

Preference Construction and Persistence in Digital Marketplaces: The Role of Electronic Recommendation Agents

Gerald Häubl and Kyle B. Murray

*School of Business
University of Alberta*

This article examines the role of electronic recommendation agents in connection with consumers' construction of multiattribute preferences. We propose that such digital agents have the potential to influence consumers' preferences in a systematic fashion. Our key hypothesis is that, everything else being equal, the inclusion of an attribute in a recommendation agent renders this attribute more prominent in consumers' purchase decisions. The results of a controlled agent-assisted shopping experiment provide strong support for this hypothesis. We also demonstrate that this preference-construction effect may persist beyond the initial shopping experience and into subsequent choice settings in which no recommendation agent is available. Finally, we propose three possible explanations of the effect and discuss each of these in light of our results.

Against the background of the rapid growth of business-to-consumer electronic commerce (e.g., Ernst & Young, 2001; Forrester, 2001), it becomes increasingly important to develop an understanding of how consumers process product information and make purchase decisions in digital marketplaces, that is, in networked computer-mediated information environments for online shopping. An important property of such digital marketplaces is that they allow for a high degree of information personalization and for the implementation of sophisticated decision aids designed to assist customers in making their purchase decisions (e.g., Alba et al., 1997; Häubl & Trifts, 2000). Although these characteristics of electronic information environments may be highly beneficial to consumers in some respects, they also provide marketers with powerful, yet subtle, new opportunities to influence consumers' preferences and, ultimately, consumers' purchase decisions.

Consumers often do not have well-defined pre-existing preferences that are merely revealed when they make choices among available products or services. Instead, they tend to construct their preferences on the spot when the external environment prompts them to make a decision (Payne, Bettman, & Johnson, 1992). As a result, preferences are sensitive to the particular way in which a decision

problem is framed and to the format in which pertinent information is presented (Slovic, 1995). This is consistent with the general notion that human behavior is determined by the interaction between properties of the individual's information-processing system and properties of the task environment (Lynch & Srull, 1982; Simon, 1990). Because an increasing proportion of consumers' purchase decisions are made in digital marketplaces and given the central role that the properties of the decision environment play in consumers' construction of preference (Bettman, Luce, & Payne, 1998), it is important to examine the potential of key characteristics of such electronic information environments to influence consumers' preferences in a systematic fashion.

In this article, we focus on the role of an important and increasingly prevalent feature of digital shopping environments, a recommendation agent, in connection with the construction of preferences by decision makers. We conceptualize an electronic recommendation agent as a software tool that (a) attempts to understand a human decision maker's multiattribute preference with respect to a particular domain or product category based on a learning (or "calibration") phase during which the human reveals subjective preference information to the agent and (b) makes recommendations in the form of a sorted list of alternatives to the human based on its understanding of that individual's preference structure (see also Häubl & Trifts, 2000).

Consistent with the notion that preferences are constructive (Bettman et al., 1998), we propose that the characteristics of a recommendation agent may systematically influence decision makers' multiattribute preferences. Because real-world recommendation agents are almost inevitably selective in the sense that they consider only a subset of the pertinent product attributes, the particular set of attributes that is included in an agent—that is, used during its calibration phase and considered by its sorting algorithm—is a key characteristic of a recommendation agent and of the digital shopping environment that it is embedded in. We hypothesize that whether or not a particular attribute is included will affect the subjective importance of that attribute to the decision maker. More specifically, we predict an inclusion effect, such that an attribute will be rendered more important in preferential choice merely as a result of its inclusion in a recommendation agent. The existence of such a preference-construction effect is examined in an experiment. Furthermore, we investigate the possibility that this inclusion effect may persist beyond the agent-assisted shopping experience and carry forward into subsequent preferential choice scenarios in which no recommendation agent is available.

The remainder of the article is organized as follows. First, we discuss the role of recommendation agents in connection with electronic shopping environments. We then briefly review the notion of constructive consumer preferences and discuss the potential of electronic recommendation agents to influence consumers' preferences in digital marketplaces. This is followed by a discussion of the method and results of a laboratory experiment aimed at enhancing our understanding of how attribute inclusion in a recommendation agent may influence preferences. Our results provide support for the existence of the predicted inclusion effect. In addition, we find that the constructed preferences persist beyond agent-assisted shopping experiences. However, the evidence indicates that the inclusion effect is not ubiquitous—it is moderated by the interattribute correlation structure of the marketplace and the perceived rationale for the selective inclusion of attributes. The article concludes with a discussion of the potential mechanisms underlying the inclusion effect, and a general discussion of how the findings of this research enhance our understanding of consumer decision making in digital marketplaces.

RECOMMENDATION AGENTS FOR ELECTRONIC SHOPPING

Unlike in traditional retail environments, the constraints of physical space no longer dictate the organization of information in digital marketplaces. One consequence of this is that online vendors are able to offer a very large number of products due to their virtually infinite "shelf space," that is, the lack of physical constraints with respect to product display. Combined with the fact that the cost of searching for product information across merchants is substantially lower in digital mar-

ketplaces than in the physical world (Bakos, 1997; Lynch & Ariely, 2000), this results in the availability of a potentially vast amount of information about market offerings to consumers.

Easy access to large amounts of product information is both a blessing and a curse. It is a blessing in the sense that more information may allow consumers to make better purchase decisions (e.g., to select products that better match their personal preferences) than they would otherwise. However, the curse of having access to vast amounts of information is that consumers, due to their limited cognitive capacity, may be unable to adequately process this information. The idea that human decision makers have limited resources for information processing—whether those limits are in memory, attention, motivation, or elsewhere—has deep roots in the literature of both marketing and psychology (e.g., Payne, Bettman, & Johnson, 1993; Shugan, 1980; Simon, 1955). In a digital shopping environment, consumers are less constrained by the availability of product information, yet they remain bounded by the cognitive limitations of human information processing.

A response to the problem of information overload in digital marketplaces is the emergence of electronic decision aids for consumers. The latter represent a technology that takes advantage of an important and unique characteristic of digital shopping interfaces—the potential for real-time personalization of an information environment based on explicit input by, or other information about, a consumer (see Häubl & Trifts, 2000). Software tools that generate personalized product recommendations in the form of a list in which alternatives are sorted by their predicted attractiveness to an individual shopper, thus allowing the latter to screen a large set of alternatives in a systematic and efficient manner, are particularly valuable to consumers and have become highly prevalent in real-world digital marketplaces (e.g., Amazon.com, AOL's shopping site, Microsoft's MSN eShop, and the Yahoo! shopping site). We refer to electronic decision aids of this type as recommendation agents. Our focus is on recommendation agents that attempt to understand a consumer's preference in terms of a multiattribute preference model (based, e.g., on a weighted additive evaluation rule) that is calibrated using subjective preference information revealed to the agent by the consumer.¹ Such attribute-based recommendation agents are an integral component of many online shopping sites, and this technology can be licensed by vendors for inclusion in their digital stores (e.g., the virtual sales assistants developed by Active Decisions, Inc.)

Almost inevitably, real-world attribute-based recommendation agents are selective in the sense that only a subset of all the relevant product attributes can be used in their calibration

¹Whereas the majority of real-world recommendation agents for online shopping are attribute based, it is worth noting that another important class of recommendation systems uses so-called community-based approaches that do not involve a multi-attribute model of a consumer's preference (e.g., Ansari, Essegai, & Kohli, 2000).

and, thus, in the algorithm used to generate the recommendations. This is apparent in the implementation of many commercial recommendation systems for online shopping (e.g., Active Buyer's Guide and Yahoo.com's shopping guides.). The reasons for such selectivity in recommendation agents include (a) the very large number of attributes that exist in many product categories; (b) the substantial amount of data about, or interaction with, a consumer that would be required to develop an accurate understanding of the consumer's subjective preference in a high-dimensional attribute space; (c) an inclination to use only those attributes that are common to most or all available products; and (d) a tendency to include only attributes that are quantitative in nature (i.e., the levels of which can be represented numerically). Apart from the previous reasons, the selective inclusion of attributes in a recommendation agent may also be driven by strategic objectives (e.g., to de-emphasize specific attributes) on the part of whoever controls the design of the agent.

