

Consumer Learning and Brand Equity

STIJN M. J. VAN OSSELAER
JOSEPH W. ALBA*

A series of experiments illustrates a learning process that enhances brand equity at the expense of quality-determining attributes. When the relationship between brand name and product quality is learned prior to the relationship between product attributes and quality, inhibition of the latter may occur. The phenomenon is shown to be robust, but its influence appears sensitive to contextual variations in the learning environment. Tests of process are inconsistent with attentional explanations and popular models of causal reasoning, but they are supportive of associative learning models that portray learners as inherently forward looking.

Purchase decisions are based on predictions of product performance. Consumers base their predictions in part on product cues and are accurate to the extent that they have properly learned the relationship between the cues and performance. Consumer research has devoted little attention to this learning process despite its fundamental importance (Hutchinson and Alba 1991; Meyer 1987). In the present research we examine consumer learning of product cues as predictors of product quality with particular emphasis on the distinction between brand and attribute cues.

To illustrate, consider the cases in which consumers rely strictly on either brand or attribute cues to predict quality. If consumers learn the relationship between product attributes and quality, they will differentiate among brands that possess different attributes and treat as commodities those brands that share the same attributes. Once the predictive rule is learned, it may be applied to any new brand that possesses the attributes. In contrast, consumers who rely strictly on brand cues will ignore the underlying attributes and may incorrectly differentiate physically identical brands. The latter case is important because it can be costly and is not uncommon (such as when consumers pay high premiums for branded drugs that are chemically identical to their generic counterparts). An appealing explanation of this phenomenon is that consumers are unaware of the attributes of these brands. Indeed, firms attempt to foster such ignorance by making attribute information difficult to find or process (Bergen, Dutta, and Shugan 1996; Hoch and Deighton 1989). The present research investigates whether consumers will routinely learn the determinants of product quality when attribute cues are

freely available and processing is unconstrained. We suggest that learning can be suppressed even under these relatively favorable conditions due to the learning phenomenon known as blocking.

BLOCKING

As originally documented in the study of classical conditioning, the learning of one predictive cue can “block” the learning of subsequently encountered predictive cues. Specifically, once the relationship between a conditioned stimulus (e.g., a tone) and an unconditioned stimulus (e.g., a shock) is fully learned, consistent copresentation of a second stimulus (e.g., a light) with the original conditioned and unconditioned stimuli will fail to produce any evidence of learning of the relationship between the second conditioned stimulus and the unconditioned stimulus (Kamin 1969). That is, the second stimulus evokes no response from the organism despite being perfectly predictive of the unconditioned stimulus. Analogous effects have been reported in human learning tasks (e.g., Dickinson, Shanks, and Evenden 1984; Waldmann and Holyoak 1992).

This phenomenon is pertinent to consumer learning because consumers often do not encounter all predictive cues simultaneously, and it is likely that exposure to brand cues will precede exposure to information about the attributes that truly determine quality. In such instances, consumer learning mimics the blocking paradigm, with brands blocking the learning of attributes. Insofar as blocking is observed, a previously unidentified source of brand equity is uncovered.

The notion that brands may inhibit learning of attributes may seem nonintuitive in light of the causal status of the former. Brands signal quality by virtue of their correlation with quality but brands do not “cause” quality. Brand names are merely surrogates for the intrinsic attributes and processes that truly determine quality. Consumers should treat a brand name as a useful heuristic—or as proxy for quality-

*Stijn M. J. van Osselaer is assistant professor of marketing, University of Chicago. Joseph W. Alba is distinguished professor of marketing, University of Florida. This work was supported by the Beatrice Foods Co. Faculty Research Fund at the University of Chicago Graduate School of Business. The authors thank Alan Cooke, Joffre Swait, Chris Janiszewski, and Bart Weitz for their helpful comments.

determining attributes—but should abandon it when heuristic reasoning is inappropriate. Evidence of blocking, however, would demonstrate that a correlational cue can suppress learning of a causal feature, even when the latter is available and equally predictive of quality.

In the following experiments, we examine the possibility of this outcome while also assessing the merits of two popular but competing theoretical explanations of the blocking phenomenon. Consistent with its roots in animal conditioning, one view relies on associative learning processes. Briefly, associative and connectionist theories hold that humans and animals predict outcomes based on associations with cues (see, e.g., Gluck and Bower 1988; Shanks 1994; Shanks, Medin, and Holyoak 1996; Young 1995). Co-occurrence of a cue and an outcome strengthens their association via a simple learning process. A critical assumption is that no additional learning occurs once perfect prediction is achieved. Thus, additional cues that later appear contemporaneously with the original predictive cue gain no association strength. Colloquially, these additional cues are merely redundant, lacking the ability to improve predictability and therefore lacking value to a system designed to predict the environment.

In contrast, replication of the effect with human subjects and meaningful stimuli has inspired a more cognitive or logic-based explanation (e.g., Cheng 1997; Cheng and Novick 1992; Waldmann and Holyoak 1992). According to this view, people assess predictiveness and causality through normative causal-reasoning processes. Just like scientists, people search for data that provide unconfounded contrasts free of ceiling and floor effects. In the case of blocking, the first cue associated with the outcome is unconfounded with any other cue and therefore is likely to be viewed as causal via normal reasoning processes. The second cue occurs only in the presence of the first cue and therefore never uniquely predicts the outcome. Thus, people should be agnostic about its causal status. We examine the merits of these competing views in more detail as our data become pertinent. Based on our empirical outcomes, we also entertain a third explanation.

We begin the research with an investigation of the basic applicability of the blocking paradigm to consumer learning and the role that prior beliefs may play in it. In experiment 1 (and throughout), the effect of preexposure to a predictive cue on learning of subsequent predictive cues is examined. To examine the relevance of the procedure to brand equity as well as the generalizability of the effect to other contexts, the cue to which subjects are preexposed is manipulated. In some conditions, preexposure is to predictive brand names; in other conditions, preexposure is to an equally predictive intrinsic product attribute. The remaining experiments are conducted within the context of predictive brand cues. Experiment 2 examines the moderating effects of prior beliefs by manipulating the a priori relationship between the subsequently provided predictive cues and the to-be-predicted outcome. Experiments 2a and 2b examine alternative ex-

planations for the results. Experiments 2b, 3, and 4 explore the merits of currently competing models of the basic effect.

EXPERIMENT 1

The most basic goal of experiment 1 was to determine whether the blocking effect can be replicated with product features in a consumer context. The second and more important goal was to explore consumers' causal beliefs about brands versus attributes as determinants of product quality. To the extent that brands are regarded as noncausal signs, they should yield to predictive intrinsic attributes encountered later. We examine this prediction by manipulating the initially encountered predictive cue. Some subjects first learned that brand name was predictive of quality; other subjects first learned that an attribute was predictive of quality. All subjects were then exposed to a second predictive (attribute) cue. If brands are viewed as signals rather than causes, subjects in the brand condition should exhibit less blocking; that is, they should be more willing to treat the subsequently learned cue as the true cause of quality.

In addition to the manipulation of the preexposed feature (brand name vs. intrinsic attribute), the consumer learning context used in experiment 1 differed from previous blocking experiments in at least two ways: (1) it employed a very small number of preexposures (four), and (2) subjects did not make an explicit outcome prediction (e.g., a quality prediction) on each trial. Both differences made experiment 1 a more conservative and externally valid test of the blocking phenomenon than any study to date.

Method

Design and Subjects. The basic paradigm in experiment 1 consisted of showing subjects a number of product descriptions that included information about product features as well as the overall quality level. Subjects in the experimental conditions received two series of product profiles. In the first set, only one product cue was predictive of quality. In the second set, a second predictive cue was added. Following Waldmann and Holyoak (1992), we will refer to the first predictive feature as the "Predictive" feature and to the second predictive feature as the "Redundant" feature. The Predictive feature referred to either brand name or an intrinsic attribute and was manipulated between subjects. The Redundant feature was always an intrinsic attribute. The control conditions differed from the experimental cells only by the absence of the first set of profiles. Thus, in the control conditions, subjects learned about both predictive cues simultaneously and were not preexposed to information about one cue.¹ The basic experimental design in experiment 1 was a 2 (experimental vs. control) × 2 (type of Predictive feature: brand vs. attribute) completely ran-


¹Note that the Redundant feature in the control conditions is not any more or less "redundant" than the Predictive feature. In the control conditions, both features are always presented together. Hence, they are introduced simultaneously and appear equally often.

domized factorial design (see Fig. 1 for a schematic representation of stimuli; see the Appendix for design summaries of this and the following experiments). Forty undergraduate subjects participated in experiment 1 in return for extra course credit.

