

Personalized Product Presentation: The Influence of Electronic Recommendation Agents on Consumer Choice

Gerald Häubl

Banister Professor of Electronic Commerce and Associate Professor of Marketing
School of Business
University of Alberta
Edmonton, AB; Canada T6G 2R6
Voice: 780-492-6886
Fax: 780-492-3325
E-mail: Gerald.Haeubl@ualberta.ca

Kyle B. Murray

Marketing Ph.D. Student
School of Business
University of Alberta
Edmonton, AB; Canada T6G 2R6
E-mail: kbmurray@ualberta.ca

Valerie Trifts

Marketing Ph.D. Student
School of Business
University of Alberta
Edmonton, AB; Canada T6G 2R6
E-mail: vtrifts@ualberta.ca

To appear in:
The Power of One — Leverage Value From Personalization Technologies
(Editors: Arvind Rangaswamy and Nirmal Pal)

First version: November 14, 2001

This version: January 9, 2002

© Gerald Häubl, Kyle B. Murray, Valerie Trifts
Please do not reproduce or quote without the authors' permission. Comments are welcome.

Introduction

Personalized Product Presentation

Imagine that a retailer is able to organize its store in such a way that, when a customer looking for a tent enters the store, the first products she sees are tents. In addition, imagine that the customer wants a backpacking tent for about \$250 and that, therefore, the retailer has displayed the tents of this type and in this price range right by the door. Moreover, the retailer has organized the tents from lightest to heaviest because this customer is more concerned about the weight of the tent than about its durability. If the customer moves beyond the first tent, which — according to the vendor's understanding of her personal preference — most closely matches what she is looking for, she will come to the tent that is the next closest match to the retailer's estimate of her preference. In essence, the entire store has been arranged to suit this one customer's preference. The next customer who enters the store wants to buy a stereo, so the retailer quickly moves the tents to the back of the store and stocks the shelves closest to the door with stereos, arranging them to match that shopper's personal preference for stereos.

Such a scenario is unthinkable in a traditional retail setting, where customers with very different goals and preferences continuously pass through the store. However, this degree of personalization *is* attainable in an electronic store, where the traditional bricks-and-mortar constraints on space and organization do not apply. In fact, with the advent of online decision aids, personalization of product presentation is available to any consumer with access to the Internet who is willing to answer a few questions. For example, if you are in the market for a

running shoe, you may use Nike.com's online shoe advisor. It asks you how competitive you are, where you like to run, what sort of features you prefer, and how much you are prepared to spend. After a brief discussion with this advisor, it gives you an initial recommendation that it has personalized for you based on your expressed preference (Figure 1). In fact, it has rated Nike's entire line of running shoes for you based on the information you provided. The artificial intelligence of the electronic advisor has relabeled all of the shoes on Nike.com's virtual shelves, leaving you to browse a product line personalized specifically for you.

This type of personalized shopping assistant is often referred to as a shopbot, which is shorthand for a shopping robot. A shopbot does the work of shopping: it searches through the marketplace to find the products that best suit the subjective preference of the consumer for whom it is working. Often, such robots simply search for the lowest price for the item the consumer wants. Examples of this type of shopbot include allbookstores.com (finds the lowest price on a book) and destinationrx.com (finds the lowest price on prescription drugs and other health products). Another such example, dealtime.com, will even page shoppers when it finds a good deal on the merchandise that they are interested in. However, price-comparison robots are helpful only after a consumer knows what he or she wants to buy, and, therefore, they do only part of the work for the shopper. The consumer must consider the product's nonprice attributes before the robot can go to work. Another class of shopbots is designed to help consumers to find the product that, given their subjective preference in terms of product attributes, is right for them. In some cases, such digital decision aids may search across a number of stores (e.g., ActiveBuyersGuide.com), or they may search the database of products offered by a particular retailer or manufacturer (e.g., Nike's shoe advisor). We refer to shopbots that gather information

on a consumer's personal preference or taste in a particular product category and then base product recommendations on this information as *recommendation agents*. This type of shopbot, which actively seeks to understand a consumer's preference and makes personalized product recommendations based on this understanding, is the focus of this chapter.

The Trade-Off between Effort and Accuracy

The amount of thought that a human decision maker devotes to making a particular choice depends largely on the degree of difficulty, or thinking cost, associated with the decision (Shugan 1980). This cost of thinking is positively related to both the complexity of the decision (in terms of the number of relevant dimensions) and the desired level of confidence in having made the best possible choice, and inversely related to the difference in the decision maker's preference between the available options. As a result, complex and important decisions are more costly in terms of cognitive effort than simple and routine decisions. Individuals often settle for less accurate decisions in return for a reduction in effort (Bettman, Johnson, and Payne 1990). Because of this trade-off between effort and accuracy, decision makers often choose options that are satisfactory but would be suboptimal if decision costs were zero (Simon 1955). This is particularly true when alternatives are numerous and/or difficult to compare (Payne, Bettman, and Johnson 1993).

