Recommending or Persuading? The Impact of a Shopping Agent's Algorithm on User Behavior

Gerald Häubl
School of Business
University of Alberta
Edmonton, AB
Canada, T6G 2R6
Gerald.Haeubl@ualberta.ca

Kyle B. Murray
School of Business
University of Alberta
Edmonton, AB
Canada, T6G 2R6
kbmurray@ualberta.ca

ABSTRACT
This paper investigates the potential of recommendation agents for electronic shopping to influence human decision making by shaping user preferences. Specifically, we examine how the type of information that is elicited by a shopping agent for use in its recommendation algorithm may affect consumers' preference for product features and ultimately their product choice in an electronic marketplace. A recommendation agent is defined as a software tool that (a) calibrates a model of a user's preference based on his/her input and (b) uses this model to make personalized product recommendations. We report the results of a controlled experiment that demonstrates that, everything else being equal, the inclusion of a product feature in a recommendation agent renders this feature more prominent in shoppers' purchase decisions. In addition, we find that this effect is moderated by an important property of the marketplace — the correlation structure among the features of available products. We conclude that electronic shopping agents, through the design of their recommendation algorithms, have the potential to influence user preferences in a systematic fashion.

Categories and Subject Descriptors
H.1.2 [**Models and Principles**]: User/Machine Systems – *human information processing, human factors, software psychology.*

H.5.2 [**Information Interfaces and Presentation (e.g., HCI)**]: User Interfaces – *evaluation/methodology, interaction styles, screen design, theory and methods.*

General Terms
Algorithms, Management, Design, Economics, Experimentation, Human Factors, Theory.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
EC'01, October 14-17, 2001, Tampa, Florida, USA.
Copyright 2001 ACM 1-58113-387-1/01/00010…$5.00.

Keywords
Recommendation Systems, Shopping Agents, Personalization, Human Decision Making, Consumer Behavior, Product Choice, Online Shopping, User Preferences, Persuasion, Influence.

1. INTRODUCTION
The constraints of physical space no longer dictate the organization of information in electronic shopping environments [6]. 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," i.e., 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 electronic marketplaces than in the physical world [1, 7], 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 [9, 11, 12]. In electronic shopping environments, 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 user [5]. 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, MySimon, Microsoft's MSN eShop, and the Yahoo! shopping site). We refer to electronic decision aids of this type as recommendation agents.

Following Häubl and Trifts [5], we conceptualize an electronic *recommendation agent* as a software tool that (a) attempts to understand a human decision maker's multi-attribute 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 provided to the human in a decision task based on its understanding of that individual's subjective preference structure.

Our focus is on recommendation agents that attempt to understand a consumer's preference in terms of a multi-attribute preference model (based, e.g., on a weighted additive evaluation rule) that is calibrated using subjective preference information revealed to the agent by the human user. Such *feature-* or *attribute-based* recommendation agents are an integral component of many of the major online shopping sites (e.g., MySimon), and they represent a standardized technology that can be licensed by vendors for inclusion in their electronic stores (e.g., Active Buyer's Guide and Frictionless Commerce's PurchaseSource).

Almost inevitably, real-world attribute-based recommendation agents are *selective* in the sense that only a subset of all the relevant product features can be used in their calibration and, thus, in the algorithm used to generate the recommendations. This is apparent in the implementation of many commercial recommendation systems for online shopping (see, e.g., MySimon, Active Buyer's Guide, or Nike's online product recommender). 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 high-dimensional attribute space, (c) an inclination to use only those attributes that are common to most or all alternatives, 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 these 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., Nike's online store), in order 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), in order 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. For example, Nike's product recommender does not include any features that are not available in Nike products. Similarly, Active Buyer's Guide typically does not include all possible product attributes, as this would defeat one of the main purposes of recommendation agents, namely to simplify consumers' product search and purchase decisions. The present 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 paper, 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' pre-purchase information search [5]. This finding suggests that, due to the limited information-processing capacity of the human mind, users tend to rely heavily upon an electronic agent's product recommendation in order 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*.

