiment 7 demonstrates that the dilution effect reverses when consumers set out to test the hypothesis that the product will not deliver the benefit.

EXPERIMENT 1

The objectives of the first experiment are to demonstrate the diluting effect of obviously irrelevant product information and to examine whether the effect is caused by distraction of attention. Although there is considerable evidence for dilution in nonsocial and social judgments, these demonstrations do not necessarily generalize to a product judgment context. First, the abstract cues and within-subject manipulations typical of the nonsocial demonstrations of the dilution effect seem to have few parallels in product judgment contexts. Second, unlike subjects in social judgment experiments, consumers who make product judgments cannot rely on easily accessible stereotypes that may be an essential requirement for some decision mechanisms to occur (e.g., representativeness). Third, many studies have shown that adding nondiagnostic neutral or ambiguous information can polarize rather than dilute judgments (e.g., Fischhoff and Beyth-Marom 1983; Ha and Hoch 1989; Hoch and Ha 1986; Wallsten 1981). Therefore, experiment 1 provides an initial test of the dilution effect in a product judgment context.

A second objective of the first experiment is to test whether the irrelevant information is diluting consumers' beliefs by engaging resources that would otherwise be allocated to the diagnostic information. It is possible that people extract less information from supportive evidence when they are distracted by irrelevant evidence. To test this explanation, we measured subjects' recognition of the diagnostic information at the end of the experiment. If the irrelevant information did cause subjects to elaborate less on the supportive information, then subjects who had been exposed to this irrelevant information should recognize the supportive information less quickly than those who had only received the supportive information.

Subjects and Design

Subjects were 36 undergraduate students who participated in return for class credit. The design was a 2 (type of information) \times 8 (product replicates) mixed design. Each subject was presented with descriptions of eight different products. For each of the product replicates, subjects were randomly assigned to either the baseline condition or the treatment condition.

Procedure

The entire experiment was administered by personal computer. Subjects were first informed that they would receive information about eight different products (services) and that they would have to indicate whether the product would deliver a particular benefit. Subjects were told that the information they would receive “may or may not be helpful for the decision that you have to make.” Subjects were then informed of the desired benefit for the first product (e.g., “You are looking for a fast computer”). Immediately after this statement, subjects were provided with the first piece of information, which was always the supportive attribute (e.g., “very powerful processor”). In the baseline condition, subjects did not receive any additional product information. However, in the treatment condition, the supportive information was followed by three sequentially presented pieces of irrelevant information (e.g., “manufactured in the USA,” “airs commercials on NBC and CBS,” and “can be ordered on-line”). Finally, while the entire product description remained on the screen, subjects were asked to indicate their belief that the product would deliver the benefit (e.g., “Is this computer fast?”). Responses were made on a nine-point scale anchored by 1 = “Definitely NOT Fast” and 9 = “Definitely Fast.” The subjects then received the remaining seven replicates. The order of the replicates was randomized.

After all replicates had been presented, subjects were given a practice reaction time task. Subjects were shown six simple statements unrelated to the experiment and had to answer “true” or “false” as fast as possible by pressing the “1” or “0” key on their keyboard. For each subject, the average reaction time of accurate responses was recorded for use as a response latency covariate in the analysis. The practice task was followed by the actual reaction time task. Subjects were told that they would read statements about the product descriptions they had reviewed earlier and had to indicate whether the statements were true or false. For each replicate, subjects received one true statement, such as “The computer had a very powerful processor,” and one false statement, such as “The computer was loaded with games.” The true statements always concerned the supportive piece of information. The presentation order of the 16 statements was randomized.

Stimuli

We first selected eight product categories and corresponding desirable benefits: apartments (safe), package delivery service (fast), frozen dinners (healthy), airlines (superior service), toothpaste (fills cavities), car (sportive), stereo system (reliable), and computers (fast). A pretest ($n = 30$) was conducted to select three irrelevant attributes and one supportive attribute for each replicate. The pretest listed a wide range of facts for each product or service. Subjects were asked to allocate these facts to one of three categories: “suggests not [benefit],” “is not helpful for my decision,” or “suggests [benefit].” The 24 irrelevant facts selected for the main experiment were classified as “not helpful” by an

average of 90% of pretest subjects, as supportive of the benefit by 6% of the subjects, and as counterdiagnostic by only 4% of the subjects. The irrelevant information included package information (e.g., a toothpaste that comes in 6 oz. tubes), product attributes (e.g., a computer that can be ordered on-line), marketing information (e.g., an airline that sponsors the New York City Marathon), and product availability (e.g., a frozen dinner brand that is available at most grocery stores). The eight supportive facts were classified as suggesting the benefit by an average of 94% of pretest subjects.

A second pretest was conducted to examine the possibility that the facts judged as irrelevant in isolation were judged as relevant in the context of the complete product description. Thirty subjects were presented with the full product descriptions and asked to indicate the relevance of each piece of information. The 24 irrelevant facts were classified as nondiagnostic by an average of 93% of the subjects, as diagnostic of the benefit by 6% of the subjects, and as counterdiagnostic by only 1% of the subjects.

A final pretest examined the possibility that even though subjects indicated that the irrelevant facts are not diagnostic, they may still use these facts to make inferences about the desired benefit. For example, although a subject may classify a fact as irrelevant, the fact may still be informative because it is positively or negatively correlated with unstated facts that are relevant. To examine the possible direction of such an effect, 18 subjects were presented with the irrelevant facts and were asked to rate them on a six-point scale (ranging from 1 = “Will probably not [deliver benefit]” to 6 = “Will probably [deliver benefit]”), thus forcing them to classify the information as either diagnostic or counterdiagnostic. The irrelevant facts were classified as diagnostic of the benefit by an average of 65% of the subjects and as counterdiagnostic by an average of 35% of the subjects. The average rating was 3.94, which was significantly higher than 3.50, the midpoint of the scale ($t(432) = 7.22, p < .01$). Thus, if the irrelevant information would indeed lead to inferences about the benefit, these inferences would support subjects’ beliefs in the product benefit rather than counteract them.

Results and Discussion

It was expected that the addition of irrelevant information would weaken beliefs in the product benefit, thus demonstrating the dilution effect. The results confirmed that adding irrelevant information weakened subjects’ beliefs in the product’s ability to deliver the desired benefit ($F(1, 272) = 6.56, p = .01$). Subjects who received only the supportive information reported more extreme judgments ($M = 6.28$) than those who also received the irrelevant information ($M = 5.83$). The effect of irrelevant information did not depend on the specific product or service that people were judging ($F(7, 272) = 1.47, NS$).¹

¹For the sake of clarity and conciseness, nonsignificant interactions with the replicate factor will not be discussed. All interactions are nonsignificant unless explicitly stated otherwise.

If dilution was due to a diversion of resources to the irrelevant information, the addition of irrelevant information should also lead to poorer and slower recognition of the supportive information. In fact, subjects who only received supportive information were slightly less likely to recognize the supportive information ($P = 93\%$) than were subjects who also received the irrelevant information ($P = 97\%$; $\chi^2(1) = 2.91, p < .1$). Subjects who did recognize the supportive information did not show any differences in reaction time depending on the product information they had received earlier ($RT_{\text{supportive}} = 1.53$, $RT_{\text{supportive+irrelevant}} = 1.57$; $F(1, 272) = 0.23$, NS). Thus, the diluting effect of irrelevant product information was not caused by a diversion of resources to the irrelevant product information at the expense of the supportive information.