Unlike bricks-and-mortar shopping environments, digital marketplaces are not constrained by limitations in physical space in their organization of product information. Online vendors have virtually unlimited shelf space and can, therefore, offer a very large number of products to their customers. As a result, a potentially vast amount of information about market

offerings is available to consumers. Searching through a marketplace composed of many such retailers would require consumers who wish to make well-informed decisions to expend a great deal of effort.

For shoppers, easy access to large amounts of product information is both a blessing and a curse. It is a blessing in that, given more information, they may make better purchase decisions (e.g., select products that better match their personal preferences) than they would otherwise. However, it is a curse in that, given vast amounts of information but limited cognitive capacity, consumers may be unable to adequately process the information. The idea that human decision makers have limited resources for processing information — whether those limits are in memory, attention, motivation, or elsewhere — has deep roots in the fields of psychology and marketing (e.g., Payne, Bettman, and Johnson 1993; Shugan 1980; Simon 1955).

Because of the limitations of human information processing, recommendation agents may be of great value to online shoppers. While people tend to do quite well at selecting the criteria they wish to use in making a decision, computers are good at methodically searching through a problem (or product) space in order to compile and retain large amounts of information. For example, a recommendation agent may ask a consumer a set of questions in an attempt to understand his or her preference, and then do the work of searching through the products in the marketplace to find the most appropriate alternative(s) to recommend. Clearly, such a recommendation agent has the potential to assist consumers in their decision making by reducing their effort and increasing the quality of their purchase decisions.

While different types of recommendation agents exist (see, e.g., Ansari, Essegaiier, and Kohli 2000), we focus on *attribute-based* agents — i.e., on those that ask a consumer about his or her preference in terms of product attributes and then, based on this preference information, estimate a model that can be used to rate all available products for that individual. (In both of the studies reported in this chapter, a weighted additive utility model was used for this purpose.) Having estimated such a model, the agent is able to provide personalized product recommendations to the shopper. The Nike shoe advisor follows this type of process to make shoe recommendations. However, more general recommendation agents also exist. For example, *activebuyersguide.com* can assist shoppers in selecting products in categories as diverse as automobiles, family pets, online stockbrokers, beer, and belt sanders. Moreover, such an agent can recommend products across multiple manufacturers, brands, or retailers. In each case, the agent asks the consumer a series of questions, estimates a model of his or her preference, rates the products in its database (which may be compiled from a variety of vendors) based on this preference model, and makes personalized product recommendations (see Figure 2 for an example). The question is, do consumers benefit from this type of electronic assistance in deciding what to buy? In the following sections, we discuss the findings of two recent studies that, taken together, address precisely this question.

Overview of Empirical Evidence

We will review some of the findings of recent research on how electronic recommendation agents may influence shoppers' purchase decisions. In particular, we will discuss the relevant results of two major empirical studies that examine different aspects of

consumers' agent-assisted product choice behavior in personalized digital shopping environments.

The first study focuses on the impact that use of a recommendation agent has on both the *quality* and the *efficiency* of consumer decision making in an online shopping environment, i.e., how good a choice the consumer makes given the set of available products and how much effort he or she must expend to make a decision. This study provides evidence that a recommendation agent can benefit consumers, because it does much of the work of searching the product space and personalizing the information environment by presenting those alternatives likely to be most attractive to a shopper first (i.e., at the "front" of the store). However, for a recommendation agent to benefit consumers, it must do its work in an accurate and unbiased fashion. The electronic agent must be effective at determining what the consumer wants and at searching for a product that meets the consumer's needs.

The second study concerns recommendation agents that perform their search of the marketplace based on an incomplete conversation with the consumer. In this context, "conversation" refers to the dialogue between a computer-based recommendation agent and a consumer, in which the electronic agent asks the human questions designed to elicit information regarding his or her personal preference in terms of the features or attributes of a product. The second study examines the impact on consumer behavior of a recommendation agent that elicits limited preference information before making a personalized product recommendation. The findings suggest that, when an electronic recommendation agent is selective in its conversation with a shopper, it goes beyond merely eliciting preference information and, in fact, may *influence* the consumer's preference.