Based on recent theorizing in the area of preference construction, we propose that the characteristics of a recommendation agent may systematically influence decision makers' preferences for objects in multi-attribute space. Since real-world recommendation agents are almost inevitably selective in the sense that they consider only a subset of the pertinent product features, the particular set of features that is included in an agent — i.e., used at its calibration stage 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 prominent in preferential choice merely as a result of its inclusion in a recommendation agent.

The remainder of the paper is organized as follows. First, we briefly review the notion of constructive consumer preferences and discuss the potential for preference construction in digital marketplaces. Specifically, we focus on the possible role of electronic recommendation agents in consumers' construction of preferences. This is followed by a discussion of the methods and results of a controlled experiment aimed at enhancing our understanding of how attribute inclusion in an electronic agent may influence individuals' preferences. The results of this experiment provide support for the existence of the predicted inclusion effect — attributes receive greater weight in an agent-assisted shopping task when they are included in the recommendation agent than when they are not. The paper concludes with a discussion of how the findings contribute to the expansion of our understanding of human decision making in connection with electronic shopping agents.

2. THE POTENTIAL OF AN ELECTRONIC AGENT TO PERSUADE

The information-processing approach to human decision making recognizes that individuals' information-processing capacity is limited [2] 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 [13]. 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 [3]. 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 [8].

The basic idea that underlies the constructive-preferences perspective is that decision makers typically do not have a "master list" that they can refer to with regards to their preferences [3]. This perspective adheres to two major tenets: (a) expressions of preference are generally constructed at the time at which the valuation of an object is required, and (b) 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, leading to highly contingent behavior [10].

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 [14], the potential of electronic shopping environments, which are interactive and personalizable, to influence consumer preferences and, ultimately, purchase decisions is very significant [6]. In particular, since shoppers tend to be quite willing to rely on product recommendations made to them by digital agents [5], the latter may be an important determinant of how consumers construct their preferences in electronic shopping environments.

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 above), we propose that such selective recommendation agents may systematically influence consumers' preferences. More specifically, we predict that the relative importance that shoppers attach to different product attributes may be influenced by whether or not a particular attribute is used in the recommendation algorithm of an electronic agent. Our key hypothesis is that the inclusion of an attribute in an agent will, everything else being equal, render this attribute more prominent when consumers make product choices in digital marketplaces. We refer to this type of preference-construction effect as an *inclusion effect*. For a more detailed discussion of the psychological mechanisms that may underlie this effect, see [4].

The proposed inclusion effect was examined in a controlled experiment using human subjects, in which we gathered empirical evidence as to the existence of such a preference-construction effect in connection with product choice in an electronic shopping environment. In this study, we also investigated the potential moderating effect of an important property of the marketplace, the inter-attribute correlation structure across the set of available products, with respect to any inclusion effect that might exist.

3. METHOD

3.1 Overview

The objective of this experiment was to examine the possibility of preference construction due to the selective inclusion of attributes in a recommendation agent. The study was fully computer-based, and involved a simulated shopping trip in an Internet-based electronic store equipped with a recommendation agent and the subsequent completion of an online questionnaire. Subjects were informed that the purpose of the research was to test a new electronic shopping environment and its features. Their task was to shop for a backpacking tent in the Internet-based store and to complete their simulated shopping trip by selecting from the set of available tents the one that was the most attractive to them personally. A total of 347 subjects completed the study remotely, via a secure Internet site. Participants were randomly assigned to one of the treatment conditions (see below).

3.2 Selective Inclusion of Attributes in the Recommendation Agent

The backpacking tents were described within a four-dimensional attribute space. For the selective inclusion of attributes in the recommendation agent, 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 pilot study, 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 characteristics of the actual attributes, such as their ecological importance.

3.3 Available Alternatives

A total of 16 backpacking tents were available during the shopping task. These alternatives were hypothetical, but presented as actual products in the study. Each tent was identified by a fictitious model name and described in terms of four attributes. The two primary and the two secondary attributes were varied at eight and two levels, respectively. In order to allow for a clear and simple test of the predicted inclusion effect, the available alternatives were constructed such that subjects' 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. No alternative had the most attractive level of both primary attributes, and, thus, subjects had to rely relatively more heavily on one of the two primary attributes when making their choice. 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/her decision. The price of the backpacking tents was held constant — subjects were informed that the price of each tent was $249.

