

Consumer Decision Making in Online Shopping Environments: The Effects of Interactive Decision Aids

Gerald Häubl • Valerie Trifts

Faculty of Business, University of Alberta, Edmonton, Alberta, Canada T6G 2R6
gerald.haeubl@ualberta.ca • trifts@datanet.ab.ca

Abstract

Despite the explosive growth of electronic commerce and the rapidly increasing number of consumers who use interactive media (such as the World Wide Web) for prepurchase information search and online shopping, very little is known about how consumers make purchase decisions in such settings. A unique characteristic of online shopping environments is that they allow vendors to create retail interfaces with highly interactive features. One desirable form of interactivity from a consumer perspective is the implementation of sophisticated tools to assist shoppers in their purchase decisions by customizing the electronic shopping environment to their individual preferences. The availability of such tools, which we refer to as *interactive decision aids* for consumers, may lead to a transformation of the way in which shoppers search for product information and make purchase decisions. The primary objective of this paper is to investigate the nature of the effects that interactive decision aids may have on consumer decision making in online shopping environments.

While making purchase decisions, consumers are often unable to evaluate all available alternatives in great depth and, thus, tend to use two-stage processes to reach their decisions. At the first stage, consumers typically screen a large set of available products and identify a subset of the most promising alternatives. Subsequently, they evaluate the latter in more depth, perform relative comparisons across products on important attributes, and make a purchase decision. Given the different tasks to be performed in such a two-stage process, interactive tools that provide support to consumers in the following respects are particularly valuable: (1) the initial screening of available products to determine which ones are worth considering further, and (2) the in-depth comparison of selected products before making the actual purchase decision. This paper examines the effects of two decision aids, each designed to assist consumers in performing one of the above tasks, on purchase decision making in an online store.

The first interactive tool, a *recommendation agent* (RA), allows consumers to more efficiently screen the (potentially very large) set of alternatives available in an online shopping environment. Based on self-explicated information about a

consumer's own utility function (attribute importance weights and minimum acceptable attribute levels), the RA generates a personalized list of recommended alternatives. The second decision aid, a *comparison matrix* (CM), is designed to help consumers make in-depth comparisons among selected alternatives. The CM allows consumers to organize attribute information about multiple products in an alternatives \times attributes matrix and to have alternatives sorted by any attribute.

Based on theoretical and empirical work in marketing, judgment and decision making, psychology, and decision support systems, we develop a set of hypotheses pertaining to the effects of these two decision aids on various aspects of consumer decision making. In particular, we focus on how use of the RA and CM affects consumers' search for product information, the size and quality of their consideration sets, and the quality of their purchase decisions in an online shopping environment.

A controlled experiment using a simulated online store was conducted to test the hypotheses. The results indicate that both interactive decision aids have a substantial impact on consumer decision making. As predicted, use of the RA reduces consumers' search effort for product information, decreases the size but increases the quality of their consideration sets, and improves the quality of their purchase decisions. Use of the CM also leads to a decrease in the size but an increase in the quality of consumers' consideration sets, and has a favorable effect on some indicators of decision quality.

In sum, our findings suggest that interactive tools designed to assist consumers in the initial screening of available alternatives and to facilitate in-depth comparisons among selected alternatives in an online shopping environment may have strong favorable effects on both the quality *and* the efficiency of purchase decisions—shoppers can make much *better decisions* while expending substantially *less effort*. This suggests that interactive decision aids have the potential to drastically transform the way in which consumers search for product information and make purchase decisions.

(*Decision Making; Online Shopping; Electronic Commerce; Decision Aids; Recommendation Agents; Consumer Behavior; Information Search; Consideration Sets; Information Processing*)

Introduction

The popularity of interactive media such as the World Wide Web (WWW) has been growing at a very rapid pace (see, e.g., GVU 1999). From a marketing perspective, this has manifested itself primarily in two ways: (1) a drastic increase in the number of companies that seek to use the WWW to communicate with (potential) customers, and (2) the rapid adoption of the WWW by broad consumer segments for a variety of purposes, including prepurchase information search and online shopping (Alba et al. 1997). The combination of these two developments provides a basis for substantial growth in the commercial use of interactive media.