Recommendation Agents and Consumer Decision Making (Häubl and Trifts 2000)

The primary objective of the study by Häubl and Trifts (2000) was to obtain an understanding of the possible effects of using an electronic recommendation agent on both the quality and the efficiency of consumer decision making in online shopping environments. To that end, they examined the impact of the availability of such an electronic decision aid on three aspects of consumer decision making: (1) the amount of search that the consumer undertakes before making a purchase, (2) the set of products the consumer seriously considers purchasing (i.e., the consideration set), and (3) the quality of the consumer's ultimate purchase decision.

Method

Häubl and Trifts conducted a controlled experiment to examine the effects of an electronic recommendation agent on the above aspects of consumer decision making in an online shopping environment. A participant's task consisted of shopping for, and making a hypothetical purchase of, a product in each of two categories, backpacking tents and compact stereo systems, in an online store. These purchase decisions were tied to a lottery incentive that was designed to increase the validity of the findings by making the shopping task more consequential (see Häubl and Trifts 2000 for details). The availability of the recommendation agent was manipulated systematically. Half of the participants in this study completed the task with the help of the electronic agent, while the other half received no such assistance. In addition, the order in which subjects shopped for the two products was varied independently. Study participants were randomly assigned to one of the experimental conditions.

The data for this study were collected in a university computer lab in small group sessions of 15 to 20 subjects. The study was completed by a total of 249 participants. Upon arrival at the laboratory, subjects were assigned to personal computers and informed that they would be pilot-testing a new online store by shopping for two products, a backpacking tent and a compact stereo system. The experimenter then held a ten-minute practice session during which she demonstrated the features of the electronic shopping environment. Before starting their first shopping trip, participants rated their levels of knowledge about, and interest in, each of the product categories (using nine-point rating scales). They then read a detailed description of the task and of the lottery incentive.

In each product category, 54 products were available (nine models for each of six brands). Actual brand names and fictitious model names were used. The following tent attributes were varied across the 54 alternatives (number of levels in parentheses): pole material (3), warranty (3), weight (12), durability rating (12), and price (12). In addition, fly fabric and vestibule were used as filler attributes with levels that were the same for all backpacking tents. For stereos, the varied attributes were CD player type (3), tuner presets (3), output power (12), sound quality rating (12), and price (12). Cassette decks and remote control were used as additional attributes, and their levels were identical for all stereo models.

An innovative method for measuring the *quality* of shoppers purchase decisions, as well as of their consideration sets, was used in this study. Since consumer preferences are not subject to direct observation, it is impossible to accurately measure decision quality in uncontrolled real-world settings. In the Häubl and Trifts (2000) study, the sets of available products were constructed in such a way that, irrespective of an individual's subjective preference, the purchase

of particular alternatives represented a poor decision. This approach is based on the idea of an objective standard for quality and requires a combination of objectively dominated and nondominated alternatives. An alternative is *dominated* if there is at least one other alternative that is superior on at least one attribute while not being inferior on any attribute. By contrast, an alternative is *nondominated* if no other alternative is superior on an attribute without, at the same time, being inferior on at least one other attribute. For each product category, six nondominated alternatives—one for each brand—were constructed. While these six products were mutually nondominated, they did dominate all other products. Whether or not a participant purchased an objectively attractive (i.e., nondominated) alternative was used as one measure of decision quality, and the share of nondominated products in a subject's consideration set was used as the measure of consideration set quality.

Subjects in the no-recommendation-agent conditions were taken to a hierarchically structured Web site with all six brands listed at the top level and all models for a brand listed at the lower level. They could access detailed information about a product by first clicking on a brand name and then on a model name. In the conditions in which the attribute-based recommendation agent was available, participants started by providing attribute importance weights using a 100-point constant-sum scale, specifying minimum-acceptable attribute levels, and selecting the maximum number of alternatives to be included in the recommendation. Based on this information, the electronic agent produced a personalized list of recommended products. In this list, products were identified by their brand and model name, and sorted by their likely attractiveness to the shopper (in descending order). From the recommendation list, subjects were able to request detailed information about particular products. In all conditions, participants

could complete their purchase from any of the screens containing detailed information about a product via a checkout procedure that included the confirmation of the selected product.

After finalizing their purchase, participants completed a short online questionnaire. Next, they were presented with a list of the alternatives they had looked at and asked to report their consideration set (“Please indicate which of these products you considered seriously before making your purchase decision.”). Subsequently, participants completed a *switching task* in which they were given an opportunity to switch from the purchased alternative to each of several nondominated alternatives, all of which had been available during the shopping task. The number of switching opportunities depended upon whether a subject had initially chosen a dominated alternative (six switching opportunities) or a nondominated alternative (five switching opportunities). The switching task consisted of a series of pairwise comparisons. Participants were encouraged to switch whenever they saw an alternative that they preferred over their initial choice, and informed that the lottery incentive would reflect any changes to their product selection that they made during this task. Response behavior in this task was used as the second measure of decision quality, with switching to another, previously available alternative indicating poor initial decision quality.