3.4 Inter-Attribute Correlation

In order to examine whether the predicted inclusion effect might be moderated by the nature of the marketplace, we systematically manipulated the correlation structure among the attributes (in terms of attribute-level utilities) of the available products. Specifically, two different product spaces ("markets") were created. The first one was characterized by negative inter-attribute correlation. This marketplace is "efficient" in the sense that no alternative is clearly superior to another and choosing a product involves making (potentially very difficult) trade-offs among attributes. Thus, the setting in which attributes are negatively correlated resembles a typical real-world market. The second type of marketplace that we used was characterized by positive inter-attribute correlation. Such a market is "inefficient" in the sense that some alternatives are clearly superior to others and choosing a product involves few, if any, trade-offs among attributes. While this type of market structure is somewhat atypical, it provides us with an opportunity to also test the predicted inclusion effect in a setting that requires little effort on the part of the decision maker. Subjects were randomly assigned to one of the two treatments of inter-attribute correlation. Descriptions of the sets of products that were available in each of the two marketplaces are provided in Table 1 and Table 2.

3.5 Subjects' Interaction with the Recommendation Agent

At the beginning of the experiment, subjects read instructions relating to the task, including an explanation of the recommendation agent's purpose and functionality. In order to calibrate the agent, subjects were asked to indicate how important they personally considered each of the two included attributes to be on a 100-point rating scale (see Figure 1). Based on these subjective attribute-importance weights, and using (standardized) utility scale values for the different levels of each attribute, the recommendation agent computed a linear-weighted overall utility score for each available product. It then returned a personalized list of recommended products in which the alternatives were sorted by their utility score in descending order (see Figure 2). For each product, this list contained the model name and the levels of the two attributes that were included in the agent. From the recommendation list, subjects were able to request a detailed description of a particular product (i.e., the levels of all four attributes) by clicking on a hyperlink. From the screen containing the detailed description, subjects were able to either return to the personalized list of recommended products for further search or proceed to complete their "purchase."

3.6 Experimental Design

The experimental design for this study is a 2 (inclusion of attribute block in recommendation agent) × 2 (inter-attribute correlation between included and excluded attributes) between-subjects full factorial, yielding four different treatment conditions.

Table 1: Set of Available Products — Negative Inter-Attribute Correlation

	Model Name	Durability Rating	Fly Fabric	Weight (kilograms)	Warranty (years)
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	4
8	Drifter	85	1.9 oz Polyester	3.3	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	2.3 oz Nylon	3.7	4
16	Freestyle	97	1.9 oz Polyester	3.7	3

Note: The most attractive level of each of the two primary attributes is indicated by gray shading.

The two experimental factors were manipulated as follows:

- The *inclusion of attributes* in the electronic agent was implemented by having either attribute block 1 or attribute block 2 included in the agent's algorithm for generating product recommendations.
- The *inter-attribute correlation* was manipulated by constructing the available alternatives such that the correlation (in terms of attribute-level utilities) 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 ($\rho = +0.71$). Descriptions of the sets of products for the two treatments of inter-attribute correlation are provided in Table 1 and Table 2.

Each of the 347 study participants was randomly assigned to one of the four treatment conditions implied by the 2×2 full factorial design.