The focus of this paper is on one specific type of commercial use of interactive media: *shopping in online environments*. We conceptualize this behavior as a shopping activity performed by a consumer via a computer-based interface, where the consumer's computer is connected to, and can interact with, a retailer's digital storefront (implemented on some computer) through a network (e.g., the WWW). A consumer can engage in online shopping in any location, but our conceptualization is based on the assumptions that the products of interest are not physically present at the time and that no face-to-face assistance is available to the shopper.

A unique characteristic of online shopping environments is that they allow for the implementation of very high degrees of *interactivity*. The latter is a multidimensional construct, the key facets of which include reciprocity in the exchange of information, availability of information on demand, response contingency, customization of content, and real-time feedback (Alba et al. 1997, Ariely 2000, Zack 1993). In the context of computer-mediated communication, a distinction has been made between person interactivity and machine interactivity. While the former describes the ability to communicate with other individuals, the latter refers to the ability to interactively access information in an online database (Hoffman and Novak 1996, p. 53). Given our conceptualization of shopping in online environments, the concept of machine interactivity is of particular interest.

While it has been hypothesized that consumers' shopping behavior in online stores may be fundamentally different from that in traditional retail settings

(Alba et al. 1997, Winer et al. 1997), theorizing about the nature of these differences has been sparse. We propose that consumer behavior in an online shopping environment is determined largely by the degree and type of machine interactivity that is implemented in such a setting. Specifically, we hypothesize that the way in which consumers search for product information and make purchase decisions is a function of the particular interactive tools available in an online shopping environment. We refer to such tools as *interactive decision aids* for consumers.

In this paper, we identify two types of interactive decision aids that, in light of established notions about purchase decision processes, seem particularly valuable to consumers. Based on theoretical and empirical work in marketing, judgment and decision making, psychology, and decision support systems, we develop a set of hypotheses pertaining to the effects of each of these tools on consumers' search for product information, the size and quality of their consideration sets, and the quality of their purchase decisions in an online store. The results of a controlled experiment indicate that each of the interactive decision aids has a substantial impact on consumer decision making, thus providing a demonstration of how the availability of such tools may transform the way in which individuals search for information and make purchase decisions in online environments.

The paper is organized as follows. First, we briefly discuss the relevant literature on human decision making and decision aids. Next, we provide an overview of interactive decision aids for online shopping and motivate the choice of the two particular tools investigated in our study. We then develop a set of hypotheses pertaining to how we expect each of these decision aids to affect different aspects of consumer decision making in online shopping environments. This is followed by a description of the method used to test these hypotheses. We then report the results of our empirical study. The paper concludes with a general discussion of the findings.

Human Decision Making and Decision Aids

Humans adapt their decision making strategies to specific situations and environments (see, e.g., Payne

1982). They can be described as “cognitive misers” who strive to reduce the amount of cognitive effort associated with decision making (Shugan 1980). The notion that individuals are typically willing to settle for imperfect accuracy of their decisions in return for a reduction in effort is well supported (Bettman et al. 1990, Johnson and Payne 1985) and consistent with the idea of bounded rationality (Simon 1955). Because of this trade-off between effort and accuracy, decision makers frequently choose options that are satisfactory but would be suboptimal if decision costs were zero. This is particularly common when alternatives are numerous and/or difficult to compare, i.e., when the complexity of the decision environment is high (Payne et al. 1993).