Key Results

As to the extent of information search, shoppers who had access to a recommendation agent looked at far fewer products in detail than did those who were shopping without agent assistance. Across their two shopping trips, subjects requested detailed information for an average of 6.57 alternatives when the electronic agent was available, compared to 11.78

alternatives when it was not available (Figure 3). Furthermore, although the availability of a recommendation agent had no substantial effect on the number of products that participants considered seriously for purchase, agent-assisted shoppers had a much higher percentage of objectively desirable products in their consideration set. In particular, the share of nondominated alternatives in subjects' consideration sets doubled as a result of using an electronic recommendation agent when making purchase decisions (Figure 4). The level of statistical significance for both of these effects is very high ($p < 0.001$).

The recommendation agent also had a strong positive effect on the quality of shopper's purchase decisions. First, participants who had the assistance of the electronic agent were much more likely to select a product that was objectively of high quality than unassisted shoppers. Specifically, while only 65 percent of subjects purchased a nondominated product when no recommendation agent was available, this share increased to 93 percent when shoppers had the assistance of the electronic agent (Figure 5). In addition, consumers who were able to use a recommendation agent on their digital shopping trip were significantly less likely to abandon their initial choice during the subsequent switching task than were those who had no such assistance. The share of participants who switched to another, previously available product was 59.5 percent of those who shopped without agent assistance and only 21.5 percent of those who did use the electronic recommendation agent during their shopping experience (Figure 6). The effects on both measures of decision quality are highly significant ($p < 0.001$).

In sum, these results indicate that the personalization of product presentation through an attribute-based recommendation agent allowed consumers to engage in less search, while improving the average quality of the products they considered and, most important, the quality of

their ultimate purchase decisions. In sum, use of the electronic recommendation agent enabled consumers to make *better decisions with less effort*.

The findings of Häubl and Trifts (2000) show how a recommendation agent implemented in an online shopping environment can transform the way in which consumers search for product information and make purchase decisions. Given that the trade-off between effort and accuracy has been demonstrated consistently in the offline world (Payne, Bettman, and Johnson 1993), it is remarkable that an increase in decision quality would not require an increase in effort. This study provides strong evidence that consumers can greatly benefit from the personalization of product recommendations. However, it is important to note that the recommendation agent available to shoppers in this study was fully cooperative and was carefully designed to effectively screen the marketplace on behalf of the consumer based on preference information provided by the consumer. Real-world recommendation agents may not be as altruistic or as complete in their design. For example, they may not cover all available products but instead represent only the products of a particular vendor. They may fail to ask questions or elicit information about some important product attribute, or they may be biased in the way they process the information that they do elicit, either of which may result in recommendations that do not reflect the true preference structure of the consumer. While Nike's shoe advisor is overtly a tool for recommending only Nike shoes, a recommendation agent could be much more covert about its algorithm and its objectives.

Recommendation Agents and Consumer Preference Construction (Häubl and Murray, forthcoming)

Almost inevitably, real-world attribute-based recommendation agents are selective in that they consider only a subset of all the relevant attributes in a product category. This is apparent in the implementation of many commercial recommendation systems for online shopping (e.g., Active Buyer's Guide or Nike's shoe advisor). The reasons for such selectivity in electronic recommendation agents include (1) the large number of attributes that exist in many product categories, (2) 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 preferences for products with many attributes, (3) an inclination to use only those attributes that are common to most or all available products, and (4) a tendency to include only attributes that are quantitative in nature (i.e., whose levels can be represented numerically). Apart from these reasons, the attributes to include in a recommendation agent may be chosen for strategic reasons (e.g., to de-emphasize specific attributes) by the designer of the agent.

An electronic recommendation agent may be made available either by a particular online vendor (e.g., Nike's online store) to help shoppers choose one of the products in its own assortment or by a third-party provider (e.g., Active Buyer's Guide) to help consumers choose a product from those of various vendors. The two types of providers may have different motivations for including certain attributes in these decision aids. Häubl and Murray's forthcoming work pertains equally to the two provider scenarios (vendor and third-party provider), if the recommendation agent is selective in the attributes it includes.