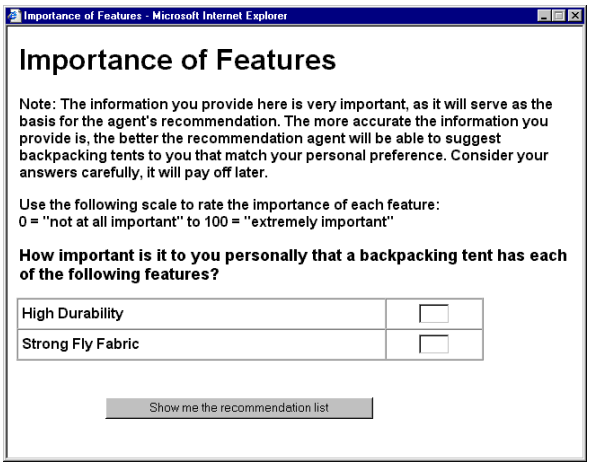

Figure 1: Calibration of Recommendation Agent

4. RESULTS

As a test of the predicted inclusion effect, we examine the relative choice shares in the shopping task for alternatives that have the most attractive level of the primary included attribute, i.e., the primary attribute that was considered by the recommendation agent. Our directional prediction is that alternatives that are superior on the included attribute are more likely to be chosen than ones that are superior on the excluded attribute. The corresponding null hypothesis is that the extent to which an attribute drives subjects' choices of products is independent of whether that attribute was included in the recommendation agent or not, i.e., that half of the subjects select an alternative that has the most attractive level of the included attribute and the other half choose a product that has the most attractive level of the excluded attribute (when controlling for potential differences in the ecological importance of the actual attributes through counterbalancing). A significant departure from such a fifty-fifty split in choice shares in the predicted direction (i.e., greater importance of an attribute when it is included in the recommendation agent) would provide support for the inclusion effect. Since attribute-specific characteristics were controlled for through the counterbalancing of the two blocks of attributes, any such departure would be independent of the relative importance of the actual attributes used.

Table 2: Set of Available Products — Positive Inter-Attribute Correlation

	Model Name	Durability Rating	Fly Fabric	Weight (kilograms)	Warranty (years)
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	4
14	Wanderer	94	1.9 oz Polyester	3.3	3
15	Mountain	97	2.3 oz Nylon	3.4	4
16	Outfitter	97	1.9 oz Polyester	3.4	3

Note: The most attractive level of each of the two primary attributes is indicated by gray shading.