EXPERIMENT 1A

While the first experiment demonstrated the dilution effect and provided evidence against a distraction explanation, experiment 1A tests the robustness of the phenomenon. First, we examine whether the dilution effect also occurs when consumers compare multiple products and select the product that is most likely to deliver the benefit. For example, consumers could choose between two computers, one described with supportive information and a second described with supportive and irrelevant information. Second, we also manipulate the order in which the information is presented. Consistent with most previous demonstrations of the dilution effect (e.g., Fein and Hilton 1992; Hilton and Fein 1989; Nisbett et al. 1981), we presented the supportive information first in experiment 1. However, one could argue that inserting the irrelevant information between the supportive information and the belief measure reduced the salience of the supportive information at the time of the belief measure. To test this possibility, we manipulated the order of the product information. In the supportive-first condition, the supportive information was always presented before the irrelevant information, while in the supportive-last condition, the supportive information was always presented after the irrelevant information.

Subjects and Design

Subjects were 131 undergraduate students who participated in return for class credit. The design was a 2 (type of information) \times 2 (order of information) \times 2 (counterbalancing of supportive information) \times 8 (product replicates) mixed design. Each subject was presented with descriptions of eight different products. For each of the product replicates, subjects were randomly assigned to either the baseline condition or the treatment condition. The order of information and the counterbalancing factor were manipulated between subjects.

Stimuli and Procedure

We used the same product information as was used in experiment 1. However, we needed eight extra pieces of

supportive information to describe the alternative product in each choice pair. These pieces of supportive information were pretested with 30 undergraduate students. The eight additional pieces of information were classified as supportive, rather than irrelevant or counterdiagnostic, by an average of 93% of the pretest subjects.

For each product category, subjects were first informed of the desired benefit and then received the description of the target product, which contained either one piece of supportive information or one piece of supportive information and three pieces of irrelevant information. The information for the alternative product, which always consisted of one piece of supportive information, was then displayed below the target product description. The supportive information used for the two products was counterbalanced. After the information for both products had been displayed, subjects indicated which of the two products was more likely to deliver the desired benefit.

Results and Discussion

When both product descriptions contained only supportive information, 53% of subjects indicated that the target product was more likely to deliver the benefit as compared with 47% who selected the alternative product ($Z = 1.59$, NS). However, when the irrelevant information was added to the description of the target product, the proportion of subjects who selected the target product dropped significantly to 38% ($\chi^2(1) = 26.26, p < .01$). While subjects were, on average, indifferent between the target product and the alternative product when both were only described using supportive information, the majority of subjects perceived the alternative product as more likely to deliver the benefit when the target product description contained additional irrelevant information ($Z = 5.56, p < .01$). The effect of the irrelevant information did not depend on the order in which the product information was presented ($\chi^2(1) = 1.81$, NS).

In summary, experiments 1 and 1A demonstrate that adding irrelevant product information not only weakens beliefs about the product but that it also reduces the likelihood that the product will be selected in a choice task. Moreover, the dilution effect does not depend on the information presentation order, thus refuting the hypothesis that dilution is caused by the reduced salience of the supportive information at the time of the decision.

EXPERIMENT 2

While the first experiments demonstrated the dilution effect, experiments 2–4 examine the biased hypothesis testing account of the effect by systematically testing it against the existing explanations. Experiment 2 uses two strategies to test between the biased hypothesis testing account and the averaging explanation. First, we examine the effect of adding less supportive, rather than irrelevant, information to the supportive information. Less supportive information is not as strong as supportive information, but it still suggests that the product will deliver the benefit. According to the av-

eraging model proposed by Lichtenstein et al. (1975), adding less supportive product information to strongly supportive information should dilute product beliefs because less supportive information should result in less extreme judgments than supportive information. Averaging these separate predictions should result in an overall prediction that is less extreme than a prediction based only on the supportive information. Alternatively, the biased hypothesis testing account predicts that adding less supportive information to strongly supportive information should lead to more extreme judgments. Since the less supportive information suggests that the product will deliver the benefit, it will be classified by consumers as confirming, instead of not confirming, and it will strengthen consumers' belief in the hypothesis that the product will deliver the benefit.

The averaging explanation can also be tested by examining how the addition of irrelevant information affects consumers' sensitivity to the diagnostic information. According to the averaging model, "since the weights must sum to one, adding a new relevant stimulus to a set will cause the weights of the old stimuli to decrease" (Anderson 1974, p. 239). Since the irrelevant information must have a weight that is significantly greater than zero to account for the dilution effect, adding irrelevant information should decrease the weight of the supportive information. Therefore, the addition of irrelevant information should reduce consumers' sensitivity to changes in the strength of the supportive information. In other words, an averaging model predicts that consumers should be less sensitive to the difference between supportive and less supportive product descriptions if both descriptions also contain irrelevant information. In contrast, the biased hypothesis testing account predicts that the irrelevant information will have an independent diluting effect on consumers' beliefs. The irrelevant information should be classified as not confirming and directly weaken consumers' beliefs in the hypothesis that the product will deliver the benefit, without affecting consumers' sensitivity to the supportive information.

Subjects and Design

Subjects were 58 undergraduate students who participated in return for class credit. The design was a 5 (type of information) \times 8 (product replicates) mixed design. Each subject was presented with descriptions of eight different products or services. For each of the product replicates, subjects were randomly assigned to the supportive information condition, the supportive + irrelevant information condition, the less supportive information condition, the less supportive + irrelevant information condition, or the supportive + less supportive information condition.

Stimuli and Procedure

The stimulus set used in experiment 2 included all stimuli used in experiment 1 plus three pieces of less supportive information for each replicate. A pretest ($n = 30$) was conducted to select the 24 less supportive pieces of product

information. On average, the less supportive information was classified as suggesting the benefit by 71% of the pretest subjects. The procedure included a manipulation check to test whether the less supportive information was indeed weaker than the original supportive information.

The procedure was identical to experiment 1, with the exception that three information conditions were added to the design. In the supportive + less supportive condition, the product description contained supportive information (e.g., "very powerful processor"), as well as three pieces of less supportive information (e.g., "well-known brand name," "64 Mbyte working memory," "32-speed CD-Rom"). In the less supportive condition, the product description only consisted of a single piece of less supportive information, which was randomly drawn from the three pieces of less supportive information selected for that replicate. Finally, in the less supportive + irrelevant information condition, the product description contained one randomly selected piece of less supportive information, as well as three pieces of irrelevant information.