One form of coping with highly complex decision environments is to use decision support systems. The latter are computer-based technologies designed to assist an individual (or a group) in making a decision or choosing a course of action in a nonroutine situation that requires judgment (Kasper 1996). Decision support systems contain one or more tools, or *decision aids*, that perform distinct information processing tasks or functions (e.g., search a database or sort objects by some criterion). The motivating principle underlying decision aids is that resource-intensive, but standardizable, information processing tasks are performed by a computer-based system, thus freeing up some of the human decision maker’s processing capacity. Determining an adequate “division of labor” between human and computer is crucial. Human decision makers are typically good at selecting variables that are relevant in the decision process, but weak at integrating and retaining large amounts of information. Effective decision aids should be designed to capitalize on the strengths and compensate for the inherent weaknesses of their users (Hoch and Schkade 1996).

A standard assumption in past research on decision support systems, most of which has focused on managerial decisions (e.g., Pearson and Shim 1994), is that decision makers who are provided with decision aids that have adequate information processing capabilities will use these tools to analyze problems in greater depth and, as a result, make better decisions (cf. Hoch and Schkade 1996). However, behavioral decision theory suggests that because feedback on effort expenditure tends to be immediate while feedback on accuracy

is subject to delay and ambiguity, decision makers may be inclined to focus more on reducing cognitive effort than on improving decision accuracy (Einhorn and Hogarth 1978, Kleinmuntz and Schkade 1993). Thus, decision aids may lead individuals to merely reduce effort without improving the quality of their decisions. In fact, there is empirical evidence that the use of decision aids does not necessarily enhance decision making performance (cf. Benbasat and Nault 1990), and that the latter may even be reduced as a result (Todd and Benbasat 1992, p. 373). Given this mixed evidence, it cannot be assumed that a consumer’s use of interactive decision aids in an online shopping context will lead to increased decision quality. Rather, this represents an open question, which is addressed in this paper.

In the following section, we first provide a general overview of interactive decision aids available to consumers for the purpose of online shopping. Based on established notions about purchase decision making and on characteristic features of online shopping environments, two decision aids are selected for inclusion in our empirical study. These two tools are then discussed in detail.

Interactive Decision Aids for Online Shopping

Overview of Tools

The technology available for implementing machine interactivity in online shopping environments has the potential to provide consumers with unparalleled opportunities to locate and compare product offerings (Alba et al. 1997, p. 38). Such capabilities are particularly valuable given that online stores cannot offer physical contact with products, do not allow face-to-face interaction with a salesperson, and may offer a very large number of alternatives because of their virtually infinite “shelfspace,” i.e., the lack of physical constraints with respect to product display.

Interactive decision aids that may be of use to consumers who wish to shop online include a wide variety of software tools, ranging from general-purpose search engines (e.g., www.infoseek.com, www.lycos.com) to sophisticated agent-mediated electronic commerce

systems (e.g., compare.net, www.jango.com). A common classification of interactive shopping agents is based on whether a tool is designed to help a consumer determine (1) *what* to buy or (2) *whom* to buy from. These two tasks may be referred to as *product brokering* and *merchant brokering*, respectively (see Guttman et al. 1998). For the purpose of this paper, we confine our attention to the former.

Among tools for product brokering, a distinction can be made between decision aids that operate *within* a particular merchant's online store (e.g., www.personalogic.com) and ones that operate *across* merchants (e.g., www.shopper.com). The primary focus of this paper is on the former. The decision aids we investigate are implemented within an online store (see below). However, this research *also* pertains to those cross-merchant decision aids that allow shoppers direct access to a common product database (e.g., www.jango.com), provided that these tools do not discriminate between products on the basis of which vendor they are associated with.