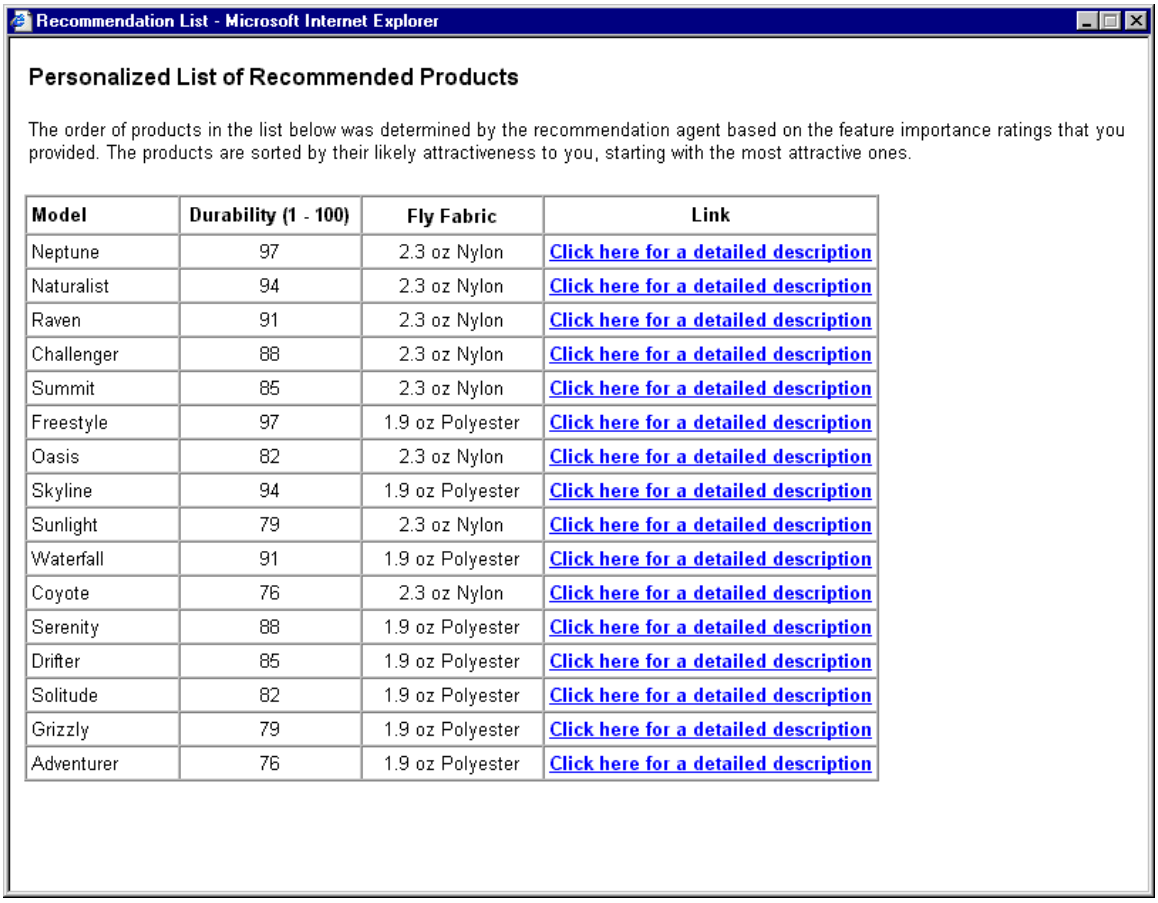

Figure 2: Personalized List of Recommended Products

We find that 60.7 percent of subjects chose an alternative that had the most desirable level of the primary included attribute and only 39.3 percent (see Figure 3) chose an alternative that had the most desirable level of the primary excluded attribute. Based on a binomial test (using equal choice probabilities as the null hypothesis), this departure from a fifty-fifty split in choice shares is statistically significant ($p < 0.0001$) and provides strong support for the predicted inclusion effect. This demonstrates that, everything else being equal, the weight that a user attaches to a particular product attribute when making a purchase decision in an agent-assisted shopping environment may in fact be enhanced as a result of the inclusion of that attribute in the recommendation agent. Thus, the empirical evidence obtained in this experiment illustrates the potential of electronic agents for online shopping to influence consumer preferences.

This effect of attribute inclusion in the recommendation agent is moderated by the correlation between included and excluded attributes in the set of available alternatives. A strong inclusion effect was observed when the primary attribute included in the agent was negatively correlated (in terms of utility) with the excluded primary attribute. In conditions with negative inter-attribute correlation, 71 percent of subjects purchased an alternative that was superior on the primary included attribute and only 29 percent chose a product that was superior on the primary attribute that was excluded from the agent. According to a binomial test, these relative choice shares depart significantly from the base case of equal choice shares ($p < 0.0001$). By contrast, when inter-attribute correlation was positive, we do not find evidence of such preference construction. The choice shares for the two types of alternatives (superior on the included versus on the excluded attribute) are not significantly different from each other in this case ($p > 0.75$). A graphical representation of the moderating effect of inter-attribute correlation with respect to the inclusion effect, which is highly significant ($\chi^2 = 12.234$, $df = 1$, $p < 0.0001$), is provided in Figure 4. We observe a strong agent-induced preference-construction effect when the marketplace is efficient in the sense that no alternative is clearly superior to another and choosing a product involves making trade-offs among attributes (i.e., market with negative inter-attribute correlation), but not when the marketplace is inefficient (i.e., positive inter-attribute correlation). Thus, the inclusion of an attribute in an electronic agent tends to influence user preferences only when decision making involves some degree of difficulty (e.g., when decision makers are forced to trade off one feature for another), and not when arriving at a decision is easy (e.g., when an attractive level of one attribute tends to be associated with attractive levels of other attributes and, thus, some products are