Häubl and Trifts (2000) found that an attribute-based recommendation agent in an electronic shopping environment can result in a substantial reduction in the amount of consumers' pre-purchase information search. This finding suggests that, due to the limited information-processing capacity of the human mind, consumers rely heavily upon an electronic agent's recommendations to reduce the effort required to make a purchase decision. Given this tendency to rely on the recommendations of these agents and given the rapidly increasing prevalence of such decision aids in digital marketplaces, it is important to examine whether and how electronic recommendation agents may *influence consumers' preferences*.

The information-processing approach to decision making recognizes that human 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, Luce, and Payne 1998). That is, in a domain (e.g., a product category) in which the alternatives have multiple attributes, individuals typically do not have specific pre-formed strategies pertaining 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 to evaluate alternatives or to make a decision (Payne, Bettman, and Johnson 1993).

The constructive preferences perspective adheres to two major tenets: (1) that expressions of preference are generally constructed when individuals are required to evaluate an object, and (2) that the process of preference construction is shaped by the interaction between the properties of the human information-processing system and the properties of the decision task (Payne, Bettman, and Schkade 1999). In a similar vein, Slovic (1995) notes that preferences appear to be remarkably labile, i.e., sensitive to the way in which a choice problem is described or framed and to the mode of response used to express the preference (see also Bettman, Luce, and Payne 1998).

Given the large amount of empirical evidence suggesting that the characteristics of the decision environment play a central role in individuals' construction of preference (e.g., Slovic 1995), digital shopping environments, which are interactive (rather than static) and personalizable (rather than standardized), have great potential to influence consumer preferences and, ultimately, purchase decisions (Johnson, Lohse, and Mandel 1999). In a recent study, Häubl and Murray (forthcoming) examined this possibility by looking at the choice behavior of consumers shopping online with the assistance of an electronic recommendation agent that is selective in its conversation with consumers, i.e., that elicits preference information in terms of only a subset of the relevant product attributes.

Method

The main task for participants in this study was to shop for a backpacking tent in an experimental online store. All participants used an electronic recommendation agent, which asked shoppers to specify their preferences for particular tent attributes. However, unlike the tool used in the Häubl and Trifts (2000) study, this recommendation agent was *selective* in the

information it asked for from different shoppers. All available backpacking tents were described on four quality attributes, and price was the same for all tents. Participants were randomly assigned to one of two agent conditions. In one, the electronic agent asked shoppers to indicate (on a 100-point scale) how important weight and warranty were to them when choosing a tent. In the second condition, the agent asked subjects how important durability and fly fabric were to them when choosing a tent (see Figure 7). The recommendation agent used in this study was selective, eliciting preference information from a shopper in terms of only two of the four relevant attributes.

The electronic recommendation agent then searched the product space based on this fragmentary preference information and provided the shopper with a list of backpacking tents, sorted based on the attribute preferences the shopper had expressed. As a result, although all available products were displayed, their presentation was personalized based on consumers' subjective preferences in terms of a subset of the relevant attributes. From the recommendation list, subjects were able to request detailed descriptions (i.e., on all four attributes) of individual tents. Participants could complete their hypothetical tent purchase from any of the screens containing detailed information about a product via a checkout procedure that included the confirmation of the selected product.

Given that preferences are often constructed on the fly rather than pre-formed, Häubl and Murray were interested in examining whether and how consumers' preferences would be affected by the selective inclusion of product attributes in a recommendation agent. They expected to find an *inclusion effect*: all else being equal, included attributes would be more important in a consumer's decision process simply because they had been included by the recommendation

agent. If this is the case, then it is also important to examine whether such a preference-construction effect may *persist* over time, especially into situations where no recommendation agent is available. To investigate the possibility of the effect's persistence, Häubl and Murray asked subjects to perform a follow-up choice task after the initial shopping experience.

The experiment 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 subjects accessed via a standard Web browser. Participants entered all of their responses via this Web interface. In addition, subjects' interaction behavior with the experimental environment was recorded electronically. Data were collected in group sessions with 10 to 15 participants per session. A total of 347 subjects completed the study. Participants were informed that the overall purpose of the study was to test a new electronic shopping environment. The main task was taking an online shopping trip for a backpacking tent in an electronic store equipped with a recommendation agent.

To allow for a clear and simple test of the predicted inclusion effect, the set of available tents was constructed such that a subject's product choice was informative as to which attribute was the most important one in making his or her decision. Shoppers had to choose a product that had the most desirable level of one attribute, but not of the other attributes. As a result, which of the attributes the selected alternative was superior on served as an indicator of the relative importance of the attributes in a subject's purchase decision. We refer interested readers to the detailed description of this method in Häubl and Murray (forthcoming).