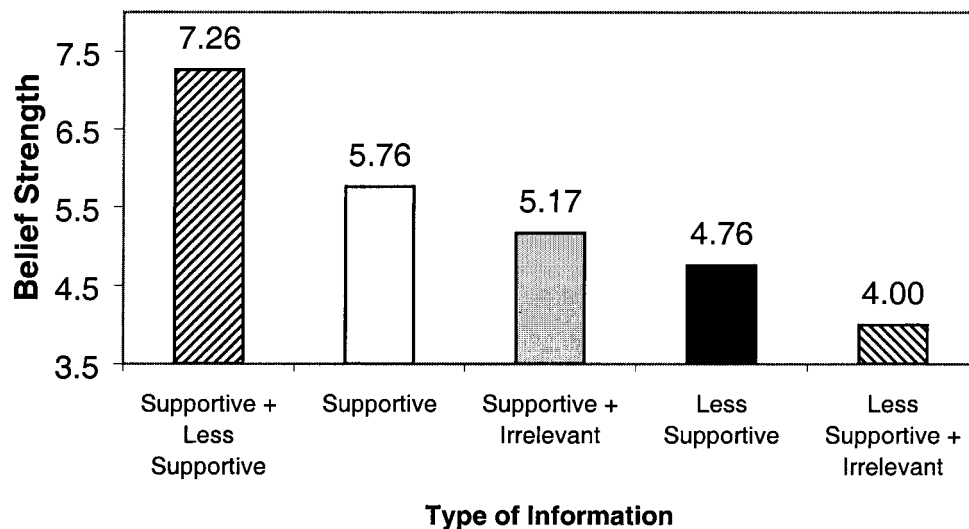
Results and Discussion

The results are summarized in figure 1. First, subjects' beliefs in the product benefit were weakened when irrelevant information was added to either supportive information ($M_{\text{supportive}} = 5.76$, $M_{\text{supportive+irrelevant}} = 5.17$; $F(1, 424) = 6.23$, $p < .05$) or less supportive information ($M_{\text{less supportive}} = 4.76$, $M_{\text{less supportive+irrelevant}} = 4.00$; $F(1, 424) = 5.87$, $p < .05$). Thus, the dilution effect observed in the first experiments was replicated in experiment 2. Second, subjects' beliefs in the product benefit were stronger when the product information they had received was supportive ($M = 5.76$) rather than less supportive ($M = 4.76$; $F(1, 424) = 12.91$, $p < .01$). This confirms that the manipulation of the degree of support was successful. Third, subjects receiving both supportive and less supportive information reported more extreme judgments ($M = 7.26$) than subjects receiving only supportive information ($M = 5.76$; $F(1, 424) = 22.42$, $p < .01$). This polarization effect is inconsistent with the averaging model but consistent with the biased hypothesis testing explanation. Fourth, the difference between the supportive information condition and the less supportive information condition did not become less pronounced when the irrelevant information was added ($D_{\text{no irrelevant}} = 1.00$, $D_{\text{irrelevant}} = 1.17$; $F(1, 424) = 0.04$, NS). This result is inconsistent with the averaging model prediction that the irrelevant information should reduce the impact of the diagnostic information, but it is consistent with the biased hypothesis testing account prediction that the irrelevant information has an independent influence on product beliefs.

The results of experiment 2 not only contradict that subjects are averaging separate predictions based on each piece of information (e.g., Lichtenstein et al. 1975) but are also inconsistent with more complex averaging models that include an initial impression (e.g., Anderson 1967; Lopes

FIGURE 1

EXPERIMENT 2: EFFECT OF THE TYPE OF INFORMATION ON SUBJECTS' BELIEFS ABOUT THE PRODUCT BENEFIT



1987). For example, it could be argued that when subjects' initial impressions are very low, the three pieces of less supportive information could lessen the relative impact of the low initial impression and create a polarization effect. Yet, even averaging models that include an initial impression predict that adding information that receives a nonzero weight in the decision should reduce people's sensitivity to the original information. The fact that the difference in belief strength between the supportive + irrelevant and less supportive + irrelevant conditions was not smaller than the corresponding difference between the supportive and less supportive conditions is also inconsistent with these more complex averaging models.

EXPERIMENT 3

The previous experiments demonstrated the robustness of the dilution effect and provided evidence that was inconsistent with the distraction and averaging accounts of the effect but consistent with the biased hypothesis testing explanation. Experiments 3 and 3A will pit the biased hypothesis testing explanation against the conversational norms account of the dilution effect. According to the conversational norms account, the dilution effect may result from subjects' misguided reliance on the maxim of relation, which states that communications have to be relevant to the goal of the conversation (Grice 1975). Subjects may assume that all the information provided by the experimenter is relevant for the judgment they are asked to make and therefore rely on all available information. To discourage this assumption, subjects in experiment 3 are asked whether the additional information is relevant before they state their be-

lief in the product benefit. If consumers are relying on conversational norms, information that is acknowledged to be irrelevant should not affect their belief in the benefit. On the other hand, the biased hypothesis testing account predicts that, even when subjects are aware that the information is irrelevant, they will automatically classify it as not supportive of their hypothesis and weaken their beliefs.

Subjects and Design

Subjects were 47 undergraduate students who participated in return for class credit. The design was a 2 (type of information) \times 12 (product replicates) mixed design. Each subject was presented with six different products or services. For each of the product replicates, subjects were randomly assigned to one of two information conditions, either the supportive information condition or the supportive + irrelevant information condition.

Stimuli and Procedure

The stimulus set used in experiments 1 and 2 was expanded with four additional replicates: hotel (luxurious), mountain bike (sturdy), movie (action-packed), and printer (high graphic quality). A pretest ($n = 30$) was conducted to select the product information for each additional replicate. The 12 irrelevant facts were classified as "not helpful" by an average of 87% of the pretest subjects. The four supportive facts were classified as suggesting the benefit by an average of 96% of the pretest subjects. Each subject was only exposed to six randomly selected replicates. Two filler descriptions were inserted in the second and fifth position

and served to make the preprogrammed structure of the stimuli less obvious.

The procedure was identical to the procedure followed in the first experiment, except for the presence of intermediate questions. After each piece of information had been presented, the following question was displayed: "What does this particular piece of information tell you about this [product]?" Subjects responded by clicking on one of three buttons labeled "That it is [benefit]," "That it is NOT [benefit]," and "This information is not helpful here." After all pieces of information had been presented and evaluated, subjects rated their belief in the product's ability to deliver the benefit.

Results and Discussion

Subjects clearly perceived the additional information as irrelevant. The irrelevant information was classified as not helpful in 92% of the cases, as diagnostic in 4% of the cases, and as counterdiagnostic in the remaining 4% of the cases. Yet, subjects' belief ratings were more extreme when the product description only contained supportive information ($M = 6.84$) than when it also contained irrelevant information ($M = 5.54$; $F(1, 252) = 51.89$, $p < .01$). This result is consistent with biased hypothesis testing but not with a reliance on conversational norms. Even when the analysis was restricted to trials on which all three pieces of information were classified as irrelevant, the additional information still weakened subjects' beliefs about the product benefit ($M_{\text{supportive}} = 6.84$, $M_{\text{irrelevant}} = 5.49$; $F(1, 228) = 50.79$, $p < .01$).

EXPERIMENT 3A

The results of experiment 3 clearly indicate that even when subjects acknowledge the irrelevance of the information, they still display the dilution effect. Although these findings cast doubt on the conversational norms explanation of the observed dilution effect, they cannot rule it out completely. Indeed, the fact that someone chooses to express an irrelevant assumption may itself be highly relevant (Sperber and Wilson 1986, p. 121). Subjects may assume that the experimenter must have provided the irrelevant information for some reason and therefore use the information despite its apparent irrelevance. To address this interpretation, we conducted an additional study ($n = 22$) in which we presented the information as randomly selected by a computer. This procedure is similar to manipulations used in studies supporting the conversational norms explanation (e.g., Tetlock et al. 1996). Subjects were told that the information was being randomly sampled by the computer and that, consequentially, some information would be helpful while other information would not. To reinforce this guise, a rapidly filling clock and the message "Randomly Drawing Information" was displayed before each piece of information was retrieved. Otherwise, the procedure was identical to the procedure used in experiment 1, except for the expanded

stimulus set of experiment 3. The results showed that the additional information still decreased subjects' beliefs in the product's ability to deliver the benefit ($M_{\text{supportive}} = 6.95$, $M_{\text{supportive+irrelevant}} = 5.94$; $F(1, 152) = 8.67$, $p < .01$). Together, the results from studies 3 and 3A cast doubt on a conversational norms explanation of the dilution effects observed in the current research.

EXPERIMENT 4


Although the results of the preceding experiments are inconsistent with averaging and conversational norms accounts of the dilution effect, they are consistent with the biased hypothesis testing explanation. However, these results can also be accounted for by consumers' reliance on a representativeness heuristic. Consumers may assess their belief in the product benefit by relying on the similarity between the described product and the typical desired product. Irrelevant information may reduce the similarity with this prototype, while less supportive information may enhance it. Furthermore, the representativeness explanation also predicts that the dilution effect will persist when subjects are aware of the irrelevance of the information and when the information appears to be randomly selected.