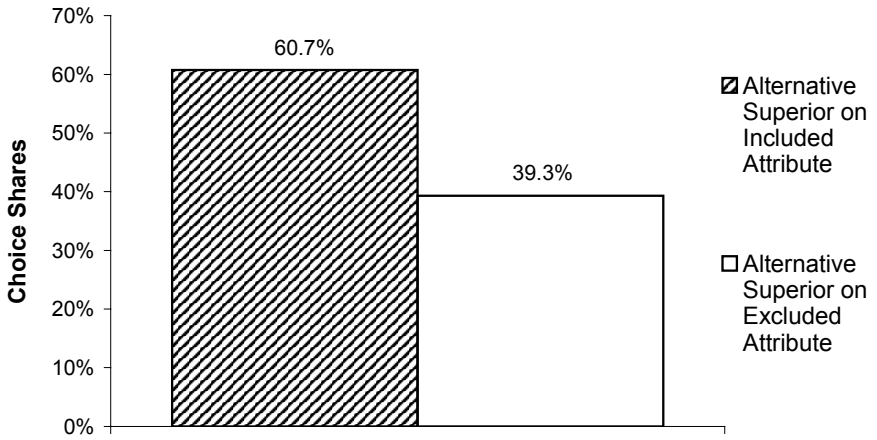

Figure 3: Attribute Inclusion in the Agent and Choice Shares in Agent-Assisted Shopping

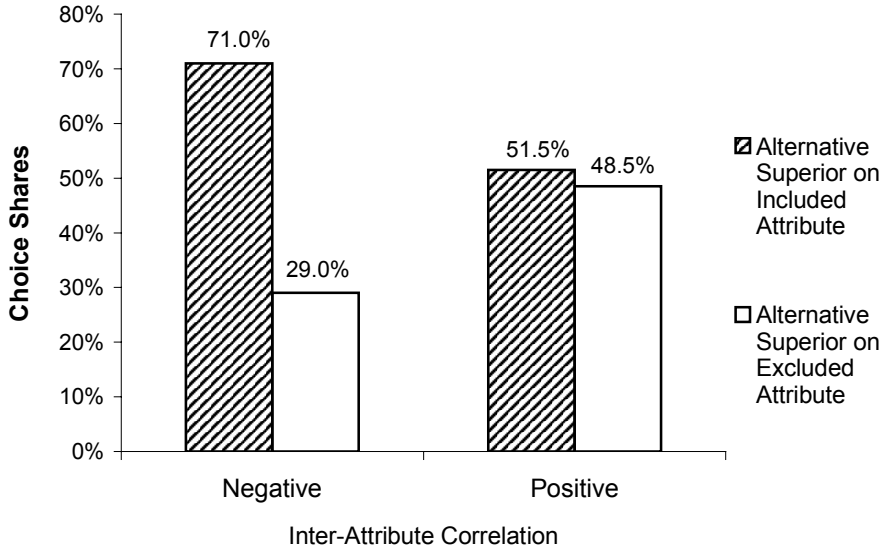

Figure 4: Moderating Effect of Inter-Attribute Correlation

clearly superior to others). This suggests that the agent-induced preference-construction effect may require that a certain level of cognitive effort on the part of the user be invoked by the shopping task or decision environment. Finally, it is worth pointing out that, since most real-world markets tend to be competitive (and therefore efficient), the setting in which we do find the predicted inclusion effect is the one that most closely resembles reality.

5. DISCUSSION

Although electronic shopping environments are not subject to the space constraints of bricks-and-mortar stores, consumers remain bounded by 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 human users, thereby assisting shoppers in making better decisions with limited cognitive effort [5]. However, for such a decision aid to be effective, the consumer must place some confidence in the product recommendations made by the agent, as well as in the process by which these recommendations are being generated. This required level of confidence, or trust, raises the potential for an electronic agent to not only assist the user in the decision-making process given his/her subjective preference, but also to *influence* this preference. The objective of this paper has been to investigate the potential of recommendation agents to systematically influence user preferences.

Our key hypothesis has been that, everything else being equal, the inclusion of an attribute in a selective recommendation agent renders this attribute more prominent in consumers' purchase decisions in an electronic shopping environment. The results of our controlled experiment provide strong support for the existence of such an inclusion effect under typical market conditions where no alternative is clearly superior to another and choosing a product involves making trade-offs among attributes (i.e., negative inter-attribute correlation). Our findings suggest that, in addition to providing a recommendation, an electronic agent has the potential, whether intentionally or unintentionally, to *persuade* users that certain alternatives are preferable to others.

