

CONSUMERS' BELIEFS IN PRODUCT BENEFITS:
THE EFFECT OF IRRELEVANT PRODUCT INFORMATION

By
TOM MEYVIS

A DISSERTATION PRESENTED TO THE GRADUATE SCHOOL
OF THE UNIVERSITY OF FLORIDA IN PARTIAL FULFILLMENT
OF THE REQUIREMENTS FOR THE DEGREE OF DOCTOR OF PHILOSOPHY

UNIVERSITY OF FLORIDA

2001

ACKNOWLEDGEMENTS

This dissertation has benefited, directly and indirectly, from interactions with many people. I would like to thank these persons for their discussions, comments, and seminars, and for supporting and encouraging my research. First of all, I would like to thank Chris Janiszewski, my dissertation chair, for the many discussions, for his mentorship in general, and for making this entire process so much fun. This dissertation has also benefited from the extensive comments I have received from Joe Alba and Alan Cooke and from discussions with the other members of the dissertation committee: Steve Shugan and Ira Fischler. I would also like to thank all the students, former students, and faculty at the University of Florida, both for stimulating discussions and for guiding me through the Ph.D. experience, especially Amitav Chakravarti, Michel Pham, John Pracejus, Stijn van Osselear, Joel Cohen, Barry Schlenker, Joffre Swait, Bart Weitz, and Jinhong Xie. This research has also improved significantly thanks to comments from participants in interviews at the 2000 AMA conference in Chicago and in seminars at INSEAD, MIT, Harvard University, UCLA, the University of Colorado, Stanford University, Duke University, the University of Pennsylvania, the University of Chicago, Northwestern University, Columbia University, New York University, UC-Berkeley, and Tilburg University. I especially would like to thank Bob Wyer and John Lynch for their very helpful comments.

In a less direct, yet not less important, manner, this dissertation has also benefited from the continuing support I have received from my family and parents. I would especially like to convey my gratitude to my parents for supporting me throughout my education and for always giving me the freedom to explore my interests, even though these interests tended to change rather unpredictably. Finally, I would like to thank Els for bringing so much fun to my life that I almost forgot to write this dissertation.

TABLE OF CONTENTS

	<u>page</u>
ACKNOWLEDGEMENTS	ii
ABSTRACT	vi
CHAPTERS	
1 INTRODUCTION	1
The Effect of Irrelevant Information	2
Overview	4
2 THEORETICAL BACKGROUND	6
Irrelevant Product Information in the Consumer Research Literature	6
The Dilution Effect	9
Existing Accounts of the Dilution Effect	10
A Biased Hypothesis Testing Account	15
3 EXPERIMENT 1: DEMONSTRATION OF THE DILUTION EFFECT	24
Method	26
Results	30
Discussion	31
Experiment 1A	32
4 EXPERIMENT 2: A TEST OF THE AVERAGING ACCOUNT	36
Method	39
Results and Discussion	40
5 EXPERIMENT 3: A TEST OF THE CONVERSATIONAL NORMS ACCOUNT	44
Method	47
Results and Discussion	48
Experiment 3A	49

6	EXPERIMENT 4: A TEST OF THE REPRESENTATIVENESS ACCOUNT	51
	Method	55
	Results	57
	Discussion	59
7	EXPERIMENT 5: PROCESSING THE INFORMATION WITHOUT THE BENEFIT IN MIND	60
	Method	62
	Results and Discussion	63
	Experiment 5A	65
8	EXPERIMENT 6: CONSIDERING BOTH THE FOCAL AND ALTERNATIVE HYPOTHESIS	69
	Method	70
	Results	71
	Discussion	72
9	EXPERIMENT 7: EXPECTING THAT THE PRODUCT WILL NOT DELIVER THE BENEFIT	73
	Method	74
	Results and Discussion	76
10	GENERAL DISCUSSION	79
11	LIMITATIONS AND FUTURE RESEARCH	83
	REFERENCES	87
	APPENDICES	
	A INSTRUCTIONS TO SUBJECTS	93
	B PRODUCT INFORMATION	98
	C EXAMPLES OF COMPUTER DISPLAYS	110
	BIOGRAPHICAL SKETCH	112

Abstract of Dissertation Presented to the Graduate School of the University of Florida in
Partial Fulfillment of the Requirements for the Degree of Doctor in Philosophy

CONSUMERS' BELIEFS IN PRODUCT BENEFITS:
THE EFFECT OF IRRELEVANT PRODUCT INFORMATION

By

Tom Meyvis

May 2001

Chairman: Chris Janiszewski
Major Department: Marketing

When consumers try to assess the performance of a product on a key benefit, they often encounter irrelevant product information while searching for diagnostic information. Normatively, this irrelevant information should not influence consumers' belief in the product's ability to deliver the benefit. However, this dissertation presents evidence that adding irrelevant product information to supportive benefit information systematically *weakens* consumers' belief that the product will deliver the desired benefit. This dilution effect persists when subjects are forced to acknowledge the irrelevance of the additional information prior to stating their beliefs, when the information is allegedly randomly selected by a computer, and when the irrelevant information increases the similarity of the product description to the typical desired product. Furthermore, the irrelevant information does not dilute the impact of the supportive information, but instead has a direct, independent effect on consumers' beliefs in the benefit.

These findings are inconsistent with previous accounts of the dilution effect. The diluting effect of irrelevant information observed in these studies is not caused by an averaging strategy, a misguided reliance on conversational norms, or the use of a representativeness heuristic. Instead, I suggest that this dilution effect is caused by consumers' reliance on a *biased hypothesis testing* strategy. It is proposed that consumers test the hypothesis that the product will deliver the benefit, selectively search for information that confirms this hypothesis, and categorize all other information (be it irrelevant or disconfirming) as "not confirming." When information is classified as "not confirming," it weakens belief in the hypothesis, even when it does not confirm the alternative hypothesis. As a consequence, irrelevant information weakens consumers' beliefs in the product's ability to deliver the benefit.

Consistent with this explanation, the dilution effect is shown to disappear when consumers first process the information without the benefit in mind, and when consumers are forced to consider both the focal hypothesis and the alternative hypothesis. Moreover, the dilution effect even reverses when the product description concerns a brand with a poor reputation, suggesting that, in this case, consumers may be testing the hypothesis that the product will *not* deliver the benefit.

CHAPTER 1 INTRODUCTION

When consumers evaluate products or services, they often focus on one or more critical product benefits. To find out if the product under consideration will indeed deliver these important benefits, consumers will usually search for diagnostic product information. For instance, a consumer who wonders if a particular computer model is fast may be primarily interested in the computer's processing speed or in its working memory capacity. Managers understand this concern and will usually prominently display these important features.

However, while searching for this diagnostic information, consumers will almost inevitably also be exposed to *obviously irrelevant* information, i.e., information which consumers perceive as obviously irrelevant with respect to the benefit they are looking for. This irrelevant information may be encountered as a natural consequence of the search process, or it may be provided as part of a deliberate management strategy. For example, a consumer who is searching a web site for information on the processing speed of a computer, may also notice that the computer has high quality speakers or that the company sponsors an important art exhibition. Managers may have good reasons to provide this information. The high quality speakers may appeal to those consumers who value the sound quality of their computer, while the sponsorship of the arts may create general positive affect. However, for a consumer who is only trying to find out whether

this computer is fast, this information is completely irrelevant: it does not suggest that the computer is fast, but it does not indicate that it is slow either. Thus, normatively, the presence of this irrelevant information should not change consumers' assessment of the product's ability to deliver the desired benefit. A computer with a powerful processor should not be perceived as faster or slower than a computer with a powerful processor and high quality speakers, which is manufactured by a company that sponsors art exhibits.

The Effect of Irrelevant Information

Yet, many studies suggest that objectively irrelevant information may actually influence consumers' decisions. A first set of studies indicate that irrelevant product information may *polarize* consumers' beliefs in the product benefit. Stated differently, a computer with a powerful processor and high quality speakers may be perceived as *faster* than a computer with only a powerful processor. This is consistent with studies on the configural primacy effect (e.g., Asch 1946; Peterson and DuCharme 1967; Wallsten 1981), which have shown that people interpret additional information as consistent with information they have received earlier. Thus, the additional, irrelevant information may be interpreted as consistent with the supportive benefit information. Similarly, the literature on the hypothesis confirmation bias (e.g., Beyth-Marom and Fischhoff 1983; Ha and Hoch 1989; Hoch and Ha 1986; Nickerson 1998) has demonstrated that people often interpret evidence in a manner that favors the hypothesis being tested. Therefore, if the supportive information creates the hypothesis that the product will deliver the benefit,

the irrelevant information may be interpreted as supportive of the benefit. Finally, the literatures on belief polarization (e.g., Tesser 1978) and schema-congruent processing (e.g., Cantor and Mischel 1979) also suggest that the additional irrelevant information will be interpreted as consistent with the schema created by the supportive information and will lead to more extreme judgments. In sum, adding irrelevant information to supportive benefit information may strengthen consumers' beliefs that the product will deliver the benefit.

Several studies in the previously cited literatures have indeed demonstrated that beliefs can become more extreme after confrontation with neutral evidence (e.g., Beyth-Marom and Fischhoff 1983; Wallsten 1981), ambiguous evidence (e.g., Asch 1949; Ha and Hoch 1989; Hoch and Ha 1986), and, in some cases, mixed evidence (e.g., Lord, Ross, and Lepper 1979; but see Kuhn and Lao 1996). However, while these studies demonstrated that *objectively irrelevant* information can polarize beliefs, they did not examine the effect of information that is *obviously irrelevant* to decision makers, which is the focus of this research project. In fact, those studies that have explicitly addressed the impact of obviously irrelevant information have observed that adding this information to diagnostic information actually *dilutes* people's judgments. Most of the demonstrations of this dilution effect can be found in the social judgment literature (e.g., De Dreu, Yzerbyt, and Leyens 1995; Fein and Hilton 1992; Nisbett, Zukier, and Lemley 1981; Tetlock and Boettger 1989). For instance, Tetlock and Boettger (1989) observed that a student who studied 31 hours a week received a higher GPA estimate than another student who studied 31 hours a week and played tennis 3 times a month. A similar

dilution effect has also been observed in some studies using nonsocial stimuli (e.g., Shanteau 1975; Troutman and Shanteau 1977). For instance, Troutman and Shanteau (1977) showed subjects draws of beads from one of two boxes. The boxes contained different proportions of red and white beads, but an equal proportion of blue beads. They observed that when an initial draw with predominantly red beads was followed by an irrelevant draw with only blue beads, subjects reduced their confidence in the assumption that the red box was being sampled. In sum, these studies suggest that obviously irrelevant product information may weaken consumers' belief in a product's ability to deliver a desired benefit. In other words, a computer with a powerful processor and high quality speakers may be perceived as *slower* than a computer with only a powerful processor.

Overview

The results reported in this dissertation are consistent with the findings of the dilution effect literature. Across ten studies, adding irrelevant product information to supportive information systematically weakens consumers' beliefs in the product's ability to deliver the desired benefit. More importantly, these studies also provide insight into the mechanism behind this dilution effect, and show under which conditions this effect may disappear or even reverse. Previous studies have argued that the dilution effect may be caused by an averaging strategy, the use of the representativeness heuristic, or the reliance on conversational norms. This dissertation proposes a new explanation for the dilution effect based on the literature on biased hypothesis testing. It is argued that

consumers may be selectively testing the hypothesis that the product will actually deliver the desired benefit. They may selectively search for evidence that supports this hypothesis and classify all information with respect to their search goal: information that is supportive of the benefit is classified as “confirming,” while any additional information is classified as “not confirming.” Consumers’ belief in the hypothesis will be stronger to the extent that their search for confirmatory evidence produced more “confirming” evidence and will be weaker to the extent that this search produced more “not confirming” evidence. This alternative explanation will be systematically tested against the existing accounts of the dilution effect. To construct these tests, I examine a number of important factors which have not yet received attention in the dilution literature, such as the manner in which the information is processed and consumers’ awareness of the irrelevance of the information. Furthermore, I will explore when irrelevant information may in fact *strengthen* consumers’ belief in the product benefit, thus effectively reversing the dilution effect.

Before presenting the studies, I will first discuss previous research on the influence of irrelevant product information and review prior demonstrations of the dilution effect as well as the prior accounts of the effect. I will then propose a new explanation based on the biased hypothesis testing literature, which will be tested in ten experiments. Finally, I will provide a general discussion of the findings, as well as a discussion of the limitations of this dissertation and possible avenues for future research.

CHAPTER 2 THEORETICAL BACKGROUND

Irrelevant Product Information in the Consumer Research Literature

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 Nakomoto (1994) report that a brand with a distinguishing, but irrelevant attribute receives a higher preference rating than the same brand without the attribute. On the other hand, 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. Similarly, Simonson, Nowlis, and Simonson (1993) report that consumers who notice that others prefer a product for a reason which they find irrelevant are less likely to choose that product. 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 effect partially depends on the size of the choice set.

All these studies demonstrate that irrelevant information influences consumer decisions. Yet, they may not tell us much about the influence of obviously irrelevant information on consumers' beliefs in product benefits, because the situations examined in these studies differ in two important ways from the subject of this dissertation. First,

while the information presented in the preceding studies was objectively irrelevant, subjects did not *perceive* the information as irrelevant for their decision and knowingly relied on it. Subjects in Hoch and Ha (1986) perceived the ambiguous product experience as diagnostic information and interpreted the experiences as confirming the advertising claims. On the other hand, while subjects in Carpenter et al. (1994) were given information from which they could infer that the differentiating attribute did not influence the quality of the product, they did seem to make quality inferences based on this information and thus perceived it as relevant for their decision. Consistent with this assertion, the irrelevant information only increases preference ratings when the target brand is offered at a high price (study 2), and the effect of the irrelevant information even reverses when the labels have a negative rather than positive connotation (Broniarczyk and Gershoff 1997). Similarly, in the experiments by Simonson and colleagues (1993, 1994), subjects are probably aware that an unwanted feature does not change the quality of the product, but they do perceive it as unattractive and state it as a reason for not choosing the target brand. Likewise, the majority of subjects in Brown and Carpenter (2000) refer to the trivial information to justify their choices. Thus, while these previous studies indicate that consumers can be influenced by objectively irrelevant product information, they do not inform us whether consumers can be influenced by product information that they perceive to be irrelevant for their decision.

A second important distinction concerns the type of consumer decisions. None of the previous studies examined specific product judgments, such as beliefs in a product's ability to deliver a particular benefit. With the exception of Hoch and Ha (1986), all

these studies examine the influence of trivial information in a *choice* or *preference* context in which consumers have to make comparisons between brands. While this distinction may seem trivial, it is in fact of great importance, as the choice context is essential for the explanation of the effects. The irrelevant information influences consumer decisions because it differentiates the target brand from the other brands (Carpenter et al. 1994) and thus 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. 1993, 1994) and the moderating effect of the choice context (Brown and Carpenter 2000) provide strong support for these interpretations. However, these mechanisms certainly do not operate in situations in which consumers make independent assessments of a brand's ability to deliver a specific benefit. In sum, while these 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 which consumers *perceive as obviously irrelevant* with respect to these benefits. However, the psychological literature on the dilution effect does look at the influence of obviously irrelevant information on specific predictions and may thus offer a more promising perspective on this research problem. I will first review the general findings from this research stream, followed by an overview of the existing accounts of the effect.