A well-known phenomenon regarding decision making in complex environments is that individuals are often unable to evaluate all available alternatives in great depth prior to making a choice (Beach 1993). Instead, they tend to use *two-stage* processes to reach their decisions, where the depth of information processing varies by stage (Payne 1982, Payne et al. 1988). In the context of purchase decision making, a typical two-stage process may unfold as follows. First, the consumer screens a large set of relevant products, without examining any of them in great depth, and identifies a subset that includes the most promising alternatives. Subsequently, s/he evaluates the latter in more depth, performs comparisons across products on important attributes, and makes a purchase decision. Given the different tasks to be performed in the course of such two-stage purchase decision processes, interactive tools that provide support to consumers in the following two respects seem particularly valuable: (1) the initial screening of available products to determine which ones are worth considering further, and (2) the in-depth comparison of selected products before making the actual purchase decision. We focus on two decision aids, each designed to assist consumers in performing one of these key tasks. The two interactive tools are discussed in turn.

Recommendation Agent: A Tool for Screening Alternatives

We conceptualize a *recommendation agent* (RA) as an interactive decision aid that assists consumers in the initial screening of the alternatives that are available in an online store. Based on information provided by the shopper regarding his/her own preference, an RA "recommends" a set of products that are likely to be attractive to that individual. Elementary forms of this type of decision aid are currently implemented on a number of online retail sites (e.g., www.macys.com, www.netmarket.com). A real-world tool that corresponds very closely to our conceptualization of RA is the consumer decision guide developed by PersonaLogic™ (www.personalogic.com).

The RA used in the present study generates a personalized list of recommended alternatives, in which alternatives are described by their brand and model name.¹ This recommendation is based on three types of parameters provided by the consumer. First, a consumer's self-explicated attribute importance weights are used to compute a summary score for each alternative as the sum over all products of (standardized) attribute level scale value and corresponding importance weight.² This score determines the order of alternatives in the RA's output. Thus, the RA is effectively an automated implementation of a weighted additive evaluation rule (Payne et al. 1993). Second, the RA allows consumers to specify minimum acceptable attribute levels, and only alternatives that meet all such specifications are included in the personalized list. This corresponds to an automated implementation of a conjunctive decision rule (see Wright 1975). Finally, the RA allows shoppers to impose a quota cut-off (Feinberg and Huber 1996), i.e., to limit the number of products to be included in the list.

Comparison Matrix: A Tool for Organizing Product Information

The second decision aid we examine, a *comparison matrix* (CM), is conceptualized as an interactive tool that assists consumers in making in-depth comparisons

¹From there, detailed information about a product may be requested by clicking on its model name.

²These scores are imperfect, approximate indicators of an alternative's (unknown) true utility to a consumer.

among those alternatives that appear most promising based on the initial screening. The CM allows shoppers to organize attribute information about multiple products. Very basic forms of this type of decision aid, usually referred to as a shopping cart or basket, are implemented on many online retail sites (e.g., www.amazon.com, www.shopping.com). Most of these tools do not currently allow for side-by-side comparisons of products in terms of their attributes. However, one real-world comparison aid that does closely match the above definition of CM is available on CompareNet's site (compare.net).

The CM used in the present study is implemented as an interactive display format in which product information is presented in an alternatives (rows) \times attributes (columns) matrix. It is designed to enable shoppers to compare products more efficiently and accurately. While viewing detailed information about an alternative in the online shopping environment, a consumer can choose to have the product added to his/her personal CM. (Once included, alternatives may also be deleted from the CM.) The display format is interactive in that a shopper can have all products in the CM sorted by any attribute. Use of this decision aid should result in a shift in emphasis from memory-based to stimulus-based purchase decisions in the sense that retaining specific attribute information about relevant alternatives in memory becomes less important (see Alba et al. 1997).

Hypotheses

Dependent Variables

We expect that use of the RA and the CM will have an impact on three general aspects of consumer decision making in an online shopping environment: (1) amount of information search, (2) consideration sets, and (3) decision quality.

Amount of search for product information is conceptualized as the number of alternatives for which detailed information is acquired (Moorthy et al. 1997). In our study, this coincides with the number of pages containing attribute information about a particular product that are viewed. This is an indicator of the effort an individual expends to screen available alternatives.