Two different market scenarios were used for this shopping task: the inter-attribute correlations across the set of available products were either positive or negative. In a market with positive inter-attribute correlations, an alternative that is favorable on one attribute tends to also be favorable on other attributes. By contrast, in a market characterized by negative inter-attribute 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. Subjects were randomly assigned to one of these two market conditions.

After finishing their shopping trip by selecting their subjectively most-preferred tent, participants completed an extensive online questionnaire. The final part of the study involved a series preferential-choice questions, whereby subjects were asked to consider six two-alternative choice sets containing new backpacking tents, i.e., ones they had not encountered in the shopping task. All alternatives were described in terms of the same four attributes used in the shopping task. The pairs of tents were personalized using a dynamic choice design (for details see Häubl and Murray, forthcoming). In addition to choosing their preferred tent from each choice set, subjects also indicated the strength of their preference on a five-point rating scale with end points 1 = “just barely prefer” and 5 = “very strongly prefer”. The choice response and the strength-of-preference rating were combined into a 10-point graded-paired-comparison response variable representing an individual’s relative preference for the two alternatives in a choice set. This task allowed for a test of whether the inclusion effect persisted into a choice environment in which no recommendation agent was available and all products were new to consumers.

Key Results

First and foremost, Häubl and Murray were interested in whether or not a less than perfect recommendation agent (which in this case was selective in its elicitation of preference information) would affect consumers' construction of preference. The results of this experiment provide strong evidence for the proposed inclusion effect. However, the authors found this effect only in connection with a market that required consumers to make trade-offs between product attributes, i.e., one with negative inter-attribute correlations. Such a marketplace is more efficient and more analogous to real-world markets than one characterized by positive inter-attribute correlations. For example, when consumers can buy a tent that is both the lightest and the most durable (where it doesn't matter which attribute they base their choices on), Häubl and Murray found no evidence of an inclusion effect. They observed such a preference-construction effect only in marketplaces in which consumers must to make trade-offs among product attributes — for example, when a more durable tent is heavier and a less durable tent is lighter.

Figure 8 shows the choice shares for a market with negative inter-attribute correlations of (1) the products with the most desirable level of an attribute that was included in the electronic recommendation agent (i.e., an attribute that the agent asked the shopper for preference information about) and (2) the products that were superior on an attribute that was excluded from the recommendation agent (i.e., an attribute that the agent did not ask about). Note that, because each attribute was included in the recommendation agent for half of the shoppers, equal choice shares (50 percent each) in Figure 8 would have indicated the absence of any effect of attribute inclusion in the electronic agent on consumer preference. However, subjects tended to prefer products that were superior on an attribute for which the agent had elicited preference

information (71 percent choice share). This inclusion effect, measured as the departure from equal choice shares, is highly significant ($p < 0.001$).

Having demonstrated that a selective recommendation agent can influence shoppers' construction of preference, Häubl and Murray also examined whether this effect would persist in future decision making. Subjects' responses to the six preferential-choice questions, which followed the agent-assisted shopping task and an online questionnaire, were used to test for such persistence. The graded-paired-comparison responses (see above) indicate that participants attached significantly greater overall importance to the attributes that had been included in the electronic agent during the earlier shopping task than to those that had not been included ($p < 0.01$). This shows that the preference-construction effect based on attribute inclusion in the recommendation agent persisted over time and into a setting in which no recommendation agent was available. Once again, we refer interested readers to Häubl and Murray (forthcoming) for a more detailed discussion of this study's findings.

Summary of Findings

The two studies described here provide some initial insights into the potential of electronic recommendation agents to affect consumer decision making. A well-designed recommendation agent can help consumers to increase the quality of their purchase decisions and, at the same time, reduce the amount of effort required to make these decisions. However, these results also suggest that the potential for systematically manipulating consumer behavior in digital marketplaces through the design of electronic decision aids is very significant. This was

demonstrated in the Häubl and Murray (forthcoming) study despite the fact that their recommendation agent was, apart from being selective in its inclusion of attributes, perfectly cooperative. For example, it considered all available products and provided product recommendations that were fully accurate (given an individual's input about his or her 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). The findings discussed in this chapter are conservative in the sense that they tend to understate the potential for influencing consumer preferences and purchase decisions through non-cooperative recommendation agents.

Conclusion

The type of personalization we have considered in this chapter is important because consumers like to make good decisions with low effort, and personalized product recommendations can be very helpful in this regard. However, when consumers rely on an electronic recommendation agent to screen the marketplace, they open the door to influence in much the same way they would by relying on a salesperson in a bricks-and-mortar store. While we know a lot about how people can influence other people (e.g., Cialdini 2001), we know very little about how electronic entities, such as recommendation agents, can influence people.