The Dilution Effect

The dilution effect has been found using social stimuli with natural judgments and nonsocial stimuli with probability judgments. Studies using social stimuli with natural judgments always use a person as target stimulus, manipulate the presence of irrelevant information between subjects, and ask subjects to make a single judgment about a stereotypic trait (De Dreu et al. 1995; Fein and Hilton 1992; Nisbett et al. 1981; Tetlock, Lerner, and Boettger 1996; Tetlock and Boettger 1989; Zukier 1982; Zukier and Jennings 1983). For example, in Zukier and Jennings (1983), students were asked to act as jurors in the trial of a man accused of murdering his aunt. In the control condition, subjects only received evidence indicative of guilt (e.g., "he was known to have argued with his aunt" and "he had no alibi"). In the treatment condition, subjects received the same diagnostic information, as well as additional information that had no implication for the defendant's guilt (e.g., "the defendant is of average height and vision"). Zukier and Jennings (1983) observed that subjects in the treatment condition were less likely to conclude that the suspect committed the murder than were subjects in the control condition.

Studies using nonsocial stimuli with probability judgments use commodity-like objects as target stimuli, manipulate the presence of irrelevant information within-subject, and ask subjects to make multiple, explicit probability judgments (Birnbaum and Mellers 1983; Lichtenstein, Earle, and Slovic 1975; Shanteau 1975; Troutman and Shanteau 1977). For example, Troutman and Shanteau (1977) showed subjects two boxes containing different proportions of red and white beads (one box had more red beads and

one box had more white beads), but equal proportions of blue beads. Subjects were then presented with successive samples of beads and, after each sample, indicated the probability that the predominantly white box was being sampled. Troutman and Shanteau (1977) found that when a diagnostic draw (a red or white bead) was followed by an irrelevant draw (a blue bead) or neutral draw (an equal number of red and white beads), subjects moderated their estimates.

Existing Accounts of the Dilution Effect

Three major explanations of the dilution effect have been proposed in the literature. The representativeness heuristic and conversational norm explanations have been proposed in the social judgment literature, while the averaging hypothesis is the most prominent explanation in the nonsocial judgment literature. I will review these existing explanations and then propose a new explanation based on the biased hypothesis testing literature.

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 and the *typical* to-be-predicted outcome to estimate the probability that the outcome will occur (Kahneman and Tversky 1972, 1973). It is assumed that diagnostic information is highly representative of the to-be-predicted behavior and that irrelevant information is not

representative of this behavior. Therefore, adding irrelevant information to diagnostic information makes the individual less representative of the behavior and attenuates the judgment.

For example, a student who studies 31 hours a week is more similar to the stereotypical "high GPA student" than a student who studies 31 hours a week and plays tennis 3 times a month. Therefore, if people base their estimate of a student's GPA on the student's similarity to the stereotypical "high GPA student," then the student who also plays tennis 3 times a month, should receive a lower GPA rating than the student who does not. This is indeed what Tetlock and Boettger (1996) observed. Similarly, a computer with a powerful processor will be perceived as more similar to a typical "fast computer" than a computer with a powerful processor manufactured by a company that sponsors an arts exhibit. Therefore, if consumers rely on a representativeness heuristic to assess product benefits, irrelevant product information should weaken consumers' belief in the product's ability to deliver the benefit.

Averaging. The *averaging* explanation is the most popular account of the dilution effect within the nonsocial judgment literature (Birnbau and Mellers 1983; Lichtenstein et al. 1975; Shanteau 1975; Troutman and Shanteau 1977). A first group of algebraic averaging models assumes that the weights of the attributes are being adjusted according to the weights of the other attributes (e.g., Anderson 1967, 1971, 1974; Birnbau and Mellers 1983):

$$R = C + \sum_{i=0} w_i s_i / \sum_{i=0} w_i + \epsilon_{ij}$$

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 suggest that people use an adjustment model in which they make separate predictions based on each piece of information, and then average the predicted outcomes (e.g., Lichtenstein et al. 1975; Lopes 1987; Shanteau 1975). An example of the latter is Shanteau's (1975) study in which subjects were shown samples from one of two boxes with different proportions of red and white beads. A single red bead indicated that there was a 60 percent chance that the red box was being sampled; a draw with equal numbers of red and white beads implied that there was a 50 percent chance that either box was being sampled. Subjects who saw the red bead sample followed by the neutral sample adjusted their predictions downward, resulting in a prediction between 50 and 60 percent, consistent with the use of an averaging rule. Troutman and Shanteau (1977) replicated this finding, but added an equally large number of blue beads to both boxes. They demonstrated that a draw consisting of only blue beads had the same diluting effect as a draw consisting of equal numbers of red and white beads. It could be argued that, since both boxes contained a large proportion of blue beads, the blue beads draw was very representative of each box. Therefore, Troutman and Shanteau's (1977) findings suggest that reliance on the representativeness heuristic cannot explain all occurrences of the dilution effect.

However, there is considerable debate regarding the status of the averaging explanation. Anderson (1974) argues that the averaging model is an as-if model, a data model, rather than a process model: "The work on the averaging hypothesis only allows

the conclusion that the subject acts *as if* he were averaging. . . . It seems unlikely that the subject carries out the steps indicated by equation 1." (Anderson 1974, p 253). Likewise, Lichtenstein et al. (1975) propose their averaging model as an adequate description of individuals' observed behavior, but suggest that the representativeness heuristic may be a more accurate description of the actual decision process. On the other hand, Shanteau (1975) and Troutman and Shanteau (1977) present their averaging model as a process explanation and an alternative to the representativeness account. Finally, Lopes (1987) distinguishes between the algebraic averaging models, which she perceives as data models, and the adjustment models, which she presents as a possible description of the decision process.

Conversational norms. A third account of the dilution effect argues that dilution is an experimental artifact resulting from subjects' mistaken reliance on *conversational norms*. This account has recently become very prominent in the social judgment literature (e.g., Schwarz, Strack, Hilton, and Naderer 1991, Slugoski and Wilson 1998; Tetlock et al. 1996). In intentional communication, a number of conversational norms are assumed to be respected, one of which is that all information that is being provided 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 ("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 were deactivated by mentioning that the computer had randomly selected the information that was presented. However, non-accountable subjects still showed a dilution effect when conversational norms had been deactivated, suggesting that conversational norms may contribute to the dilution effect, but are not necessary for the effect to occur.

While the conversational norms account explains why subjects relied on the information despite its irrelevance, this account does not explain why adding the irrelevant information results in less extreme rather than more extreme judgments. However, the conversational norms explanation can easily be combined with the representativeness or averaging accounts to address this problem. These latter explanations describe decision mechanisms that can lead to less extreme responses, whereas the conversational norms account explains why the irrelevant information is actually incorporated into these decision mechanisms. If subjects would not assume that all the information provided by the experimenter has to be relevant, they may not incorporate this information in their similarity judgments nor use it as a basis for separate predictions which will be averaged with predictions based on the diagnostic information.

In any case, regardless of whether the reliance on conversational norms is a sufficient cause of the dilution effect, this perspective suggests that the assumption that the information provided is relevant is a necessary condition for the dilution effect to occur. Does this imply that irrelevant product information may only dilute product judgments in an experimental setting? Surely not. When consumers encounter product

information in an advertisement, they are also likely to make certain assumptions. They may, for instance, assume that the advertiser will only release information that will encourage consumers to purchase the product. Thus, consumers may assume that all the information provided is relevant. However, an advertisement does not constitute an intentional face-to-face communication. The advertiser may have good reasons to assume that the information will appeal to a certain consumer segment, but this does not imply that the information is relevant for *all* consumers, regardless of the benefits they are personally interested in. Thus, consumers may assume that the irrelevant product information is intended for a different consumer segment and may therefore not rely on this information. In sum, if the reliance on the relevance principle is a necessary condition for the dilution effect to occur, the dilution effect may not apply to many realistic advertising situations.

A Biased Hypothesis Testing Account

All three preceding mechanisms can account for a diluting effect of irrelevant product information on consumers' beliefs in product benefits. However, I propose an alternative, fourth account which will be tested against the three existing explanations within a product judgment context. It is proposed that consumers who are assessing product benefits, as well as subjects in the dilution experiments, are faced with a hypothesis testing task. For instance, consumers are testing the hypothesis that the product will deliver the benefit, while subjects in Zukier and Jennings (1983) are testing the hypothesis that the defendant is guilty. Thus, the extensive literature on biased

hypothesis testing may provide additional insight into the effect of irrelevant information in this context. The proposed *biased hypothesis testing* explanation is based on four assumptions which will be discussed next.

First, it is assumed 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. Consumers are more likely to benefit from finding out which products deliver the desired benefit than by identifying all the products that will not. Moreover, advertising claims can explicitly create the hypothesis that the product will indeed deliver the benefit (Ha and Hoch 1989; Hoch and Ha 1986). Furthermore, this preference for positively framed hypotheses may generalize beyond consumers' search for product benefits. Evidence suggests that the acceptance of an idea is part of the automatic comprehension of that idea, occurs before the rejection of the idea, and is also less effortful than the rejection (Gilbert 1991; Gilbert, Tafarodi, and Malone 1993). In sum, consumers may start out with assuming that the product will deliver the benefit, and set out to test this hypothesis rather than its complement.

Second, it is assumed that, when testing this hypothesis, consumers will search for evidence in a biased fashion. Consumers will selectively *search for confirming evidence*, i.e., evidence that suggests that the product will indeed deliver the benefit. There have been many demonstrations of a biased search for confirming evidence (e.g., Shaklee and Fischhoff 1982; Snyder and Cantor 1979; Snyder and Swann 1978) or, in rule testing, positive test cases (e.g., Klayman and Ha 1987). While early research demonstrated that this biased search strategy will often increase the likelihood that the hypothesis will be

accepted (e.g., Snyder and Swann 1978; Wason 1960), later studies indicated that a biased strategy does not necessarily lead to a biased outcome (e.g., Klayman and Ha 1987). Consistent with this latter perspective, I propose that consumers' biased search for product information that confirms the hypothesis that the product will deliver the benefit does not necessarily strengthen consumers' belief in the product benefit.

Third, consistent with the literature on goal-based classification (e.g., Barsalou 1983; Ratneshwar, Pechmann, and Shocker 1996), it is assumed that consumers will *classify the product information with respect to their search goal*. In other words, consumers will 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). While ambiguous and obviously supportive information can be interpreted as supportive and will therefore be classified as "confirming," counterdiagnostic and obviously irrelevant information can not be interpreted as supportive and will be classified as "not confirming."

Finally, it is assumed that when consumers evaluate the product information, they will only consider its implications for the focal hypothesis and ignore the consequences for the alternative hypothesis. Thus, consumers' hypothesis testing is not only confirmatory, but it is also *selective*. Several studies on selective or pseudo-diagnostic hypothesis testing have demonstrated that people indeed only consider the focal hypothesis when interpreting evidence (Beyth-Marom and Fischhoff 1983; Sanbonmatsu, Posavac, Kardes, and Mantel 1998; Sanbonmatsu, Posavac, and Stasney 1997; 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. On the other hand, 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 support the hypothesis that the product will *not* deliver the benefit. While the first process leads to a confirmation bias, the latter results in dilution.

Similar to the literature on biased information search, most demonstrations of selective hypothesis testing have shown that it can lead to an unwarranted *increase* in confidence in the focal hypothesis when the evidence is very likely given either hypothesis (e.g., Beyth-Marom and Fischhoff 1983). However, several researchers (Beyth-Marom and Fischhoff 1983; Sanbonmatsu et al. 1997, 1998) have indicated that the same strategy may actually lead to an unwarranted decrease in confidence in the hypothesis when the evidence is unlikely given any of the hypotheses. Beyth-Marom and Fischhoff (1983) have even suggested that this strategy may explain the dilution findings reported by Nisbett et al. (1981). They propose that people only consider the likelihood of the evidence given the focal hypothesis, i.e., $P(D|H)$, and do not consider the likelihood of the evidence given the alternative hypothesis, i.e., $P(D|\text{not } H)$, thus ignoring the denominator of the likelihood ratio in Bayes' theorem. Beyth-Marom and Fischhoff (1983) argue that, because Nisbett et al. (1981) designed their stimuli so that the irrelevant information did not fit with either of the possible classification categories, a selective focus on $P(D|H)$ may have decreased subjects' confidence in the focal

hypothesis. For example, some subjects in Nisbett et al. were asked to estimate the probability that a student, who was either a music or an engineering major, was in fact a music major. The irrelevant, highly personal information about the student was not very likely given that the student was a music major (i.e., $P[D|H]$ was low), but was also not very likely given that the student was an engineering major (i.e., $P[D|\text{not } H]$ was also low). Thus, the personal information may have decreased subjects' belief that the student was a music major despite its irrelevance, because students only considered the low likelihood of this information given that the student was a music major.

Yet, while this process, by itself, can indeed explain the dilution effect when the irrelevant information is unlikely given either hypothesis, it cannot account for the many demonstrations of the dilution effect using irrelevant information that does not fit this restriction. In fact, Zukier and Jennings (1983) even demonstrated that the dilution effect was more pronounced when the irrelevant information was typical rather than atypical. Therefore, while the selective focus on a single hypothesis is a *necessary* assumption for a biased hypothesis testing account of the dilution effect, it is not *sufficient*. However, when it is combined with the previous three assumptions, it may indeed account for a wide range of dilution effects.

In sum, it is proposed that consumers are testing the hypothesis that the product will deliver the benefit. They selectively search for information that suggests that the product will indeed deliver the benefit. Consumers will then classify product information with respect to this search goal. Obviously supportive and ambiguous information will be classified as "confirming," while counterdiagnostic and obviously irrelevant

information will be classified as “not confirming.” When information is classified as “confirming,” it will strengthen consumers’ belief in the benefit, even when it also supports that the product will not deliver the benefit. On the other hand, when information is classified as “not confirming,” it will weaken consumers’ belief in the benefit, even when it also does not support that the product will not deliver the benefit.

For instance, suppose that a consumer is testing an advertising claim and is confronted with an *ambiguous* product experience. According to the process described above, the consumer will search for information that confirms the advertising claim and, since the ambiguous experience can easily be interpreted as consistent with the claim, she will classify it as “confirming” and strengthen her belief in the claim. This is indeed what Hoch and Ha observed in their studies (Ha and Hoch 1989; Hoch and Ha 1986). However, what would happen if this consumer would encounter *obviously irrelevant* instead of ambiguous product information? In that case, she cannot interpret this information as confirming the claim and she will classify it as “not confirming.” Moreover, her belief in the advertising claim will be weakened, even though the information does not confirm that the claim is not true either. Thus, while this biased hypothesis testing process may produce a confirmation bias when the product information is ambiguous, the same process may produce a dilution effect when the product information is obviously irrelevant.

The generality of the biased hypothesis testing explanation. While the biased hypothesis testing explanation was formulated in a product judgment context, this same mechanism may also apply to prior demonstrations of the dilution effect in other domains

(although I do not argue that one single mechanism is responsible for all observations of the dilution effect). Consider, for instance, Shanteau's (1975) experiment in which subjects draw from two boxes with different proportions of red and white beads. If someone first draws predominantly red beads, this may create the hypothesis that the beads are being sampled from the "red box." The person will now be looking for information that confirms this hypothesis, i.e., draws with predominantly red beads. If the second sample contains equal numbers of red and white beads, the sample will be perceived as "not confirming" and confidence in the hypothesis will be reduced.

The four assumptions on which this alternative explanation is based will not hold in all situations. For instance, consumers may not start out with the hypothesis that the product will deliver the benefit, when they have a very negative prior opinion of that product. In that case, they may test the hypothesis that the product will *not* deliver the benefit. Similarly, consumers may sometimes consider the implications of the product information for both the focal and alternative hypotheses (i.e., "Does this suggest that the product will deliver the benefit, or does this suggest that the product will not deliver the benefit?"). As will be demonstrated, the dilution effect will disappear, or even reverse, in these situations.