A recommendation agent may be made available, either by a particular online vendor (e.g., Sony's online product recommendation tool or FirstBank's online mortgage advisor), to assist shoppers in choosing one of the products in its own assortment, or by a third-party provider (e.g., Active Buyer's Guide or ConsumerReports.org), to help consumers in selecting a product from among different vendors. This distinction might be associated with different motivations for including certain attributes in the decision aid. This work pertains equally to both provider scenarios (vendor and third-party provider), as long as the recommendation agent is selective in terms of attribute inclusion. Throughout the remainder of this article, we refer to attribute-based recommendation tools in general, regardless of who might have control over the specific aspects of their design.

Recent empirical research shows that the availability of an attribute-based recommendation agent in an electronic shopping environment may result in a substantial reduction in the amount of consumers' prepurchase information search (Häubl & Trifts, 2000). This finding suggests that, due to the limited information-processing capacity of the human mind, consumers tend to rely heavily to an electronic agent's product recommendations to reduce the amount of effort required to make a purchase decision. Given this tendency to rely on suggestions made by recommendation agents, and given the rapidly increasing prevalence of such decision aids in digital marketplaces, it is critically important to develop an understanding of whether and how electronic recommendation agents may influence consumers' preferences.

THE ROLE OF ELECTRONIC RECOMMENDATION AGENTS IN PREFERENCE CONSTRUCTION

The information-processing approach to human decision making recognizes that individuals' information-processing

capacity is limited (e.g., Bettman, 1979) and that most decisions are consistent with the notion of bounded rationality in that decision makers seek to attain some satisfactory, although not necessarily maximal, level of achievement (Simon, 1955). As a result of these constraints, individuals typically do not have well-defined preferences that are stable over time and invariant to the context in which decisions are made (Bettman et al., 1998). That is, in a domain (e.g., product category) involving alternatives that are characterized in terms of multiple attributes, individuals typically do not have specific pre-formed strategies with respect to exactly how important each of several attributes is to them personally, what kind of integration rule they should use to combine different pieces of attribute information into overall assessments of alternatives, or precisely how they wish to make trade-offs between attributes. Instead, decision makers tend to construct their preferences on the spot when they are prompted either to express an evaluative judgment or to make a decision (Payne et al., 1992).

The constructive-preferences perspective adheres to two major tenets: (a) that expressions of preference are generally constructed at the time at which the valuation of an object is required and (b) that this construction process will be shaped by the interaction between the properties of the human information-processing system and the properties of the decision task (Payne, Bettman, & Schkade, 1999). In a similar vein, Slovic (1995) noted that preferences appear to be remarkably labile, sensitive to the way a choice problem is described or framed, and to the mode of response used to express the preference.²

Given the large amount of empirical evidence suggesting that the particular characteristics of the decision environment may play a central role in individuals' construction of preference (see, e.g., Slovic, 1995), the potential of electronic shopping environments, which are interactive (rather than static) and personalizable (rather than standardized), to influence consumer preferences and, ultimately, purchase decisions is very significant. In particular, because consumers tend to be quite willing to rely on product recommendations made to them by electronic agents (Häubl & Trifts, 2000), the latter may be an important determinant of how consumers construct their preferences in digital marketplaces.

In view of the fact that real-world recommendation tools for online shopping almost inevitably base their product suggestions on only a subset of the pertinent attributes (see previously), we propose that such selective recommendation agents may influence consumers' preferences in a systematic fashion. More specifically, we predict that the relative importance weight that shoppers attach to different attributes may be influenced by whether a particular attribute is included in the calibration and sorting algorithm of such an electronic agent. Our key hypothesis is that the mere inclusion of an attribute in a

²For a detailed discussion of constructive consumer choice processes, the reader is directed to Bettman et al. (1998).

recommendation agent will, everything else being equal, render this attribute more important when consumers make product choices in digital marketplaces.

Such an inclusion effect could be due to several mechanisms. First, it might be a direct consequence of the format of information presentation (see Jarvenpaa, 1989, 1990; Slovic, 1972). The inclusion of an attribute in a selective recommendation agent enhances that attribute's processability, which may, in turn, lead to an increase in the relative weight that consumers attach to the included attributes when making a purchase decision. Second, this type of preference-construction effect could be the result of feature-based priming, such that the inclusion of an attribute may cause the latter to become temporarily more salient in the consumer's mind (see Herr, 1989; Srull & Wyer, 1979). Finally, the inclusion effect may be the result of inferences about relative attribute importance that consumers make based on the notion that the inclusion of a product attribute in a recommendation agent conveys some information about its relevance or importance (see Grice, 1975; Levinson, 1983). These three possible explanations of the proposed preference-construction effect are discussed in greater detail when we interpret the results of the experiment (see the following).

We now turn to a discussion of the method and results of a laboratory experiment aimed at enhancing our understanding of the role of selective recommendation agents in connection with consumers' construction of preference in digital marketplaces. The primary objectives of this study were (a) to test for the existence of the proposed inclusion effect under a number of systematically varied circumstances and (b) to examine the effect's persistence over time and into digital environments in which no recommendation agent is available. In addition, the results provide some evidence with respect to the different mechanisms that might underlie such a preference-construction effect.

EXPERIMENT

General Method

The experiment consisted of three primary tasks. After a brief overview, a detailed description of the method is presented separately for each of the tasks, with the results that pertain to a particular task following immediately after the discussion of the relevant method.

The study, which was fully computer based, was conducted in a research laboratory equipped with state-of-the-art networked personal computers. All stimuli were embedded in a dynamic Web environment, which was accessed via a standard Web browser. Participants entered all of their responses via this Web interface. In addition, their interaction behavior with the experimental environment was recorded. Data were collected in sessions of 10 to

15 participants, with an experimenter present throughout. The study was completed by a total of 347 respondents.

Participants were informed that the overall purpose of the study was to test a new electronic shopping environment. The first task was a simulated shopping trip in an Internet-based electronic store that was equipped with an attribute-based recommendation agent. After finishing their shopping trip, and following a set of filler tasks, participants completed an extensive online questionnaire that contained, among other measures, two series of preferential choice questions, which we refer to as the first and second choice task. Once all participants in a session had completed the study, the group was dismissed. The duration of the experimental sessions was approximately 40 min. About 1 week after the experiment, a written debriefing was sent to participants by e-mail.

AGENT-ASSISTED SHOPPING TASK

Method

The first task was designed to test our prediction that the inclusion of an attribute in a recommendation agent will, everything else being equal, render this attribute more important when consumers make product choices during an agent-assisted shopping trip. The participants' task was to shop for a backpacking tent in an electronic store that was equipped with a selective attribute-based recommendation agent and to complete this simulated shopping trip by selecting from the set of available tents the subjectively most attractive one.³ Sixteen backpacking tents were available during this task. Each of the tents was identified by a fictitious model name and described in terms of four attributes (durability, weight, fly fabric, and warranty).⁴

In terms of the experimental procedure, participants first read instructions relating to the task, including an explanation of the recommendation agent's purpose and functionality. To calibrate the agent, participants were asked to indicate how important they personally considered each of the included attributes to be on a 100-point rating scale (see Figure B1 in Appendix B). Based on these subjective attribute-importance weights, and using (standardized) scale values for the different levels of each attribute, the recommendation agent computed a linear-weighted overall utility score for each available product.⁵ It then returned a personalized list of recommended products in which the alterna-

³For about half of the sample (182 of the 347 participants), a lottery incentive was used to make the shopping task more involving. Participants in the lottery condition were informed that two winners, selected at random from the group of study participants, were to receive the product they "purchased" during this shopping trip plus the difference between \$500 and the price of that product in cash. Because the products used in the study were constructed (see later), the two winners received real-world products that differed slightly from their "purchased" alternatives. The presence of the lottery incentive was counterbalanced with all experimental factors.

⁴Price was held constant at \$249 across the available tents.

tives were sorted by their utility score in descending order (see Figure B2 in Appendix B). For each product, this list contained the model name and the levels of the two attributes that were included in the agent. From the list of recommended products, participants were able to request the detailed description of a particular product (i.e., the levels of all four attributes) by clicking on a hyperlink. From each of the screens containing detailed product descriptions, participants were able to either return to the personalized list of recommended products for further search or proceed to complete their “purchase.”⁶

To allow for a test of the hypothesized preference-construction effect, the inclusion of attributes in the recommendation agent had to be manipulated systematically. For this purpose, the four attributes were divided into two subsets, Block 1 and Block 2, each containing two attributes. Each block contained one attribute that, according to a pretest, is of high importance (“primary attribute”) and one that is of only moderate importance (“secondary attribute”). Block 1 included durability (primary attribute) and fly fabric (secondary attribute), and Block 2 included weight (primary attribute) and warranty (secondary attribute). The inclusion of attributes was manipulated by using either Attribute Block 1 or Attribute Block 2 in the agent’s calibration interface and sorting algorithm. Through this counterbalancing, it was possible to manipulate attribute inclusion in the recommendation agent independently of the inherent characteristics of the actual attributes, such as their ecological importance.