Unlike flesh-and-blood salespeople, electronic agents can control the choice environment. Although a retail salesperson at Niketown may come to understand a particular customer's product preference over time, she or he cannot rearrange the store and personalize the presentation of products in the same way that Nike's online shoe advisor can. A human sales assistant may change his or her behavior and advice for different customers, but the electronic agent can alter the entire online shopping environment in response to each individual consumer.

While we are just beginning to develop an understanding of how this level of personalization may affect consumer behavior and consumer preferences, the research findings reported in this chapter show that the personalization of electronic shopping interfaces through recommendation agents can not only improve consumer decision making, but also systematically influence consumer preferences.

References

- Ansari, A.; S. Essegaier; and R. Kohli, 2000. Internet Recommendation Systems. *Journal of Marketing Research*, 37 (August): 363-375.
- Bettman, J. R., 1979. *An Information Processing Theory of Consumer Choice*. Reading, MA: Addison-Wesley.
- Bettman, J. R.; E. J. Johnson; and J. W. Payne, 1990. A Componential Analysis of Cognitive Effort in Choice. *Organizational Behavior and Human Decision Processes*, 45: 111-139.
- Bettman, J. R.; M. F. Luce; and J. W. Payne, 1998. Constructive Consumer Choice Processes. *Journal of Consumer Research*, 25 (December): 187-217.
- Cialdini, R. B., 2001. *Influence: Science and Practice (4th edition)*. Needham Heights, MA: Allyn & Bacon.
- Häubl, G.; and K. B. Murray, forthcoming. Preference Construction and Persistence in Digital Marketplaces: The Role of Electronic Recommendation Agents. *Journal of Consumer Psychology*.
- Häubl, G.; and V. Trifts, 2000. Consumer Decision Making in Online Shopping Environments: The Effects of Interactive Decision Aids. *Marketing Science*, 19, 1: 4-21.
- Johnson, E. J.; G. L. Lohse; and N. Mandel, 1999. Designing Marketplaces of the Artificial: Four Approaches to Understanding Consumer Behavior in Electronic Environments. *Working Paper*. Columbia University, New York, NY.
- Payne, J. W.; J. R. Bettman; and E. J. Johnson, 1993. *The Adaptive Decision Maker*. New York, NY: Cambridge University Press.
- Payne, J. W.; J. R. Bettman; and D. Schkade, 1999. Measuring Constructed Preferences: Towards a Building Code. *Journal of Risk and Uncertainty*, 19: 243-271.
- Shugan, S. M., 1980. The Cost of Thinking. *Journal of Consumer Research*, 7: 99-111.
- Simon, H. A., 1955. A Behavioral Model of Rational Choice. *Quarterly Journal of Economics*, 69 (Feb): 99-118.
- Slovic, P., 1995. The Construction of Preference. *American Psychologist*, 50 (May): 364-371.

Figure 1
Nike.com's Online Shoe Advisor



The screenshot shows the Nike.com online shoe advisor interface. The browser window title is "Shoe Advisor -- Selecting the Right Shoe! - Microsoft Internet Explorer". The address bar contains the URL: <http://nikev2.an.ask.com/advisor/nikeadvisor.asp?pqNums=q1%2Cq2%2Cq3%2Cq4%2Cq5%2Cq6%2Cq7%2Cq8%2Cq9&paTypes=r>. The page features a navigation menu with options like "ASK NIKE", "POLICIES", "STORE LOCATOR", "PRODUCT SEARCH", "SHOE ADVISOR", and "SIZING". The "SHOE ADVISOR" section displays "Results" based on user input, suggesting three shoes: "Air Durham (men)", "Air Structure Triax (men)", and "Air Span Triax (men)". Each shoe is accompanied by an image, price, and a "Buy It!" button. Below the shoe recommendations is a comparison table with various attributes.