Experiment 4 used four conditions to test between the representativeness and biased hypothesis testing accounts of the dilution effect. Our strategy was to use irrelevant information that would increase, as opposed to decrease, the similarity of the product description to the typical, desired product. In the control condition, subjects were presented with the product descriptions and asked to state their belief in the product benefit. In the biased hypothesis testing condition, subjects were first asked to classify each piece of information as either supportive or not supportive of the benefit, then asked to rate their belief in the product benefit. In the representativeness condition, subjects were first asked to rate the similarity between the product description and the typical desired product, then asked to rate their belief in the product benefit. In the perceived similarity condition, subjects did not have to rate their belief in the product benefit but were instead asked to rate the similarity between the product description and the typical desired product. This condition was used to confirm that adding the specially selected irrelevant information did in fact increase the perceived similarity of the product to the typical desired product.

If subjects are relying on a representativeness heuristic and if the irrelevant information does indeed increase the perceived typicality of the product, then adding irrelevant information in the control condition should lead to polarization instead of dilution. Moreover, when subjects in the representativeness condition are encouraged to rely on their typicality judgments to assess their belief in the product benefit, they should produce belief ratings that are similar to those observed in the control condition. However, if the biased hypothesis testing process applies, adding irrelevant information in the control condition should still dilute product beliefs. Furthermore, when subjects in the biased hy-

FIGURE 2

EXPERIMENT 4: EFFECT OF STRATEGY AND THE TYPE OF INFORMATION ON SUBJECTS' BELIEFS ABOUT THE PRODUCT BENEFIT



hypothesis testing condition are first asked to classify the product information as either supportive or not supportive, they should produce belief ratings that are similar to those observed in the control condition.

Subjects and Design

Subjects were 83 undergraduate students who participated in return for class credit. The design was a 2 (type of information) \times 4 (strategy conditions) \times 8 (product replicates) mixed design. Each subject was presented with eight different products or services. For each of the product replicates, subjects were randomly assigned to either the supportive information condition or the supportive + irrelevant information condition. The strategy factor was manipulated between subjects.

Stimuli and Procedure

To ensure that adding the irrelevant product information would increase the perceived typicality of the product description, we selected atypical supportive information (e.g., a computer with “revolutionary triple processors”) and typical irrelevant information (e.g., “includes DVD player”). Although triple processors suggest that the computer is fast, they are not part of the representation of a typical fast computer. On the other hand, the DVD player does not affect the speed of the computer, but it does enhance the similarity to the typical fast computer. A pretest ($n = 40$) showed that the eight supportive facts were classified as suggesting the benefit by an average of 93% of pretest subjects, while the 24 irrelevant facts were classified as “not helpful” by an average of 93% of the subjects.

The procedure in the control condition was identical to

the one used in experiment 1. In the biased hypothesis testing condition, after each piece of information appeared, subjects were asked whether “this particular piece of information indicates that this [product] is [benefit].” They could respond by clicking on buttons labeled “yes” and “no.” After the entire product description had been displayed, they indicated their belief that the product would deliver the benefit. In the representativeness condition, subjects received the same product information, but before the information appeared, the statement “Try to imagine a typical [desired product]” appeared at the top of the screen (e.g., “Try to imagine a typical fast computer”). After all the product information had been displayed, subjects were asked to rate the similarity between the product description and the typical desired product (e.g., a typical fast computer) on a nine-point scale anchored by “not similar at all” and “very similar.” Subjects then indicated their belief that the product would deliver the benefit. In the perceived similarity condition, subjects followed a procedure identical to the representativeness condition, except that they were not asked to provide belief ratings.

Results and Discussion

The results are summarized in figure 2. The perceived similarity condition was used as a manipulation check. As expected, adding the typical irrelevant information to the supportive information increased the perceived similarity between the product description and the typical desired product ($M_{\text{supportive}} = 5.16$, $M_{\text{supportive+irrelevant}} = 6.51$; $F(1, 600) = 15.83$, $p < .01$). Yet, in the control condition, adding the typical irrelevant information still weakened subjects’ beliefs about the product benefit ($M_{\text{supportive}} = 6.59$, $M_{\text{supportive+irrelevant}} = 5.55$; $F(1, 600) =$

9.34, $p < .01$. This is inconsistent with the prediction of the representativeness account, but it is consistent with the prediction of the biased hypothesis testing perspective.

The two theories can also be tested by comparing the forced strategy conditions to the control condition. If subjects in the control condition were indeed using the proposed biased hypothesis testing strategy, then they should behave similarly to subjects who were explicitly encouraged to follow this strategy. The results show that adding irrelevant product information weakened subjects' beliefs in the biased hypothesis testing condition ($M_{\text{supportive}} = 6.18$, $M_{\text{supportive + irrelevant}} = 5.32$; $F(1, 600) = 7.56$, $p < .01$), a dilution effect that is not significantly different from the effect observed in the control condition ($F(1, 600) = 0.17$, NS). Alternatively, if subjects in the control condition were relying on the representativeness heuristic, they should behave similarly to subjects who were explicitly encouraged to follow this strategy. The results show that the addition of irrelevant information did not affect subjects' belief ratings in the representativeness condition ($M_{\text{supportive}} = 6.24$, $M_{\text{supportive+irrelevant}} = 6.30$; $F(1, 600) = 0.04$, NS), a result that is significantly different from the dilution effect observed in the control condition ($F(1, 600) = 4.50$, $p < .05$).

EXPERIMENT 5

The results observed in the previous experiments are inconsistent with averaging, conversational norms, and representativeness accounts of the dilution effect. However, while the results are consistent with the biased hypothesis testing explanation, the experiments did not test the essential characteristics of the proposed mechanism. Therefore, the following experiments will directly test the implications of the biased hypothesis testing explanation. In experiments 5 and 5A, we manipulate whether consumers initially process the information with the hypothesis in mind. Experiment 6 examines if the dilution effect disappears when consumers consider both the focal and alternative hypotheses. Finally, in experiment 7, we test if the effect reverses when consumers consider the hypothesis that the product will not deliver the benefit.

In the previous studies, as well as in previous demonstrations of the dilution effect, subjects knew the outcome that had to be predicted prior to processing the evidence. They could therefore engage in goal-oriented, top-down processing of the information. The fifth experiment replicates this top-down scenario, but it also adds a condition in which the desired benefit is only revealed after subjects have read the product description, thus forcing subjects to first process the product information without the benefit in mind (i.e., in a bottom-up fashion). According to the biased hypothesis testing explanation, subjects in the top-down condition (1) process the product description while searching for information that supports the hypothesis that the product will deliver the benefit, (2) classify the irrelevant information as

not confirming, and (3) weaken their belief in the hypothesis. In contrast, subjects in the bottom-up condition should (1) initially process the product description without a specific hypothesis in mind, (2) learn about the hypothesis, and (3) search the screen for information that supports the hypothesis. However, because subjects in the bottom-up condition have already processed all the product information, their search for supportive evidence can be more selective. They can immediately focus on the supportive evidence and ignore information that cannot possibly be supportive. Although they have initially processed the irrelevant information, they do not use it to evaluate the hypothesis; hence dilution should not occur.

A second objective of experiment 5 was to examine the robustness of the polarizing effect of the less supportive information observed in the second experiment. To this end, the experimental design also contained conditions with less supportive information. We expected to replicate the polarization effect in both the top-down and bottom-up conditions. The less supportive information suggests that the product will deliver the benefit and should therefore be revisited in the search for confirmatory evidence once the desired benefit has become known. It should be classified as confirming and strengthen consumers' beliefs in the hypothesis.