Stimuli. The stimuli are most easily described by condition. Each of the four cells of the design involved slight variations on the others. In the brand/control condition, subjects were presented with 12 profiles of whitewater rafts. Each profile described a particular whitewater raft, listing its brand name (Hypalon or Riken), type of compartments (airecell or closed-cell), type of hull (always polyurethane), type of floor (tubular or I-beam), and its quality level (high or low). Six of the rafts were high quality and six were low quality. Two features, brand and floor, were perfectly predictive of quality. All high-quality rafts carried the Hypalon brand name and had a tubular floor. All low-quality rafts

were Riken brand rafts and had an I-beam floor. The position of the two predictive features in the product descriptions (at the top of the list of features vs. at the bottom) was counterbalanced between subjects. Thus, for half the subjects, brand information was presented at the top of the profile and floor information near the bottom; for the other half, brand information was presented near the bottom and floor information at the top. Hull and compartments were un-predictive filler features. All high-quality rafts and all low-quality rafts had a polyurethane hull. Half of the high-quality rafts and half of the lower-quality rafts had airecell compartments. The other half of the high-quality rafts and the other half of the lower-quality rafts had closed-cell compartments. Thus, hull was a “constant” filler feature that was not predictive of quality, and compartments was a varying but “uncorrelated” filler feature that was equally un-predictive of quality. In sum, subjects in the brand/control con-

FIGURE 1
SCHEMATIC REPRESENTATION OF STIMULI (EXPERIMENT 1, BRAND CONDITIONS)



NOTE.—In the attribute conditions, brand was replaced by an attribute label (frame, bow, steering system, or outer skin).

dition were simultaneously exposed to information about two predictive features, brand (Predictive feature) and floor (Redundant feature).

Subjects in the brand/experimental condition were presented with the same 12 profiles but were first preexposed to four extra profiles. In these four profiles, information was given about brand name, compartments, hull, and quality but not about the type of floor. Two of these additional rafts were high quality and two were low quality. As in the subsequent 12 trials, the high-quality rafts carried the Hypalon brand name and the low-quality rafts were Riken brand rafts. Compartments and hull again were not predictive. Thus, the experimental group was identical to the control group with the exception that it was preexposed to four profiles that indicated a brand-quality association but were mute about a floor-quality association.

Stimuli in the attribute/control and attribute/experimental conditions differed from corresponding brand conditions only in terms of the Predictive feature. Brand name was replaced by an intrinsic attribute (frame, bow, steering system, or outer skin). Thus, high-quality rafts had hypalon frames, bows, steering systems, or outer skins and low-quality rafts had riken frames, bows, steering systems, or outer skins. Consequently, subjects in the attribute/control condition were simultaneously exposed to information about two predictive intrinsic attributes; subjects in the attribute/experimental condition were preexposed to information about one predictive attribute (frame, bow, steering system, or outer skin) before receiving information about another predictive attribute (floor). The four different attribute labels (frame, bow, steering system, and outer skin) were counterbalanced across subjects and used to enhance generalizability beyond any specific attribute label.

Procedure. Subjects were run in groups of one to five and were randomly assigned to conditions. Each subject was seated in front of a personal computer shielded from other subjects. Instructions and stimuli were presented via the Micro Experimental Laboratory software system (Schneider 1988). Instructions informed subjects about the type of stimulus information they would receive and requested that they try to “learn how to predict the quality level of whitewater rafts” and to “determine what really matters in whitewater rafts.” Subjects then were sequentially presented with 16 (experimental conditions) or 12 (control conditions) profiles of whitewater rafts. Presentation was self-paced. The profiles in the experimental conditions were separated into two learning phases. Phase 1 contained the four preexposure profiles that indicated a Predictive feature-quality association but were mute about the floor-quality association. Phase 2 followed phase 1 and contained 12 profiles with information about both cues. Phases 1 and 2 were separated by at least 10 seconds. Subjects in the control conditions received only the phase 2 stimuli.

Measures. Several dependent measures were collected after phase 2. The primary dependent measure consisted of eight additional profiles of rafts containing information

about all four product cues. The eight profiles were constructed as a $2 \times 2 \times 2$ factorial combination of the two levels of the Predictive cue, the Redundant cue (floor), and the uncorrelated filler feature (compartments). The constant filler feature (hull) was also added to each profile, as in the stimulus profiles of the learning phase(s). For each of the eight rafts, subjects indicated whether the quality level of the raft was high or low.²

Results

To summarize briefly, results showed a strong blocking effect on the Redundant feature that was not qualified by the nature of the Predictive feature (brand vs. attribute). Results also showed a positive effect of preexposure on the value subjects placed on the Predictive feature. No statistically significant effects were found of the Predictive and Redundant features' position in the list of features. Hence, the data were collapsed across the levels of this counterbalancing variable.

The results for the primary measure were analyzed using different logistic regression models that each included the levels of one of the two predictive features and the independent variables. Because all hypotheses are at the level of the feature dimension (e.g., “floor” instead of “tubular floor”) the conceptual dependent variable is the difference in evaluation between products that had one level of a feature and products that had the other level of the feature. For example, to assess the value placed on the floor feature in the brand/control cell, a “mean feature effect” (MFE) of floor was computed by subtracting the percentage of I-beam floors in that cell that were judged to be high quality (12.5 percent) from the percentage of profiles with a tubular floor judged to be high quality (85 percent, leading to a mean feature effect of floor of $.850 - .125 = .725$). For practical purposes, this mean feature effect can be interpreted as a floor part-worth or a floor weight. Means for each condition are provided in Table 1.

Statistical analyses revealed, unsurprisingly, that preexposure to a predictive cue raises its weight relative to a control condition (asymptotic $t = 2.28, p < .01$). These analyses also revealed that this effect occurs equally strongly for brand and attribute cues (asymptotic $t = -.09$) and that

²In addition to the binary quality judgments, willingness-to-pay judgments and predictiveness ratings were collected. For the willingness-to-pay measure, subjects were presented with some additional four-feature profiles. This time, the levels of both filler features were held constant. Nine profiles were constructed as a 3×3 factorial combination of the two levels of the Predictive cue and Redundant cue presented in the learning phase(s) plus a not previously presented level (“ensolite” brand name, steering system, bow, frame, or outer skin, and “batten” floor). Subjects were asked to guess the price that they thought would be reasonable for each raft. They were told that prices ranged from \$1,500 to \$2,500 with an average of \$2,000. For the predictiveness measure, subjects were asked to rate each of the product dimensions (brand, steering, bow, frame, or outer skin; compartments; hull; and floor) independently as to whether it predicted the quality level of rafts. A 21-point scale ranging from -10 (definitely does not predict quality level) to 0 (don't know) to +10 (definitely predicts quality level) was used.

brand and attribute cues received similar weights overall (asymptotic $t = .48$). More important, such preexposure resulted in a strong blocking effect. When subjects received just four preexposures to one predictive (Predictive) cue before being exposed to information about a second predictive (Redundant) feature, the second feature had a much smaller effect on product evaluations (MFE = .21) than when subjects were exposed to both predictive features simultaneously (MFE = .60; asymptotic $t = -3.51, p < .001$). Thus, blocking can be obtained in a consumer context using a relatively conservative procedure.

Also important is the test of the causal beliefs hypothesis. If brands are treated as signals rather than causes, the blocking effect should be smaller when the first predictive cue refers to a brand than when it refers to an intrinsic attribute. The data, however, do not support this hypothesis (asymptotic $t = 1.39, p > .16$). The effect of the Redundant feature (floor) on quality judgments after preexposure was as small when the first predictive cue referred to an intrinsic attribute (MFE = .23) as when it referred to brand name (MFE = .19; asymptotic $t = .19$). The main effect of the type of Predictive feature (brand vs. intrinsic attribute) on the effect of the equally predictive Redundant feature also was not statistically significant (asymptotic $t = -1.62, p > .10$).

With regard to the issue of brand equity, the results indicate that there was no tendency for subjects in the brand/experimental condition to discard the brand once a predictive attribute was encountered. In fact, the impact of the latter was strongly suppressed. One might argue that the brand retained its dominance over the Redundant attribute simply because brands enjoy special status. However, results from the brand/control condition do not support this explanation. When the predictive brand and attribute cues were presented simultaneously, subjects placed a much higher weight on the attribute—consistent with the hypothesis that attributes are viewed as more causal of quality than is a brand name.³

³Results from the other dependent measures were essentially identical with those observed from the binary quality measure. These additional measures are important because they refute an alternative explanation of the blocking effect. The low weight placed on the Redundant feature (floor) may have been caused not by a decrease in the perceived importance of this feature but indirectly by an increase in the perceived importance of the Predictive feature due to the four extra exposures to that feature in the experimental conditions. Thus, it is possible that subjects did learn that floors were important predictors of quality in whitewater rafts but did not use this information because another, more accessible or diagnostic cue (e.g., a known brand name) was present. This explanation is not consistent with the strong blocking effect found on predictiveness ratings of the Redundant feature only (i.e., independent of the Predictive feature; $F(1, 36) = 7.31, p < .02$) and also is not consistent with the results of a separate analysis on the WTP judgments for profiles containing a new, unknown level of the Predictive feature. In these cases, subjects should have relied on what they learned about the Redundant feature (floor), because no other diagnostic information was available. Results again revealed a significant blocking effect ($F(1, 36) = 5.71, p < .05$). In absolute dollar terms, the raw dollar premiums subjects were willing to pay for a tubular floor versus an I-beam floor were three times larger in the control conditions than in the experimental conditions. Thus, even in the absence of the preexposed Predictive cue, subjects in the experimental condition failed to place much credence in the Redundant attribute.