The two primary and the two secondary attributes were varied at eight and two levels, respectively, across the 16 available tents. To allow for a clear and simple test of the predicted inclusion effect, the alternatives were constructed such that participants’ choice of their most-preferred product would be informative with respect to which attribute was the most important one to them in making their decision. Specifically, this was accomplished by combining (in products) the most attractive level of each primary attribute with a level of the other primary attribute that is not the most attractive one. Two alternatives had the best level of durability, and two others had the most attractive level of weight. There was no objectively dominating product, that is, no alternative had the most attractive level of both primary attributes, and, thus, participants had to rely relatively more heavily on one of the two primary attributes when making their choice. It was expected that the majority of participants would choose a product with the most attractive level of one of the two primary attributes. Which of the

two primary attributes the selected alternative was superior on was an indicator of the relative importance of these attributes to a respondent in making his or her decision.

In addition to investigating the existence of the predicted preference-construction effect, which requires the manipulation of attribute inclusion in the recommendation agent, we are also interested in examining the effect’s robustness to variations in two important factors, the structure of the marketplace and the perceived rationale for the inclusion of attributes in the recommendation agent, both of which were varied experimentally (see the following). These two factors will also help shed some light on the different mechanisms that may underlie the proposed inclusion effect.

We examine the possibility that the predicted inclusion effect might be moderated by an important property of the marketplace, the interattribute correlation across the set of available products. In a market that is characterized by negative interattribute correlations, a more attractive level of one attribute tends to be associated with a less attractive level of another attribute and, therefore, purchase decisions require trade-offs among attributes. In such an environment, which may be characterized as an “efficient” marketplace, the relative importance that is attached to different attributes tends to be highly consequential with respect to the decision outcome, and even very small differences in relative attribute importance may affect which of a set of alternatives is chosen. By contrast, in a market in which interattribute correlations are positive, an alternative that is favorable on one attribute tends to also be favorable on other attributes. In this type of scenario, which corresponds to a relatively “inefficient” marketplace, the relative importance that a consumer places on different attributes has much less of an impact on which product is chosen, and the potential cost of making a subjectively poor decision is significantly lower than in the case of negative inter-attribute correlations.⁷

For the purpose of the agent-assisted shopping task, the interattribute correlation was manipulated by constructing the set of products such that the correlation (in terms of utility) between the primary attribute that was included in the recommendation agent and the one that was not was either negative ($\rho = -0.71$) or positive ($\rho = +0.71$). Complete descriptions of the sets of available alternatives for the two treatments of inter-attribute correlation are provided in Appendix A (Tables A1 and A2).

In addition to the interattribute correlation structure in the market, the inclusion effect may also be sensitive to what consumers perceive to have been the reason(s) for the selection of the particular set of attributes that are included in an electronic recommendation agent. In particular, if shoppers believe that the included attributes were deliberately selected because of their high importance and/or relevance in connection with

⁵Each of these scores is an approximate measure of an alternative’s (unobservable) utility to the participant.

⁶On completion of the shopping task, participants performed a set of filler tasks. This was followed by a recall task in which participants were asked to state the levels of the four attributes for the product they had “purchased” during the agent-assisted shopping task.

⁷Of the two market scenarios considered here, the one with negative interattribute correlation more closely resembles typical real-world markets, which tend to be fairly competitive and efficient.

their task of selecting a product for purchase, they may have greater confidence in—and rely more heavily on—the electronic agent's recommendations and also engage in less pre-purchase information search than they would otherwise.

To examine its possible moderating role with respect to the predicted preference-construction effect, the perceived rationale for attribute inclusion was manipulated at three levels. Participants were informed that the attributes included in the recommendation agent had been selected (a) because of their high importance to a relevant group of consumers who participated in a recent study (*strong rationale*); (b) randomly from the set of pertinent attributes (*neutral rationale*); or (3) because, although they have been found to be of low importance to other consumers who participated in a recent study, they should be given some attention in the decision process (*weak rationale*). This manipulation was embedded (as one sentence) in the task instructions that participants were asked to read at the beginning of the experiment.

In sum, the experimental design was a 2 (attribute block 1 vs. attribute block 2 included in the recommendation agent) \times 2 (negative vs. positive interattribute correlation) \times 3 (weak vs. neutral vs. strong perceived rationale for attribute inclusion in the agent) between-subject full factorial. The 347 study participants were randomly assigned to one of 12 treatment conditions.

Results

As a first test of preference construction in the form of the predicted inclusion effect, we examine the relative choice shares in the shopping task for those alternatives that have the most attractive level of the primary included attribute, that is, the primary attribute that was used by the recommendation agent. The directional prediction implied by the inclusion effect is that alternatives that are superior on the included primary attribute are more likely to be chosen than ones that are superior on the excluded primary attribute. The corresponding null hypothesis is that the extent to which an attribute drives participants' product choices is independent of whether or not it was included in the agent, that is, that half of the participants select an alternative that had the most attractive level of the included primary attribute and the other half choose a product that had the most attractive level of the excluded primary attribute (when controlling for differences in the ecological importance of the actual attributes through counterbalancing).

Of the 347 participants who completed the agent-assisted shopping task, 308 chose a product that had the most attractive level of either of the primary attributes. The choices of these participants were used in the analysis.⁸ Overall, 60.7%

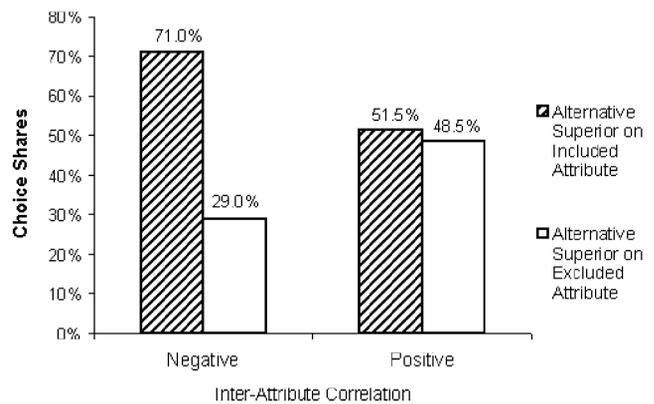


FIGURE 1 Inclusion effect: Moderating role of interattribute correlation.

of participants chose an alternative that had the most desirable level of the primary included attribute, and only 39.3% chose an alternative that had the most desirable level of the primary attribute that was not included in the recommendation agent. Based on a binomial test (using equal choice probabilities as the null hypothesis), these relative choice shares depart significantly from the base case of equal choice shares ($p < .0001$). This provides strong support for the predicted inclusion effect. The inclusion of an attribute in the recommendation agent rendered it more important in participants' product choices.⁹

This effect is, however, moderated by the interattribute correlation in the set of available products. A strong inclusion effect is observed in the conditions with negative interattribute correlation. When the included primary attribute was negatively correlated with the excluded primary attribute, 71% of participants purchased an alternative that had the most desirable level of the included primary attribute, and only 29% chose a product that was superior on the primary attribute that was not included in the recommendation agent. According to a binomial test, this departure from a 50–50 split in choice shares is significant ($p < .0001$). By contrast, when the interattribute correlations are positive, we find no evidence of such an effect. In this case, the choice shares for the two types of alternatives (superior on the included vs. on the excluded attribute) are not significantly different from each other in these conditions ($p > .75$). Figure 1 provides a graphical representation of the moderating effect of interattribute correlation with respect to the inclusion effect.¹⁰ The descriptive finding in terms of choice shares is consistent with the result of a logistic regression analysis, according to which the probability that the chosen product has the most at-

⁸For those participants who chose an alternative that did not have the most attractive level of either of the two primary attributes, no simple measure of relative attribute importance is available.