Subjects and Design

Subjects were 57 undergraduate students who participated in return for class credit. The design was a 4 (type of information) \times 2 (processing mode) \times 12 (product replicates) mixed design. Each subject was presented with eight different product descriptions out of a total of 12 product replicates. For each of the replicates, subjects were randomly assigned to either the supportive, supportive + irrelevant, less supportive, or supportive + less supportive information condition. The processing strategy (bottom-up or top-down) was manipulated between subjects.

Stimuli and Procedure

Experiment 5 used the same 12 replicates that were used in experiment 3. A pretest ($n = 30$) was conducted to find less supportive information for the four replicates not used in experiment 2. The 12 additional pieces of less supportive information were perceived as suggesting the benefit by 87% of respondents. The procedure in the top-down condition was identical to the procedure used in experiment 1, while the procedure in the bottom-up condition differed in that the desired benefit was not displayed before the product description appeared. Instead, the experiment instructions informed subjects that they would have to evaluate each product "on a certain dimension." As in the previous experiments, all information remained on the screen when the belief measure appeared.

Results and Discussion

The results are summarized in figure 3. The biased hypothesis testing account predicts that when the irrelevant information has first been processed without the benefit in mind, consumers will not revisit this information in their search for supportive evidence. Hence, the irrelevant information will not be used to test the hypothesis. Consistent with this prediction, the dilution effect did depend on the manner in which the information was processed ($F(1, 360) = 4.59, p < .05$). The irrelevant information significantly diluted subjects' beliefs in the top-down condition ($M_{\text{supportive}} = 6.56, M_{\text{supportive+irrelevant}} = 5.68; F(1, 360) = 7.62, p < .01$), but it did not influence beliefs in the bottom-up condition ($M_{\text{supportive}} = 6.18, M_{\text{supportive+less supportive}} = 6.45; F(1, 360) = 0.02, \text{NS}$). It was also predicted that consumers would revisit additional information when it is less supportive. As a consequence, the processing strategy should not influence the strength of the polarization effect. As predicted, the polarization effect did not depend on the manner in which the information was processed ($F(1, 360) = 0.63, \text{NS}$). Adding less supportive information to supportive information increased the strength of the subjects' product beliefs, regardless of the manner in which the information was presented ($M_{\text{supportive}} = 6.37, M_{\text{supportive+less supportive}} = 7.33; F(1, 360) = 13.10, p < .01$). A manipulation check confirmed that product descriptions consisting only of supportive information ($M = 6.37$) led to stronger beliefs than descriptions consisting only of less supportive information ($M = 5.57; F(1, 360) = 15.92, p < .01$).

EXPERIMENT 5A


The results from experiment 5 suggest that the dilution effect in the bottom-up condition failed to occur because subjects' preprocessing of the irrelevant information allowed them to subsequently ignore this information when selectively searching for supportive information. To directly test this assumption, we created a situation in which the initial processing of the irrelevant information suggested that the information could be supportive of the yet-to-be-revealed benefit. The obviously irrelevant information of study 5 was replaced with pseudorelevant information. Pseudorelevant information is information that is relevant in similar situations but is not relevant for the actual decision (Hilton and Fein 1989; Yzerbyt, Leyens, and Schadron 1997). For example, the fact that a computer has high quality speakers, is loaded with games, and has a flat screen monitor is not relevant for assessing the speed of the computer, but it is often relevant in typical computer purchase decisions. As a consequence, we expected subjects would have to reconsider pseudorelevant information after the benefit was revealed, classify this information as not confirming, and lessen their belief that the product would deliver the benefit.

Subjects and Design

Subjects were 51 undergraduate students who participated in return for class credit. The design was a 2 (type of information) \times 2 (processing mode) \times 8 (product replicates) mixed design. Each subject was presented with eight different product descriptions. For each replicate, subjects were randomly assigned to either the supportive information condition

FIGURE 3

EXPERIMENT 5: EFFECT OF PROCESSING MODE AND THE TYPE OF INFORMATION ON SUBJECTS' BELIEFS ABOUT THE PRODUCT BENEFIT



or the supportive + pseudorelevant information condition. The processing mode, top-down or bottom-up, was manipulated between subjects.

Stimuli and Procedure

The stimulus set was a subset of the products used in experiment 5: hotel, movie, car, apartments, package delivery service, frozen dinners, toothpaste, and computers. In a first pretest ($n = 29$), 24 pseudorelevant facts were classified as "not helpful" for the benefit judgment by an average of 90% of the subjects. In a second pretest ($n = 24$), subjects were shown both the pseudorelevant facts and the irrelevant information used in experiment 5 and were asked to indicate whether the information was typically helpful for evaluating the product (1 = not helpful at all, to 7 = very helpful). The results showed that the pseudorelevant facts were perceived as more helpful for evaluating the products ($M = 5.65$) than the irrelevant information ($M = 4.19$; $F(1, 1040) = 202.74$, $p < .01$). The procedure was identical to the procedure used in experiment 5.

Results and Discussion

Adding pseudorelevant information significantly diluted product beliefs ($M_{\text{supportive}} = 6.65$, $M_{\text{supportive+pseudorelevant}} = 5.98$; $F(1, 376) = 16.34$, $p < .01$). This effect did not interact with the manner in which the information was being processed ($F(1, 376) = 0.91$, NS). The irrelevant information weakened product beliefs whether the subjects processed the information with the benefit in mind ($M_{\text{supportive}} = 6.49$, $M_{\text{supportive+pseudorelevant}} = 5.87$; $F(1, 376) = 5.19$, $p < .05$) or without the benefit in mind ($M_{\text{supportive}} = 6.82$, $M_{\text{supportive+pseudorelevant}} = 6.09$; $F(1, 376) = 11.55$, $p < .01$).

The results from experiments 5 and 5A suggest that processing product information prior to having a benefit in mind may inhibit dilution but only under certain conditions. When additional information is obviously irrelevant, people know the information is unlikely to be supportive of any benefit, and they can restrict their subsequent search for supportive evidence to the supportive information. In contrast, when additional information is pseudorelevant, subjects must reconsider the information when they subsequently learn about the desired benefit because the information may be supportive. Since the pseudorelevant information is not supportive, reconsidering this information results in dilution.

EXPERIMENT 6

In experiment 6, we test a second implication of the proposed biased hypothesis testing mechanism. One assumption of this account is that irrelevant information dilutes product beliefs because consumers only consider whether the information supports the focal hypothesis, while ignoring whether the information supports the alternative hypothesis. Whereas obviously irrelevant information does not support the hypothesis that the product will deliver the benefit, it

also does not support the hypothesis that the product will not deliver the benefit. Therefore, the dilution effect should not occur when consumers consider the implications of the irrelevant information for both hypotheses. Consistent with this argument, McKenzie (1998) observed that decision makers rely less on a nondiagnostic cue when they simultaneously consider the focal and alternative hypotheses. Experiment 6 tests this prediction by manipulating the number of questions subjects are asked about each product description. In the single hypothesis condition, subjects are only asked to rate their belief in the benefit, whereas in the dual hypotheses condition, subjects are also asked to rate their belief that the product will not deliver the benefit.

Subjects and Design

Subjects were 112 undergraduate students who participated in return for class credit. The design was a 2 (type of information) \times 2 (number of hypotheses) \times 8 (product replicates) mixed design. Each subject was presented with descriptions of nine different products, one practice product and eight target products. For each product, subjects were randomly assigned to either the supportive condition or the supportive + irrelevant condition. The number of hypotheses was manipulated between subjects.

Stimuli and Procedure

The stimulus set was identical to the information used in experiment 1, with the exception that one additional practice category was included as the first description for each subject. This practice category was included to ensure that subjects in the dual hypotheses condition would realize that they would have to rate the products on both dimensions.