Comparison with existing explanations. The proposed biased hypothesis testing explanation is not incompatible with previous accounts of the dilution effect. As was mentioned earlier, the averaging account may provide a description of people's behavior rather than a description of their decision process (e.g., Lichtenstein et al. 1975). Thus, it is plausible that when people engage in biased hypothesis testing, they may generate

responses that can often be perfectly described by an averaging model of information integration. However, there are situations in which the predictions of the hypothesis testing account can not be adequately represented by an averaging model. One such situation will be examined in experiment 2.

The representativeness explanation, on the other hand, is more similar to the biased hypothesis testing account. In fact, one could interpret the representativeness heuristic as a heuristic hypothesis testing strategy: people may test hypotheses by relying on the similarity between the evidence and the typical hypothesized outcome. However, there are two important distinctions between the proposed explanation and the representativeness account. On the one hand, the representativeness explanation is *less specific* than the biased hypothesis testing account. The biased hypothesis testing explanation explicitly states that, for the dilution effect to occur, the product information needs to be processed with the hypothesis (i.e., the desired benefit) in mind. While this constraint can be added to the representativeness explanation (i.e., the information needs to be processed with the desired prototype in mind), it is not essential. If consumers first form an impression of the described product and only then compare the product to the desired prototype, the dilution effect could still occur, since the irrelevant product information would still make consumers' impression of the product less similar to the typical desired product. On the other hand, the representativeness explanation is also *more specific* than the biased hypothesis testing account. The biased hypothesis testing account does not specify how consumers decide whether to classify a piece of product information as “confirming” or “not confirming.” They may make inferences based on

prior theories, or they may rely on heuristics such as similarity judgments. However, while the representativeness account states that this reliance on similarity judgments is essential for the explanation of the dilution effect, it is not a necessary requirement for the biased hypothesis testing account.

The following chapters will present ten experiments in which the biased hypothesis testing account is tested against each of the three existing explanations. First, experiments 1 and 1A demonstrate the effect, show its robustness across product categories, presentation orders, and belief measures, and provide evidence against a distraction of attention account. Then, experiment 2 demonstrates that the diluting effect of irrelevant product information can not be accounted for by an averaging model, while experiments 3 and 3A provide evidence against a conversational norms explanation of the effect, and experiment 4 shows evidence that is inconsistent with the representativeness account. While the findings of these first experiments are all consistent with the biased hypothesis testing explanation, the remaining experiments will more systematically manipulate the assumptions of the proposed 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 also disappears when consumers consider the implications of the evidence for both the focal and alternative hypothesis. Finally, experiment 7 demonstrates that the effect even reverses when consumers set out to test the hypothesis that the product will *not* deliver the benefit.

CHAPTER 3

EXPERIMENT 1: DEMONSTRATION OF THE DILUTION EFFECT

The objective of this first experiment is to demonstrate the diluting effect of irrelevant product information and provide an initial test of a distraction explanation of this effect. Although there is considerable evidence for dilution in social and non-social judgments, these demonstrations will not necessarily generalize to a product judgment context. Unlike subjects in social judgment experiments, consumers who make product judgments can not rely on easily accessible stereotypes, which may be an essential requirement for some decision mechanisms to occur (e.g., representativeness). Similarly, the abstract cues and within-subject manipulations, typical of the non-social demonstrations of the dilution effect, seem to have few parallels in product judgment contexts. Furthermore, many studies have shown that adding nondiagnostic neutral or ambiguous information can *polarize* rather than dilute judgments (e.g., Asch 1949; Beyth-Marom and Fischhoff 1983; Ha and Hoch 1989; Hoch and Ha 1986; Wallsten 1981). Therefore, given the limited overlap between the product judgment context and contexts in which dilution effects have been observed, and given the evidence that nondiagnostic information can also lead to polarization, experiment 1 provides an initial test of the dilution effect in a product judgment context.

A second objective of this first experiment is to test whether the irrelevant information is diluting consumers' beliefs by distracting resources otherwise allocated to

the diagnostic information. While this distraction explanation has not been tested previously, it can account for the dilution effect by assuming that people extract less information from the supportive evidence when their attention is distracted by the irrelevant evidence. To test this explanation, I measured subjects' recognition of the diagnostic information at the end of the experiment. If the irrelevant information did indeed 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 easily than those who had only received the supportive information.

Each of the ten experiments reported in this dissertation used a very similar procedure in which subjects received descriptions of different products or services and were asked to indicate their belief that this product or service would deliver a specific benefit. In experiment 1, subjects were presented with eight different products or services. For each category, they were first given a specific desirable benefit (e.g., "You are looking for a fast computer"). Subjects then received the product description. In the baseline condition, the product description only contained one piece of supportive information that strongly suggested that the product would deliver the benefit (e.g., "Very Powerful Processor"), whereas in the treatment condition, the description also contained three pieces of irrelevant information (e.g., "Assembled in the USA," "Airs Commercials on NBC & CBS," and "Can Be Ordered On-Line"). After subjects had received the entire product description, they were asked to indicate whether the product would deliver the desired benefit. When all descriptions had been shown, subjects were given a brief reaction time training, which provided a baseline measure of subjects' response latencies.

Finally, subjects were given a recognition task, in which they were presented with the eight pieces of supportive information, as well as eight new pieces of information. They were asked to indicate, as fast as possible, if the information had been presented earlier in the experiment.

First, it was expected that the addition of irrelevant information would weaken beliefs in the product benefit, thus demonstrating the dilution effect. In other words, subjects who had received both supportive and irrelevant information were expected to provide lower belief ratings than subjects who had only received supportive information. Second, if this dilution effect was due to a distraction of resources by the irrelevant information, the addition of irrelevant information should impair the recognition of the supportive information.

Method

Subjects and design. Subjects were 36 undergraduate students who participated in return for class credit. The design was a 2 (type of information) by 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.

Stimuli. Eight different products or services were selected and a desirable benefit was specified for each category: apartments (safe), package delivery service (fast), frozen dinners (healthy), airlines (superior service), toothpaste (fights cavities), car (sportive), stereo system (reliable), and computers (fast). For each replicate, three irrelevant

attributes had to be selected, as well as one attribute that strongly suggested that the product would deliver the benefit, i.e., the supportive attribute (see Appendix B for a complete list of the product information presented in this and other experiments). A pretest ($n = 30$) was conducted to select these attributes. 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. Each of the irrelevant facts was classified as "not helpful" by at least 80 % of pretest 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 NYC Marathon), and product availability (e.g., a frozen dinner brand that is available at most grocery stores). The eight supportive facts selected for the main experiment were classified as suggesting the benefit by an average of 94% of pretest subjects.

While the first pretest demonstrated that subjects perceived the information as irrelevant in isolation, it was possible that this apparently irrelevant information became relevant in the context of the complete product description. Therefore, a second pretest was conducted to measure subjects' perception of the information in the context of the complete 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 irrelevant 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. The eight supportive facts were classified as suggesting the benefit by an average of 93 % of pretest subjects, as irrelevant by 6 % of the subjects, and as counterdiagnostic by less than 1 % of the subjects.

A final pretest examined the possibility that even though pretest subjects indicated that the irrelevant facts were not diagnostic, subjects in the actual experiment may still use these facts as a basis for forming beliefs about the product 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 6-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 ($t(432) = 7.22, p < .001$). Thus, if the irrelevant information would indeed lead to inferences in the experimental context, these inferences would support subjects' beliefs in the product benefit, rather than counteract them.

Procedure. The entire experiment was administered by personal computer. Subjects entered the experimental lab and were asked to sit at a computer. All

instructions were provided by the computer program. Subjects were first informed that they would receive information about eight different products (services) and that they would have to indicate their belief that the described 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" (see Appendix A for the complete instructions used in this and other experiments). Subjects then received the information for the first replicate. First, subjects were informed of the benefit they were looking for (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 *treatment* condition, this information was followed by three additional pieces of irrelevant information that were presented sequentially (e.g., "Assembled 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 specify their belief that the product would deliver the benefit (e.g., "Is this computer fast?"). Responses were made on a 9-point scale (e.g., anchored by 1 = "Definitely Not Fast" and 9 = "Definitely Fast"). Examples of the complete screen displays in the *baseline* and *treatment* conditions can be found in Appendix C. After subjects had rated their belief in the benefit, they received the information for the next replicate, until all eight product categories had been presented. The order of the replicates was randomized.

After all replicates had been presented, subjects were given a reaction time training task. Subjects were shown six simple statements unrelated to the experiment

(e.g., “Paris is the capital of France”) and had to answer “True” or “False” as fast as possible by pressing the “1” or “0” key on their keyboard. Subjects received feedback about their response time and accuracy after each statement. For each subject, the average reaction time of the accurate responses was recorded for use as a response latency covariate in the analysis. Finally, subjects were shown the actual reaction time task. They were told that they would receive statements about the product descriptions they had received earlier and would have to indicate whether they were true or false using the same technique as in the training task. For each statement, subjects were first shown a warning screen, which mentioned the product for which they would receive a statement (e.g., "The next statement concerns: The Computer"), and asked subjects to place their fingers on the “0” and “1” keys. The screen then displayed a true statement, such as “The computer had a very powerful processor,” or a false statement, such as “The computer was loaded with games.” For each replicate, consumers received one true and one false statement. The true statements always concerned the supportive piece of information. The false statements contained randomly chosen positive features that were unrelated to the desired benefit. The presentation order of the 16 statements was randomized.

Results

The addition of the irrelevant information significantly weakened subjects' beliefs in the product's ability to deliver the desired benefit ($F(1,272) = 6.56, p = .01$). Subjects who only received the supportive information reported more extreme judgments ($\bar{X} = 6.28$) than those who also received the irrelevant information ($\bar{X} = 5.83$). 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. These results indicate that the diluting effect of irrelevant information does indeed generalize to product judgments.

Subjects' recognition of the supportive information at the end of the experiment was very high overall. Interestingly, subjects who only received supportive information were slightly *less* likely to recognize the supportive information ($\hat{\pi} = 93\%$) than were subjects who also received the irrelevant information ($\hat{\pi} = 97\%$; $\chi^2 = 2.91, df = 1, p < .1$). Those subjects who recognized the supportive information did not show any differences in reaction time depending on the product information they had received earlier ($RT_{SUPP} = 1.53, RT_{SUPP+IRREL} = 1.57; F < 1, ns$).

Discussion

The results of experiment 1 demonstrate that the diluting effect of irrelevant information generalizes to product judgments. Across eight different products and services, adding irrelevant product information to supportive benefit information systematically weakened subjects' beliefs in the product's ability to deliver the desired benefit. Furthermore, the findings of the recognition task indicate that this dilution effect is not due to the fact that the irrelevant information diverts resources which otherwise would have been used to process the supportive information. In fact, those subjects who only received supportive information were marginally less likely to recognize the

¹ For the sake of clarity and conciseness, non-significant interactions with the replicate factor will not be discussed for the remaining experiments. All interactions are non-significant, unless explicitly stated otherwise.

supportive information than were those who also received the irrelevant information. Furthermore, there were no differences between the average reaction times in the two information conditions. These results suggest that the supportive information was not processed to a lesser degree when additional irrelevant information was present, which is inconsistent with a distraction of attention account of the dilution effect.

Experiment 1A

While the first experiment demonstrated the diluting effect of irrelevant information and provided evidence against a distraction of resources explanation, experiment 1A tested the robustness of this phenomenon. First, the experiment examined whether the dilution effect also occurs when consumers are comparing multiple products and select the product that is most likely to deliver the benefit. In experiment 1, the addition of irrelevant information could not only have influenced consumers' perception of the described product, but could also have changed consumers' perception of the belief measure. In other words, similar scores on the 9-point belief scale may have reflected different perceptions depending on the information condition. This problem was addressed in experiment 1A by using a selection task rather than a belief measure.

Furthermore, this experiment also manipulated the order in which the information was presented. Consistent with most previous demonstrations of the dilution effect (e.g., Fein and Hilton 1992; Hilton and Fein 1989; Nisbett et al. 1981), the supportive information was always presented 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 elicitation. While the first experiment demonstrated that, at the end of the experiment, the accessibility of the supportive information was not reduced by the addition of irrelevant information, it is still possible that the specific ordering of the information temporarily reduced its accessibility when the belief measure was presented. To test this possibility, experiment 1A manipulated the order in which the information was presented. 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. If the dilution effect was due to the reduced salience of the supportive information at the time of the belief elicitation, then the irrelevant information should not influence subjects' beliefs when it is presented before the supportive information.

Subjects and design. Subjects were 131 undergraduate students who participated in return for class credit. The design was a 2 (type of information) by 2 (order of information) by 2 (counterbalancing factor) by 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. The product information used was the same information as used in experiment 1, with the exception of eight additional pieces of supportive information. 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. The presentation of the target product was identical to the procedure used in experiment 1. For each product category, subjects were first informed of the desired benefit, and then received the description of the target product (“product A”), which contained either only one piece of supportive information or also three pieces of irrelevant information. The information for the alternative product (“product B”), which always consisted of one piece of supportive information, was then displayed below the description of the target product. The supportive information used for the two products was counterbalanced. For half of the subjects, the target product descriptions listed the supportive information used in experiment 1, while the descriptions of the alternative product used the eight new pieces of supportive information. For the other half of the subjects, this assignment was reversed. When the information for both products had been displayed, subjects were asked to indicate which of the two products was more likely to deliver the desired benefit. Subjects indicated their choice by clicking on a button with the corresponding product label.

Results and discussion. When both product descriptions only contained supportive information, 53% of subjects indicated that the target product was more likely to deliver the benefit, compared to 47% who selected the alternative product. This difference was not significant ($Z = 1.59$, *ns*), as would be expected since the supportive information was counterbalanced. 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 less than 38% ($\chi^2 = 26.26$, $df = 1$, $p < .001$). While subjects were, on average, indifferent between the target product and the alternative product when both were described using only 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 ($Z = 5.56$, $p < .001$). The effect of the irrelevant information did not depend on the order in which the product information was being presented ($\chi^2 = 1.81$, $df = 1$, ns). Adding the irrelevant information decreased the proportion of subjects selecting the target product both when the supportive information was presented first ($\chi^2 = 5.38$, $df = 1$, $p < .05$), and when it was presented last ($\chi^2 = 24.13$, $df = 1$, $p < .001$).

Experiment 1A thus demonstrates that adding irrelevant product information not only weakens beliefs about isolated products, but also reduces the likelihood that a product is perceived as more likely to deliver a benefit than an alternative product. Moreover, the dilution effect does not depend on the order in which the product information is presented, thus refuting that the effect is caused by the reduced salience of the supportive information at the time of the decision.

CHAPTER 4

EXPERIMENT 2: A TEST OF THE AVERAGING ACCOUNT

The first experiments demonstrated the robustness of the dilution effect and provided evidence against explanations of the effect based on the distraction of attention by the irrelevant information. The following experiments will examine the biased hypothesis testing account of the effect by systematically testing it against each of the existing explanations. Experiment 2 provides such a test between the biased hypothesis testing account and the *averaging* explanation.

A first way to test between these two explanations is to examine the effect of adding *less supportive*, rather than irrelevant information. Less supportive information is information that suggests that the product will deliver the benefit, but is not as strong as the original information. For example, the fact that a computer has a well-known brand name suggests that it may be fast, but it is not as convincing as the fact that it has a powerful processor. According to the averaging model proposed by Lichtenstein et al. (1975), adding less supportive product information to strongly supportive information should also dilute product beliefs. The less supportive information by itself should result in judgments that are less extreme than judgments based on the strongly supportive information. Averaging these separate judgments should result in an overall judgment that lies between these two judgments and is less extreme than a judgment based only on the supportive information.

H1: Averaging: Adding less supportive information to strongly supportive benefit information will *weaken* consumers' beliefs in the product benefit.