The research presented here demonstrates that the preferences of human decision makers can be influenced in a systematic and predictable manner by merely altering the composition of the set of product attributes that are included in a recommendation agent for online shopping. In combination with the results from Häubl and Murray [4], which demonstrate that the inclusion effect may persist over time and into settings where an electronic agent is no longer present, this stream of research illustrates the considerable potential for systematically manipulating consumer behavior and consumer preferences in digital marketplaces through the design of electronic decision aids.

This paper extends the existing body of literature on constructive consumer preferences by proposing and demonstrating a new type of preference-construction effect that, given the rapidly increasing prevalence of electronic decision aids for online shopping, is of growing importance. In addition, this research also makes a contribution to the emerging literature on consumer behavior in the context of electronic commerce, in that it represents a step towards a more complete understanding of human decision making in agent-assisted electronic shopping environments.

6. ACKNOWLEDGMENTS

The authors gratefully acknowledge the support provided by the Social Sciences and Humanities Research Council of Canada (SSHRC grant 410-99-0677), the University of Alberta, Macromedia Inc., and the Institute for Online Consumer Studies (www.iocs.org). This research also benefited from the Banister Professorship in Electronic Commerce, the Petro-Canada Young Innovator Award, and the Southam Faculty Fellowship awarded to Gerald Häubl, as well as from the Poole PhD Endowment Fellowship, the University of Alberta School of Business PhD Fellowship, and the Province of Alberta Graduate Fellowship awarded to Kyle Murray. The authors wish to thank Terry Elrod and Jim Bettman for their valuable input. Correspondence should be addressed to the first author at University of Alberta, School of Business, Edmonton, Alberta, T6G 2R6, Canada, E-mail: Gerald.Haeubl@ualberta.ca, Voice: 780-492-6886.

7. REFERENCES

[1] Bakos, Yannis (1997), "Reducing Buyer Search Costs: Implications for Electronic Marketplaces," *Management Science*, 43, 12, 1676-1692.

[2] Bettman, James R. (1979), *An Information Processing Theory of Consumer Choice*, Reading, MA: Addison-Wesley.

[3] Bettman, James R., Mary Frances Luce, and John W. Payne (1998), "Constructive Consumer Choice Processes," *Journal of Consumer Research*, 25 (December), 187-217.

[4] Häubl, Gerald and Kyle B. Murray (2002), "Preference Construction and Persistence in Digital Marketplaces: The Role of Electronic Recommendation Agents," *Journal of Consumer Psychology*, forthcoming.

[5] Häubl, Gerald and Valerie Trifts (2000), "Consumer Decision Making in Online Shopping Environments: The Effects of Interactive Decision Aids," *Marketing Science*, 19, 1, 4-21.

[6] Johnson, Eric J., Gerald L. Lohse, and Naomi Mandel (2001), "Computer-Based Choice Environments: Four Approaches to Designing Marketplaces of the Artificial," working paper, Columbia University.

[7] Lynch, John G. and Dan Ariely (2000), "Wine Online: Search Costs and Competition on Price, Quality, and Distribution," *Marketing Science*, 19, 1, 83-103.

[8] Payne, John W., James R. Bettman, and Eric J. Johnson (1992), "Behavioral Decision Research: A Constructive Processing Perspective," *Annual Review of Psychology*, 43, 87-131.

[9] Payne, John W., James R. Bettman, and Eric J. Johnson (1993), *The Adaptive Decision Maker*, New York: Cambridge University Press.

[10] Payne, John W., James R. Bettman, and David Schkade (1999), "Measuring Constructed Preferences: Towards a Building Code," *Journal of Risk and Uncertainty*, 19, 243-271.

[11] Shugan, Steven M. (1980), "The Cost of Thinking," *Journal of Consumer Research*, 7 (September), 99-111.

[12] Simon, Herbert A. (1955), "A Behavioral Model of Rational Choice," *Quarterly Journal of Economics*, 69 (February), 99-118.

[13] Simon, Herbert A. (1990), "Invariants of Human Behavior," *Annual Review of Psychology*, 41, 1-19.

[14] Slovic, Paul (1995), "The Construction of Preference," *American Psychologist*, 50 (May), 364-371.