The procedure in the single hypothesis condition was identical to the procedure used in experiment 1. In the dual hypotheses condition, subjects were also told that they were looking for a particular benefit, followed by a description of the product, and the measure of subjects' belief in the benefit (e.g., "Is this computer fast?"). However, this measure then was followed by a measure of subjects' belief in the reverse of the benefit (e.g., "Is this computer slow?") on a similar nine-point scale (e.g., anchored by "definitely not slow" and "definitely slow").

Results and Discussion

The effect of the irrelevant information depended on the number of hypotheses subjects were evaluating ($F(1, 864) = 4.07$, $p < .05$). When subjects only rated their belief in the benefit, the irrelevant information again diluted subjects' beliefs ($M_{\text{supportive}} = 6.59$, $M_{\text{supportive+irrelevant}} = 6.09$; $F(1, 864) = 10.60$, $p < .01$). However, when subjects also indicated their belief in the opposite of the benefit, the irrelevant information did not influence their belief in the product benefit ($M_{\text{supportive}} = 6.44$, $M_{\text{supportive+irrelevant}} = 6.38$; $F(1, 864) = 0.13$, NS). Thus, consistent with the biased hypothesis

testing account, consumers' selective focus on the hypothesis that the product will deliver the benefit seems to be a necessary condition for the dilution effect to occur.

EXPERIMENT 7

The experiments presented so far have shown that adding irrelevant product information will usually weaken consumers' beliefs in the product benefit. Yet, this does not imply that brands should always avoid the communication of irrelevant information if they want to emphasize a product benefit. If consumers follow the proposed biased hypothesis testing process, then confronting irrelevant information may sometimes strengthen consumers' beliefs in the product's ability to deliver the benefit. This may happen when consumers have a strong reason to suspect that the product will not deliver the benefit. For instance, when the product carries a brand name that has a very poor reputation on the critical dimension, consumers may set out to test the hypothesis that the product will not deliver the benefit. They may then search for information that confirms this hypothesis and classify information with regard to this search goal as confirming (i.e., counterdiagnostic of the benefit) or not confirming (i.e., not counterdiagnostic of the benefit). Irrelevant information will be classified as not confirming and weaken consumers' belief in the hypothesis that the product will not deliver the benefit. In other words, irrelevant information will strengthen the belief that the product will deliver the benefit. Experiment 7 tests this prediction by manipulating the presence of a negatively perceived brand name.

Subjects and Design

Subjects were 68 undergraduate students who participated in return for class credit. The design was a 2 (type of information) \times 2 (presence of brand names) \times 7 (product replicates) mixed design with a brand name only control condition. Each subject was presented with three different product descriptions from a total of seven product replicates. For each of the replicates, subjects were randomly assigned to either the supportive condition or the supportive + irrelevant condition. The presence of brand names was manipulated between subjects. Subjects in the no brand name condition only received product descriptions, while subjects in the brand name condition received both brand names and product descriptions. Subjects in the brand name only control condition received only brand names and no product descriptions.

Stimuli and Procedure

The stimulus set consisted of seven target categories (products or services) and five filler categories. Only two of the target categories were taken from previous experiments (hotel room and car). The other five categories were either completely new or required changes in the product information: beer (great taste), apartments (safe), clothing store (trendy), shampoo (high quality hair care), and res-

taurant (healthy). A pretest ($n = 36$) was conducted to select one supportive fact and three irrelevant facts for each of the new categories. The 15 irrelevant facts were classified as "not helpful" by an average of 87% of the pretest subjects, while the five supportive facts were classified as suggesting the benefit by an average of 92% of respondents.

The procedure used in the no brand name condition was similar to the one used in experiment 1. The procedure in the brand name condition differed in some important ways. First, subjects were asked to rate a set of brands in the seven target categories and five filler categories. In each product category, subjects were presented with four to seven brand names and asked to indicate whether each brand would deliver a particular benefit on a scale ranging from -3 (definitely not [benefit]) to $+3$ (definitely [benefit]). After a filler task, subjects were exposed to information for the five filler categories and for three target categories for which they had indicated strong negative beliefs for at least one of the brands.² Subjects first received the negatively perceived brand name, followed by the product description. The filler descriptions confirmed subjects' positive or negative priors, thus reducing suspicion about the accuracy of the information. The target descriptions either contained only supportive information or both supportive and irrelevant information. For instance, if the pretest indicated that a subject thought K-Mart was not trendy, the subject would be presented with the instruction, "You are looking for a trendy store. The store you are considering is K-Mart." This instruction could be followed by supportive information (e.g., "Has announced the opening of a Tommy Hilfiger section") or by both supportive information and three pieces of irrelevant information (e.g., "Closes at 9 PM," "Major credit cards accepted," and "Airs commercials on CBS and NBC"). In the brand name only condition, subjects did not receive a product description and had to base their judgment on the brand name.


Results and Discussion

The results are summarized in figure 4. First, a manipulation check showed that adding supportive information to the brand name strengthened subjects' belief that the brand would deliver the benefit ($M_{\text{brand name only}} = 1.61$, $M_{\text{brand name+supportive}} = 3.75$; $F(1, 129) = 35.29$, $p < .01$). The remainder of the analyses will concentrate on the 2 \times 2 design manipulating the type of information and presence of the brand name. First, there was a main effect of brand name. As expected, the belief ratings were higher when subjects did not receive the negatively perceived brand name ($M_{\text{no brand name}} = 5.11$, $M_{\text{brand name}} = 4.11$; $F(1, 128) = 9.58$, $p < .01$). Second, there was no main effect of adding irrelevant information ($F(1, 128) =$

²To maintain comparability, assignments of replicates in the no brand name condition depended on the replicate selection in the brand name condition. This guaranteed that the proportion of categories selected did not differ between conditions.

FIGURE 4

EXPERIMENT 7: EFFECT OF A NEGATIVELY PERCEIVED BRAND NAME AND THE TYPE OF INFORMATION ON SUBJECTS' BELIEFS ABOUT THE PRODUCT BENEFIT



0.14, NS). Third, the effect of the irrelevant information depended on the presence of the brand name ($F(1, 128) = 6.52, p < .05$). When subjects only received the product information, the irrelevant information weakened subjects' beliefs in the product benefit from 5.54 to 4.61 ($D = -0.93$; $F(1, 128) = 5.65, p < .05$). However, when subjects were also given the brand name, the irrelevant product information strengthened product beliefs from 3.75 to 4.44 ($D = 0.69$; $F(1, 128) = 4.22, p < .05$).

These results demonstrate that providing irrelevant information in addition to supportive information will not always hurt product perceptions. When a brand has a strong, negative image, consumers' beliefs will become more favorable after they encounter both supportive and irrelevant information rather than only supportive information. These findings are consistent with a biased hypothesis testing account of the dilution effect. When consumers process information regarding a negatively perceived brand, they search for counterdiagnostic information that confirms the brand will not deliver the benefit. The irrelevant information does not confirm this hypothesis, reduces confidence in the hypothesis, and results in more favorable product beliefs.

GENERAL DISCUSSION

Research on social judgment suggests, and this research confirms, that obviously irrelevant information can have a negative impact on consumers' product perceptions. In 10 different studies, and across 17 different products and services, the addition of irrelevant information to supportive benefit information weakened consumers' beliefs in the product's ability to deliver the benefit. This dilution effect

did not depend on the order in which the information was presented or the manner in which the belief was measured (experiment 1A). Moreover, the dilution effect persisted when consumers acknowledged the irrelevance of the information prior to stating their beliefs (experiment 3), when they believed that the information was being randomly sampled by a computer (experiment 3A), and when the irrelevant information made the product description more similar to the typical desired product (experiment 4).