TABLE 1

MEAN FEATURE EFFECTS (EXPERIMENT 1)

Type of Predictive feature/feature	Control	Experimental	Difference
Brand:			
Brand (P)	.28	.69	.41
Floor (R)	.73	.19	-.54
Attribute:			
Attribute (P)	.33	.73	.40
Floor (R)	.48	.23	-.25

NOTE.—Abbreviations used are “P” for “Predictive feature” and “R” for “Redundant feature.”

Discussion

Experiment 1 suggests that providing consumers with information about a predictive intrinsic attribute does not necessarily cause them to value the attribute at the expense of the brand name. When subjects had experience with brand names before they were exposed to equally predictive attribute information, the attribute information gained little predictive value, and a very strong brand preference was observed. This strong brand preference and weak attribute preference would be unsurprising if subjects accorded special status to brand names. However, results in the control conditions of this and following experiments are not consistent with such an explanation. For example, the binary quality judgments in the brand/control condition showed that brand names retained little predictive value when subjects learned about brands and attributes simultaneously. Thus, subjects appeared sensitive to the causal status of quality cues, and exposure to information about a predictive intrinsic attribute in this situation caused subjects to value the attribute at the expense of the brand name.

The large differences observed between the brand/control and brand/experimental conditions suggest that the blocking phenomenon is robust with respect to consumers’ preexperimental beliefs about which features should be more or less important for a particular consumption outcome. It is also possible, however, that subjects’ general causal beliefs about brands versus attributes were rather weak and that stronger, more specific prior beliefs relating the Redundant feature to the outcome would decrease the blocking effect. Broniarczyk and Alba (1994) have shown that strong a priori beliefs about product attributes can drive predictions of quality, even in the face of disconfirming evidence. Thus, a stronger test of the blocking potential of a brand name would be to use a Redundant attribute that is more explicitly related to the outcome. In the next experiment we retain brand as the Predictive cue and manipulate the conceptual relationship between the Redundant feature and the outcome.

EXPERIMENT 2

Experiment 2 was designed to explore the robustness of the blocking phenomenon by increasing subjects’ opportunity to learn and understand the significance of the Redundant intrinsic attribute after preexposure to brand infor-

TABLE 2
MEAN FEATURE EFFECTS (EXPERIMENT 2)

A priori relatedness/ feature	Control			Experimental			Difference		
	12	24	All	12	24	All	12	24	All
Low (outcome: quality):									
Brand (P)	.31	.43	.36	.77	.75	.76	.46	.32	.39
Rudder (R)	.69	.43	.57	.10	.14	.12	-.59	-.29	-.45
High (outcome: steering performance):									
Brand (P)	.19	.21	.20	.41	.57	.49	.22	.36	.29
Rudder (R)	.81	.71	.76	.38	.43	.40	-.43	-.28	-.36

NOTE.—“12” refers to conditions with 12 trials in phase 2. “24” refers to conditions with 24 trials in phase 2. “All” refers to results collapsed across “12” and “24” conditions.

mation. This objective was accomplished in two ways. First, we manipulated the conceptual relationship between the outcome and the to-be-blocked Redundant attribute. Second, we manipulated the number of exposures to the Redundant attribute, doubling it from 12 to 24 exposures in half the conditions.

Method

A total of 111 subjects participated in groups of one to eight. The procedure and stimulus format for experiment 2 were similar to those used in experiment 1 with the following exceptions. In all conditions, the Predictive feature referred to brand name (Hypalon or Riken) and the Redundant feature was “rudder” (with fourcase as the positive level and backbar as the negative level). To make the Redundant feature stand out further as an intrinsic attribute, the unpredictable filler features in experiment 2 were extrinsic attributes (warranty and country of origin).

For half the subjects (low priors conditions), the consumption outcome information was information about the raft’s general quality level (high quality vs. low quality), an outcome with little particular relatedness to the Redundant attribute (rudder). For subjects in the high priors conditions, the outcome pertained to the steering quality of the rafts (good steering performance versus bad steering performance). In these conditions, a very strong a priori relationship should exist between the Redundant attribute (rudder) and the consumption outcome (steering). At a minimum, this manipulation should be a stronger manipulation of causal beliefs than the brand-versus-attribute manipulation in experiment 1, inasmuch as the type of rudder should have a stronger and more direct causal influence on steering performance than on overall quality.

The other major difference between experiments 1 and 2 was the manipulation of the number of product exposures in the second learning phase. In the second learning phase, half the subjects received 12 learning trials and the other half received 24 learning trials. Thus, the basic design of experiment 2 was a 2 (experimental vs. control) \times 2 (number of phase 2 trials) \times 2 (conceptual relatedness) completely randomized factorial design.

The primary dependent measure was a binary quality measure equivalent to the one used in experiment 1.⁴

Results and Discussion

Overall, the blocking effect was robust with respect to the conceptual relationship between the Redundant cue and outcome as well as the number of exposures to the Redundant cue–outcome relationship. Table 2 shows the mean feature effects for the binary quality judgment measure. The results for the effect of the Redundant feature (rudder) showed a significant blocking effect of brand name preexposure on quality judgments (asymptotic $t = -2.41$, $p < .02$). That is, rudder received a lower weight in the experimental than control conditions. The analysis also yielded a significant effect of the relatedness manipulation (asymptotic $t = -2.46$, $p < .05$), indicating that subjects put more weight on the rudder attribute when the outcome was steering performance than when the outcome was overall quality. This effect can be interpreted as a successful manipulation check on the relatedness manipulation. Interestingly, no support was found for the hypothesis that the blocking effect would be moderated by specific prior beliefs (asymptotic $t = .47$). In addition, this analysis yielded no significant effects of the number of trials (12 vs. 24) in phase 2 (all $ps > .25$).

The results for the effect of brand (the Predictive feature) on quality judgments generally mirrored the results for rudder (the Redundant feature). Just four preexposures to brand information caused a large increase in the mean feature ef-

⁴In addition, a simplified willingness-to-pay measure was used in which subjects were presented with three pairs of raft profiles. In each pair, one of the rafts was a Riken brand raft that was made in Canada, with a one-year warranty, a backbar rudder, and a price of \$2,000. The second raft was identical except that it had the alternative level of one of the features (e.g., Hypalon brand, made in Canada, one-year warranty, backbar rudder) and lacked a price. After being informed that prices on the market range from \$1,500 to \$2,500, subjects were asked to indicate for each pair how much more or less they were willing to pay for the second raft compared to the first raft. Thus, subjects were asked to directly indicate the dollar premium for a raft that differed from another raft by only one feature. Results for the willingness-to-pay judgments showed the same statistical pattern of results found with the binary quality judgments and will not be discussed further.

fect of brand (asymptotic $t = 2.79$, $p < .01$). The preexposure effect did not interact with relatedness (asymptotic $t = .36$) nor were there any other higher-order interactions. No statistically significant effects were found of the manipulation of the number of phase 2 trials (all $ps > .30$). The effect of the relatedness manipulation on the weight subjects placed on brand names was not statistically significant (asymptotic $t = 1.60$, $p > .10$).

An individual-level analysis of the data not surprisingly revealed that a minority of subjects in the experimental conditions failed to show blocking. Inasmuch as learning of the brand-quality relationship in phase 1 is a prerequisite for blocking, a follow-up study was conducted to assess the degree of phase 1 learning. An independent group of 23 subjects experienced phase 1 only of the low priors/experimental condition and then were asked to make binary quality judgments. Results showed that 74 percent of the subjects based their judgments exclusively on brand name, indicating learning of the brand-quality relationship. This figure corresponds almost precisely to the 73 percent of subjects in the corresponding condition of the main experiment who based their binary quality judgments exclusively on the brand name at the end of phase 2. This result suggests that blocking obtained for all subjects who initially learned the brand-quality relationship and that any acquisition of value by the Redundant attribute in the second learning phase by subjects in the experimental condition was due to a lack of learning in phase 1. In this sense, the results reported in the preceding experiments are very conservative. It appears that all subjects who learn the brand-quality relationship in phase 1 exhibit blocking.

EXPERIMENT 2A

The blocking effect found in experiments 1 and 2 could be caused by several underlying processes. One explanation is the "encoding account" (e.g., Price and Yates 1995). According to this account, subjects quickly learn which cue is most predictive of an outcome and then direct most of their attention to that cue. In a blocking context, this explanation holds that learning about the Predictive cue in phase 1 causes subjects to direct attention away from the Redundant cue in phase 2. As a result, subjects do not encode the occurrences of the Redundant cue, cannot reproduce the pattern of co-occurrence between that cue and the other cues or outcome, and do not learn that it is predictive of the outcome. Empirical support for the encoding account would include evidence that consumers not only fail to search for additional information once a sufficient cause or predictor has been found (cf. Shaklee and Fischhoff 1982) but that they even fail to attend to additional product features when available.