On the other hand, the biased hypothesis testing account predicts that adding less supportive information to strongly supportive information should lead to more extreme judgments. This account assumes that consumers are looking for information that confirms the hypothesis that the product will deliver the benefit. Since the less supportive information suggests that the product will deliver the benefit, it will be classified as “confirming,” rather than “not confirming,” and it will strengthen consumers’ belief in the hypothesis that the product will deliver the benefit.

H2: Biased Hypothesis Testing: Adding less supportive information to strongly supportive benefit information will *strengthen* consumers' beliefs in the product 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

reduce consumers' sensitivity to the supportive information. Or, in other words, the irrelevant information dilutes the impact of the diagnostic information. This can be tested by comparing the difference in consumers' belief ratings based on "supportive" versus "less supportive" information to the difference in belief ratings based on "supportive + irrelevant" versus "less supportive + irrelevant" product information. If consumers rely on an averaging strategy, the manipulation of the degree of support should have less influence on consumers' belief in the benefit when the irrelevant information is present than when it is not.

H3: Averaging: The difference in consumers' responses to strongly supportive versus less supportive benefit information will be less pronounced when irrelevant information is added to the product descriptions.

On the other hand, the biased hypothesis testing account predicts that, instead of diluting the impact of the diagnostic information, the irrelevant information will have an *independent* diluting effect on consumers' beliefs. According to this explanation, the irrelevant information is classified as "not confirming," and this classification directly weakens consumers' belief in the hypothesis that the product will deliver the benefit. Therefore, the irrelevant information will not decrease consumers' sensitivity to the supportive information.

H4: Biased Hypothesis Testing: The difference in consumers' responses to strongly supportive versus less supportive benefit information will not be affected by the addition of irrelevant information to the product descriptions.

Method

Subjects and design. Subjects were 58 undergraduate students who participated in return for class credit. The design was a 5 (type of information) by 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 (i.e., the baseline condition in experiment 1), the *supportive + irrelevant* information condition (i.e., the treatment condition in experiment 1), 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 and added three pieces of less supportive information for each replicate. A pretest was conducted to select the 24 less supportive facts. Thirty subjects were asked to classify different pieces of information for each replicate as either supportive of the benefit, counterdiagnostic of the benefit, or not helpful for the decision. The objective was to select information that was perceived as supportive of the benefit by the majority of subjects, yet was not as strong as the original supportive information. On average, the less supportive information was classified as suggesting the benefit by 71% of the pretest subjects. The design of the actual experiment included a manipulation

check to test whether the less supportive information was indeed weaker than the original supportive information.

The procedure was identical to the one used in experiment 1, with the exception that three information conditions were added to the design. In the *supportive + less supportive* condition, the product description contained one piece of supportive information (e.g., "Very Powerful Processor") followed by 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. This randomization was performed independently for each subject. Finally, in the *less supportive + irrelevant* information condition, the product description contained one randomly selected piece of less supportive information, which was followed by 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 ($\bar{X}_{\text{SUPP}} = 5.76$, $\bar{X}_{\text{SUPP+IRR}} = 5.17$; $F(1,424) = 6.23$, $p < .05$) or less supportive information ($\bar{X}_{\text{LESS_SUPP}} = 4.76$, $\bar{X}_{\text{LESS_SUPP+IRR}} = 4.00$; $F(1,424) = 5.87$, $p < .05$). Thus, the dilution effect observed in the first experiments was replicated in experiment 2. Moreover, these results demonstrate that irrelevant product information also dilutes

consumers' belief when the original product information is only weakly supportive of the product benefit.

Second, subjects' beliefs in the product benefit were stronger when the product information they had received was supportive ($\bar{X} = 5.76$) rather than less supportive ($\bar{X} = 4.76$; $F(1,424) = 12.91, p < .001$). This confirms that the manipulation of the degree of support was successful.

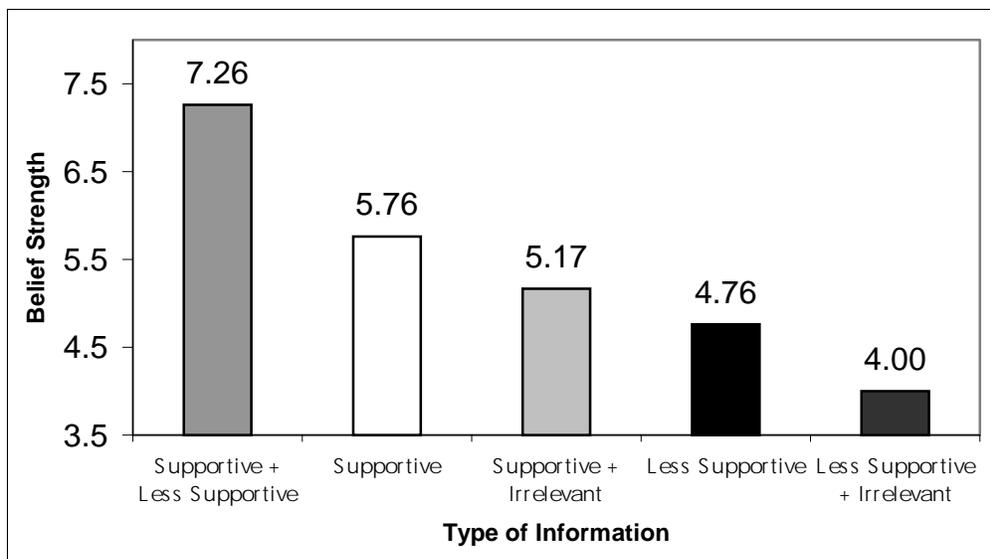


FIGURE 1
EXPERIMENT 2: EFFECT OF THE TYPE OF INFORMATION
ON SUBJECTS' BELIEF IN THE PRODUCT BENEFIT

Third, subjects receiving both supportive and less supportive information reported more extreme judgments ($\bar{X} = 7.26$) than subjects receiving only supportive information ($\bar{X} = 5.76$; $F(1,424) = 22.42, p < .001$). While adding irrelevant information to supportive information weakened subjects' belief in the product benefit, adding less supportive information to supportive information actually strengthened subjects' belief in

the benefit. This polarization effect is inconsistent with the averaging model (H1), but consistent with the biased hypothesis testing explanation (H2), which predicts that consumers will strengthen their belief in the benefit after classifying the less supportive information as “confirming” the hypothesis. However, while this result clearly indicates that subjects are not simply averaging separate predictions based on each piece of information, it can be accounted for by averaging models that include an initial impression (e.g., Anderson 1967; Lopes 1987). These latter models propose that people’s averaging calculus does not only include separate predictions based on each piece of explicitly provided information, but also includes a “baseline prediction” which reflects the decision maker’s initial impression prior to receiving any information. If subjects’ initial impression is very low, then averaging this impression and the prediction based on the supportive evidence could indeed result in weaker beliefs than would averaging this impression, the prediction based on the supportive evidence, and the three predictions based on the pieces of less supportive evidence.

To address this alternative account, I conducted a second test of the averaging explanation by examining how the addition of irrelevant information influenced subjects’ sensitivity to the supportive information. Even averaging models that include an initial impression (e.g., Anderson 1967) predict that adding information that receives a non-zero weight in the decision should reduce people’s sensitivity to the original information. Since the irrelevant information weakens subjects’ belief, it should receive a non-zero weight in the averaging model, and thus reduce the weight of the other, diagnostic information. However, the results demonstrate that the addition of irrelevant information

does not reduce the impact of the degree of support manipulation. Interestingly, the difference between the supportive information condition and the less supportive information condition becomes even slightly more pronounced when the irrelevant information is also present ($\bar{D}_{\text{NO_IRR}} = 1.00$, $\bar{D}_{\text{IRR}} = 1.17$; $F < 1$, *ns*). In sum, these results are inconsistent with an averaging account of the dilution effect, according to which the irrelevant information reduces the impact of the diagnostic information (H3), but is consistent with the biased hypothesis testing account which states that the irrelevant information directly affects the belief measure by being classified as “not confirming” (H4).

This last result not only provides evidence against an averaging account of the dilution effect, it also provides additional evidence against the distraction of resources account, and, more importantly, provides some insight into the nature of the dilution effect. While many researchers have assumed that the irrelevant information dilutes the impact of the diagnostic information (e.g., Fein and Hilton 1992; Hilton and Fein 1989), these findings suggest that the irrelevant information has an independent, direct effect on consumers’ beliefs. The irrelevant product information does not change the impact of the diagnostic information, but directly reduces consumers’ belief that the product will deliver the desired benefit.

CHAPTER 5
EXPERIMENT 3: A TEST OF THE CONVERSATIONAL NORMS ACCOUNT

The previous experiments demonstrated the robustness of the dilution effect and provided evidence that was inconsistent with 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 for the goal of the conversation (Grice 1975). Subjects may assume that all the information provided by the experimenter has to be relevant for the judgment they are asked to make and may therefore use all the available information to arrive at a decision. Thus far, evidence for this alternative hypothesis is mixed. First, conversational norms explain why people rely on the irrelevant information, but do not explain why they use it to dilute rather than polarize their judgments. In fact, the third pretest of the first experiment showed that the majority of subjects classified the irrelevant information as diagnostic of the benefit when forced to regard it as relevant. Thus, if subjects in the previous experiments perceived the irrelevant information as relevant, they were more likely to perceive it as supportive than counterdiagnostic. Since experiment 2 demonstrated that the addition of less supportive information tends to result in more extreme judgments, this suggests that treating

the irrelevant information as relevant should have lead to polarization rather than dilution. Furthermore, subjects in each of the preceding experiments were explicitly told that the information they received “may or may not be helpful for the decision that you have to make”. This should have indicated to subjects that the common principle of relevance did not apply to the experimental setting. While these arguments are far from conclusive, they do cast doubt on a conversational norms account of the observed dilution effects.

Nevertheless, it is important to provide a direct test of the conversational norms explanation in this product judgment context. First, many researchers have endorsed a conversational norms explanation of the dilution effects observed in the social judgment literature (e.g., Schwarz et al. 1991; Simonson et al. 1994; Slugoski and Wilson 1998; Tetlock and Boettger 1989; Tetlock et al. 1996). Moreover, there is evidence that subjects in these experiments indeed rely on conversational norms and that this contributes to the dilution effect (e.g., Schwarz et al. 1991; Tetlock and Boettger 1989; Tetlock et al. 1996). Finally, and most importantly, no research to date has directly tested this explanation by examining whether subjects in dilution experiments perceive the additional information as irrelevant. Experiment 3 will provide such a test by asking subjects whether the additional information is relevant before they state their belief in the product benefit.

On the one hand, one could argue that, if subjects in the dilution studies rely on the relevance principle, they should classify the additional information as supportive or counterdiagnostic. However, by asking subjects to specify the relevance of the

information, the procedure also reveals that the traditional guarantee of relevance does not apply. Hence, subjects are likely to behave as subjects in the pretests and classify the information as irrelevant. How would the dilution effect be influenced by subjects' acknowledgement of the irrelevance of the information immediately before stating their belief in the benefit? According to the conversational norms explanation, subjects who have just acknowledged the irrelevance of the information should realize that the relevance principle does not hold and simply ignore the irrelevant information, thus eliminating the dilution effect.

H5: Conversational Norms: Adding irrelevant product information to supportive benefit information will not affect consumers' belief in the benefit when consumers acknowledge the irrelevance of the information.

According to the biased hypothesis testing account, on the other hand, subjects in the earlier experiments were also aware that the information was irrelevant, but automatically classified it as “not confirming” in their search for hypothesis confirming evidence. Thus, having consumers acknowledge the irrelevance of the information should not influence the dilution effect.

H6: Biased Hypothesis Testing: Adding irrelevant product information to supportive benefit information will weaken consumers' beliefs in the benefit, even when consumers acknowledge the irrelevance of the information.

Method

Subjects and design. Subjects were 47 undergraduate students who participated in return for class credit. The design was a 2 (type of information) by 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 4 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 twelve 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 added to these six target replicates. The fillers were inserted in the second and fifth position and served to make the pre-programmed structure of the stimuli less obvious.

Subjects received the same instructions as those who participated in experiment 1. 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 could respond 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 on the same 9-point scale that was used in experiments 1 and 2.

Results and Discussion

Consider first the classification of the product information. The irrelevant information was classified as "not helpful" on 360 of 390 occasions (92%). On 15 occasions (4%) it was classified as diagnostic and on 15 other occasions (4%) it was classified as counterdiagnostic. This indicates that, on average, subjects clearly perceived the additional information as irrelevant. While this result does not necessarily imply that the additional information was also perceived as irrelevant in the previous dilution studies, it does allow us to determine the effect of the irrelevant information when its irrelevance is made very explicit immediately prior to the belief ratings.

The analysis of the belief ratings reveals that, even after subjects have acknowledged the irrelevance of the additional information, the irrelevant information still dilutes their belief in the product benefit. Subjects' belief ratings were more extreme when the product description only contained supportive information ($\bar{X} = 6.84$) than when it also contained irrelevant information ($\bar{X} = 5.54$, $F(1,252) = 51.89$, $p < .01$), consistent with the biased hypothesis testing account (H6), but inconsistent with the conversational norms explanation (H5). Most importantly, 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 in the product's ability to deliver the benefit ($\bar{X}_{\text{SUPPORTIVE}} = 6.84$, $\bar{X}_{\text{IRRELEVANT}} = 5.49$, $F(1,228) = 50.79$, $p < .01$).

Experiment 3A

The results of experiment 3 clearly indicate that subjects perceive the additional information as irrelevant and that they still display the dilution effect when this irrelevance is highly salient. 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, an additional study ($n = 21$) was conducted in which subjects were informed that the information was randomly selected by a computer. This procedure is similar to manipulations used in previous studies that have found support for 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. Moreover, before each piece of information was retrieved, the message "Randomly Drawing Information" was displayed, as well as a rapidly filling clock, thus reminding subjects that the computer was sampling information. Otherwise, the procedure was identical to the procedure used in experiment 1, but using the expanded stimulus set of experiment 3.

The results showed that adding irrelevant product information to supportive information again weakened subjects' beliefs in the product's ability to deliver the benefit ($\bar{X}_{\text{SUPP}} = 6.95$, $\bar{X}_{\text{SUPP+IRR}} = 5.94$, $F(1,321) = 8.83$, $p < .01$). Together, the results from studies 3 and 3A rule out a conversational norms explanation of the dilution effects observed in the current research. The dilution effect still occurs when subjects are explicitly aware of the irrelevance of the additional information and when they assume that the information is being randomly sampled by a computer. This does not only indicate that the dilution effect is not caused by a reliance on conversational norms, but also informs us about the range of situations in which the dilution effect may be observed. These results suggest that irrelevant product information may still weaken consumers' product beliefs when it is encountered in a mass advertising situation, in which consumers realize that the information is not necessarily intended for them and that the maxims of intentional communication may therefore not apply.

CHAPTER 6
EXPERIMENT 4: A TEST OF THE REPRESENTATIVENESS ACCOUNT

While the results of the preceding experiments are inconsistent with averaging and conversational norms accounts of the dilution effect, they do confirm the predictions of 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, similar to biased hypothesis testing, a reliance on the representativeness heuristic also predicts that the dilution effect will persist when subjects are very aware of the irrelevance of the information, and when the information seems to be randomly selected. Moreover, the addition of the irrelevant information directly reduces the similarity of the product description to the typical desired product. Therefore, the representativeness account also predicts that the irrelevant information directly affects consumers' beliefs in the product benefit, without reducing consumers' sensitivity to the diagnostic information. Since both explanations can account for this entire pattern of results, experiment 4 was designed to test between these two competing accounts.