Despite the strong evidence for the dilution effect, the data cannot be fully accounted for by the existing explanations of the dilution effect. The polarizing effect of less supportive information (experiments 2 and 5) and the failure of the irrelevant information to influence consumers' sensitivity to the supportive information (experiment 2) are not consistent with an averaging account of the dilution effect. The conversational norms account cannot explain why the dilution effect persists when consumers first acknowledge the irrelevance of the product information (experiment 3) and when the information is allegedly randomly sampled by a computer (experiment 3A). Furthermore, the representativeness explanation predicts that irrelevant information that significantly increases the product's similarity to the typical desired product should lead to more favorable product judgments, while this information in fact dilutes consumers' beliefs in the product benefit (experiment 4). Finally, the distraction account is inconsistent with the observation that the irrelevant information does not affect consumers' recognition of the supportive information (experiment 1) nor consumers' sensitivity to the supportive information (experiment 2). Thus, although the previous mechanisms may contribute to the dilution effect, none of

these accounts can explain the complete pattern of results observed in these studies.

Instead, the results are more consistent with a biased hypothesis testing procedure that has four critical characteristics. First, consumers test the hypothesis that the product will deliver the benefit, rather than the hypothesis that it will not deliver the benefit. Second, consumers selectively search for information that confirms the hypothesis. Third, consumers classify all information with regard to their search goal, either as confirming or not confirming. Finally, consumers rely on this classification to determine their belief in the product's ability to deliver the benefit. When consumers encounter irrelevant information, they classify it as not confirming and weaken their belief in the product benefit, without taking into account that the information does not confirm the alternative hypothesis either.

This proposed mechanism explains the direct diluting effect of irrelevant information on product beliefs, as well as the polarizing effect of less supportive information. It also predicts that the dilution effect will persist when consumers are aware of the irrelevance of the information, when the information is allegedly randomly sampled, and when the information increases the typicality of the product description. Moreover, the last four experiments confirm the predictions of the biased hypothesis testing perspective regarding the boundary conditions of the dilution effect. When the irrelevant information is not processed with the hypothesis in mind, the information cannot be classified as not confirming and the dilution effect disappears (experiment 5). Similarly, the dilution effect does not occur when consumers consider the implications of the information for both the focal hypothesis and the alternative hypothesis (experiment 6). Finally, when consumers have negative priors about a brand, they may test the hypothesis that the product will not deliver the benefit, leading to a reversal of the dilution effect (experiment 7).

Limitations and Future Research

Our studies did not examine situations in which desirable benefits are spontaneously generated by consumers. A benefit can be generated because it has a high habitual salience for a certain consumer or because it is primed by a certain usage situation (Ratneshwar et al. 1997). In these situations, the activated benefit may lead to similar outcomes as the explicit instructions used in our experiments. However, consumers can also derive the desired benefit from the product information itself. While some product information may seem irrelevant at first, consumers may rely on the relevance principle (Grice 1975; Sperber and Wilson 1986) and infer that the information has to convey some value. The irrelevant information may prime a benefit consumers had not considered earlier, thereby increasing the appeal of the product, rather than diluting it.

Although the reported studies are most consistent with a biased hypothesis testing explanation, our results do not imply that other processes cannot contribute to the dilution effect or even be the unique cause of dilution in situations in which

this proposed mechanism does not apply. For instance, many studies have shown that people often do rely on a representativeness heuristic when making predictions (e.g., Andreasen 1988; Kahneman and Tversky 1972, 1973). It is plausible that subjects in the social judgment dilution studies indeed relied on the similarity between the description of the individual and the stereotypical murderer or child abuser. Although it may be hard to assess the similarity between a product description and an abstract benefit, it is easier to assess the similarity between a person and an accessible stereotype. In fact, consumers may also rely on a representativeness heuristic when predicting product benefits, provided that the desired benefit is connected to a well-defined subcategory or a prototypical brand.

Finally, we would like to emphasize that we only examined the influence of one type of irrelevant information. The concept of irrelevance used in this article has three essential characteristics. First, the information was irrelevant with respect to a context. The information we used was only irrelevant with respect to one specific benefit. It was not necessarily irrelevant with respect to product choice or with respect to the overall evaluation of the product. Second, the irrelevance of the information was subjective rather than objective. We label the information as obviously irrelevant because the great majority of subjects classified the information as "not helpful." Finally, the irrelevance of the information was not absolute. Even for those who classified the information as irrelevant, the information may still have had a minimal diagnostic value. It was this third characteristic of irrelevant information that encouraged us to demonstrate that the irrelevant information tends to be categorized as supportive rather than counterdiagnostic (pretest 3, experiment 1) and that less supportive information tends to polarize, rather than dilute, judgments (experiment 2).

How does this conceptualization compare with other interpretations of relevance? It is clearly different from a consequentialist perspective, which states that information is irrelevant for a decision when it does not influence the decision. A more related perspective is that of Sperber and Wilson (1986), who argue that an assumption is relevant in a context when it has some (subjective) contextual effect, the amount of which determines the degree of (ir)relevance. It is interesting that they argue that the degree of irrelevance also depends on the effort required to process the information and obtain the contextual effect. One could indeed argue that all product information can have some diagnostic value for any desired benefit but that the effort required to extract this value is so great that it makes the information irrelevant for most consumers. On the other hand, Sperber and Wilson (1986) also indicate that the context is not determined before processing the utterance but is selected so as to maximize the possible relevance of the statement (since people assume it is relevant). This indicates that the apparently irrelevant information may suggest new benefits to the consumer, as mentioned earlier in this section. However, this assertion is based on the relevance principle, which governs intentional communication. This principle may not hold in

many advertising situations, since consumers may assume that the additional information is intended for another consumer segment. In mass advertising communications, the information has to convey value to some consumer but not necessarily to the individual processing the message.

[Received June 2000. Revised July 2001. David Glen Mick served as editor, and Frank R. Kardes served as associate editor for this article.]