Experiment 2a addressed the encoding account by asking subjects in a standard two-phase blocking condition to recall the pattern of co-occurrence of the Redundant feature and the outcome, in addition to making binary quality judgments. The encoding account predicts a negative relationship between blocking and ability to recall the pattern of co-

occurrence between the Redundant feature and quality. Further, because the encoding account holds that attention to one cue leads consumers to turn attention away from subsequently encountered cues, the encoding account also predicts a negative relationship between consumers' ability to recall the pattern of co-occurrence between the Predictive feature and quality and their ability to recall the pattern of co-occurrence between the Redundant feature and quality.

Method

Fifty-two subjects received a treatment similar to the 12 trials/low priors/experimental cell of experiment 2. The only difference was that in addition to the eight binary quality judgments, subjects estimated the co-occurrence between product features and quality. The two tasks were counterbalanced. The syntax of the co-occurrence questions was: "In the second set of rafts, how many of the (*high/low*) quality rafts had a (*level feature*)?" Responses were taken on a five-point scale anchored by 0 percent (None, all [*high/low*] quality rafts had an [*other level feature*]) and 100 percent (All [*high/low*] quality rafts had a [*level feature*]). The scale points were multiples of 25 percent.

Results and Discussion

Co-occurrence perceptions between the Redundant feature (rudder) and quality were computed by subtracting the perceived percentage of the high-quality rafts that had a fourcase rudder from the perceived percentage of the low-quality rafts that had a fourcase rudder. Subjects' co-occurrence perceptions were classified into three categories. If subjects correctly identified that 100 percent of the high-quality rafts had a fourcase rudder and that 0 percent of the low-quality rafts had a fourcase rudder, they were classified as "fully accurate." Subjects who merely identified a higher percentage of fourcase rudders for high-quality rafts than for low-quality rafts were classified as "partially accurate." Subjects claiming that the rudder in the co-occurrence question was equally often paired with high quality as with low quality were classified as "inaccurate." The latter category also was used for subjects indicating a lower percentage of fourcase rudders for high-quality rafts than for low-quality rafts.

Results for the two dependent measures (quality and co-occurrence judgments) were first tested for order effects. These analyses showed no effect of order on co-occurrence judgments (asymptotic $t = .51$), but they did show a difference in the number of subjects displaying a "blocked" pattern of quality judgments when those judgments were made after having judged co-occurrence (asymptotic $t = -1.81$, $p = .07$).⁵ We therefore concentrate our discussion on the group that made quality judgments prior to co-occurrence judgments.

Table 3 shows that of the 19 subjects demonstrating an

⁵One of the subjects failed to answer some of the dependent measures. Therefore, most analyses relied on answers from 51 instead of 52 subjects.

TABLE 3
 FREQUENCIES (PROPORTIONS) OF SUBJECTS SHOWING QUALITY AND CO-OCCURRENCE JUDGMENT PATTERNS
 (EXPERIMENT 2A)

Order/quality judgments depending on	Co-occurrence judgments			
	Fully accurate	Partially accurate	Inaccurate	All
Quality first:				
Brand	9 (.47)	6 (.32)	4 (.21)	19 (1.00)
Other	1 (.14)	1 (.14)	5 (.71)	7 (1.00)
All	10 (.38)	7 (.27)	9 (.35)	26 (1.00)
Co-occurrence first:				
Brand	3 (.25)	4 (.33)	5 (.42)	12 (1.00)
Other	7 (.54)	4 (.31)	2 (.15)	13 (1.00)
All	10 (.40)	8 (.32)	7 (.28)	25 (1.00)

NOTE.—Not all proportions add up to 1.00 due to rounding after the second digit.

entirely brand-driven pattern of quality judgments, 15 showed at least partially accurate perceptions of the co-occurrence between the Redundant attribute and quality. Of these 15, nine were fully accurate. Logistic regression analyses showed that entirely brand-driven subjects were significantly more likely than not to have at least partially accurate memory for the co-occurrence of the Redundant cue and quality level (asymptotic $t = 2.35$, $p < .02$).

Logistic regression analyses also showed that entirely brand-driven subjects were significantly *more* likely to show at least partially accurate co-occurrence recall than subjects whose quality judgments were not entirely brand driven (asymptotic $t = 2.24$, $p < .03$). Thus, evidence was found against a negative relationship between placing a low weight on rudder information (as evidenced by a “blocked” pattern of quality judgments) and ability to recall the pattern of co-occurrence between different types of rudders and quality. Finally, contrary to the encoding explanation, a correlational analysis of the rudder-quality and brand-quality co-occurrence judgments did not yield a negative correlation between the accuracy of co-occurrence judgments for brand and rudder ($r = .15$). Together, the data do not support an attentional explanation of the blocking effect. That is, blocking of the Redundant attribute obtains even when its relationship to the outcome is soundly encoded. Of course, failures to attend to the Redundant attribute surely occur in the world and provide another avenue to the blocking effect. Experiment 2a demonstrates, however, that encoding failure is not a necessary condition.

Finally, we should note that blocking was weaker among subjects who estimated the co-occurrence of the Redundant cue and the outcome prior to making quality judgments. Although such an outcome is consistent with a “demand” to be consistent across responses, it is also possible that blocking is reduced when the Redundant cue’s relationship with the outcome is made salient and consumers are compelled to consider its implications.

EXPERIMENT 2B

As noted, cognitive explanations of blocking are based on normative causal reasoning about the presence and ab-

sence of specific cues and outcomes. In terms of a previous example, the Predictive cue (tone) varies independently (and in the absence) of the second predictive cue (light) during phase 1. Because the tone is always followed by the outcome (shock) and the tone is unconfounded with the light, normative causal reasoning should lead to the conclusion that the tone causes shock. For the Redundant cue (light), the only unconfounded contrast is provided by comparing the “light present, tone present” situation with the “light absent, tone present” situation. Because the shock always appears in the “light absent, tone present” situation, any “light effect” is masked by a ceiling effect. Thus, subjects should conclude that the Predictive cue has a strong effect on the outcome but that the role of the Redundant feature is indeterminate (Cheng 1997).

Although results from the experimental groups in experiments 1 and 2 conform to these predictions, results from the control conditions do not. In the control conditions, the Predictive and Redundant cues were completely confounded. Thus, subjects should have been just as agnostic about their influence on the outcome as they were about the Redundant feature’s influence in the experimental conditions. Our findings are clearly inconsistent with this prediction.⁶ Experiment 2b was designed to assess another aspect of the cognitive account and to address an important concern regarding the generalizability of the blocking effect.

According to the causal reasoning view, people retrospectively access their memory of specific stimuli, code and summarize the stimuli, and then perform the normative analysis that produces their assessments of causality and predictiveness. The normative causal reasoning explanation of the blocking effect depends on specific assumptions about the retrospective coding of events that took place in the first learning phase. Consider again the classic demonstration. In Kamin’s (1969) experiment, the status of the Redundant cue (light) in phase 1 was clear. It was not on. Also, the status of that cue in phase 1 was the same as in trials in which

⁶For example, the average predictiveness rating for the Redundant feature in the control condition of experiment 1 did not indicate agnosticism (8.60 on a scale ranging from -10 [definitely does not predict quality level] to 0 [don’t know] to +10 [definitely predicts quality level]) and was significantly higher than the average rating in the experimental condition (4.80).

the outcome (shock) did not occur in phase 2. It was not on. Thus, the appropriate conditional contrast could be computed for the Predictive cue (tone) in phase 1, but not for the Redundant cue (light). In Kamin's experiment it can be assumed that subjects normatively code the first phase trials as "Redundant cue absent." However, in experiments 1 and 2, floors or rudders were not absent or present in phase 2 but, like many product features, were present in two different forms. The coding issue in the case of such features is more complex. A normative causal reasoning account predicts that a subject who received information in the second learning phase showing that rafts had either a tubular floor or an I-beam floor would be uncertain about the floors possessed by the rafts in phase 1. If so, subjects should be uncertain about the causal status of the first predictive feature (e.g., brand) because it is now potentially confounded with the second predictive feature (e.g., floor). Our data, however, show little sign of uncertainty about the Predictive feature in the experimental conditions.