Consideration set is conceptualized as the set of alternatives that a consumer considers seriously for purchase (Hauser and Wernerfelt 1990).³ We use both the *size* and the *quality* of this set as dependent variables. The former is simply the number of products considered seriously, which can be viewed as an indicator of a shopper's relative product uncertainty when making a purchase decision. Consideration set quality is conceptualized as the share of considered products that are "non-dominated," i.e., not objectively inferior to any alternative (see the Method section for details).

Decision quality is measured using both objective and subjective indicators. This concept can be defined by basic principles of coherence, such as not selecting dominated alternatives (Payne et al. 1993). Thus, one indicator of objective decision quality is whether or not a consumer *purchases a nondominated alternative*.⁴ Our second measure of objective decision quality is whether or not a shopper, after making a purchase decision, changes his/her mind and *switches to another alternative* when given an opportunity to do so. Switching indicates poor initial decision quality (see the Method section for details). Finally, subjective decision quality is conceptualized as the consumer's *degree of confidence* in a purchase decision.

³An alternative conceptualization of *consideration set* is to view it as a dynamic construct that evolves over time as products are being added to and dropped from the set (e.g., Nedungadi 1990). Within this dynamic framework, our conceptualization of consideration set corresponds to the *final* consideration set, i.e., the set of alternatives considered at the time the actual purchase decision is made.

⁴Measuring the *quality* of purchase decisions and consideration sets is a very ambitious endeavor. In this context, quality is conceptualized as the degree of match or fit between heterogeneous consumer preferences and differentiated products. Because an individual's preferences are not subject to direct observation, it is impossible to accurately measure decision quality in uncontrolled real-world settings. The measurement approach used in the present study is based on the idea of an objective standard for quality and requires a combination of objectively *dominated* and *nondominated* alternatives. The sets of available products were constructed in such a way that, irrespective of an individual's utility function, the purchase of particular alternatives indicates (with certainty) that s/he made a poor decision. Choices of dominated alternatives in our controlled study are the equivalent of real-world purchase decisions that are suboptimal given an individual's utility function at the time of purchase and the set of available products, irrespective of whether or not any of the alternatives are objectively dominated.

We present a set of hypotheses about how these six aspects of consumer decision making are affected by use of each of the two interactive decision aids, RA and CM. All hypotheses are stated in terms of the expected difference, everything else being equal,⁵ between a scenario in which one of these tools is used in an online shopping encounter and a case where it is not. The base case in which neither tool is available to a shopper corresponds to a typical, "bare-bones" online store.

Effects of Using the Recommendation Agent

Amount of Search for Product Information. The amount of search for product information is determined by consumers' uncertainty about the absolute utility associated with an alternative and about the relative utility of alternatives in a set (Moorthy et al. 1997, Ratchford and Srinivasan 1993). In an online shopping environment, the amount of information search is also dependent upon the consumer's ability to screen information effectively (Alba et al. 1997, Bakos 1997). Because the RA automatically sorts available products based on criteria provided by the shopper, the latter is better able to determine the relative utility of alternatives, and this should, in turn, lead to a reduction in the amount of search (Moorthy et al. 1997). Thus, we hypothesize that individuals who have this tool to assist them in their shopping task will view attribute information about fewer products than those who do not.

HYPOTHESIS H1. *Use of the recommendation agent leads to a reduction in the number of alternatives for which detailed product information is viewed.*

Consideration Set Size. Models of consideration set size are typically based on the notion of a trade-off between the marginal benefits and costs of considering an additional alternative (e.g., Hauser and Wernerfelt 1990, Roberts and Lattin 1991). These models assume that a product's utility is unknown to the consumer prior to evaluation. However, because the RA screens and ranks alternatives based on consumer-specified criteria, it provides information about the relative utility of available products *prior to* processing specific

product information. As a result, the marginal benefits of including additional products in the consideration set diminishes much more rapidly than in a situation where the consumer has no prior information about the relative utility of alternatives. Therefore, we expect that individuals who use the RA will have smaller consideration sets than those who do not.