⁹Additional results from a separate experiment, which was identical to this one except that the primary attributes were uncorrelated and no rationale for attribute inclusion was provided, corroborate the evidence in favor of an inclusion effect. These results are available from Gerald Häubl on request.

tractive level of the included (and therefore not of the excluded) attribute is significantly greater under negative than under positive interattribute correlation (Wald = 12.02, $p < .001$).¹¹

Two measures were used to examine the amount of information search that participants engaged in during the shopping task: (a) the total amount of time spent searching and (b) the number of alternatives for which a detailed description, in terms of the complete set of attributes, was viewed. On average across the experimental conditions, participants spent a total of 76.51 sec examining detailed information about available products and did so for an average of 8.81 alternatives. Whereas the manipulation of interattribute correlation structure had no effect on the extent of information search, the manipulation of the perceived rationale for the inclusion of attributes in the recommendation agent did. Participants who were informed that the included attributes had been selected because of their high importance to a relevant group of consumers who had participated in a recent study (i.e., strong rationale) engaged in significantly less information search than did those who received one of the two other treatments of perceived rationale. This was the case in terms of both the amount of time spent searching ($p < .05$) and the number of alternatives for which a detailed description was viewed ($p < .05$).¹² The cell means of these two indicators for the three levels of perceived rationale for attribute inclusion are provided in Table 1.

First Paired Choice Task

Method. In addition to testing for the existence of the inclusion effect during the agent-assisted shopping task, we also examine the persistence of this effect into other choice situations in which no recommendation agent is available. The first paired choice task involved choosing the preferred alternative from each of six pairs of backpacking tents. The alternative that a participant had selected during the shopping task was included in each of these six choice sets. Complete attribute information was presented for the new alternative in each choice set, but only the model name was available for the previously purchased product (see Figure B3 in Appendix B for a sample choice set). Therefore, to make attribute-by-attribute comparisons, participants had to retrieve at-

TABLE 1
Effect of Perceived Rationale for Attribute Inclusion on Amount of Information Search

<i>Perceived Rationale for Attribute Inclusion in Recommendation Agent</i>	<i>Mean Amount of Time Spent Searching (Seconds)</i>	<i>Mean Number of Alternatives Searched</i>
Strong (Attributes were selected because they are important.)	62.71	7.75
Neutral (Attributes were selected arbitrarily.)	81.91	9.00
Weak (Attributes were selected because they are unimportant.)	83.60	9.64

tribute information about the purchased alternative from memory.¹³

The stimuli for this task were created dynamically. In particular, the descriptions of the new product were personalized for each participant such that the tent selected by this individual during the shopping task was used as the base case and the attribute levels of all other alternatives were constructed relative to the previously chosen product. Only the primary attributes were manipulated across alternatives, and the secondary attributes were held constant at the levels of the purchased product. For a given participant, three levels were used for each of the primary attributes, with the middle level always being the one of the product previously chosen by that individual. The more desirable and the less desirable level of an attribute were constructed by adding/subtracting a fixed increment to/from the middle level. This increment was 3 points (on a 100-point scale) for the durability rating and 0.1 kilograms for weight.¹⁴ Table 2 provides an overview of the dynamic design used in constructing the alternatives for this choice task.

In addition to choosing one of the two products in each choice set, participants also indicated the strength of their preference for the alternative that they chose. A 5-point rating scale ranging from 1 (*just barely prefer*) to 5 (*very strongly prefer*) was used for this purpose. The strength-of-preference measures complement the observed choices by providing more fine-grained information about an individual’s relative preference for a given pair of alternatives. We combined the observed choice response and the strength-of-preference rating into a 10-point graded-paired-comparison (GPC) re-

¹⁰It is worth pointing out that, of the two market scenarios considered here, we find evidence of the inclusion effect in the one (i.e., negative interattribute correlations) that more closely resembles typical real-world markets.

¹¹We also estimated a logistic regression model that, in addition to the interattribute correlation, included the rationale for attribute inclusion (coded as weak = -1, neutral = 0, and strong = +1) and an interattribute correlation × rationale interaction as predictors. Neither of these additional variables has a significant effect on the probability that the chosen product has the most attractive level of the included attribute.

¹²The differences between the weak-rationale and neutral-rationale conditions are not statistically significant ($p > 0.05$).

¹³In addition to allowing a test of the persistence of the inclusion effect, the first choice task also enables us to gain some insight into whether this type of preference-construction effect tends to carry forward in an attribute-based (i.e., through consumers’ relative importance weights for the relevant attributes) or alternative based (i.e., in the form of an overall preference specifically for the previously selected product) manner.

¹⁴Participants were not aware that these stimuli were created dynamically.

TABLE 2
Dynamic Design for First Paired Choice Task

Choice Set	Alternatives	Included Primary Attribute	Excluded Primary Attribute
1	Purchased product	i	e
	New alternative	i - 1	e + 1
2	Purchased product	i	e
	New alternative	i + 1	e - 1
3	Purchased product	i	e
	New alternative	i	e
4	Purchased product	i	e
	New alternative	i	e + 1
5	Purchased product	i	e
	New alternative	i + 1	e
6	Purchased product	i	e
	New alternative	i + 1	e + 1

Note. The levels of the included and excluded primary attributes of the product that a participant had selected during the shopping task are denoted “i” and “e,” respectively. Attribute levels that are less or more attractive than the corresponding attribute level of the purchased product are indicated by “-1” and “+1,” respectively.

sponse variable, which lends itself more readily to a quantitative representation of a person’s relative preference for the two alternatives in a choice set.¹⁵

To assess the relative importance, as revealed during this preference task, of the included and excluded primary attributes to participants, we ran individual-level OLS regression models with a participant’s GPC score for each of the six choice sets as the dependent variable and the difference between the alternatives in a pair in terms of each primary attribute as independent variables. The latter were coded such that that an estimated model coefficient for a given attribute indicates an individual’s predicted reaction, expressed as the number of points on the GPC scale, to an improvement in that attribute by one level (i.e., an increase in durability rating by 3 points or a reduction in weight by 0.1 kg).

Given our primary goal of examining the persistence of the inclusion effect into the first paired choice task, we focus on two individual-level coefficients, one for the primary attribute that was included in the recommendation agent during the shopping task (β_{in}) and one for the primary attribute that was excluded from the agent (β_{ex}), which were estimated along with a unique intercept for each participant. The coefficients β_{in} and β_{ex} indicate the overall importance that an individual attached to the (previously) included primary attribute and the excluded primary attribute, respectively, when completing the first paired choice task.¹⁶

Results

Given the memory-based character of this choice task, we start by examining participants’ recall of the attribute information for the backpacking tent that they had selected during the agent-assisted shopping task. The shares of participants who correctly recalled the attribute levels of the product that they had “purchased” are 93.9% for durability rating, 81.8% for fly fabric, 85.3% for weight, and 98.6% for warranty. Recall performance was largely invariant to whether an attribute had been included in the recommendation agent.¹⁷ In addition, 80.7% of participants correctly recalled both of the primary attributes, durability rating and weight, of the tent that they had purchased. Thus, the majority of participants were able to accurately recall attribute information for their selected product to make comparisons between it and alternative products during the first paired choice task.

We now turn to the relative importance that participants attached to the primary attributes during this preference task. In particular, we are interested in whether the preference-construction effect due to attribute inclusion in the recommendation agent persisted into a scenario in which no such agent was present. We focus on those participants ($n = 145$) who completed the earlier agent-assisted shopping task in connection with a set of products characterized by negative interattribute correlation, as these were the conditions under which the inclusion effect was observed during the shopping task. The results of the individual-level regression analyses provide evi-

¹⁵For a choice set consisting of alternatives A and B, the GPC variable was coded as a zero-centered 10-point scale ranging from -4.5 (*very strongly prefer A*) to +4.5 (*very strongly prefer B*). This coding preserves the association between the scale midpoint of 0 and a participant’s state of indifference between the two alternatives in a choice set.

¹⁶Note that we controlled for inherent differences in importance between the actual attributes by counterbalancing which of the two primary attributes was used as the included versus the excluded primary attribute.

¹⁷The shares of participants with accurate attribute recall, broken down by attribute inclusion in the recommendation agent, are as follows. Durability rating: 93.1% when included and 94.7% when excluded. Fly fabric: 80.5% when included and 83.1% when excluded. Weight: 89.6% when included and 81.0% when excluded. Warranty: 98.4% when included and 98.8% when excluded.

dence for the persistence of the inclusion effect beyond the agent-assisted shopping experience. The mean coefficient for the primary attribute that had been included in the recommendation agent (β_{in}) is 2.181 and the mean coefficient for the excluded primary attribute (β_{ex}) is 1.752. This mean difference between β_{in} and β_{ex} is highly significant ($p < .01$).¹⁸ Thus, the inclusion of an attribute in the recommendation agent during the earlier shopping task resulted in a greater importance of that attribute during the first paired choice task. This indicates a substantial degree of persistence of the preference-construction effect beyond the task in which the recommendation agent was available.¹⁹

Next, we examine whether the magnitude of the preference-construction effect that was carried forward to the first choice task is moderated by the perceived rationale for the inclusion of attributes in the recommendation agent. For this purpose, we use the difference between the individual-level coefficient for the included primary attribute and that for the excluded one, $\beta_{diff} = \beta_{in} - \beta_{ex}$, as an indicator of the magnitude of this effect.