Although intuitively unlikely, it is also possible that subjects inferred that the brands in phase 1 either did not have any level of the Redundant feature or had levels different from the levels in phase 2. If so, causal reasoning models could still make the same predictions as in the classic case in which the Redundant feature is clearly absent in phase 1. This raises the possibility that the blocking effect is dependent on particular inferences about missing information. It also suggests that blocking should not occur in the common situation in which consumers naturally assume that the products they received information about first were the same as the ones they received more detailed information about later. (For example, it is unlikely that consumers assume that the Tylenol headache medicine they saw advertised on TV had different ingredients than the Tylenol found at the store.) When consumers appropriately infer that brands have the same attributes during the first learning phase as during the second learning phase, brand and the Redundant attribute are equally confounded in both phases and no preexposure effects should obtain. Thus, a stronger test of both external validity and the causal reasoning hypothesis involves the situation in which there is no uncertainty about the implicit status of the Redundant feature in phase 1 learning trials.

Method

Sixty-eight subjects, in groups of one to six persons, participated in a two-group design. In the control condition, subjects were presented with 12 profiles of whitewater rafts that were purportedly taken from a catalog. As in previous experiments, these profiles contained two product cues (brand and floor) that were predictive of quality as well as two unpredictable intrinsic filler features (rudder and fabric). In the experimental condition, subjects received the same 12 profiles but were preexposed to 12 other profiles. As in previous experiments, these phase 1 profiles contained information about quality, brand, and two filler features but were mute about floors. Brand was again predictive of quality, and the filler features were not. However, unlike in the

previous experiments, the phase 1 profiles were said to come from a different catalog than those in phase 2. In addition, the phase 1 profiles contained extrinsic filler features (country of origin and warranty) that were different from the intrinsic filler features used in phase 2.

In the answer booklet, subjects in both conditions received additional information about whitewater rafts as part of the introduction to the dependent measure. Embedded in this information were statements saying that because whitewater rafts have many different features, most catalogs can only provide part of the information and that different catalogs might offer the same products but provide information about different types of features. A sentence was then inserted for the experimental group saying that the rafts they saw in the first catalog were exactly the same as those they saw in the second catalog. These instructions effectively eliminated any ambiguity about the status of the (implicit) Redundant cue in phase 1 trials. The manipulation was inserted at the time of test (instead of earlier) to maximize its temporal proximity to the retrospective coding process that should take place at test.

Following this introduction, subjects made binary quality judgments. All eight measurement profiles contained the intrinsic filler features; no mention was made of the extrinsic filler features used in the first learning phase. Finally, subjects in the experimental group responded to a question designed to check whether they had read and attended to the critical part of the instruction saying that the rafts in the two learning phases were the same.

In sum, the main difference between experiment 2b and the previous experiments was that subjects in the experimental condition were explicitly told that the products in phase 1 and phase 2 were the same. Thus, the two predictive features were totally confounded in both phases for these subjects. According to the causal reasoning hypothesis, no blocking should occur.

Results and Discussion

The binary quality judgments showed effects of brand preexposure on the mean feature effects of brand ($MFE_{\text{control}} = .41$, $MFE_{\text{experimental}} = .63$; asymptotic $t = 3.00$, $p < .01$) and floor ($MFE_{\text{control}} = .56$, $MFE_{\text{experimental}} = .37$; asymptotic $t = -2.51$, $p < .02$). Of the 34 subjects in the experimental group, only four did not correctly remember the "same rafts" instruction. Excluding these subjects from the analysis did not change the results (asymptotic $t = 2.99$, $p < .01$ for brand, and asymptotic $t = -2.56$, $p < .02$ for floor). The results do not support the normative causal reasoning account of blocking. A significant blocking effect was found despite explicit instructions that the products in the two learning phases were identical.

The present experiment provides another conservative test of the blocking effect. The explicit realization that information was missing should have caused subjects to be more uncertain about their judgments, which, in turn, should have led to more equal weighting of the Predictive and Redundant

cues in both conditions. To the contrary, a strong blocking effect was again observed.

Taken together, experiments 1–2b demonstrate a robust blocking effect in a consumer context. Moreover, the effect cannot be attributed to relatively uninteresting factors, such as lack of attention to critical stimuli or idiosyncratic inferences about implicit information. From a theoretical perspective, our results are inconsistent with a normative causal reasoning explanation that requires retrospective consideration of brand-related experiences. Rather, subjects appear to engage in a forward-looking process aimed at establishing accurate prediction of future outcomes, with little importance accorded to information that does not help to improve prediction. The findings are consistent with associative or connectionist explanations of blocking (e.g., Kruschke 1996; Pearce 1994; Rescorla and Wagner 1972). In the remaining experiments we broaden the investigation to other realistic scenarios in an attempt to examine moderating conditions. To our knowledge, these scenarios have not been examined in prior research. Thus, another primary objective of the remaining experiments is to shed additional light on the learning process.

EXPERIMENT 3

In experiments 1–2b, we have investigated only scenarios in which consumers experience just one high- (low-) quality brand prior to exposure to redundant attribute information. Of course, consumers sometimes experience multiple high- (low-) quality brands before receiving information about the attribute that really drives product performance. For example, consumers may first learn that two brands of pain killers (e.g., Advil and Motrin) work well before they learn that both contain ibuprofen. We investigated this scenario in experiment 3.

Investigating the shared attribute scenario is important, not just to test generalizability to a common consumer learning situation, but because it allows us to investigate the role of parsimony in consumer learning. Note that in the previous experiments, both the Predictive feature and the Redundant feature provided equally parsimonious explanation of performance. However, when multiple Predictive brands share the same level of the Redundant attribute, the Redundant attribute becomes the more parsimonious cue. It is possible that consumers decide to attribute performance to the first predictive cue only when it is at least equally as parsimonious as predictive cues encountered later. Thus, it is possible that the learning system is aimed not at establishing accurate prediction but at establishing the most parsimonious accurate prediction. In that case, no blocking should occur when consumers experience multiple brands that share the same level of an attribute. In contrast to this prediction, associative models predict that the Predictive cue should still interfere with the acquisition of predictive value by the Redundant attribute, even though this Redundant attribute provides a more parsimonious explanation. As will be explained in more detail in the “General Discussion” section below, simple associative models do predict a main effect

of parsimony. However, they could not account for a decrease in the effect of preexposure on the value of the Redundant attribute as it becomes more parsimonious.

Method

Design. In experiment 3, we manipulated preexposure (experimental vs. control) and the number of brand names per attribute level (one vs. two) in a 2×2 completely randomized factorial design.

Stimuli and Procedure. Eighty-two subjects participated in groups of one to six using a procedure identical to that in the previous studies. In the one/control condition, subjects received 16 profiles of whitewater rafts with information about quality, a Predictive feature (brand), a Redundant feature (floor), a constant filler feature (warranty), and a varying but unpredictable filler feature (country of origin). Eight rafts were high quality, carried the Hypalon brand name, and had a tubular floor. The eight remaining rafts were low quality, carried the Riken brand name, and had an I-beam floor. Thus, brand and Redundant attribute information were learned simultaneously, and each level of the Redundant attribute was coupled with just one brand name. This condition was structurally identical to the basic control conditions in the previous experiments. In the one/experimental condition, subjects received the same 16 profiles but were preexposed to eight profiles. Four of these phase 1 rafts were high quality and carried the Hypalon brand name. The remaining four rafts were low quality and carried the Riken brand name. The filler features were again mentioned but unpredictable of quality, and no mention was made of the Redundant attribute (floor). This condition was structurally identical to the basic experimental conditions in the previous experiments.

The two-brand conditions differed from the one-brand conditions through the addition of a second brand at each level of quality. In the two/control condition, subjects also were exposed to 16 profiles of whitewater rafts. However, instead of all eight high-quality rafts carrying the Hypalon brand name, only four high-quality rafts carried the Hypalon brand name. Four other high-quality rafts carried the Leaf-field brand name. All eight rafts had a tubular floor. Similarly, only four low-quality rafts carried the Riken brand name; the four remaining low-quality rafts carried the Bering brand name. All eight of these rafts possessed an I-beam floor. Thus, subjects learned simultaneously about brand names and floor, and the number of exposures to each level of the Redundant attribute was kept constant. However, each level of the Redundant attribute was shared by two brand names. Thus, floor type provided a more parsimonious “explanation” of quality than the four brand names.

In the two/experimental condition, subjects received the same 16 phase 2 profiles as subjects in the two/control con-

dition but were preexposed to 16 profiles.⁷ Eight of these profiles were the same profiles as the eight phase 1 profiles in the one/experimental condition, that is, four high-quality rafts carrying the Hypalon brand name and four low-quality rafts carrying the Riken brand name. The eight remaining profiles were four high-quality rafts carrying the Leaffield brand name and four low-quality rafts carrying the Bering brand name.

Measure. The dependent measure consisted of the binary quality judgments used in the previous experiments. In the conditions with one brand name per attribute level, eight profiles were used that were constructed as a 2 × 2 × 2 factorial combination of brand (Hypalon vs. Riken), floor (tubular vs. I-beam), and country of origin (Sweden vs. Canada). In the conditions with two brand names per attribute level, subjects made the same eight binary quality judgments plus eight additional binary quality judgments created by a 2 × 2 × 2 factorial combination of the second pair of brands (Leaffield vs. Riken), floor, and country of origin.