To test between these explanations, the product information for experiment 4 was selected so that adding the irrelevant information to the supportive information would actually *increase* the similarity to the typical desired product, rather than reduce it. In the previous studies, the addition of irrelevant information tended to reduce the similarity between the product description and the typical desired product. For instance, an apartment with 24 hour on-site security is more similar to the typical safe apartment, than an apartment with 24 hour on-site security, a 40-year old manager, 1 and 2 bedroom apartments, and named “Haywood Park.” However, in this experiment, the product information was selected so that adding irrelevant information would increase the similarity with the typical desired benefit. For instance, an apartment with surveillance cameras above all doors is less similar to the typical safe apartment than an apartment with surveillance cameras above all doors, a fully equipped club house, a great fitness center, and monthly pest control included in the rent. If adding this irrelevant product information indeed increased the perceived typicality of the product description, *and* if consumers are relying on a representativeness heuristic to predict the product’s ability to deliver the desired benefit, then adding irrelevant information should strengthen rather than weaken consumers’ beliefs in the product benefit.

Furthermore, the experiment also contained two conditions in which subjects were encouraged to either follow a biased hypothesis testing strategy or a representativeness strategy. By comparing these conditions to a condition in which subjects were not directed to one particular strategy, it could be tested whether subjects spontaneously behaved as if they were engaging in biased hypothesis testing, or if they

spontaneously behaved as if they were relying on a representativeness heuristic. This resulted in four strategy conditions: a *typicality* condition, a *free strategy* condition, a *biased hypothesis testing* condition, and a *representativeness* condition.

In the *typicality* 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. It was assumed that adding the specially selected irrelevant information would increase the perceived similarity with the typical desired product. In the *free strategy* condition, subjects were presented with the same product descriptions and, as in previous experiments, 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 “confirming” or “not confirming” that the product will deliver the benefit, similar to the process proposed by the biased hypothesis testing account. After classifying all pieces of information, subjects rated their belief in the product benefit. Finally, in the *representativeness* condition, subjects were first presented with the entire product description, and were then asked to rate the similarity between this description and the typical desired product. Following the similarity rating, these subjects also rated their belief in the product benefit.

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 *free strategy* condition should lead to polarization instead of dilution. Moreover, when subjects 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 *free strategy* condition. Thus, if subjects are spontaneously relying on a representativeness heuristic, there should be no difference between the effect of irrelevant information in the *free strategy* condition and its effect in the *representativeness* condition.

H7: Representativeness: Adding irrelevant product information that enhances the product's similarity to the typical desired benefit will *strengthen* consumers' beliefs in the product benefit, regardless of whether consumers are encouraged to use similarity ratings as a basis for their belief in the product's ability to deliver the benefit.

However, if the proposed biased hypothesis testing mechanism applies, adding irrelevant information should dilute product beliefs, regardless of its effect on the perceived typicality of the product description. Even the highly typical irrelevant information will be classified as “not confirming,” resulting in weaker beliefs in the product benefit. Furthermore, when subjects are asked to classify all information as either “confirming” or “not confirming,” they should produce belief ratings that are similar to those observed in the *free strategy* condition. Thus, if subjects are spontaneously using a biased hypothesis testing strategy, there should be no difference between the effect of irrelevant information in the *free strategy* condition and its effect in the *biased hypothesis testing* condition.

H8: Biased Hypothesis Testing: Adding irrelevant product information that enhances the product's similarity to the typical desired benefit will *weaken* consumers' beliefs in the product benefit, regardless of whether consumers are encouraged to use a biased hypothesis testing strategy to assess the product's ability to deliver the benefit.

Method

Subjects and design. Subjects were 83 undergraduate students who participated in return for class credit. The design was a 2 (type of information) by 4 (strategy conditions) by 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, I selected atypical supportive information (e.g., a computer with "revolutionary triple processors") and typical irrelevant information (e.g., "includes DVD player"). While the 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 does enhance the similarity to the typical fast computer. Subjects' perception of the new information was measured using the same pretest as used in experiment 1 ($n = 40$). The eight supportive facts were classified as suggesting the benefit by an average of 93%

of pretest subjects. The 24 irrelevant facts were classified as "not helpful" by an average of 93% of the subjects, as suggesting the benefit by 5% of subjects, and as counterdiagnostic of the benefit by only 2% of subjects. The *typicality* condition in the actual experiment tested whether this irrelevant information did indeed increase the perceived similarity to the typical desired product.

The procedure in the *free strategy* condition was identical to the one used in experiment 1. For each replicate, subjects were told that they were looking for a specific benefit, followed by the product information, which was sequentially presented on the screen. When the entire product description had been displayed, subjects were asked to indicate their belief in the benefit on a 9-point scale. In the *typicality* condition, subjects received the same product information, but before the product description appeared, the statement "try to imagine a typical [desired product]" appeared at the top of the screen. After all the product information had been displayed, subjects were asked to rate the similarity between the product description and the typical desired product on a 9-point scale anchored by "not similar at all" and "very similar". The *representativeness* condition was identical to the *typicality* condition, with the exception that subjects were also asked to indicate their belief in the product benefit after providing the similarity rating. Finally, in the *biased hypothesis testing* condition, subjects were asked, for each piece of information that appeared on the screen, 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,

subjects indicated their belief in the product benefit, but did not make any similarity judgments.

Results

The results are summarized in Figure 2. As intended, adding the irrelevant product information to the supportive information increased the perceived similarity between the product description and the typical desired product ($\bar{X}_{\text{SUPP}} = 5.16$, $\bar{X}_{\text{SUPP+IRR}} = 6.51$; $F(1,600) = 15.83$, $p < .001$). However, while adding the irrelevant information increased the perceived typicality of the product description in the *typicality* condition, it still weakened subjects' beliefs in the product benefit in the *free strategy* condition ($\bar{X}_{\text{SUPP}} = 6.59$, $\bar{X}_{\text{SUPP+IRR}} = 5.55$; $F(1,600) = 9.34$, $p < .01$). This is inconsistent with the prediction of the representativeness account (H7), but does confirm the prediction of the biased hypothesis testing perspective (H8). An additional test between these two theories is provided by comparing the forced strategy conditions to the *free strategy* condition.

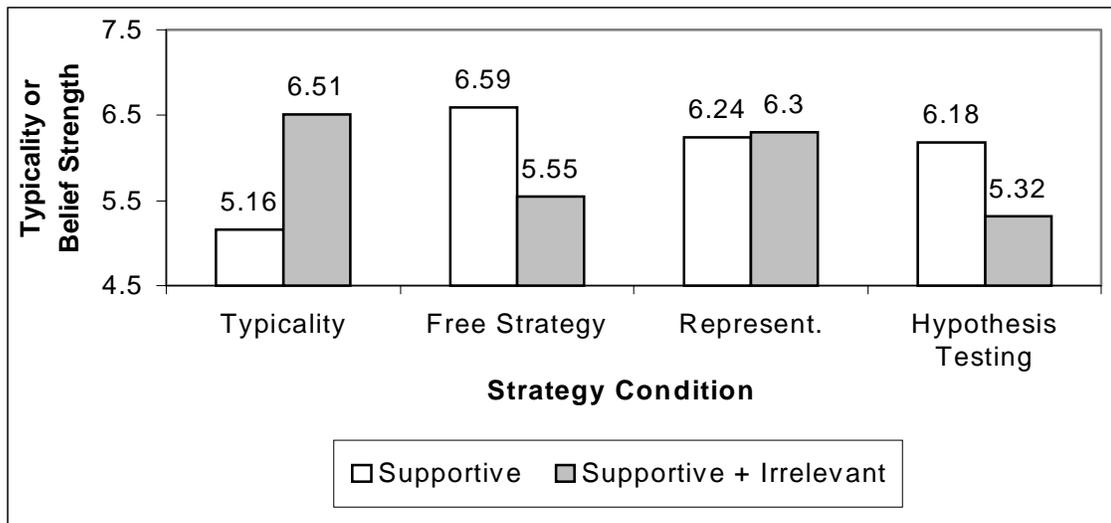


FIGURE 2
 EXPERIMENT 4: EFFECT OF STRATEGY AND TYPE OF INFORMATION
 ON SUBJECTS' BELIEF IN THE PRODUCT BENEFIT OR THE PERCEIVED
 TYPICALITY OF THE PRODUCT DESCRIPTION

If subjects in the *free strategy* condition are indeed using the proposed biased hypothesis testing strategy, then they should behave similarly to subjects who are explicitly encouraged to follow this strategy. On the other hand, if subjects in the *free strategy* condition are relying on the representativeness heuristic, they should behave similarly to those subjects who are making explicit similarity judgments before rating their belief in the benefit. The results show that the addition of irrelevant product information weakened subjects' beliefs ratings in the *biased hypothesis testing* condition ($\bar{X}_{\text{SUPP}} = 6.18$, $\bar{X}_{\text{SUPP+IRR}} = 5.32$; $F(1,600) = 7.56$, $p < .01$), and that this dilution effect was not significantly different from the effect observed in the *free strategy* condition ($F < 1$, *ns*). However, the addition of irrelevant information did not affect subjects' belief ratings in the *representativeness* condition ($\bar{X}_{\text{SUPP}} = 6.24$, $\bar{X}_{\text{SUPP+IRR}}$

= 6.30; $F < 1$, *ns*), resulting in an effect that was significantly different from the dilution effect observed in the *free strategy* condition ($F(1,600) = 4.50$, $p < .05$). Thus, subjects in the *free strategy* condition seemed to react to the irrelevant information in a similar way as those in the *biased hypothesis testing* condition, but differently from those in the *representativeness* condition, indicating that, while these subjects may have been using a biased hypothesis testing strategy, they were most likely not relying on the representativeness heuristic.

Discussion

The findings of experiment 4 demonstrate that adding irrelevant information still weakens consumers' belief in the product benefit when this information actually increases the perceived similarity between the product description and the typical desired product. Moreover, consumers who choose their own strategy for assessing whether a product will deliver a benefit behave similarly to those who are encouraged to follow a biased hypothesis testing strategy, but different from those who are encouraged to rely on similarity ratings to determine their beliefs in the product benefit. These findings therefore indicate that the diluting effect of irrelevant product information reported in these studies is not due to subjects' reliance on a representativeness heuristic. Instead, these findings are consistent with the biased hypothesis testing account of the dilution effect.

CHAPTER 7
EXPERIMENT 5: PROCESSING THE INFORMATION
WITHOUT THE BENEFIT IN MIND

The results observed in the previous experiments are inconsistent with distraction, averaging, conversational norms, and representativeness accounts of the dilution effect, while they are consistent with the biased hypothesis testing explanation. On the other hand, the previous experiments failed to directly test some of the essential characteristics of the biased hypothesis testing mechanism. The following experiments will therefore examine some direct implications of the proposed explanation. In experiments 5 and 5A, I will manipulate whether consumers initially process the information with the hypothesis in mind. Experiment 6 will examine whether the dilution effect disappears when consumers consider the implications of the product information for both the focal and the alternative hypothesis. Finally, experiment 7 will test whether the effect reverses when consumers start out with the hypothesis that the product will not deliver the desired benefit, instead of the hypothesis that the product will deliver the benefit.

In the previous studies, as well as in virtually all 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 will replicate this *top-down* scenario, but will also add a condition in which the desired benefit is only revealed after subjects have read the product description, 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 will (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 with regard to their search goal as "not confirming", and (3) weaken their belief in the hypothesis. However, subjects in the *bottom-up* condition will (1) process the product description without a specific search goal, (2) learn about the hypothesis, and (3) redirect their attention to those pieces of information that may confirm the hypothesis. Because subjects in the *bottom-up* condition have already processed the product information, they can engage in a more efficient search for supportive evidence. They can immediately focus on the supportive evidence and ignore information that is obviously irrelevant. While they have still processed the irrelevant information, they have not used it to evaluate any hypothesis, hence dilution should not occur.

H9: Biased Hypothesis Testing: Adding irrelevant product information to supportive benefit information will weaken consumers' beliefs in the product benefit when the information is first processed with the benefit in mind, but not when it is first processed without the benefit in mind.

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. It was expected that the polarization effect would be replicated in both the *top-down* and *bottom-up* conditions. Even when the less supportive information has first been processed without the benefit in mind, it should be revisited in the search for confirming evidence once the desired benefit has become known. It will then be classified as "confirming," and strengthen consumers' belief in the hypothesis that the product will deliver the desired benefit.

Method

Subjects and design¹. Subjects were 57 undergraduate students who participated in return for class credit. The design was a 4 (type of information) by 2 (processing mode) by 12 (product replicates) mixed design. Each subject was presented with 8 different product descriptions out of a total of 12 product replicates. For each of the product 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. The same 12 replicates used in experiment 3 were also used in experiment 5. However, an additional pretest ($n = 30$) was conducted to select less supportive information for the four replicates not used in experiment 2. The twelve additional pieces of less supportive information were perceived as suggesting the benefit

¹ The original design of the experiment also contained a *top-down reversed* condition. This condition was identical to the *top-down* condition, with the exception that the irrelevant information was presented first, followed by the supportive information. The results obtained in this condition did not differ from those in the *top-down* condition, demonstrating that the effects do not depend on the order in which the information is presented.

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 instructions at the beginning of the experiment informed subjects that they would have to evaluate each product “on a certain dimension”. As in the previous experiments, the separate pieces of product information appeared sequentially and all information remained on the screen when the belief measure appeared.

Results and Discussion

The results are summarized in Figure 3. First, it can be seen that, across both processing modes, product descriptions consisting only of supportive information lead to stronger beliefs than descriptions consisting only of less supportive information ($F(1,408) = 15.68, p < .001$). This effect was significant in the *top-down* condition ($F(1,408) = 15.87, p < .001$), and marginally significant in the *bottom-up* condition ($F(1,408) = 2.24, p = .13$). Thus, the manipulation of the degree of support was successful. Second, the addition of less supportive information increased the strength of subjects' product beliefs, regardless of the manner in which the information was presented ($F(1,408) = 13.00, p < .001$). There was a significant polarization effect of the less supportive information in both the *top-down* condition ($F(1,408) = 4.42, p < .05$) and the *bottom-up* condition ($F(1,408) = 8.83, p < .01$). The strength of this polarization effect did not depend on the manner in which the information was being processed ($F < 1, ns$). These results attest to

the robustness of the polarizing effect of less supportive information observed in experiment 2.

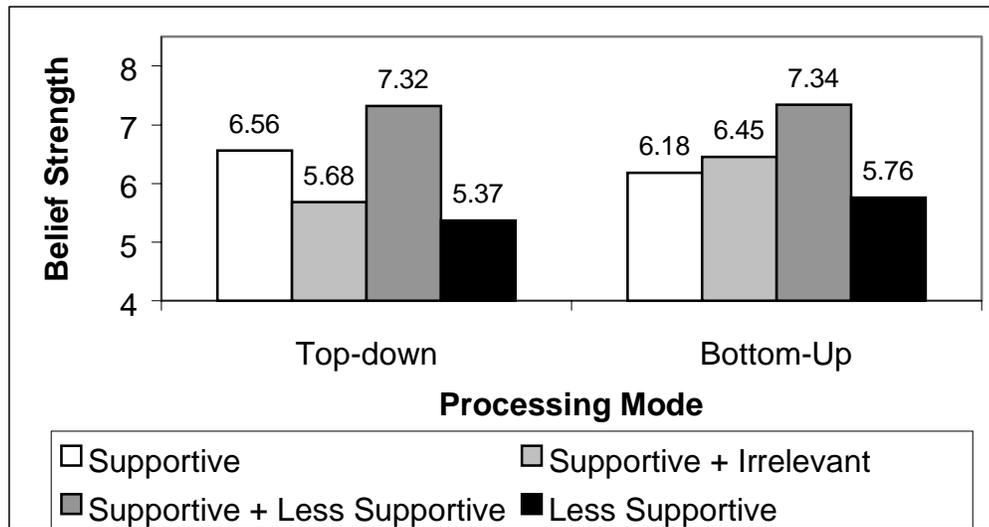


FIGURE 3
EXPERIMENT 5: EFFECT OF PROCESSING MODE AND
TYPE OF INFORMATION ON SUBJECTS' BELIEF IN THE PRODUCT BENEFIT

However, the main objective of this experiment is to examine how the dilution effect is affected by the manner in which the product information is being processed. A general dilution effect was observed when irrelevant information was added to the supportive information ($F(1,408) = 4.38, p < .05$). However, this effect did depend on the manner in which the information had been processed ($F(1,408) = 4.46, p < .05$). The irrelevant information significantly diluted subjects' beliefs in the *top-down* condition ($F(1,408) = 7.42, p < .01$), but not in the *bottom-up* condition ($F < 1, ns$). Thus, consistent with the prediction of the biased hypothesis testing account, irrelevant information only affects product beliefs when it is initially processed with the desired benefit in mind. When consumers have already processed the irrelevant information

without the benefit in mind, they can selectively focus on the supportive information when testing the hypothesis. Thus, these consumers never classify the irrelevant information as "not confirming," and consequentially do not weaken their belief in the hypothesis that the product will deliver the desired benefit.