The results of an analysis of variance show that the perceived rationale for the inclusion of attributes in the recommendation agent during the earlier shopping task had a very substantial impact on the extent to which the preference-construction effect persisted into the first choice task, $F(2, 142) = 6.719, p < .01$. The nature of this moderating effect is exactly as expected. The strongest inclusion effect was observed when the perceived rationale for attribute inclusion in the recommendation agent had been strong. In these conditions, the mean of β_{diff} is 0.817. In the neutral-rationale conditions, the mean difference is 0.562, and in the weak-rationale conditions it is -0.155 (see Figure 2). All pair wise differences in the magnitude of this effect among the three types of rationale are statistically significant with at least $p < .01$.²⁰ This provides strong evidence that the perceived rationale for the inclusion of attributes in the electronic agent during the earlier shopping task affected the extent to which the agent-induced preference-construction effect persisted into the first choice task.

Second Paired Choice Task

Method. The second choice task also involved a series of choices from pairs of backpacking tents. However, un-

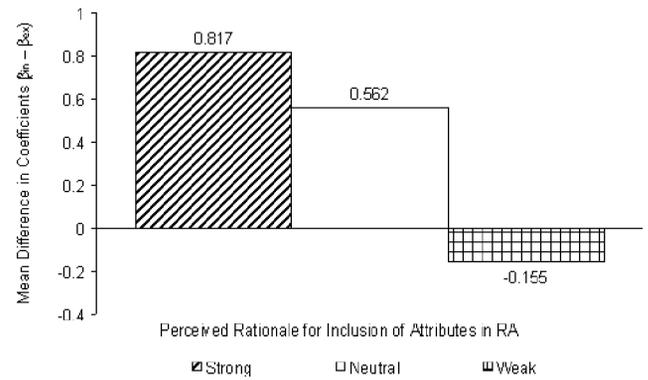


FIGURE 2 Persistent inclusion effect (first choice task): Moderating role of perceived rationale for attribute inclusion.

like the first choice task, participants were now informed that both backpacking tents in each choice set represented new alternatives, i.e., ones they had not encountered in any of the preceding tasks. Consequently, both alternatives in a choice set were described on all four attributes in this task. The tents in each pair were identified as Alternative A and Alternative B (see Figure B4 in Appendix B for a sample choice set). Because full attribute information was provided for both products in a set, participants were not required to retrieve any information from memory to make attribute-by-attribute comparisons. As a result, this task allows us to examine whether the inclusion effect may persist into a choice environment in which no recommendation agent is available, all alternatives are new to the decision maker, and all relevant attribute information is available for all products.

As in the first choice task, the stimuli were personalized using a dynamic choice design. The attribute levels of all alternatives were specified relative to those of the backpacking tent selected by an individual during the shopping task, and the actual attribute levels for the stimulus products were generated dynamically for each participant. Whereas the secondary attributes were held constant at the levels of the previously purchased product, the primary attributes were manipulated at three levels. The procedure used for selecting the three actual levels for each of the primary attributes was identical to the one used in connection with the first choice task (mentioned previously). A dynamic version of a basic utility-balanced efficient choice design (Huber & Zwerina, 1996) was used to construct the six choice sets for this task (see Table 3). Measures of extent of preference were also obtained along with each choice, using the same 5-point rating scale as in the first choice task. Finally, the observed choice response and the strength-of-preference rating were again combined into a 10-point GPC response variable representing an individual's relative preference for the two alternatives in a choice set.

¹⁸As expected given the results of the shopping task, there was no difference between β_{in} and β_{ex} for those participants who had completed the shopping task in one of the positive-interattribute-correlation conditions ($n = 163$).

¹⁹The mean intercept for these individual-level models is not significantly different from zero. Thus, participants' responses during this task do not reflect any bias in favor of the previously purchased product. This suggests that the way in which the preference construction effect carried forward was tied to the product attributes, rather than to a specific product.

²⁰The mean of β_{diff} in connection with the weak rationale is not significantly different from zero.

TABLE 3
Dynamic Design for Paired Choice Task

Choice Set	Alternatives	Included Primary Attribute	Excluded Primary Attribute
1	Alternative 1	$i - 1$	$e + 1$
	Alternative 2	i	e
2	Alternative 1	$i + 1$	e
	Alternative 2	$i - 1$	$e + 1$
3	Alternative 1	$i - 1$	$e + 1$
	Alternative 2	$i + 1$	$e - 1$
4	Alternative 1	$i + 1$	e
	Alternative 2	i	$e + 1$
5	Alternative 1	i	$e + 1$
	Alternative 2	$i + 1$	$e - 1$
6	Alternative 1	i	$e - 1$
	Alternative 2	$i - 1$	e

Note. The levels of the included and excluded primary attributes of the product that a participant had selected during the shopping task are denoted “ i ” and “ e ,” respectively. Attribute levels that are less or more attractive than the corresponding attribute level of the purchased product are indicated by “ -1 ” and “ $+1$,” respectively.

Results

The second choice task provides an opportunity to examine whether the preference-construction effect due to the inclusion of attributes in a recommendation agent during the shopping task persisted over a significant period of time and into a situation that involved choices among, new, completely described alternatives. Using the same approach as for the first choice task, individual-level coefficients for the included (β_{in}) and for the excluded primary attribute (β_{ex}) were estimated using a participant’s GPC score for each of the six sets in the second choice task as the dependent variable. These coefficients indicate the relative importance of a primary attribute as a function of whether it had been included in the recommendation agent during the shopping task while controlling for inherent differences in importance between the actual attributes used through attribute counterbalancing.

The individual-level coefficients for the second choice task provide evidence for the persistence of the preference-construction effect well beyond the agent-assisted shopping experience. In the conditions with negative interattribute correlations in the shopping task (i.e., the ones where the basic mere-inclusion effect was observed), the mean coefficient for the primary attribute that had been included in the recommendation agent (β_{in}) is 1.671 and that for the excluded primary attribute (β_{ex}) is 1.490. The difference between β_{in} and β_{ex} is in the direction predicted by our preference-construction effect and statistically significant ($p < 0.05$).²¹ Everything else being equal, the importance of an attribute during the second choice task was enhanced by having been included in the recommendation agent during an earlier shopping experience.²²

Once again, we examine whether the magnitude of the agent-induced preference-construction effect that persisted into the second choice task is moderated by the perceived rationale for the inclusion of attributes in the recommendation agent. As in the case of the first choice task, we use the difference between the individual-level coefficient for the included primary attribute and that for the excluded one, $\beta_{diff} = \beta_{in} - \beta_{ex}$, as an indicator of the magnitude of this effect. Due to our experimental design, this indicator is not contaminated by any inherent differences in the actual attributes.

As was the case in connection with the first choice task, the perceived rationale for the inclusion of attributes in the recommendation agent during the earlier shopping task had a substantial impact on the extent to which the preference-construction effect persisted into the this new setting, $F(2, 142) = 4.149, p < .05$. Once again, the nature of this moderating effect is exactly as expected (see Figure 3). The greatest persistence of the inclusion effect was observed when the rationale for attribute inclusion had been strong. The mean of β_{diff} for these conditions is 0.363. In the neutral-rationale conditions, it is 0.218. Both of these are significantly greater than zero ($p < .05$). In the weak-rationale scenarios, the mean of β_{diff} is not significantly different from zero. The magnitude of this effect is greater in both the strong-rationale and neutral-rationale conditions than in the weak-rationale treatments ($p < .05$). Finally, the difference between the strong-rationale

²¹As expected, no difference between β_{in} and β_{ex} was found for those participants who were in one of the positive-interattribute-correlation conditions for the initial shopping task.

²²Although our findings indicate that the inclusion effect persisted from the agent-assisted shopping task into situations where the recommendation agent is no longer present and new products are being evaluated, the magnitude of this preference-construction effect does appear to diminish between the first and the second choice task. To what extent this might be due to the different nature of the second choice task (completely stimulus-based) or to the greater amount of time that had passed because participants’ interaction with the recommendation agent cannot be determined on the basis of this experiment.