HYPOTHESIS H2. *Use of the recommendation agent leads to a reduction in the number of alternatives considered seriously for purchase.*

Consideration Set Quality. Because the RA uses self-explicated attribute importance weights and minimum acceptable attribute levels to produce a personalized list of recommended alternatives, the products with the highest subjective utility will tend to appear towards the top of a shopper's list. Therefore, consumers should be less likely to consider inferior alternatives for purchase. In addition, consideration sets are more likely to be composed of products with similar utility values than products with dissimilar ones (Lehmann and Pan 1994). Therefore, we expect that the share of alternatives included in the consideration set that are nondominated will be greater when the RA is used than when it is not.

HYPOTHESIS H3. *Use of the recommendation agent leads to a larger share of nondominated alternatives in the set of alternatives considered seriously for purchase.*

Decision Quality. The RA enables shoppers to screen products using complex decision rules with very low effort. Research on decision support systems indicates that decision aids designed to screen large numbers of alternatives may reduce decision makers' cognitive effort (Todd and Benbasat 1994) and improve decision quality by enabling individuals to make complex decisions with high accuracy (Singh and Ginzberg 1996). By applying decision rules in an automated fashion, such tools can reduce the amount of superfluous information to be processed and, thus, augment human information processing capabilities. In addition, the ability to screen alternatives in an efficient manner enhances the "quality" of the information that is processed, which, combined with reduced information quantity, should have a positive impact on decision quality (Keller and Staelin 1987, 1989). Finally,

⁵For brevity, we suppress the statement "ceteris paribus" in all our hypotheses.

Widing and Talarzyk's (1993) findings suggest that electronic decision formats based on weighted average scores for alternatives lead to less switching after initial choice. Thus, we hypothesize that consumers' use of the RA will have the following effects on the three indicators of decision quality.

HYPOTHESIS H4. *Use of the recommendation agent leads to an increased probability of a nondominated alternative being selected for purchase.*

HYPOTHESIS H5. *Use of the recommendation agent leads to a reduced probability of switching to another alternative (after making the initial purchase decision).*

HYPOTHESIS H6. *Use of the recommendation agent leads to a higher degree of confidence in purchase decisions.*

Effects of Using the Comparison Matrix (CM)

Amount of Search for Product Information. Through its capability for organizing information, the CM allows consumers to more efficiently compare and determine the relative attractiveness of alternatives. When searching for detailed product information, shoppers who have access to the CM will anticipate being able to subsequently use this tool to make accurate side-by-side comparisons of products and, therefore, tend to initially acquire information about a larger number of alternatives. If a product appears attractive at first glance, the consumer can add it to the CM, evaluate it in direct comparison with other alternatives, and then decide whether or not to retain it in the matrix. Because the CM facilitates stimulus-based, as opposed to memory-based, comparisons (see Alba et al. 1997), it reduces the combined marginal cost of acquiring and processing attribute information about an alternative. Therefore, we expect that individuals who use this tool will view information about more products than those who do not.

HYPOTHESIS H7. *Use of the comparison matrix leads to an increase in the number of alternatives for which detailed product information is viewed.*

Consideration Set Size. While availability of the CM should increase the amount of search (see H7), we expect that once consumers actually take advantage of this decision aid's comparison-facilitating capabilities, they will be able to more quickly and more

accurately eliminate unwanted products from their consideration set. Decision aids that help organize information have been found to reduce the number of alternatives considered by decision makers (Goslar et al. 1986). The CM improves consumers' ability to both determine their personal efficient frontiers and identify dominated alternatives (Winer et al. 1997). As a result, use of the CM reduces relative product uncertainty and, thus, the marginal benefit of including an additional product in the consideration set (Hauser and Wernerfelt 1990, Roberts and Lattin 1991). Thus, we hypothesize that consumers who use the CM will seriously consider fewer alternatives for purchase than those who do not.