REFERENCES

- Anderson, Norman H. (1967), "Averaging Model Analysis of Set-Size Effect in Impression Formation," *Journal of Experimental Psychology*, 75 (October), 158-165.
- (1971), "Integration Theory and Attitude Change," *Psychological Review*, 78 (May), 171-206.
- (1974), "Information Integration Theory: A Brief Survey," in *Contemporary Developments in Mathematical Psychology*, Vol. 1, *Learning, Memory and Thinking*, ed. David H. Krantz et al., 236-270.
- Andreassen, Paul B. (1988), "Explaining the Price-Volume Relationship: The Difference between Price Changes and Changing Prices," *Organizational Behavior and Human Decision Processes*, 41 (June), 371-389.
- Barsalou, Lawrence W. (1983), "Ad Hoc Categories," *Memory and Cognition*, 11 (May), 211-227.
- Beyth-Marom, Ruth and Baruch Fischhoff (1983), "Diagnosticity and Pseudodiagnosticity," *Journal of Personality and Social Psychology*, 45 (December), 1185-1195.
- Birnbaum, Michael H. and Barbara A. Mellers (1983), "Bayesian Inference: Combining Base Rates with Opinions of Sources Who Vary in Credibility," *Journal of Personality and Social Psychology*, 45 (October), 792-804.
- Broniarczyk, Susan M. and Andrew D. Gershoff (1997), "Meaningless Differentiation Revisited," in *Advances in Consumer Research*, Vol. 24, ed. Merrie Brucks and Deborah J. MacInnis, Provo, UT: Association for Consumer Research, 223-228.
- Brown, Christina L. and Gregory S. Carpenter (2000), "Why Is the Trivial Important? A Reasons-Based Account for the Effects of Trivial Attributes on Choice," *Journal of Consumer Research*, 26 (March), 372-385.
- Carpenter, Gregory S., Rashi Glazer, and Kent Nakamoto (1994), "Meaningful Brands from Meaningless Differentiation: The Dependence on Irrelevant Attributes," *Journal of Marketing Research*, 31 (August), 339-350.
- De Dreu, Carsten K. W., Vincent Y. Yzerbyt, and Jacques-Philippe Leyens (1995), "Dilution of Stereotype-Based Cooperation in Mixed-Motive Interdependence," *Journal of Experimental Social Psychology*, 31 (November), 575-593.
- Fein, Steven and James L. Hilton (1992), "Attitudes toward Groups and Behavioral Intentions toward Individual Group Members: The Impact of Nondiagnostic Information," *Journal of Experimental Social Psychology*, 28 (March), 101-124.
- Fischhoff, Baruch and Ruth Beyth-Marom (1983), "Hypothesis Evaluation from a Bayesian Perspective," *Psychological Review*, 90 (July), 239-260.
- Gilbert, Daniel T. (1991), "How Mental Systems Believe," *American Psychologist*, 46 (February), 107-119.
- Gilbert, Daniel T., Romin W. Tafarodi, and Patrick S. Malone (1993), "You Can't Not Believe Everything You Read," *Journal of Personality and Social Psychology*, 65 (August), 221-233.
- Grice, H. Paul (1975), "Logic and Conversation," in *Syntax and Semantics*, Vol. 3, *Speech Acts*, ed. Peter Cole and Jerry L. Morgan, New York: Academic Press, 41-58.
- Ha, Young-Won and Stephen J. Hoch (1989), "Ambiguity, Processing Strategy, and Advertising-Evidence Interactions," *Journal of Consumer Research*, 16 (December), 354-360.
- Hilton, James L. and Steven Fein (1989), "The Role of Typical Diagnosticity in Stereotype-Based Judgments," *Journal of Personality and Social Psychology*, 57 (August), 201-211.
- Hoch, Stephen J. and Young-Won Ha (1986), "Consumer Learning: Advertising and the Ambiguity of Product Experience," *Journal of Consumer Research*, 13 (September), 221-233.
- Kahneman, Daniel and Amos Tversky (1972), "Subjective Probability: A Judgment of Representativeness," *Cognitive Psychology*, 3 (July), 430-454.
- (1973), "On the Psychology of Prediction," *Psychological Review*, 80 (July), 237-251.
- Klayman, Joshua and Young-Won Ha (1987), "Confirmation, Disconfirmation, and Information in Hypothesis Testing," *Psychological Review*, 94 (April), 211-228.
- Lichtenstein, Sarah, Timothy C. Earle, and Paul Slovic (1975), "Cue Utilization in a Numerical Prediction Task," *Journal of Experimental Psychology: Human Perception and Performance*, 104 (February), 77-85.
- Locksley, Anne, Christine Hepburn, and Vilma Ortiz (1982), "Social Stereotypes and Judgments of Individuals: An Instance of the Base-Rate Fallacy," *Journal of Experimental Social Psychology*, 18 (January), 23-42.
- Lopes, Lola L. (1987), "Procedural Debiasing," *Acta Psychologica*, 64 (February), 167-185.
- McKenzie, Craig R. M. (1998), "Taking into Account the Strength of an Alternative Hypothesis," *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 24 (May), 771-792.
- Nisbett, Richard E., Henri Zukier, and Ronald E. Lemley (1981), "The Dilution Effect: Nondiagnostic Information Weakens the Implications of Diagnostic Information," *Cognitive Psychology*, 13 (April), 248-277.
- Ratneshwar, S., Cornelia Pechmann, and Allan D. Shocker (1996), "Goal-Derived Categories and the Antecedents of Across-Category Consideration," *Journal of Consumer Research*, 23 (December), 240-250.
- Ratneshwar, S., Luk Warlop, David Glen Mick, and Gail Seeger (1997), "Benefit Salience and Consumers' Selective Attention to Product Features," *International Journal of Research in Marketing*, 14 (July), 245-259.
- Sanbonmatsu, David M., Steven S. Posavac, Frank R. Kardes, and Susan P. Mantel (1998), "Selective Hypothesis Testing," *Psychonomic Bulletin and Review*, 5 (June), 197-220.
- Sanbonmatsu, David M., Steven S. Posavac, and Randon Stasney (1997), "The Subjective Beliefs Underlying Probability Overestimation," *Journal of Experimental Social Psychology*, 33 (May), 276-295.
- Schwarz, Norbert, Fritz Strack, Denis Hilton, and Gabi Naderer (1991), "Base Rates, Representativeness and the Logic of Conversation: The Contextual Relevance of 'Irrelevant' Information," *Social Cognition*, 9 (Spring), 67-84.
- Shaklee, Harriet and Baruch Fischhoff (1982), "Strategies of Information Search in Causal Analysis," *Memory and Cognition*, 10 (November), 520-530.
- Shanteau, James (1975), "Averaging versus Multiplying Combi-

- nation Rules of Inference Judgment," *Acta Psychologica*, 39 (February), 83–89.
- Simonson, Itamar, Ziv Carmon, and Suzanne O' Curry (1994), "Experimental Evidence on the Negative Effect of Product Features and Sales Promotions on Brand Choice," *Marketing Science*, 13 (Winter), 23–41.
- Slugoski, Ben R. and Anne E. Wilson (1998), "Contribution of Conversation Skills to the Production of Judgmental Errors," *European Journal of Social Psychology*, 28 (July), 575–601.
- Snyder, Mark and Nancy Cantor (1979), "Testing Hypotheses about Other People: The Use of Historical Knowledge," *Journal of Experimental Social Psychology*, 15 (July), 330–342.
- Snyder, Mark and William B. Swann (1978), "Hypothesis-Testing Processes in Social Interaction," *Journal of Personality and Social Psychology*, 36 (November), 1202–1212.
- Sperber, Dan and Deirdre Wilson (1986), *Relevance: Communication and Cognition*, Cambridge, MA: Harvard University Press.
- Tetlock, Philip E. and Richard Boettger (1989), "Accountability: A Social Magnifier of the Dilution Effect," *Journal of Personality and Social Psychology*, 57 (September), 288–398.
- Tetlock, Philip E., Jennifer S. Lerner, and Richard Boettger (1996), "The Dilution Effect: Judgmental Bias, Conversational Convention, or a Bit of Both?" *European Journal of Social Psychology*, 26 (November), 915–934.
- Trope, Yaacov and Akiva Liberman (1996), "Social Hypothesis Testing: Cognitive and Motivational Mechanisms," in *Social Psychology: Handbook of Basic Principles*, ed. Edward Tory Higgins and Arie W. Kruglanski, New York: Guilford Press, 239–270.
- Troutman, C. Michael and James Shanteau (1977), "Inferences Based on Nondiagnostic Information," *Organizational Behavior and Human Performance*, 19 (June), 43–55.
- Wallsten, Thomas S. (1981), "Physician and Medical Student Bias in Evaluating Diagnostic Information," *Medical Decision Making*, 1 (October), 145–164.
- Yzerbyt, Vincent Y., Jacques-Philippe Leyens, and Georges Schadrone (1997), "Social Judgeability and the Dilution of Stereotypes: The Impact of the Nature and Sequence of Information," *Personality and Social Psychology Bulletin*, 23 (December), 1312–1322.
- Zukier, Henri (1982), "The Dilution Effect: The Role of the Correlation and the Dispersion of Predictor Variables in the Use of Nondiagnostic Information," *Journal of Personality and Social Psychology*, 43 (December), 1163–1174.
- Zukier, Henri and Dennis L. Jennings (1983), "Nondiagnosticity and Typicality Effects in Prediction," *Social Cognition*, 2 (July), 187–198.