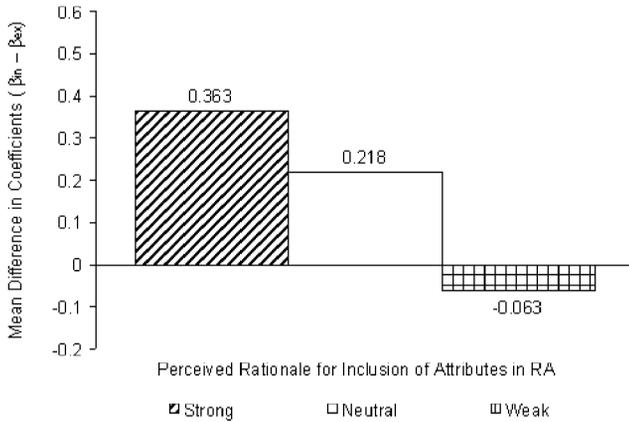


FIGURE 3 Persistent inclusion effect (second choice task): Moderating role of perceived rationale for attribute inclusion.

and neutral-rationale conditions is significant at $p < .1$ (but not at $p < .05$). In sum, these results suggest that the strength of the perceived rationale for the inclusion of attributes in the recommendation agent during the earlier shopping task had a positive effect on the degree to which the agent-induced preference-construction effect persisted into the second choice task.²³

Discussion

Taken together, the findings obtained from the three tasks provide strong evidence for the ability of electronic recommendation agents to influence consumers' preferences. Specifically, the results of this research suggest that the inclusion of an attribute in a recommendation agent will, everything else being equal, render this attribute more important when consumers make product choices in digital marketplaces. However, our findings also indicate that this preference-construction effect depends on an important property of the marketplace—the interattribute correlation structure across the set of available products. We find a strong inclusion effect in a scenario characterized by negative interattribute correlation (i.e., when the efficiency of the market is high), but not in the case of positive interattribute correlation (i.e., when the efficiency of the market is low). Moreover, the inclusion effect is moderated by the perceived rationale for the selective inclusion of attributes. Consumers who are provided with a strong rationale as to why certain attributes are included in the recommendation agent will rely more on the latter's recommendations and engage in less information search than consumers who are not. In addition, we find that this type of preference-con-

struction effect persists into subsequent choice tasks where no electronic decision aid is present and that the extent of such persistence is greater if a stronger rationale for the selective inclusion of attributes in a recommendation agent was provided during an earlier shopping trip.

The focus of this article has been on the potential for recommendation agents to influence consumers' preferences through the selection of attributes that are used during an agent's calibration phase and considered by its sorting algorithm, as well as on the persistence of such preference-construction effects over time and into subsequent decision scenarios in which no electronic agent is available. Although our experiment was designed to test these outcomes rather than the process by which they occur, the results do provide some insight into the mechanisms that may underlie the preference-construction effect that we have observed. In particular, we consider three potential explanations for the inclusion effect. We now discuss each of these in turn and then proceed to evaluating the extent to which each of the explanations appears to be consistent with the results of our study.

First, a substantial amount of effort has been devoted to studying the effects of the format of information presentation on how individuals acquire and process information. The results of this body of research suggest that the manner in which information is presented tends to have a substantial impact on human information processing and decision making (Bettman & Kakkar, 1977; Jarvenpaa, 1989, 1990; MacGregor & Slovic, 1986). The notion that the format of information presentation may influence decision making is formalized in Slovic's (1972) principle of concreteness. The latter suggests that human decision makers tend to (a) use only that information which is explicitly displayed in a stimulus environment and (b) process this information in the particular form in which it is presented. The information presentation format created by a recommendation agent—a list of products rank-ordered by their likely attractiveness to an individual in terms of the included attributes—should render the included attributes relatively more processable. In turn, this enhanced processability may lead to an increase in the relative weight that consumers attach to the included attributes when making a purchase decision. However, individuals do tend to depart from the particular type of processing that is encouraged by the format of information presentation if the cost of the potential inaccuracy (e.g., of making an inferior choice) that may result from such processing is significant (see Coupey, 1994, p. 97). Consumers will use the information in the format in which it is provided when doing so results in acceptable decision quality, but not when this would lead to a vastly inferior decision outcome.

A second possible explanation of this type of preference-construction effect is feature-based priming. The selective inclusion of attributes in a recommendation agent may act as a prime that causes the included product attrib-

²³It is worth noting that we found no evidence that familiarity, knowledge, or interest with respect to the product category had any moderating influence on the inclusion effect during the agent-assisted shopping task or on the effect's persistence into the two choice tasks.

utes to become temporarily more salient in the consumer's mind. This form of priming implies a spreading-activation mechanism that renders one or more cognitive concepts more easily accessible in memory (see, e.g., McNamara, 1994; Srull & Wyer, 1979;). In turn, this enhanced accessibility may lead to an increase in the likelihood that the primed concept (e.g., a category, decision rule, or feature) is used by an individual, as long as the concept is relevant to the cognitive task at hand and no other, competing concept is chronically more accessible (see Herr, 1989, p. 74). Such feature-based priming is evident in recent work by Mandel and Johnson (2002) that demonstrates the possibility of influencing individuals' preferences by merely altering the background of a Web page that participants view prior to completing a product choice task. In our experiment, the initial exposure to the included attributes during the calibration phase of the recommendation agent may have primed these attributes and therefore increased the likelihood that they would be relied on in the subsequent shopping task.

A third mechanism by which a recommendation agent may influence consumer preferences is based on the potential information value of attribute inclusion, that is, the notion that the inclusion of an attribute in the agent may be informative with respect to the importance of that attribute. Consumers might assume that the particular attributes that are used in the calibration of a recommendation agent have been included because they are relevant and important aspects of the alternatives in the product category of interest. In line with the theory of conversational logic, which suggests that exchanges of information are generally guided by a cooperative principle (Grice, 1975; Levinson, 1983), consumers are likely to believe that the electronic agent has been designed so as to be a meaningful tool that is capable of assisting them in their decision-making processes. In this study, the inclusion of attributes in the recommendation agent may have been perceived to convey information about other shoppers' preferences in terms of the set of attributes that are deemed important by a relevant population of consumers. Therefore, participants may have inferred that the attributes that were included in the recommendation agent are of high importance when buying a tent.

Although our experimental design was not intended to disentangle the three alternative explanations of the inclusion effect, our results do provide some insight as to how consistent each of these three mechanisms is with our results. The most important finding in this regard is the absence of an inclusion effect in connection with the marketplace characterized by positive interattribute correlations (i.e., the less efficient marketplace). This result is entirely consistent with the information-value explanation, but much less so with the other two explanations. According to the format-of-information- presentation

explanation, the inclusion effect should be stronger in an environment with positive inter-attribute correlations, where using the information in the format in which it is presented has limited costs in terms of decision accuracy (i.e., is unlikely to result in a significant loss of decision quality). When interattribute correlations are negative, the cost of using the information as it is presented may be high, because trade-offs between the primary attributes are required. Yet, it is precisely in the conditions with negative interattribute correlations where we observe a strong agent-induced preference-construction effect.

In feature-based priming (Mandel and Johnson, 2002; Yi, 1990), an individual attaches greater weight to the feature or attribute that was primed (i.e., made more salient in the consumer's mind) immediately prior to a choice or evaluation task. However, this form of priming cannot explain why we observe a strong inclusion effect in a marketplace where the cost of relying on the primed attribute in decision making may be high, whereas finding no evidence of such an effect in a setting where the cost of doing so is insignificant. A priming effect should not be moderated by the inter-attribute correlation structure in the marketplace, and even if it were, the effect should be stronger when these correlations are positive (i.e., lower cost associated with making a poor decision) than when they are negative (i.e., higher cost).

The finding that the preference-construction effect is stronger when interattribute correlations are negative is, however, consistent with the predictions based on the information-value-of-attribute-inclusion explanation. In these conditions, the marketplace is more efficient, and purchase decisions require trade-offs among attributes. To make decisions that involve trading off one attribute against another, an individual must consider the subjective importance of the different attributes to him or her. Moreover, in a marketplace characterized by negative interattribute correlations, the relative importance that is attached to different attributes tends to be highly consequential with respect to the decision outcome—even very small differences in relative attribute importance may affect which of a set of alternatives is chosen. Therefore, the information-value explanation suggests that the agent-induced preference-construction effect will be stronger when the inter-attribute correlations are negative than when they are positive. The choice shares that we observed in the shopping task are highly consistent with this explanation. In addition, the moderating effect of the perceived rationale for attribute inclusion on the extent to which the inclusion effect persists into subsequent choice tasks also supports the information-value explanation.