	<u>Air Durham (men)</u> \$100 Buy It!	<u>Air Structure Triax (men)</u> \$90 Buy It!	<u>Air Span Triax (men)</u> \$85 Buy It!
Surface	Road Shoe	Road Shoe	Road Shoe
Midsole	PU midsole	Phylon midsole	Phylon midsole
Width	Standard	Standard	Standard
Motion Control	★★★★	★★★★★	★★★★★
Feel: Impact Protection	★★★★★	★★★★★	★★★★★
Feel: Responsiveness	★★★★★	★★★★★	★★★★★
Breathability	★★★★★	★★★★★	★★★★★
Water Resistance	No	No	No
Weight	14.5 ounces	12.6 ounces	12 ounces

Figure 2
Example of Personalized List of Recommended Products (Active Buyer's Guide)

The screenshot shows a Microsoft Internet Explorer window titled "Active Buyer's Guide". The address bar contains a complex URL. The main content area displays a list of products sorted by "YOUR FIT". The list includes five items, each with a fit percentage, a product name, a camera icon, and a list of attributes. A "Compare" button is visible at the top right of the list area.

your fit	Beer Type	Compare
90% fit	Specialty Beer Lindemans Kriek Lambic Ale 10 Palate - Belgium Brewery Country - 8 Presentation - 9 Appearance - 10 Bouquet - Bottle Container - National or International Brewery Type more features , reviews	<input type="checkbox"/> compare
89% fit	Stout St. Ambroise Oatmeal Stout 8 Palate - Canada Brewery Country - 9 Presentation - 9 Appearance - 7 Bouquet - Bottle Container - National or International Brewery Type more features , reviews	<input type="checkbox"/> compare
89% fit	Ale Maudite 9 Palate - Canada Brewery Country - 8 Presentation - 8 Appearance - 8 Bouquet - Bottle Container - National or International Brewery Type more features , reviews	<input type="checkbox"/> compare
89% fit	Specialty Beer Gordon Biersch Marzen (Draught) 9 Palate - United States Brewery Country - 9 Presentation - 9 Appearance - 8 Bouquet - Draught Container - Brewpub Brewery Type more features , reviews	<input type="checkbox"/> compare
89% fit	Specialty Beer San Andreas Apricot Ale 9 Palate - United States Brewery Country - 8 Presentation - 9 Appearance - 10 Bouquet - Bottle Container - Microbrewery Brewery Type more features , reviews	<input type="checkbox"/> compare



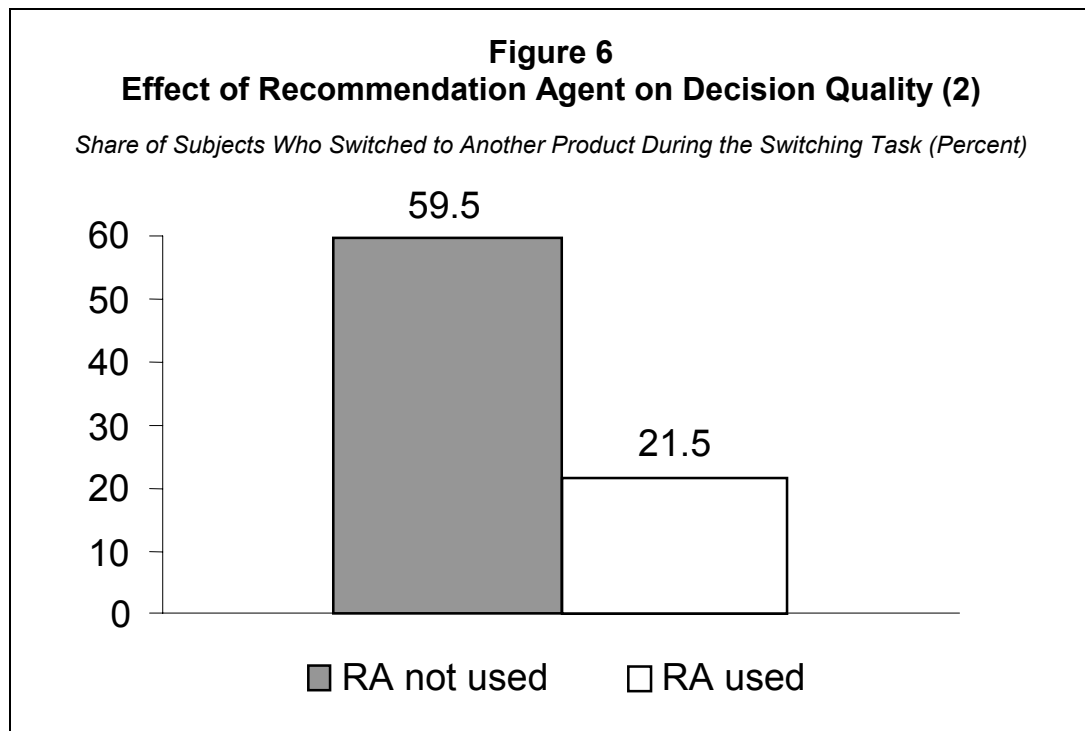



Figure 7
Preference Elicitation by the Recommendation Agent

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="text"/>
Strong Fly Fabric	<input type="text"/>

Show me the recommendation list