Results and Discussion

As in the previous experiments, brand preexposure lowered the value subjects placed on the Redundant attribute and increased the value they placed on the Redundant name (asymptotic $t = -3.38, p < .001$ for the Redundant attribute; and asymptotic $t = 3.87, p < .001$ for brand name; see Table 4). In addition, however, results supported the hypothesis that subjects would place more value on a Redundant attribute and less on brand name when the attribute is a more parsimonious predictor of quality (asymptotic $t = 3.79, p < .001$ for the Redundant attribute; asymptotic $t = -2.35, p < .02$ for brand name). Interaction effects between the parsimony manipulation and the preexposure manipulation were not statistically significant (asymptotic $t = -1.82, p > .05$ for the Redundant attribute; asymptotic $t = -1.22, p > .22$ for brand name). Thus, the results suggest that a blocking effect occurs even when the Redundant attribute is the more parsimonious predictor, as when consumers experience multiple brands before receiving causal attribute information. Consumers are sensitive to parsimony, however.

The results reinforce the conclusion of experiment 2a that our data are inconsistent with an attentional explanation. Subjects exposed to predictive brand cues in phase 1 nonetheless placed a high weight on the Redundant cue encountered in phase 2. Clearly, subjects attended to the Redundant attribute even after some learning of predictive brand names.

More important, the results fail to refute a purely associative account. Nonetheless, the two-brand experimental

⁷Two additional conditions were run to check for differences between the experimental conditions in learning of the brand-quality relationships in phase 1. The additional conditions were similar to the two experimental conditions but omitted phase 2. Results for both conditions indicated imperfect but similar levels of learning of brand-quality relationships.

TABLE 4
MEAN FEATURE EFFECTS (EXPERIMENT 3)

Number of brands per attribute level/feature	Control	Experimental	Difference
One:			
Brand (P)	.30	.69	.39
Floor (R)	.63	.29	-.34
Two:			
Brand (P)	.04	.34	.30
Floor (R)	.93	.62	-.31

NOTE.—Brand effects are based on the two brand names (Hypalon and Riken) that were present in both the one and two brands per attribute level conditions. Brand effects including the two other brand names (Leaffield and Bering) in the two brands conditions were almost identical (.03 in the control condition and .35 in the experimental condition).

condition is the first instance in which we find a higher weight being placed on the Redundant cue than the Predictive cue—a result that is consistent with subjects’ preference for a parsimonious account of product performance. In the next experiment, we conduct a stronger test of the attractiveness of parsimony. Specifically, we examine the extent to which subjects rely on a parsimonious brand cue when retrospective processing is required to detect parsimony.

EXPERIMENT 4

The focal condition in the present experiment involved pairing a predictive brand cue with a predictive attribute cue in phase 1 (predictive extra attribute condition). In phase 2 the brand cue was paired with a different predictive attribute. At the end of both phases, the high-quality brand had been paired separately with two different predictive attributes (and the low-quality brand had been paired with the alternative levels of the same two attributes). Thus, parsimony favors the brand as the causal cue.

Method

Procedure and Stimuli. Fifty-one subjects participated in groups of one to six. Procedure and format of this experiment were similar to those in the previous experiments. In both of the two conditions in this experiment, subjects were exposed to 24 profiles of whitewater rafts, divided over two learning phases of 12 profiles each. Both conditions involved the inclusion of an extra attribute in phase 1. The extra attribute was either predictive or unpredictable of quality.

In the first learning phase of the unpredictable extra attribute condition, six high-quality rafts carried the Hypalon brand name, were made either with polyurethane fabric in Sweden or with neoprene fabric in Canada, and had a one-year warranty. Six low-quality rafts carried the Riken brand name, were also made either with polyurethane fabric in Sweden or with neoprene fabric in Canada, and also had a one-year warranty. Thus, as in earlier experiments, brand was used as the Predictive feature; warranty and country of

origin were used as constant and varying filler features, respectively. Unlike in earlier experiments, an extra attribute, fabric, was added that was unpredictable of quality, effectively making it another varying filler feature. In phase 2, six high-quality rafts carried the Hypalon brand name (Predictive feature), had tubular floors (Redundant feature), had a one-year warranty (constant filler feature), and were made either in Sweden or in Canada (varying filler feature). Six low-quality rafts carried the Riken brand name, had I-beam floors, and also had a one-year warranty and were made in Sweden or Canada. The extra attribute from phase 1 (fabric) was not mentioned in the second learning phase. In sum, the unpredictable extra attribute condition was structurally the same as the standard experimental conditions in the previous experiments.

The first learning phase in the predictive extra attribute condition was identical to the first learning phase in the unpredictable extra attribute condition, with the exception that the extra attribute was now predictive of quality. Thus, all six high-quality rafts carried the Hypalon brand name and were made of polyurethane fabric. All six low-quality rafts carried the Riken brand name and were made of neoprene fabric. The phase 2 profiles were identical to those in the unpredictable extra attribute condition.

Measures. The main dependent measure was the same binary quality judgments measure used in the previous experiments. In the binary quality judgments measure, no mention was made of the extra feature. Two additional profiles with a binary quality scale were included that had identical levels of brand and the two filler features but differed in their level of the extra feature.

Consistent with the results so far, subjects in the unpredictable extra attribute condition should place a high weight on the Predictive feature (brand) and very little weight on the Redundant feature (floor). In the predictive extra attribute condition, the brand cue should again receive a high weight if parsimony drives causal judgment. However, this parsimony is not evident until phase 2 and requires subjects to integrate across both phases. The relative importance of parsimony versus retrospective processing should be revealed by the relative weights placed on brand (Predictive feature) versus floor (Redundant feature) in the predictive extra attribute condition.

Results and Discussion

The unpredictable extra attribute condition exhibited the expected blocking pattern ($MFE_{\text{floor}} = .35$, $MFE_{\text{brand}} = .59$). In contrast, subjects in the predictive extra attribute condition placed more weight on a Redundant floor attribute ($MFE = .58$) than on Predictive brand names ($MFE = .25$). This was true even though (1) subjects were exposed to brand name information before being exposed to Redundant attribute information and (2) brand names constituted a more parsimonious explanation of quality than did attributes. In addition, analyses showed that initial exposure to a predictive extra attribute led subjects to place significantly more

value on the Redundant attribute and less value on brand names than in the regular blocking scenario in which the extra attribute was not predictive of quality (asymptotic $t = 2.76$, $p < .01$ for the Redundant attribute; asymptotic $t = -3.66$, $p < .001$ for brand names). Of course, results in the separate two-profile measure showed that subjects placed more value on the extra attribute when it was predictive ($MFE = .82$) than when it was not predictive ($MFE = .13$) of quality (asymptotic $t = 2.64$, $p < .01$).⁸

In sum, the results suggest that when parsimony requires retrospective processing, the equity-preserving blocking effect observed repeatedly in the previous experiments may not obtain. This surprising result raises two questions. The first question is whether the reversal from experiment 3 is due to the fact that brand name was the parsimonious cue in experiment 4 or to the fact that the nonparsimonious cues (the extra and Redundant attribute cues) were presented consecutively instead of simultaneously (which was the case for the brand names in the two/control condition of experiment 3). To address this question we ran a replication study in which all cues were attributes (chemical ingredients of eye drops). Results again showed a relatively high weight on the Redundant feature. Although this weight was slightly lower than the weight of the Predictive feature, the results tentatively suggest that the effect of parsimony is limited more by a lack of retrospective thought than by the nature of the cues. Second, experiment 4 begs the question of process. We address this issue below.

GENERAL DISCUSSION

The present series of experiments suggests that “blocking” of consumer learning is a robust phenomenon that has implications for brand equity. Strong blocking effects were found despite a limited number of brand preexposures and extensive exposure to predictive attribute information. In addition, blocking effects were found even when attributes were conceptually related to the benefit of interest, when subjects had no reason to believe that the preexposed products were any different from the ones encountered later, and when subjects experienced several brands sharing the same attribute. Although the conceptual nature of product features (brand vs. intrinsic attribute), a priori conceptual relatedness between features and outcomes, and explanatory parsimony did not significantly interact with the effect of brand preexposure, manipulations of these factors yielded strong main effects. Finally, adding a predictive attribute at initial exposure damaged brand equity and led to a dramatic reduction of blocking of another predictive attribute encountered later. The latter effect occurred despite the fact that brand name was a more parsimonious predictor of performance than were the attributes.

⁸Due to experimenter error, data for this measure were obtained from only 46 of 51 subjects.

Process

In terms of the psychological processes that might underlie the blocking phenomenon, our results provide evidence against explanations involving truncated processing (experiments 2a and 4) or idiosyncratic inference making (experiment 2b). More important, our results are inconsistent with a currently prominent explanation that relies on normative causal reasoning processes. A central premise of the causal reasoning explanation is that consumers retrospectively retrieve instances of relevant experiences and perform a causal analysis on those experiences. The data suggest, however, that subjects did very little retrospective processing in forming product evaluations.