Whereas we cannot rule out any of the three possible explanations of the inclusion effect, the information-value-of-attribute inclusion explanation is the one that is most consistent with our empirical results. Nevertheless, given that this experiment was designed to investigate the

potential of an inclusion effect to occur, rather than the process(es) that might underlie such an effect, a more systematic examination of the latter remains an important objective for future research on this type of preference–construction effect.

CONCLUSION

Instead of having well-defined, stable preferences that are merely revealed when making purchase decisions, consumers tend to construct their preferences on the spot when they are prompted by the external environment to make a choice or express an evaluative judgment (Bettman et al., 1998). This preference–construction process is shaped by the interaction between the properties of the human information-processing system and those of the external decision-making environment. As a result of the growing importance of online shopping, the environments in which consumers make purchase decisions are increasingly electronic, rather than physical, in nature. Therefore, it is essential to develop a deeper understanding of how consumers construct their preferences in such digital marketplaces.

Although electronic shopping environments are not subject to the space constraints of physical stores, consumers who shop in such digital marketplaces are nevertheless subject to the familiar cognitive constraints in terms of their ability to process information. Electronic recommendation agents can play a key role in reducing the amount of information about available products that has to be processed by consumers, thereby assisting the latter in making better decisions with limited cognitive effort (Häubl & Trifts, 2000). However, for such a decision aid to be effective, the user must place some confidence in the agent's product recommendations, as well as in the process by which these recommendations are generated. This required level of confidence, or trust, raises the potential for an electronic agent to not only assist the consumer in the decision-making process given his or her subjective preference, but also to influence this preference.

It is apparent that real-world recommendation systems for online shopping are almost inevitably selective in the sense that they include, or consider, only a subset of the product attributes that are pertinent in a given product category. Decision aids such as Active Buyer's Guide may be selective to simplify and facilitate the recommendation process, whereas tools such as Sony's online product recommender may be selective for strategic reasons. In either case, the selective inclusion of attributes in a recommendation agent is an important aspect of an electronic shopping environment and may, therefore, play a key role in consumers' construction of preference in a digital marketplace. The objective of this article has been to investigate the potential of recommendation agents to systematically influence consumer preferences.

The key hypothesis proposed in this article is that, everything else being equal, the inclusion of an attribute in a selec-

tive recommendation agent renders this attribute more prominent in consumers' purchase decisions in an electronic shopping environment. Our research demonstrates that the preferences of human decision makers can in fact be influenced in a systematic and predictable manner by merely altering the composition of the set of product attributes that are included in an electronic recommendation agent. Moreover, our findings suggest that, rather than being confined to the particular digital shopping environment in which the selective recommendation agent is available, this preference–construction effect is likely to persist for some time and affect subsequent purchase decisions in different settings—either in other electronic shopping environments or possibly in bricks-and-mortar stores.

From a consumer-welfare standpoint, our results suggest that the potential for systematically manipulating consumer behavior in digital marketplaces, such as through the design of digital decision aids, is very significant. We find this despite the fact that we used an electronic recommendation agent that, apart from being selective with respect to the inclusion of attributes, was perfectly cooperative. For example, it considered all available products and provided product recommendations that were fully accurate (given an individual's input regarding his or her own preference). Less cooperative recommendation systems may silently omit certain products or entire classes of products (e.g., all models of certain brand) or use a biased algorithm to generate a "personalized list of recommended products" (e.g., by attenuating the importance of price in a consumer's subjective preference model, or by boosting the rank-positions of certain alternatives). We conjecture that our findings are conservative in the sense that they understate the potential for influencing consumer preferences and purchase decisions through noncooperative recommendation agents.

The research presented here suggests that electronic decision aids for online shopping can have a powerful influence on consumer choice behavior in digital marketplaces. It extends the existing body of literature on constructive consumer preferences by proposing and demonstrating the existence of a new type of preference–construction effect that, given the rapidly increasing prevalence of electronic decision aids for consumers, is of growing importance and that, based on our empirical evidence, appears to be both strong in magnitude and persistent over time. Finally, this research also makes a contribution to the emerging literature on consumer behavior in the context of the Internet, in that it represents a step toward a more complete understanding of consumer decision-making in digital marketplaces.

ACKNOWLEDGMENTS

We acknowledge the support provided by the Social Sciences and Humanities Research Council of Canada, through both its *Initiative on the New Economy* Research Alliances Program (SSHRC 538–2002–1013) and its Standard Re-

search Grants Program (SSHRC 410–99–0677), by the Institute for Online Consumer Studies (www.iocs.org), and by Macromedia, Inc. This research also benefited from the Banister Professorship in Electronic Commerce, the Petro-Canada Young Innovator Award, and the Southam Faculty Fellowship, all awarded to Gerald Häubl, as well as from the Poole Ph.D. Endowment Fellowship, the University of Alberta School of Business Ph.D. Fellowship, and the Province of Alberta Graduate Fellowship awarded to Kyle Murray.

We thank Jim Bettman and Terry Elrod for their valuable comments on an earlier version of this article.

REFERENCES

- Alba, Joseph, Lynch, John, Weitz, Barton, Janiszewski, Chris, Lutz, Richard, Sawyer, Alan, & Wood, Stacey. (1997). Interactive home shopping: consumer, retailer, and manufacturer incentives to participate in electronic marketplaces. *Journal of Marketing*, *61*, 38–53.
- Ansari, Asim, Essegai, Skander, & Kohli, Rajeev. (2000). Internet recommendation systems. *Journal of Marketing Research*, *37*, 363–375.
- Bakos, J. Yannis. (1997). Reducing buyer search costs: Implications for electronic marketplaces. *Management Science*, *43*, 1676–1692.
- Bettman, James R. (1979). *An information processing theory of consumer choice*. Reading, MA: Addison-Wesley.
- Bettman, James R., & Kassar, Pradeen. (1977). Effects of information presentation format on consumer information acquisition strategies. *Journal of Consumer Research*, *3*, 233–240.
- Bettman, James R., Luce, Mary Frances, & Payne, John W. (1998). Constructive consumer choice processes. *Journal of Consumer Research*, *25*, 187–217.
- Coupey, Eloise. (1994). Restructuring: Constructive processing of information displays in consumer choice. *Journal of Consumer Research*, *21*, 83–99.
- Ernst & Young (2001). *Global Online Retailing* (special report). New York, NY: Ernst & Young. Retrieved July 23, 2001, from http://www.ey.com/GLOBAL/gcr.nsf/US/Online_Retailing_-_Thought_Center_-_Ernst_Young_LLP
- Forrester Research. (2001). *NRF/Forrester Online Retail Index*. Cambridge, MA: Forrester Research, Inc. Retrieved July 23, 2001, from <http://www.forrester.com/NRF>
- Grice, Herbert P. (1975). Logic and conversation. In P. Cole & J. L. Morgan (Eds.), *Syntax and semantics*, *3: Speech acts* (pp. 41–58). New York: Academic Press.
- Häubl, Gerald, & Trifts, Valerie. (2000). Consumer decision making in online shopping environments: The effects of interactive decision aids. *Marketing Science*, *19*, 1, 4–21.
- Herr, Paul M. (1989). Priming price: Prior knowledge and context effects. *Journal of Consumer Research*, *16*, 67–75.
- Huber, Joel, & Zwerina, Klaus. (1996). The importance of utility balance in efficient choice designs. *Journal of Marketing Research*, *33*, 307–317.
- Jarvenpaa, Sirkka L. (1989). The effect of task demands and graphical format on information processing strategies. *Management Science*, *35*, 285–303.
- Jarvenpaa, Sirkka L. (1990). Graphic displays in decision making — the visual salience effect. *Journal of Behavioral Decision Making*, *3*, 247–262.
- Levinson, Stephen C. (1983). *Pragmatics*. Cambridge, UK: Cambridge University Press.
- Lynch, John G., & Ariely, Dan. (2000). Wine Online: Search Costs and Competition on Price, Quality, and Distribution. *Marketing Science*, *19*, 83–103.
- Lynch, John G., & Srull, Thomas K. (1982). Memory and attentional factors in consumer choice — concepts and research methods. *Journal of Consumer Research*, *9*(1), 18–38.
- MacGregor, D., & Slovic, Paul. (1986). Graphical representation of judgmental information. *Human Computer Interaction*, *2*, 179–200.
- Mandel, Naomi, & Johnson, Eric J. (2002). When Web pages influence choice: Effects of visual primes on novices. *Journal of Consumer Research*, *29*, 235–245.
- McNamara, Timothy P. (1994). Theories of priming II: Types of primes. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *20*, 3, 507–520.
- Payne, John W., Bettman, James R., & Johnson, Eric J. (1992). Behavioral decision research: A constructive processing perspective. *Annual Review of Psychology*, *43*, 87–131.
- Payne, John W., Bettman, James R., & Johnson, Eric J. (1993). *The adaptive decision maker*. New York: Cambridge University Press.
- Payne, John W., Bettman, James R., & Schkade, David. (1999). Measuring constructed preferences: Towards a building code. *Journal of Risk and Uncertainty*, *19*, 243–271.
- Shugan, Steven M. (1980). The cost of thinking. *Journal of Consumer Research*, *7*, 99–111.
- Simon, Herbert A. (1955, February). A behavioral model of rational choice. *Quarterly Journal of Economics*, *69*, 99–118.
- Simon, Herbert A. (1990). Invariants of human behavior. *Annual Review of Psychology*, *41*, 1–19.
- Slovic, Paul. (1972). From Shakespeare to Simon: Speculations — and some evidence — about man's ability to process information. *Oregon Research Institute Bulletin*, *12*.
- Slovic, Paul. (1995). The construction of preference. *American Psychologist*, *50*, 364–371.
- Srull, Thomas K., & Wyer, Robert S. (1979). The role of category accessibility in the interpretation of information about persons: Some determinants and implications. *Journal of Personality and Social Psychology*, *37*, 1660–1672.
- Yi, Youjae. (1990). The effects of contextual priming in print advertisements. *Journal of Consumer Research*, *17*, 215–222.