Experiment 5A

The results from experiment 5 suggest that the dilution effect in the *bottom-up* condition failed to occur because subjects' pre-processing of the information allowed them to ignore the irrelevant information when searching for information that supported the hypothesis. To directly test this assumption, experiment 5A examines a situation in which the initial processing of the information would not inform subjects that the information could not possibly be supportive for the to-be-revealed benefit. This would force subjects to reconsider the information after the benefit was revealed, and thus classify this information as "not confirming". In this case, the irrelevant information should still weaken consumers' beliefs in the product benefit, even when the information has initially been processed without the benefit in mind.

To create such a situation, the "obviously irrelevant" information used in the previous experiments was replaced with information that had a high degree of "typical diagnosticity." In other words, the selected information was often relevant in similar decisions, but was clearly irrelevant for the benefit desired in this specific case. For example, the fact that a computer has been assembled in the United States, can be ordered online, and airs commercials on NBC is not only irrelevant for assessing the speed of the

computer, but it does not inform consumers about any other important benefits either. On the other hand, the fact that a computer has high quality speakers, is loaded with games, and has a flat screen monitor may not be relevant for assessing the speed of the computer, but it is often relevant in typical computer purchase decisions. This latter type of information can be called "pseudo-relevant" information or "typically diagnostic" information (Hilton and Fein 1989; Yzerbyt et al. 1997). When consumers first process this pseudo-relevant information, they cannot rule out that the information may support the desired benefit. Thus, once they learn about the desired benefit, they will have to go back to this information and evaluate whether or not it implies that the product will 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) by 2 (processing mode) by 8 (product replicates) mixed design. Each subject was presented with 8 different product descriptions. For each replicate, subjects were randomly assigned to either the *supportive* information condition or the *supportive + pseudo-relevant* information condition. The processing strategy was manipulated between subjects.

Stimuli and procedure. Experiment 5A used a subset of the replicates used in experiment 5: hotel (luxurious), movie (action-packed), car (sportive), apartments (safe), package delivery service (fast), frozen dinners (healthy), toothpaste (fights cavities), and computers (fast). A first pretest ($n = 29$) was conducted to select the "pseudo-relevant" information. The 24 pseudo-relevant facts were classified as "not helpful" for the benefit judgment by an average of 90% of the pretest subjects. A second pretest ($n = 24$) was

conducted to compare the pseudo-relevant information used in this study to the irrelevant information used in experiment 3. Subjects were asked to indicate to what extent each piece of information was typically helpful for evaluating the product (endpoints 1 = “Not Helpful at All”, to 7 = “Very Helpful”). The results showed that, across product categories, the pseudo-relevant information was perceived as more helpful for evaluating the products ($\bar{X} = 5.65$) than was the obviously irrelevant information ($\bar{X} = 4.19$; $F(1,1040) = 202.74, p < .01$). The procedure was identical to the procedure used in experiment 5, with the exception that the product descriptions either contained one piece of supportive information, or one piece of supportive information and three pieces of pseudo-relevant information.

Results and discussion. The results are summarized in Figure 4. There is a significant main effect of the information manipulation. Adding pseudo-relevant information diluted product beliefs ($F(1,390) = 20.20, p < .01$). This effect did not interact with the manner in which the information was being processed ($F < 1, ns$). The irrelevant information weakened product beliefs both when subjects processed the information with the benefit in mind ($F(1,390) = 5.27, p < .05$) or when they processed the information without the benefit in mind ($F(1,390) = 11.72, p < .01$).

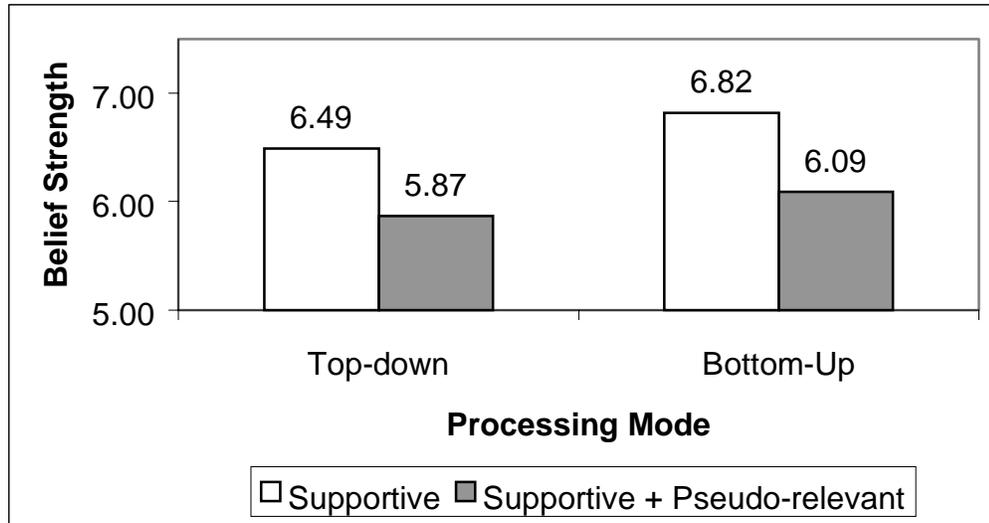


FIGURE 4
EXPERIMENT 5A: EFFECT OF PROCESSING MODE AND TYPE OF INFORMATION ON SUBJECTS' BELIEF IN THE PRODUCT BENEFIT

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 the additional information is obviously irrelevant, people know that the irrelevant information is not possibly supportive when they learn about the desired benefit, and they can restrict their search for supportive evidence to the actual supportive information. In contrast, when additional information is "pseudo-relevant", subjects must reconsider it when they learn about the desired benefit, since the information may be supportive. Even though this "pseudo-relevant" information is not diagnostic with respect to the benefit, the simple act of considering this information results in dilution. These findings indicate that the dilution effect is the result of goal-directed processing and provide further evidence for the biased hypothesis testing account of the dilution effect.

CHAPTER 8
EXPERIMENT 6: CONSIDERING BOTH THE FOCAL
AND ALTERNATIVE HYPOTHESES

Experiment 6 examines 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 confirm that the product will deliver the benefit, it also does not confirm 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. Experiment 6 tests this prediction by manipulating the number of questions subjects are asked about each product description. In the *single hypothesis* condition, they 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. For example, subjects who are looking for a fast computer do not only have to rate their belief that the described model is fast, but also have to indicate their belief that the described model is slow.

Method

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

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 understand that they would have to rate each product on both dimensions. The practice category consisted of a description of an MBA program and always contained the same four pieces of information. Subjects' responses for this replicate were not included in the analyses.

With the exception of the inclusion of this practice replicate, the procedure in the *single hypothesis* condition was identical to the procedure used in experiment 1. In the *dual hypotheses* condition, subjects were also first told that they were looking for a particular benefit (e.g., "You are looking for a fast computer."), followed by a description of the product, and the measure of subjects' belief in the benefit (e.g., "Is this computer fast?"). However, unlike in the *single hypothesis* condition, this measure was followed by a measure of subjects' belief in the opposite of the benefit (e.g., "Is this computer

slow?") on a similar 9-point scale (e.g., anchored by 1 = "definitely not slow" and 9 = "definitely slow").

Results

The results are summarized in Figure 5. The addition of irrelevant product information to supportive information again weakened subjects' belief in the product benefit ($F(1,864) = 6.42, p = .01$). However, this dilution effect did depend on the number of hypotheses subjects were evaluating ($F(1,864) = 4.07, p < .05$). When subjects only rated their belief in the benefit, irrelevant information indeed diluted subjects' beliefs ($\bar{X}_{\text{SUPP}} = 6.59, \bar{X}_{\text{SUPP+IRR}} = 6.09; F(1,864) = 10.60, p < .01$). However, when subjects also indicated their belief that the product would not produce the benefit, the irrelevant information did not influence their belief in the product benefit ($\bar{X}_{\text{SUPP}} = 6.44, \bar{X}_{\text{SUPP+IRR}} = 6.38; F < 1, ns$). Furthermore, the irrelevant information also did not influence their belief that the product would not deliver the benefit ($\bar{X}_{\text{SUPP}} = 3.13, \bar{X}_{\text{SUPP+IRR}} = 3.27; F < 1, ns$).

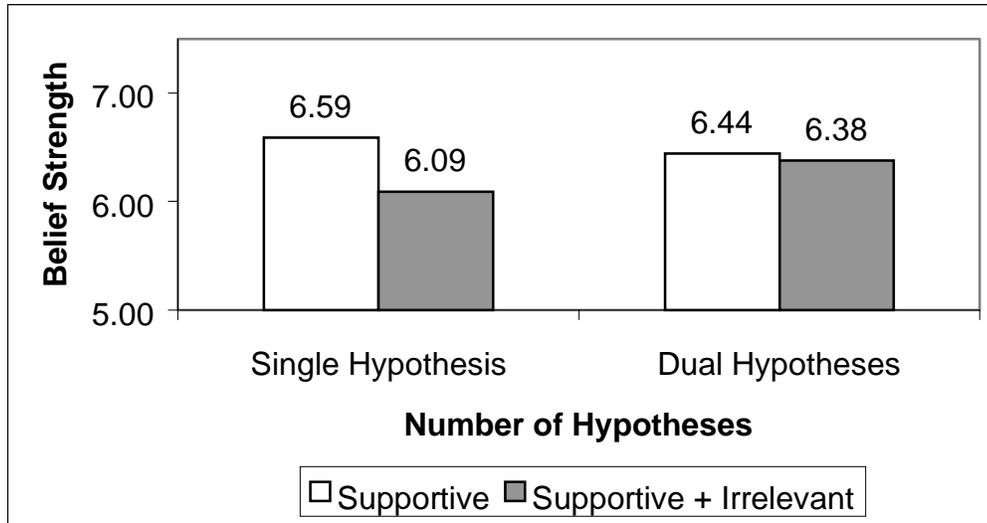


FIGURE 5
EXPERIMENT 6: EFFECT OF NUMBER OF HYPOTHESES AND TYPE OF INFORMATION ON SUBJECTS' BELIEF IN THE PRODUCT BENEFIT

Discussion

Experiment 6 demonstrates that irrelevant product information does not influence consumers' beliefs in the product benefit when consumers assess the likelihood that the product will deliver the benefit, as well as the likelihood that the product will *not* deliver the benefit. Thus, consistent with the biased hypothesis testing account, consumers' unique focus on the implications for the hypothesis that the product will deliver the benefit seems to be a necessary condition for the dilution effect to occur. The dilution effect disappears when consumers also consider whether the evidence supports the hypothesis that the product will not deliver the benefit.

CHAPTER 9
EXPERIMENT 7: EXPECTING THAT THE PRODUCT
WILL NOT DELIVER THE BENEFIT

The experiments 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 additional irrelevant information if they want to emphasize a product benefit. If consumers follow the proposed biased hypothesis testing process, then the confrontation with irrelevant information may sometimes strengthen consumers' beliefs in the product's ability to deliver the benefit.

The first assumption of the proposed explanation states that consumers will usually test the hypothesis that the product *will* deliver the benefit. However, this assumption may not hold 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 information that confirms this hypothesis (i.e., information that is counterdiagnostic of the benefit) 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 consumers' belief in the benefit.

Consider, for instance, a consumer who is looking for a trendy store and reads an ad for a K-Mart store. The consumer may assume that this store is unlikely to be trendy and search for information that confirms that the store is indeed not trendy. She may then encounter information that is irrelevant with respect to the benefit (e.g., the store accepts all major credit cards), which is not the type of information she was looking for, will be classified as "not confirming," and reduce her confidence in the initial hypothesis that the store was not trendy. In other words, the irrelevant information actually strengthened her belief that the store *is* trendy. Experiment 7 will test this prediction by manipulating the presence of a negatively perceived brand name.

Method

Subjects and design. Subjects were 68 undergraduate students who participated in return for class credit. The design was a 2 (type of information) by 2 (presence of brand names) by 7 (product replicates) mixed design with a *brand name only* baseline condition that served as a control. Each subject was presented with three different product descriptions from a total of seven product replicates. For each of the product 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)¹. 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 restaurant (healthy). A pretest ($n = 36$) was conducted to select one supportive fact and three irrelevant facts for each of the new categories. The fifteen 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 categories². 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]).

¹ The majority of the twelve replicates used in the previous experiments could not be used here for two reasons. First, many of the categories were not associated with brands for which subjects held clearly negative priors. Second, product information needed to be compatible with whatever brand name subjects perceived as least likely to provide the benefit.

² This procedure was necessary because pretests had shown that subjects not only differed in their brand beliefs within a category, but also in their brand beliefs between categories. While some subjects perceived some beer brands very negatively and gave neutral evaluations to all shampoo brands, other subjects showed an exactly reversed pattern.

After a filler task, subjects were exposed to information for five filler categories and three target categories for which they had indicated strong negative beliefs for at least one of the brands³. Subjects first received the name of the brand for which they had strong negative priors, followed by the description of the product or service. The target category descriptions either contained only supportive information or both supportive and irrelevant information⁴. 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 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 6. First, a manipulation check showed that adding supportive information made the negatively perceived brand more desirable ($\bar{X} = 3.75$) than presenting the brand name by itself ($\bar{X} = 1.61$; $F(1,129) = 35.29$, $p < .01$).

The remainder of the analyses will concentrate on the 2 X 2 design manipulating the type

³ To 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.

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 only received the product description ($\bar{X} = 5.11$), than when they also received the negatively perceived brand name ($\bar{X} = 4.11$, $F(1,128) = 9.58$, $p < .01$). Second, there was no main effect of adding irrelevant information ($F(1,128) = 0.14$, ns). Third, the effect of the irrelevant information depended on the presence or absence 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 ($\bar{D} = -0.93$; $F(1,128) = 5.65$, $p < .05$). However, when subjects were also given the brand name, the irrelevant product information *strengthened* product beliefs from 3.75 to 4.44, consistent with H9 ($\bar{D} = 0.69$; $F(1,128) = 4.22$, $p < .05$).

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.

⁴ The descriptions in the filler product categories confirmed subjects' positive or negative priors, thereby reducing suspicion about the accuracy of the product information.