The results are consistent, however, with theories that view learning as a forward-looking process. Our data suggest that evaluative learning is dominated by a satisficing process aimed at establishing accurate prediction of future consumption benefits. Such a process explains why little value accrues to information that does not directly improve prediction. The forward-looking property and the satisficing property fit at least two more specific types of process explanations: parallel associative updating and sequential hypothesis testing.

Parallel Associative Updating. According to associative or connectionist theories, humans and animals predict outcomes on the basis of associations with cues—with stronger associations implying higher evaluative weights. These theories hold that blocking and the other phenomena found in our experiments are the result of the particular way in which associations are updated. Associations between many product cues and consumption outcomes are updated simultaneously according to a learning rule that has a limited set of fundamental properties. The first of these properties is the property of “additivity,” which holds that the activation of an output node is an additive function of the association strengths of all activated input cues. In a consumer-learning context, this property holds that the predicted level of a consumption outcome (e.g., quality) is an additive function of the strengths of its associations with all presented product features. For example, if a brand and an attribute both have positive quality associations, the predicted quality of a product will be higher if it is known to possess both the brand name and the attribute than if it possesses just one of the two features. The second property is the property of “error reduction,” which holds that the amount and direction of the updating of an association’s strength depends directly on the discrepancy between the actual activation of an output node and the desired activation of an output node. Larger discrepancies lead to larger changes in association strength, and no updating takes place when the discrepancy is zero. The result is that product features will gather association strength only to the extent that predicted quality levels are different from actual quality levels. Thus, if consumers have learned to predict quality accurately on the basis of brand associations formed during initial product exposure, no further updating will take place. Attribute in-

formation provided after perfect prediction has been established will not lead to an increase in or a loss of brand association strength.

In addition to explaining the blocking phenomenon and the effect of brand preexposure on brand equity, the associative account can also explain the finding that the presence of attribute information harms brand equity when brands and attributes appear simultaneously. Note that, together, the properties of additivity and error reduction imply that the final sum of association strengths from copresent product features is limited to the level of output activation at which predicted and actual consumption outcomes are identical. Thus, simultaneously presented brand and attribute cues should develop associations with a consumption outcome (e.g., quality), but the summed strength of those associations will be limited to the same level as when only one predictive cue is present. Thus, a brand’s association with quality is lower when it is consistently presented with an attribute cue than when it is presented alone. Brands and attributes may be said to “compete” for equity.

By including a salience parameter, associative theories can also account for the finding that consumers tend to place more weight on attributes than on brands in simultaneous learning conditions. This salience parameter reflects the fact that some associations are learned faster and therefore absorb more of the limited available association strength than other cues. In the same vein, more “parsimonious” cues occur more often and therefore are updated more often than less parsimonious cues that are learned simultaneously. Thus, the more parsimonious cues should gain more of the available association strength. A similar, but somewhat attenuated process should occur when multiple brands fail to establish perfectly strong associations before exposure to more parsimonious attribute information (e.g., in the two/experimental condition of experiment 3). Finally, associative theories can account for the results in experiment 4, assuming that, during initial learning, the less salient brand and more salient attribute led the attribute information to out-compete the brand information. The result would be a weak association between brand and quality and a strong association between the attribute and quality. At the beginning of phase 2, the attribute information was replaced by a new predictive attribute. Output activation based on brand information alone would be low, leaving much opportunity for the new attribute to establish a strong association to the outcome.

Sequential Hypothesis Testing. Although parallel associative updating seems to fit our data surprisingly well, we propose another forward-looking process that might explain how consumers learn to value product features. The sequential hypothesis testing explanation holds that consumers sequentially test product cues as potential explanations of product performance. Moreover, hypothesis testing ceases once a sufficient explanation has been found. Such an explanation is consistent with related research showing that people do not conduct an exhaustive search for causal variables (Shaklee and Fischhoff 1982) and have

difficulty entertaining more than one hypothesis at a time (Mynatt, Doherty, and Dragan 1993; Van Wallendael and Hastie 1990).

The sequential hypothesis testing explanation explains the basic blocking phenomenon because initial exposure to brand information establishes a sufficient explanation of product performance before attribute information can be tested. This explanation also accounts for the “competition” between cues found in the control conditions of experiments 1–3. In those situations, sufficient brand and attribute explanations became available simultaneously. Assuming that consumers favor testing attribute explanations before brand explanations, most subjects should have tested the attribute explanation first. Because the attribute information provided a sufficient explanation of product performance, no testing was done of the brand explanation. A similar account can be applied to experiment 4, in which initial exposure was to both a predictive attribute and a predictive brand. In this case, subjects should test the attribute but not the brand explanation. When the attribute explanation is no longer viable (as in phase 2 of experiment 4), forward-looking hypothesis testers should begin testing of a new attribute explanation instead of the brand explanation. Thus, brands should receive little evaluative weight despite being the most parsimonious predictor of product performance across the sequence of experiences. The hypothesis testing account can also accommodate the parsimony effect in the two/experimental condition of experiment 3 if one assumes that the tendency to cease testing of additional explanations is reduced when the initial explanation is a multiple-sufficient explanation instead of a singular-sufficient explanation.

Future Research

Distinguishing between the associative and hypothesis testing processes is an important task for future research. In particular, future research should focus on the essential differences between the two theories in terms of sequential versus parallel processing. Once a distinction is made between the two types of theories, further research should address more detailed properties of the learning process. Several associative models share the essential properties of additivity and error reduction but differ in other important aspects (e.g., Gluck and Bower 1988; Kruschke 1996; Pearce 1994). Similarly, the sequential process may have several additional properties and boundary conditions. To generate and test predictions about such properties and conditions, it would be beneficial to develop more formal process models of evaluative learning involving sequential hypothesis testing.

Future research also should examine the generalizability and boundary conditions of the blocking phenomenon. In particular, research should assess the sensitivity of the blocking phenomenon to “disconfirming evidence” about the brand’s relationship to quality. Prior research on belief perseverance suggests that disconfirming evidence does not necessarily lead to a large change in attitude. However, disconfirming brand-quality experiences may reduce the

brand’s power to block attribute information from acquiring predictive value.

Future research on blocking could also help identify contexts in which forward- and backward-looking processes are more or less likely to play a role. For example, encouraging consumers to recall specific instances may lead to quality predictions that are less subject to blocking—as suggested by experiment 2A. In addition, forward-looking, predictive learning may be limited to characteristics that consumers are motivated to learn. For example, predictive learning may be less likely when consumers have less knowledge about the product category or when category involvement is lower. Consumers may also be more inclined to learn how to predict benefits than attributes. Thus, attribute learning may have to be recalled retrospectively from an episodic memory system that is not subject to blocking. The existence of an episodic memory system, functioning in parallel with the predictive memory system, also explains our finding in experiment 2a that subjects were able to recall co-occurrence frequencies between the Redundant attribute and quality despite blocking of the attribute information in their predictive quality judgments. Finally, consumers who are exposed only to brand information may at some point decide to learn about the causal impact of product attributes. Such consumers may then “restart” the learning process while ignoring brand names. Such a scenario would be tantamount to ignoring phase 1 at the time of phase 2. Presumably, humans can engage in hypothesis testing that is unconstrained by prior associative learning.

More pragmatically, future research should examine other instantiations of the blocking phenomenon. Brand equity is a convenient but important context in which to study consumer learning. As experiment 1 demonstrated, however, the effect applies more broadly.

Conclusion

This research supplements our understanding of why consumers value brand cues at the expense of attribute cues. Our explanation is based on learning processes not previously investigated in consumer research. Stated simply, brand cues may “block” the learning of quality-determining attribute cues. Our findings also supplement research on the blocking phenomenon by demonstrating its robustness in the face of prior beliefs and the presence of more parsimonious predictors of quality. Moreover, we cast doubt on explanations involving causal reasoning.

More generally, this research suggests an alternative approach to understanding product evaluation and choice. At one level, we show that in order to understand the value consumers place on product cues it is helpful to understand how consumers learn. Our preliminary conclusion is that learning is both forward looking and competitive in the sense that cues compete against each other for predictive value. We find little support for the notion that consumers engage in causal analysis of retrospectively considered experiences. At a broader level our evidence in favor of associative learning processes suggests that there are alter-

native views of consumer behavior that do not invoke higher level cognitive activity. This perspective has not been prominent in consumer research but seems to be reemerging as a credible approach to understanding human behavior (see Alba, in press; Janiszewski and van Osselaer, in press).