Accepted by Dawn Iacobucci.

APPENDIX A

TABLE A1
Set of Available Products: Negative Interattribute Correlation

<i>Model Name</i>	<i>Attribute Block 1</i>		<i>Attribute Block 2</i>	
	<i>Primary Attribute</i>	<i>Secondary Attribute</i>	<i>Primary Attribute</i>	<i>Secondary Attribute</i>
	<i>Durability Rating</i>	<i>Fly Fabric</i>	<i>Weight (Kilograms)</i>	<i>Warranty (Years)</i>
1. Coyote	76	2.3 oz Nylon	3.4	4
2. Adventurer	76	1.9 oz Polyester	3.4	3
3. Sunlight	79	2.3 oz Nylon	3.5	4
4. Grizzly	79	1.9 oz Polyester	3.5	3
5. Oasis	82	2.3 oz Nylon	3.6	4
6. Solitude	82	1.9 oz Polyester	3.6	3
7. Summit	85	2.3 oz Nylon	3.3 ^a	4
8. Drifter	85	1.9 oz Polyester	3.3 ^a	3
9. Challenger	88	2.3 oz Nylon	3.8	4
10. Serenity	88	1.9 oz Polyester	3.8	3
11. Raven	91	2.3 oz Nylon	3.9	4
12. Waterfall	91	1.9 oz Polyester	3.9	3
13. Naturalist	94	2.3 oz Nylon	4.0	4
14. Skyline	94	1.9 oz Polyester	4.0	3
15. Neptune	97 ^a	2.3 oz Nylon	3.7	4
16. Freestyle	97 ^a	1.9 oz Polyester	3.7	3

^aThe most attractive level of each of the two primary attributes.

TABLE A2
Set of Available Products: Positive Interattribute Correlation

<i>Model Name</i>	<i>Attribute Block 1</i>		<i>Attribute Block 2</i>	
	<i>Primary Attribute</i>	<i>Secondary Attribute</i>	<i>Primary Attribute</i>	<i>Secondary Attribute</i>
	<i>Durability Rating</i>	<i>Fly Fabric</i>	<i>Weight (Kilograms)</i>	<i>Warranty (Years)</i>
1. Traveler	76	2.3 oz Nylon	3.9	4
2. Journey	76	1.9 oz Polyester	3.9	3
3. Seabreeze	79	2.3 oz Nylon	4.0	4
4. Moonscape	79	1.9 oz Polyester	4.0	3
5. Galaxy	82	2.3 oz Nylon	3.5	4
6. Lakeside	82	1.9 oz Polyester	3.5	3
7. BackTrail	85	2.3 oz Nylon	3.6	4
8. Eagle	85	1.9 oz Polyester	3.6	3
9. Eclipse	88	2.3 oz Nylon	3.7	4
10. Daydream	88	1.9 oz Polyester	3.7	3
11. Spirit	91	2.3 oz Nylon	3.8	4
12. Westwind	91	1.9 oz Polyester	3.8	3
13. Glacier	94	2.3 oz Nylon	3.3 ^a	4
14. Wanderer	94	1.9 oz Polyester	3.3 ^a	3
15. Mountain	97 ^a	2.3 oz Nylon	3.4	4
16. Outfitter	97 ^a	1.9 oz Polyester	3.4	3

^aThe most attractive level of each of the two primary attributes.

APPENDIX B

Importance of Features

Note: The information you provide here is very important, as it will serve as the basis for the agent's recommendation. The more accurate the information you provide is, the better the recommendation agent will be able to suggest backpacking tents to you that match your personal preference. Consider your answers carefully, it will pay off later.

Use the following scale to rate the importance of each feature:
0 = "not at all important" to 100 = "extremely important"

How important is it to you personally that a backpacking tent has each of the following features?

High Durability	<input type="checkbox"/>
Strong Fly Fabric	<input type="checkbox"/>

FIGURE B1 Calibration of the recommendation agent (shopping task).

Personalized List of Recommended Products

The order of products in the list below was determined by the recommendation agent based on the feature importance ratings that you provided. The products are sorted by their likely attractiveness to you, starting with the most attractive ones.

Model	Durability (1 - 100)	Fly Fabric	Link
Neptune	97	2.3 oz Nylon	Click here for a detailed description
Naturalist	94	2.3 oz Nylon	Click here for a detailed description
Raven	91	2.3 oz Nylon	Click here for a detailed description
Challenger	88	2.3 oz Nylon	Click here for a detailed description
Summit	85	2.3 oz Nylon	Click here for a detailed description
Freestyle	97	1.9 oz Polyester	Click here for a detailed description
Oasis	82	2.3 oz Nylon	Click here for a detailed description
Skyline	94	1.9 oz Polyester	Click here for a detailed description
Sunlight	79	2.3 oz Nylon	Click here for a detailed description
Waterfall	91	1.9 oz Polyester	Click here for a detailed description
Coyote	76	2.3 oz Nylon	Click here for a detailed description
Serenity	88	1.9 oz Polyester	Click here for a detailed description
Drifter	85	1.9 oz Polyester	Click here for a detailed description
Solitude	82	1.9 oz Polyester	Click here for a detailed description
Grizzly	79	1.9 oz Polyester	Click here for a detailed description
Adventurer	76	1.9 oz Polyester	Click here for a detailed description

FIGURE B2 Personalized list of recommended products (shopping task).

Paired Choice Task - Microsoft Internet Explorer

Paired Choice Task 1a

The model of tent that you purchased: *Naturalist*

An alternative tent:	Model	Durability Rating (1 - 100)	Fly Fabric	Weight (kg)	Warranty (in years)	<input type="radio"/>
	Range	91	2.3 oz Nylon	3.7	4	

How much do you prefer the tent you have chosen? (Click one of the buttons.)

just barely prefer 1 2 3 4 5 very strongly prefer

[Click here to continue](#)

FIGURE B3 First choice task (sample set).

Paired Choice Task - Microsoft Internet Explorer

Paired Choice Task 2a

Alternative A:	Durability Rating (1 - 100)	Fly Fabric	Weight (kg)	Warranty (in years)	<input type="radio"/>
	91	2.3 oz Nylon	3.7	4	

Alternative B:	Durability Rating (1 - 100)	Fly Fabric	Weight (kg)	Warranty (in years)	<input type="radio"/>
	94	2.3 oz Nylon	4.0	4	

How much do you prefer the tent you have chosen? (Click one of the buttons.)

just barely prefer 1 2 3 4 5 very strongly prefer

[Click here to continue](#)

FIGURE B4 Second choice task (sample choice set).