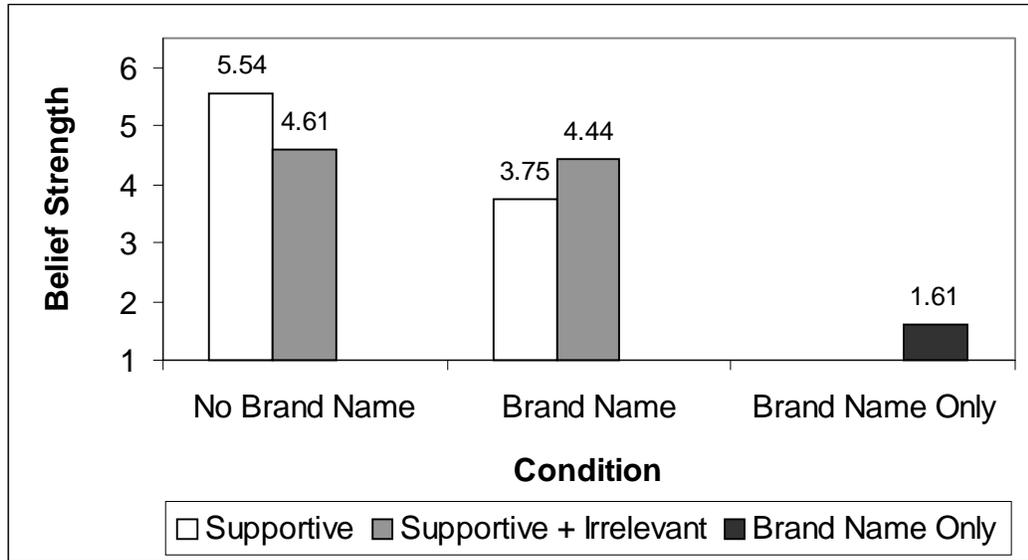


FIGURE 6
EXPERIMENT 7: EFFECT OF A NEGATIVELY PERCEIVED BRAND NAME AND
TYPE OF INFORMATION ON SUBJECTS' BELIEF IN THE PRODUCT BENEFIT

CHAPTER 10 GENERAL DISCUSSION

Marketing managers are undoubtedly aware that consumers are often looking for specific product benefits and will try to emphasize product information that suggests that their brand will deliver those benefits. A website may prominently feature a computer's powerful processor, suggesting that it is fast. An airline commercial may mention that the airline was voted #1 in a survey on airline service, suggesting that their service is superior to that of the competition. In spite of this emphasis, consumers will still encounter a myriad of product information that is not relevant for the benefit they are seeking. The website may also list that the computer features high quality speakers and the airline commercial may also mention that the airline is a proud sponsor of the Olympic Games. There are undoubtedly good reasons to mention this additional information. An additional feature may appeal to a niche in the market that highly values a specific benefit, while other product information may create generalized positive affect toward the brand. Managers often assume that these actions may only increase brand equity and cannot possibly hurt it. Indeed, if consumers are behaving normatively, they should simply ignore this information when assessing whether the product will deliver the benefit.

However, as research on social judgment (e.g., Nisbett et al. 1981) suggests, and as this research demonstrates, the irrelevant information can actually have a negative

impact on consumers' product perceptions. In ten different studies, and across seventeen 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 effect even 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).

While the findings indicate that this dilution effect is a quite robust phenomenon, they also show that irrelevant information does not always dilute consumers' beliefs in the benefit. Irrelevant information does not affect consumers' belief in the benefit when the information is processed without the benefit in mind (experiment 5) or when consumers also consider whether the information suggests the opposite of the benefit (experiment 6). In fact, adding irrelevant information to supportive benefit information may even *strengthen* consumers' beliefs in the product's ability to deliver the benefit, when the brand under consideration has a very poor reputation (experiment 7).

By examining the influence of these factors which have not been examined previously, these studies also provide new insight into the mechanism that underlies the dilution effect. None of the explanations put forth in earlier research can adequately explain the entire pattern of results observed across the ten studies. The polarizing effect

of less supportive information (experiments 2 and 3) and the fact that the irrelevant information does not influence consumers' sensitivity to the supportive information (experiment 2) provide evidence against an *averaging* account of the dilution effect. The *conversational norms* account, on the other hand, cannot explain why the 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 of resources* 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).

Instead, a new explanation is proposed that can account for the entire pattern of findings. It is argued that consumers are following a biased hypothesis testing procedure when assessing the likelihood that a product will deliver a benefit. This proposed mechanism is based on four assumptions. 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, i.e., supports the benefit. 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. Classifying information as “confirming” strengthens their belief, while classifying information as “not confirming” weakens their belief. When consumers encounter irrelevant information, they classify it as “not confirming” and weaken their belief in the product benefit. They do not consider the fact that the information does not confirm the alternative hypothesis either.

This mechanism is consistent with the findings of the first six experiments. It 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 is not classified as “confirming” or “not confirming” and the dilution effect disappears (experiments 5 and 5A). Similarly, the dilution effect also does not occur when consumers consider the implications of the information for both the focal hypothesis (that the product will deliver the benefit) and the alternative hypothesis (that the product will not deliver the benefit) (experiment 6). Finally, when consumers have very negative priors regarding the brand under consideration, they may set out to test the hypothesis that the product will *not* deliver the benefit, leading to a reversal of the dilution effect (experiment 7).

CHAPTER 11 LIMITATIONS AND FUTURE RESEARCH

Although these studies indicate that irrelevant information can weaken consumers' belief in a product's ability to deliver a benefit, they do not inform us about its effect on the overall product evaluation. Obviously, if a consumer is only interested in a single product benefit, her decreased confidence in the product's ability to deliver that benefit will reduce her overall product evaluation. However, if a consumer is interested in multiple benefits, the effect of additional information about other product benefits on the product's overall evaluation will depend on the trade-off between reduced confidence in the original benefit and increased confidence in the other benefit. Similarly, the effect of irrelevant sponsorship associations on overall product evaluations depends on the trade-off between the dilution effect and the creation of positive affect.

The reported studies are also limited in that subjects are always instructed that a particular benefit is desirable for them. This does not necessarily correspond to situations in which desirable benefits are spontaneously generated by consumers. Whether a benefit becomes salient can depend on its habitual salience for a certain consumer or because it is primed by a certain usage situation (Ratneshwar, Warlop, Mick, and Seeger 1997). In these situations, the activated benefit may lead to similar outcomes as the explicit instructions used in the experiments reported here. However, consumers can also

derive the desired benefit from the product information itself. While some product information may seem irrelevant to consumers at first, they may rely on the relevance principle (Grice 1975, Sperber and Wilson 1986) and infer that the information has to convey some value. In this manner, the irrelevant information may prime a benefit consumers had not considered earlier, thus increasing the appeal of the product, rather than diluting it.

Furthermore, though the reported studies are most consistent with a biased hypothesis testing explanation, these results do not imply that other processes cannot contribute to the dilution effect, or even be the unique cause of the effect 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., Andreassen 1988; Epstein, Donovan, and Denes-Raj 1999; Kahneman and Tversky 1972, 1973). It is plausible that subjects in the social judgment dilution studies (e.g., Nisbett et al. 1981) indeed relied on the similarity between the described person and the stereotypical murderer or child abuser. While it may be hard to assess the similarity between a product description and an abstract benefit, it is a lot easier to assess the similarity between a person description and an easily accessible stereotype. In fact, consumers may also rely on a representativeness heuristic when predicting product benefits. However, either the product information, or the desired benefit, may have to be connected to a well defined subcategory or a prototypical brand.

Finally, I would like to emphasize that I only examined the influence of one type of irrelevant information. The concept of irrelevance used in this dissertation has three

essential characteristics. First, information can only be irrelevant with respect to a *context*. The information used in these studies 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 is *subjective* rather than objective. The information was labeled as obviously irrelevant because the great majority of subjects classified the information as “not helpful” rather than “suggesting [the benefit]” or “suggesting [not the benefit]”. Finally, the irrelevance of the information is *not absolute*. I assume that there are degrees of (ir)relevance. The information is usually not classified as irrelevant by all subjects. Moreover, even for those who classify it as irrelevant, the information may still have a minimal diagnostic value. Therefore, it was essential 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 to other interpretations of relevance? It is clearly very different from a consequentialist perspective, which would state that information is irrelevant for a decision when it does not influence the decision. From this perspective, the additional product information only is irrelevant in those conditions in which the dilution effect disappears. 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. Interestingly, 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, they 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 actually intended for another consumer segment. The principle governing advertising communications may be that the communication has to convey value to *some* consumers, but not necessarily to the individual processing the message.

REFERENCES

- Anderson, Norman H. (1967), "Averaging Model Analysis of Set-Size Effect in Impression Formation," *Journal of Experimental Psychology*, 75 (2), 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: I. Learning, Memory and Thinking*, eds. David H. Krantz, Richard C. Atkinson, Duncan R. Luce, and Patrick Suppes, San Francisco: W.H. Freeman, 236-270.
- Andreassen, Paul B. (1988), "Explaining the Price-Volume Relationship: The Difference Between Price Changes and Changing Prices," *Organizational Behavior & Human Decision Processes*, 41 (June), 371-389.
- Asch, Solomon E. (1946), "Forming Impressions of Personality," *Journal of Abnormal & Social Psychology*, 41 (July), 258-290.
- Barsalou, Lawrence W. (1983), "Ad Hoc Categories," *Memory & 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," *Advances in Consumer Research*, 24, 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), p 372-385.

- Cantor, Nancy and Walter Mischel (1979), "Categorization Processes in the Perception of People," in *Advances in Experimental Social Psychology*, Vol. 12, ed. L. Berkowitz, New York: Academic Press, 3-52.
- Carpenter, Gregory S., Rashi Glazer, 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.
- Epstein, Seymour, Sean Donovan, and Veronika Denes-Raj (1999), "The Missing Link in the Paradox of the Linda Conjunction Problem: Beyond Knowing and Thinking of the Conjunction Rule, the Intrinsic Appeal of Heuristic Processing," *Personality & Social Psychology Bulletin*, 25 (February), 204-214.
- 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.
- _____, Romin W. Tafarodi, and Patrick S. Malone (1993), "You Can't Not Believe Everything You Read," *Journal of Personality & Social Psychology*, 65 (August), 221-233.
- Grice, H. Paul (1975), "Logic and Conversation," in *Syntax and Semantics: 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 & 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.
- _____ and Amos Tversky (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.
- Koriat, Asher, Sarah Lichtenstein, and Baruch Fischhoff (1980), "Reasons for Confidence," *Journal of Experimental Psychology: Human Learning & Memory*, 6 (March), 107-118.
- Kuhn, Deanna and Joseph Lao (1996), "Effects of Evidence on Attitudes: Is Polarization the Norm?" *Psychological Science*, 7 (March), 115-120.
- 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.
- Lord, Charles G., Lee Ross, and Mark R. Lepper (1979), "Biased Assimilation and Attitude Polarization: The Effects of Prior Theories on Subsequently Considered Evidence," *Journal of Personality & Social Psychology*, 37 (November), 2098-2109.
- Nickerson, Raymond S. (1998), "Confirmation Bias: A Ubiquitous Phenomenon in Many Guises," *Review of General Psychology*, 2 (June), 175-220.
- 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.

- Peterson, Cameron R. and Wesley M. DuCharme (1967), "A Primacy Effect in Subjective Probability Revision," *Journal of Experimental Psychology*, 73 (January), 61-65.
- Pitz, Gordon F., Leslie Downing, and Helen Reinhold (1967), "Sequential Effects in the Revision of Subjective Probabilities," *Canadian Journal of Psychology*, 21 (October), 381-393.
- 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.
- _____, 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 & Review*, 5 (June), 197-220.
- _____, _____, 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 & Cognition*, 10 (November), 520-530.
- Shanteau, James (1975), "Averaging versus Multiplying Combination 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.
- _____, Stephen M. Nowlis, and Yael Simonson (1993), "The Effect of Irrelevant Preference Arguments on Consumer Choice," *Journal of Consumer Psychology*, 2 (3), p 287-306.
- 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 Swann, William B. (1978), "Hypothesis-Testing Processes in Social Interaction," *Journal of Personality and Social Psychology*, 36 (November), 1202-1212.
- ____ and Nancy Cantor (1979), "Testing Hypotheses About Other People: The Use of Historical Knowledge," *Journal of Experimental Social Psychology*, 15 (July), 330-342.
- 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.
- ____, 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.
- Tesser, Abraham (1978), "Self-Generated Attitude Change," in *Advances in Experimental Social Psychology*, 11, 289-338.
- 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, NY: The 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.
- Wason, Peter C. (1960), "On the Failure to Eliminate Hypotheses in a Conceptual Task," *Quarterly Journal of Experimental Psychology*, 12 (August), 129-140.
- Yzerbyt, Vincent Y., Jacques-Philippe Leyens, and Georges Schadron (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.

_____ and Dennis L. Jennings (1983), "Nondiagnosticity and Typicality Effects in Prediction," *Social Cognition*, 2 (July), 187-198.

APPENDIX A INSTRUCTIONS TO SUBJECTS

Standard Instructions

We will now show you 8 different products. For each of these products, you are looking for a particular benefit (e.g., you want a fast computer).

For each of these products, we will provide you with some information. This information may or may not be helpful for the decision you have to make. You should look at this information as objective information, i.e., information as provided by Consumer Reports. After each product description, we will ask you to rate the product on an important dimension.

Instructions Experiment 1A

We will now show you 8 different products categories. In each category, you are looking for a particular benefit (e.g., you want a fast computer).

For each category, we will provide you with information about two different brands. The information may or may not be helpful for the decision you have to make. You should look at this information as objective information, i.e., information as provided by Consumer Reports. After you have received information for both brands, we will ask you to indicate which brand is most likely to deliver the benefit.

Instructions Experiment 3

We will now show you 8 different products. For each of these products, you are looking for a particular benefit (e.g., you want a fast computer).

For each of these products, we will provide you with some information. The information may or may not be helpful for the decision you have to make. You should look at this information as objective information, i.e., information as provided by Consumer Reports. After each piece of information, you will be asked to indicate whether or not it suggests that the brand will deliver the benefit, or whether it is not helpful for your decision.

After you have received the full product description, you will be asked to rate the product based on all the information you've received.

Instructions Experiment 3A

We will now show you 8 different products. For each of these products, you are looking for a particular benefit (e.g., you want a fast computer).

For each of these products, the computer will randomly select 1 to 4 pieces of information from a database on these products. The information may or may not be helpful to you. You should look at this information as objective information, i.e., information as provided by Consumer Reports. After you have received the full product description, you will be asked to rate the product based on all the information you've received.

Instructions Experiment 4 (Typicality Condition)

We will now show you 8 different products or services. For each of these products or services, you will receive some information. You should look at this information as objective information, i.e., information as provided by Consumer Reports.

After each product description, we will ask you to make a similarity rating. We will ask you to rate the **similarity** between the description you've received and some **typical** product.

For instance, we will ask you to imagine a typical fast computer, provide you with some product description and then ask you to rate the similarity between the description and that typical fast computer.

Instructions Experiment 4 (Biased Hypothesis Testing Condition)

We will now show you 8 different products. For each of these products, you are looking for a particular benefit (e.g., you want a fast computer).

For each of these products, we will provide you with some information. You should look at this information as objective information, i.e., information as provided by Consumer Reports. This information may or may not tell you anything about the benefit you are looking for. After each piece of information is presented, you will be asked to indicate whether or not this information indicates that the product will deliver the benefit. When the entire product description has been displayed, we will ask you to rate the product on this benefit dimension.

Instructions Experiment 4 (Representativeness Condition)

We will now show you 8 different products or services. For each of these products or services, you will receive some information. You should look at this information as objective information, i.e., information as provided by Consumer Reports.

After each product description, we will ask you to make a similarity rating. We will ask you to rate the **similarity** between the description you've received and some **typical** product.

For instance, we will ask you to imagine a typical fast computer, provide you with some product description and then ask you to rate the similarity between the description and that typical fast computer. After you have given your similarity rating, we will also ask you to make a prediction about this product or service.