HYPOTHESIS H8. *Use of the comparison matrix leads to a reduction in the number of alternatives considered seriously for purchase.*

Consideration Set Quality. The CM's alternatives \times attributes format facilitates side-by-side comparisons of products in terms of their attributes. This display format, in conjunction with the CM's capability for sorting all selected alternatives by any attribute, reduces the demand on memory and improves consumers' ability to identify suboptimal alternatives (see Payne et al. 1993, Winer et al. 1997). Because the CM allows for efficient discrimination between products with respect to their subjective overall utility, it will render consumers less likely to either eliminate excellent alternatives from or retain inferior alternatives in their consideration set. Therefore, we expect that consideration set quality will be higher for shoppers who use the CM than for those who do not.

HYPOTHESIS H9. *Use of the comparison matrix leads to a larger share of nondominated alternatives in the set of alternatives considered seriously for purchase.*

Decision Quality. The findings of numerous studies suggest that the way in which information is displayed influences decision processes by affecting the ease of carrying out different processing operations (see Kleinmuntz and Schkade 1993). Because decision makers generally try to conserve cognitive effort, they tend to use processing strategies that are facilitated by a given display format (e.g., Russo 1977). The CM enhances consumers' ability to compare products in

terms of their attributes (Alba et al. 1997). As a result, use of this tool should lead to a shift in emphasis from memory-based to stimulus-based choice. The latter has been found to result in a reduced probability of an inferior alternative being chosen (Muthukrishnan 1995, Exp. 2). In addition, information display formats that reduce task difficulty have been found to lower the frequency of preference reversals (Johnson et al. 1988). The latter may be viewed as indicators of suboptimal choice. Thus, we hypothesize that use of the CM will have the following effects on the three indicators of decision quality.

HYPOTHESIS H10. *Use of the comparison matrix results in an increased probability of a nondominated alternative being selected for purchase.*

HYPOTHESIS H11. *Use of the comparison matrix leads to a reduced probability of switching to another alternative (after making the initial purchase decision).*

HYPOTHESIS H12. *Use of the comparison matrix leads to a higher degree of confidence in purchase decisions.*

Method

A controlled experiment was conducted to test the above hypotheses about the effects of the RA and the CM on the six dependent measures of interest. The main task consisted of shopping for and making a purchase of a product in each of two categories—backpacking tents and compact stereo systems—in an online store. In this section, we discuss (1) the experimental design of the study, (2) the modeling approach, (3) the sample and incentive, and (4) the experimental procedure.

Experimental Design

A 2⁴ full-factorial experimental design was used. The manipulated factors are: RA (yes, no), CM (yes, no), product category (backpacking tent, compact stereo system), and product category order (tent first, stereo first). While product category is a within-subjects factor, RA, CM, and order were manipulated between subjects. Respondents were randomly assigned to one of the eight conditions of the 2³ between-subjects subdesign.

For each of the two product categories, a total of 54

alternatives were constructed (9 models for each of 6 brands). Six actual brand names were used in each product category. All model names were fictitious but representative of the respective category. Each alternative was described on seven attributes in addition to brand and model name. Five of these attributes were varied systematically across alternatives, while two attributes were held constant. The varied attributes that were varied are (number of levels in parentheses): pole material (3), warranty (3), weight (12), durability rating (12), and price (12). Fly fabric and vestibule were held constant across alternatives. For stereos, the varied attributes are: CD player type (3), tuner presets (3), output power (12), sound quality rating (12), and price (12). Cassette decks and remote control were held constant.⁶

The measurement of consideration set quality and of two aspects of decision quality requires alternatives that are known to be *nondominated*. For each product category, six nondominated alternatives—one for each brand—were constructed. That is, 6 of the 54 products were mutually nondominated. They did, however, dominate all remaining models.⁷ Having one nondominated alternative for each brand guaranteed that, irrespective of an individual's relative preference for brand names, one of the nondominated products was the single most preferred alternative. The six nondominated alternatives were constructed by first assigning to them the best level of the two three-level attributes. Next, for the three attributes with 12 levels, all six were assigned the best level of one, the second best of another, and the third best of the remaining attribute. All possible combinations of first, second, and third best were used. The two best levels of the 12-level attributes were reserved for the nondominated alternatives. The remaining 48 products were constructed by means of

⁶The exact descriptions of all 54 models of backpacking tents are provided in Appendix A. The corresponding information for compact stereo systems is available from the authors upon request.