APPENDIX

DESIGN SUMMARIES

TABLE A1

EXPERIMENT 1

Type of Predictive feature/experimental vs. control	Phase in experiment		
	Learning 1	Learning 2	Test
Brand:			
Experimental	$P_{br1}+$	$P_{br1}R_{at1}+$	$P_{br}?$
Control	$P_{br2}-$	$P_{br2}R_{at2}-$	$R_{at}?$
		$P_{br1}R_{at1}+$	$P_{br}?$
		$P_{br2}R_{at2}-$	$R_{at}?$
Attribute:			
Experimental	$P_{at1}+$	$P_{at1}R_{at1}+$	$P_{at}?$
Control	$P_{at2}-$	$P_{at2}R_{at2}-$	$R_{at}?$
		$P_{at1}R_{at1}+$	$P_{at}?$
		$P_{at2}R_{at2}-$	$R_{at}?$

NOTE.—P = Predictive feature, R = Redundant feature, br = brand, at = attribute, + = high quality, and - = low quality. Subscripts 1 and 2 refer to different levels of a feature.

TABLE A2

EXPERIMENT 2

A priori relatedness/experimental vs. control	Phase in experiment		
	Learning 1	Learning 2	Test
Low:			
Experimental	$P_{br1}+_{ip}$	$P_{br1}R_{at1}+_{ip}$	$P_{br}?$
Control	$P_{br2}-_{ip}$	$P_{br2}R_{at2}-_{ip}$	$R_{at}?$
		$P_{br1}R_{at1}+_{ip}$	$P_{br}?$
		$P_{br2}R_{at2}-_{ip}$	$R_{at}?$
High:			
Experimental	$P_{br1}+_{hp}$	$P_{br1}R_{at1}+_{hp}$	$P_{br}?$
Control	$P_{br2}-_{hp}$	$P_{br2}R_{at2}-_{hp}$	$R_{at}?$
		$P_{br1}R_{at1}+_{hp}$	$P_{br}?$
		$P_{br2}R_{at2}-_{hp}$	$R_{at}?$

NOTE.—+_{ip} = positive outcome, unrelated to the Redundant attribute; -_{ip} = negative outcome, unrelated to the Redundant attribute; +_{hp} = positive outcome, related to the Redundant attribute; and -_{hp} = negative outcome, related to the Redundant attribute. Half the subjects received 12 trials and the other half of the subjects received 24 trials in learning phase 2.

TABLE A3

EXPERIMENT 2B

Experimental vs. control	Phase in experiment		
	Learning 1	Learning 2	Test
Experimental	$P_{br1}R_{at1}+$	$P_{br1}R_{at1}+$	$P_{br}?$
Control	$P_{br2}R_{at2}-$	$P_{br2}R_{at2}-$	$R_{at}?$
		$P_{br1}R_{at1}+$	$P_{br}?$
		$P_{br2}R_{at2}-$	$R_{at}?$

NOTE.—Values of the Redundant feature (R) in learning phase 1 were manipulated via instructions in the test phase.

TABLE A4

EXPERIMENT 3

Number of brands per attribute level	Phase in experiment		
	Learning 1	Learning 2	Test
One:			
Experimental	$P_{br1}+$	$P_{br1}R_{at1}+$	$P_{br}?$
Control	$P_{br2}-$	$P_{br2}R_{at2}-$	$R_{at}?$
		$P_{br1}R_{at1}+$	$P_{br}?$
		$P_{br2}R_{at2}-$	$R_{at}?$
Two:			
Experimental	$P_{br1}+$	$P_{br1}R_{at1}+$	$P_{br}?$
Control	$P_{br2}-$	$P_{br2}R_{at2}-$	$R_{at}?$
	$P_{br3}+$	$P_{br3}R_{at1}+$	
	$P_{br4}-$	$P_{br4}R_{at2}-$	
		$P_{br1}R_{at1}+$	$P_{br}?$
		$P_{br2}R_{at2}-$	$R_{at}?$
		$P_{br3}R_{at1}+$	
		$P_{br4}R_{at2}-$	

TABLE A5

EXPERIMENT 4

Type of extra phase 1 attribute (X)	Phase in experiment		
	Learning 1	Learning 2	Test
Unpredictive	$P_{br1}X_1+$	$P_{br1}R_{at1}+$	P?
	$P_{br2}X_1-$	$P_{br2}R_{at2}-$	R?
	$P_{br1}X_2+$		
	$P_{br2}X_2-$		
Predictive	$P_{br1}X_1+$	$P_{br1}R_{at1}+$	P?
	$P_{br2}X_2-$	$P_{br2}R_{at2}-$	R?

NOTE.—"X" refers to the "extra attribute."

[Received December 1998. Revised December 1999. Robert E. Burnkrant served as editor, and Barbara Loken served as associate editor for this article.]

REFERENCES

Alba, Joseph W. (in press), "Dimensions of Consumer Expertise . . . or Lack Thereof," in *Advances in Consumer Research*, Vol. 27, ed. Stephen J. Hoch and Robert J. Meyer, Provo, UT: Association for Consumer Research.

Bergen, Mark, Shantanu Dutta, and Steven M. Shugan (1996),

- "Branded Variants: A Retail Perspective," *Journal of Marketing Research*, 33 (February), 9–19.
- Broniarczyk, Susan M. and Joseph W. Alba (1994), "Theory versus Data in Prediction and Correlation Tasks," *Organizational Behavior and Human Decision Processes*, 57 (January), 117–139.
- Cheng, Patricia W. (1997), "From Covariation to Causation: A Causal Power Theory," *Psychological Review*, 104 (April), 367–405.
- and Laura R. Novick (1992), "Covariation in Natural Causal Induction," *Psychological Review*, 99 (April), 365–382.
- Dickinson, Anthony, David Shanks, and John Evenden (1984), "Judgement of Act-Outcome Contingency: The Role of Selective Attribution," *Quarterly Journal of Experimental Psychology*, 36A (February), 29–50.
- Gluck, Mark A. and Gordon H. Bower (1988), "From Conditioning to Category Learning: An Adaptive Network Model," *Journal of Experimental Psychology: General*, 117 (September), 227–247.
- Hoch, Stephen J. and John Deighton (1989), "Managing What Consumers Learn from Experience," *Journal of Marketing*, 53 (April), 1–20.
- Hutchinson, J. Wesley and Joseph W. Alba (1991), "Ignoring Irrelevant Information: Situational Determinants of Consumer Learning," *Journal of Consumer Research*, 18 (December), 325–345.
- Janiszewski, Chris and Stijn M. J. van Osselaer (in press), "A Connectionist Model of Brand-Quality Associations," *Journal of Marketing Research*.
- Kamin, Leon J. (1969), "Predictability, Surprise, Attention, and Conditioning," in *Punishment and Aversive Behavior*, ed. Byron A. Campbell and Russell M. Church, New York: Appleton-Century-Crofts, 279–296.
- Kruschke, John K. (1996), "Base Rates in Category Learning," *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 22 (January), 3–26.
- Meyer, Robert J. (1987), "The Learning of Multiattribute Judgment Policies," *Journal of Consumer Research*, 14 (September), 155–173.
- Mynatt, Clifford R., Michael E. Doherty, and William Dragan (1993), "Information Relevance, Working Memory, and the Consideration of Alternatives," *Quarterly Journal of Experimental Psychology*, 46A (November), 759–778.
- Pearce, John M. (1994), "Similarity and Discrimination: A Selective Review and a Connectionist Model," *Psychological Review*, 101 (October), 587–607.
- Price, Paul C. and J. Frank Yates (1995), "Associative and Rule-Based Accounts of Cue Interaction in Contingency Judgment," *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 21 (November), 1639–1655.
- Rescorla, Robert A. and Allan R. Wagner (1972), "A Theory of Pavlovian Conditioning: Variations in the Effectiveness of Reinforcement and Nonreinforcement," in *Classical Conditioning II: Current Research and Theory*, ed. Abraham H. Black and William F. Prokasy, New York: Appleton-Century-Crofts, 64–99.
- Schneider, W. (1988), "Micro Experimental Laboratory: An Integrated System for IBM PC Compatibles," *Behavior Research Methods, Instruments, and Computers*, 20 (April), 206–217.
- Shaklee, Harriet and Baruch Fischhoff (1982), "Strategies of Information Search in Causal Analysis," *Memory and Cognition*, 10 (November), 520–530.
- Shanks, David R. (1994), "Human Associative Learning," in *Animal Learning and Cognition*, ed. N. J. Mackintosh, San Diego, CA: Academic Press, 335–374.
- , Douglas L. Medin, and Keith J. Holyoak, eds. (1996), *The Psychology of Learning and Motivation*, Vol. 34, *Causal Learning*, San Diego, CA: Academic Press.
- Van Wallendaal, Lori R. and Reid Hastie (1990), "Tracing the Footsteps of Sherlock Holmes: Cognitive Representations of Hypothesis Testing," *Memory and Cognition*, 18 (May), 240–250.
- Waldmann, Michael R. and Keith J. Holyoak (1992), "Predictive and Diagnostic Learning within Causal Models: Asymmetries in Cue Competition," *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 121 (June), 222–236.
- Young, Michael E. (1995), "On the Origin of Personal Causal Theories," *Psychonomic Bulletin and Review*, 2 (March), 83–104.