Instructions Experiment 5 and Experiment 5A (Bottom-Up Conditions)

We will now show you 8 different products. For each of these products, we will provide you with some information. You should look at this information as objective information, i.e., information as provided by Consumer Reports. After each product description, we will ask you to rate the product on a certain dimension.

Instructions Experiment 7 (Brand Name Conditions)

We will now show you information about 8 different products. For each of these products, you are looking for a particular benefit (e.g., you want a fast computer).

We will first tell you the name of the option you are considering. Then, we will provide you with some information about that option. This information may or may not be helpful for your decision. You should look at this information as objective, unbiased information, i.e., information as provided by Consumer Reports. After each product description, we will ask you to rate the product on an important dimension.

APPENDIX B
PRODUCT INFORMATION

Information Used in Experiments 1, 1A, 2, 3, 3A, 5, and 6

Product Category	Apartment
Desired Benefit	Safe
Supportive Information	24 Hours On-Site Security
Irrelevant Information	Complex Name: Haywood Park 40-Year Old Manager Both 1 & 2 Bedroom Apartments
Product Category	Car
Desired Benefit	Sportive
Supportive Information	Very Powerful Engine
Irrelevant Information	Dark Blue Color Dual Airbags Average Resale Value
Product Category	Package Delivery Service
Desired Benefit	Fast
Supportive Information	14 Hour Delivery or Money-Back Guarantee
Irrelevant Information	Sponsors Arts Events Corporate Headquarters in Chicago Company Founded in 1972

Product Category	Frozen Entree
Desired Benefit	Healthy
Supportive Information	Very Low in Fat
Irrelevant Information	Available in Most Grocery Stores Both Single and Double Serving Sizes Specializes in Italian Food
Product Category	Toothpaste
Desired Benefit	Good at Fighting Cavities
Supportive Information	Recommended by the American Dental Association
Irrelevant Information	Comes in 6 OZ Tubes Brand Owned by P & G Comes in Kid and Adult Versions
Product Category	Stereo System
Desired Benefit	Reliable
Supportive Information	Lifetime Warranty
Irrelevant Information	Original Design Includes Double Tape Deck Comes in Black and Grey
Product Category	Airline
Desired Benefit	Provide Superior Service
Supportive Information	#1 in JD Power & Associates Survey on Airline Service
Irrelevant Information	Sponsors NYC Marathon Company Founded in 1978 Corporate Headquarters in Boston

Product Category	Computer
Desired Benefit	Fast
Supportive Information	Very Powerful Processor
Irrelevant Information	Assembled in USA Airs Commercials on NBC & CBS Can Be Ordered On-Line

Additional Information Used in Experiment 1A

Product Category	Apartment
Supportive Information	In Low Crime Neighborhood
Product Category	Car
Supportive Information	Very Fast Acceleration
Product Category	Package Delivery Service
Supportive Information	Very Efficient Pick-Up and Delivery System
Product Category	Frozen Entree
Supportive Information	Very High in Vitamins and Minerals
Product Category	Toothpaste
Supportive Information	Keeps Fighting Bacteria After Brushing
Product Category	Stereo System
Supportive Information	High Consumer Reports Reliability Rating
Product Category	Airline
Supportive Information	Well Trained and Highly Motivated Personnel

Product Category	Computer
Supportive Information	High Score on PC Magazine's Performance Tests

Additional Information Used in Experiment 2

Product Category	Apartment
Less Supportive Information	Optional Panic Button Mostly Female Students Close to Campus
Product Category	Car
Less Supportive Information	Leather Interior Italian Design V8-Cylinder Engine
Product Category	Package Delivery Service
Less Supportive Information	Extended Network of Warehouses Local Personnel (Familiar with the Area) Very Reliable Package Tracking System
Product Category	Frozen Entree
Less Supportive Information	No Artificial Flavoring No Preservatives Low in Sodium
Product Category	Toothpaste
Less Supportive Information	Contains Fluoride Extensively Tested in Lab Remains Active After Brushing
Product Category	Stereo System
Less Supportive Information	Familiar Brand Name Uses Durable Components Very Popular Model

Product Category	Airline
Less Supportive Information	Multi-Lingual Personnel Efficient Customer Complaint Center Very Reliable Departure Times
Product Category	Computer
Less Supportive Information	64 Mbyte Working Memory Well-Known Brand Name 32-Speed CD-ROM

Additional Information Used in Experiments 3, 3A, and 5

Product Category	Mountain Bike
Desired Benefit	Sturdy
Supportive Information	Kevlar Reinforced Aluminum Fork
Irrelevant Information	Available in Four Colors Bikes Can Be Ordered On-Line Prices Include Pump and Bike Lock
Less Supportive Information ¹	One Piece Super-Stiff Rocker Heavy Frame Stainless Steel Pivots
Product Category	Hotel Room
Desired Benefit	Luxurious
Supportive Information	Executive Suite
Irrelevant Information	Hotel is 8 Years Old Room Located on Fourth Floor Can Be Booked On-Line
Less Supportive Information	Satin Bed Sheets Spacious Balcony Complementary Bottle of Champagne

¹ The less supportive information was only used in experiment 5.

Product Category	Movie
Desired Benefit	Action-Packed
Supportive Information	James Bond Film
Irrelevant Information	Available on Video Casting By Michael Perry Filmed in Mexico
Less Supportive Information	Very Large Budget Stars Arnold Schwarzenegger Lots of Pyrotechnics (Explosions, etc...)
Product Category	Printer
Desired Benefit	Provide High Graphical Quality
Supportive Information	Premier Photo-Quality Printing
Irrelevant Information	Robust 150-Sheet Paper Tray Print-Cancel Button Compact Size
Less Supportive Information	Smudge-Resistant Results Enhanced Color-Layering Technology High Resolution Range

Information Used in Experiment 4

Product Category	Apartment
Desired Benefit	Safe
Supportive Information	Surveillance Cameras Above All Doors
Irrelevant Information	Great Fitness Center Pest Control Included in Rent Fully Equipped Club House

Product Category	Car
Desired Benefit	Fast
Supportive Information	Engine Based on Jet Plane Engines
Irrelevant Information	Available in Red Flashy Wheel Covers Leather Seats
Product Category	Package Delivery Service
Desired Benefit	Fast
Supportive Information	14 Hour Delivery or Money-Back Guarantee
Irrelevant Information	Uniformed Employees Sponsors Sports Events Mostly Male Employees
Product Category	Frozen Entree
Desired Benefit	Healthy
Supportive Information	Great Source of Iron, Zinc, and Vitamin B12
Irrelevant Information	Can Be Found in Frozen Food Section Has Aluminum Cover Heat in Microwave or Conventional Oven
Product Category	Toothpaste
Desired Benefit	Good at Fighting Cavities
Supportive Information	Germ-Fighting Formula that Remains Active After Brushing
Irrelevant Information	Mint Flavored White Color Comes in Standard Squeeze Tube

Product Category	Stereo System
Desired Benefit	Reliable
Supportive Information	Made With Ultra-Durable Titanium
Irrelevant Information	Black Color Includes CD Player Digital Display
Product Category	Computer
Desired Benefit	Fast
Supportive Information	Uses Revolutionary Triple Processors
Irrelevant Information	Keyboard and Mouse Included Pre-Installed Software Includes DVD Player
Product Category	Beer
Desired Benefit	Great Taste
Supportive Information	Was Voted the Best Lager at the Boulder Beer Festival
Irrelevant Information	Sold in 12 Ounce Bottles Served in Restaurants Available Individually or in Six-Packs

Information Used in Experiment 5A

Product Category	Apartment
Desired Benefit	Safe
Supportive Information	24 Hours On-Site Security
Irrelevant Information	Huge Swimming Pool Laundry Facility Tennis Courts

Product Category	Car
Desired Benefit	Sportive
Supportive Information	Very Powerful Engine
Irrelevant Information	3 Year Limited Warranty 1997 Model Manufactured in the USA
Product Category	Package Delivery Service
Desired Benefit	Fast
Supportive Information	14 Hour Delivery or Money-Back Guarantee
Irrelevant Information	Uses Durable Packaging Material Shipments Fully Insured Guaranteed Confidentiality
Product Category	Frozen Entree
Desired Benefit	Healthy
Supportive Information	Very Low in Fat
Irrelevant Information	Easy to Prepare Large Variety of Dinners Italian Food
Product Category	Toothpaste
Desired Benefit	Good at Fighting Cavities
Supportive Information	Recommended by the American Dental Association
Irrelevant Information	Reasonable Breath Freshener Whitens Teeth Mint Flavored

Product Category	Hotel Room
Desired Benefit	Luxurious
Supportive Information	Executive Suite
Irrelevant Information	Close to Downtown Eight Story Building Parking Spaces Available
Product Category	Movie
Desired Benefit	Action-Packed
Supportive Information	James Bond Film
Irrelevant Information	Well-Known Director Shown at Regal Cinemas, Butler Plaza Filmed in Mexico
Product Category	Computer
Desired Benefit	Fast
Supportive Information	Very Powerful Processor
Irrelevant Information	Flat Screen Monitor Quality Speakers Loaded with Games

Information Used in Experiment 7

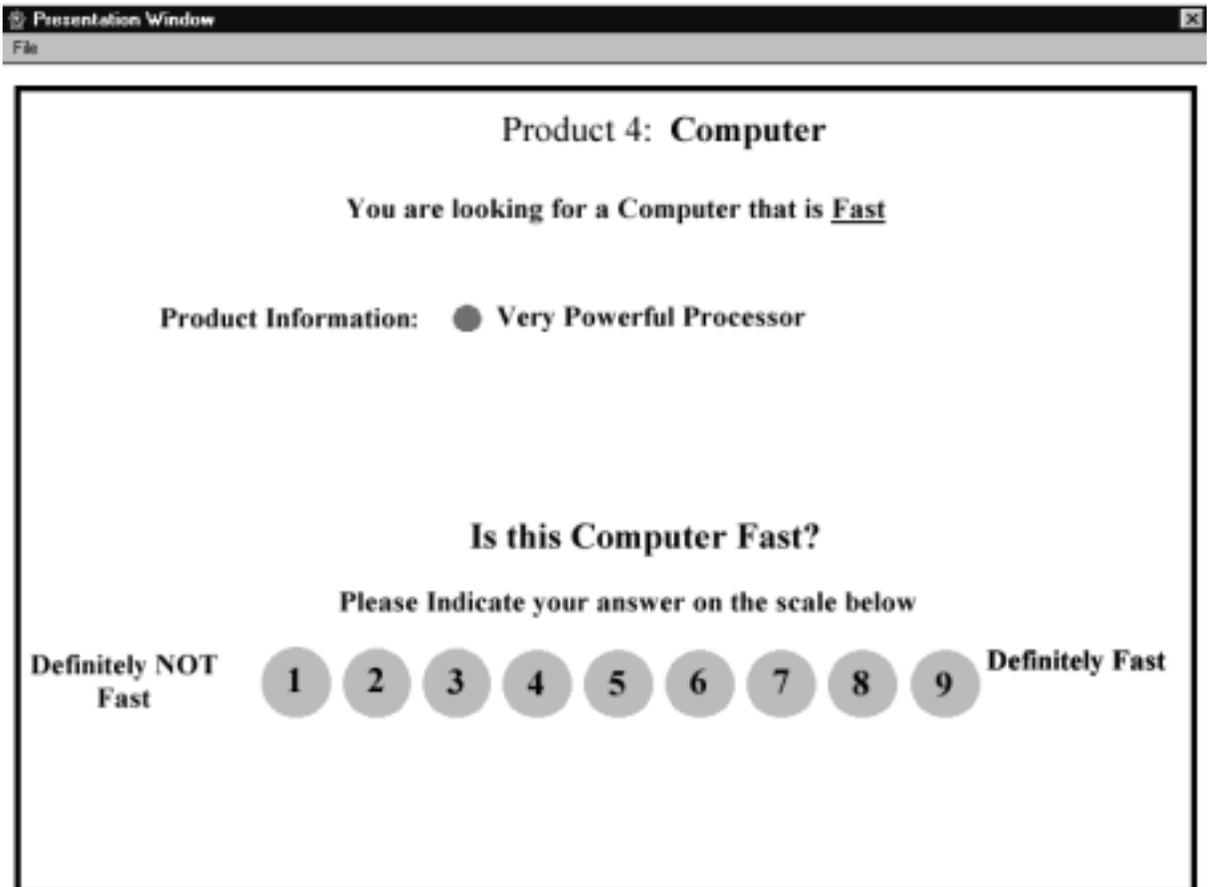
Product Category	Apartment
Desired Benefit	Safe
Supportive Information	24 Hours On-Site Security
Irrelevant Information	Laundry Facility Sewage Fee Included in Rent Walk-In Closets

Product Category	Car
Desired Benefit	Sportive
Supportive Information	Very Powerful Engine
Irrelevant Information	Dark Blue Color Dual Airbags Average Resale Value
Product Category	Hotel Room
Desired Benefit	Luxurious
Supportive Information	Executive Suite
Irrelevant Information	Hotel is 8 Years Old Room Located on Fourth Floor Can Be Booked On-Line
Product Category	Restaurant
Desired Benefit	Healthy
Supportive Information	4 Meals Available with Very Low Sodium Content
Irrelevant Information	Accepts Credit Cards Free Soda Refills Kids Menu Available
Product Category	Beer
Desired Benefit	Great Taste
Supportive Information	Honorable Mention in Great American Beer Contest
Irrelevant Information	Available in Six-Packs 12 FL OZ. Bottles Bottled on Three Locations

Product Category	Shampoo
Desired Benefit	High Quality
Supportive Information	Used by a Variety of Hair Salons
Irrelevant Information	Comes in Different Sizes Package is Recyclable Made in USA
Product Category	Clothing Store
Desired Benefit	Trendy
Supportive Information	Has Announced the Opening of a Tommy Hilfiger Section
Irrelevant Information	Closes at 9 PM Major Credit Cards Accepted Airs Commercials on NBC and CBS

APPENDIX C
EXAMPLES OF COMPUTER DISPLAYS

Example of Supportive Information Display



Example of Supportive + Irrelevant Information Display

Presentation Window

File

Product 1: Computer

You are looking for a Computer that is Fast

Product Information:

- Very Powerful Processor
- Assembled in USA
- Airs Commercials on NBC & CBS
- Can Be Ordered Online

Is this Computer Fast?

Please Indicate your answer on the scale below

Definitely NOT Fast 1 2 3 4 5 6 7 8 9 Definitely Fast

BIOGRAPHICAL SKETCH

Tom Meyvis was born in Brasschaat, Belgium, on December 5, 1972. After attending high school at the Sint-Jan Berchmans College in Merksem, Belgium, he attended the Katholieke Universiteit Leuven in Louvain, Belgium. At this university, he obtained the degrees of Kandidaat in de Psychologie (1993), Kandidaat in de Sociologie (1994), and Licentiaat in de Psychologie (1996). He wrote a thesis on intergroup discrimination under the guidance of Professor Norbert Vanbeselaere. He also spent seven months on a research internship, split between the Laboratory of Experimental Psychology at the Catholic University of Louvain, where he worked with Peter Degraef, and the Dipartimento di Psicologia dello Sviluppo e della Socializzazione at the University of Padova, Italy, where he worked with Professor Clara Casco. Encouraged by Professors Piet Vanden Abeele and Luk Warlop, Tom decided to join the Ph.D. program in marketing at the University of Florida. During his five years at the University of Florida, he has been involved in research projects with Professors Chris Janiszewski, Alan Cooke, and Jinhong Xie at the University of Florida, Michel Pham at Columbia University, and Stijn van Osselaer at the University of Chicago. In 1999, he taught an undergraduate course in international marketing at the University of Florida. Tom has accepted a position as Assistant Professor of Marketing at New York University.