⁷An 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. That is, a dominated alternative is known to be within the efficient frontier of any consumer. 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.

an iterative algorithm that approximated a target matrix of plausible across-attribute correlations, while adhering to the required pattern of (non-)dominance among alternatives.

Modeling Approach

We use *Generalized Estimating Equations* (GEE) models (Diggle et al. 1995, Liang and Zeger 1986) to test our hypotheses. This modeling technology generalizes classical linear models in two ways, both of which are essential to our study. We briefly discuss each of these extensions in turn.

First, GEE models can accommodate a variety of response distributions in addition to the common normal (Gaussian) distribution. Given the different response types used in the present study, this capability is required for the proper modeling of our dependent variables. Within the GEE framework, the relationship between a dependent variable and a set of predictors is expressed as

$$g(E(Y|\mathbf{x})) = \beta_0 + \sum_{i=1}^P \beta_i x_i, \quad (1)$$

where Y is the dependent variable and $\mathbf{x} = (x_1, \dots, x_p)$ are the values of a set of predictor variables X_1, \dots, X_p . The intercept β_0 and a coefficient β_i for each predictor are estimated. The *link function* g , which may be any monotonic differentiable function, allows nonlinear relationships between predictor and outcome variables (McCullagh and Nelder 1989).

In addition, GEE models relax the assumption that responses have independent distributions with constant variance. In particular, the variance of the dependent variable can be specified as a function of the mean response $E(Y|\mathbf{x})$ via a *variance function* V such that

$$\text{var}(Y) = V(E(Y|\mathbf{x}))\phi, \quad (2)$$

where ϕ is a scale parameter. The link and variance functions allow a wide range of non-normal response distributions, including binomial, Poisson, and gamma (Zeger and Liang 1986).

The second generalization of classical linear models reflected in GEE models is the relaxation of the assumption of independence among observations, which allows a more adequate modeling of data that follow

a hierarchical sampling pattern or are otherwise clustered by design (Liang and Zeger 1986). Because product category is a within-subject factor in our experiment, the responses for each dependent variable are clustered (by respondent) rather than independent. Thus, the ability to account for systematic relationships among multiple observations for an individual is essential to our study. For each dependent variable, a “working correlation” between the two responses of an individual is estimated as a free parameter (Zeger and Liang 1986). This is required for an adequate modeling of the data, although the working correlations are not of substantive interest here.

Two of our dependent measures—amount of search and consideration set size—are based on count data with no effective upper limit and are properly treated as following a Poisson distribution. In our GEE models, this is implemented by specifying $g(\mu) = \log(\mu)$ as the link function and $V(\mu) = \mu$ as the variance function, where $\mu = E(Y|\mathbf{x})$. Two other dependent variables—choice of a nondominated alternative and switching—are binary. In addition, consideration set quality is measured as a fraction based on a set of binary responses. A binomial distribution is adequate for these three response variables. This is handled by using $g(\mu) = \log(\mu/(1 - \mu))$ and $V(\mu) = \mu(1 - \mu)$ as the GEE model’s link and variance functions, respectively. Finally, confidence in purchase decisions is measured on a nine-point rating scale and can thus be treated as standard Gaussian (i.e., using a GEE model with identity link function and constant variance).

Sample and Incentive

Eighty undergraduate psychology students participated (for partial course credit) in a pilot study aimed at testing the validity of the experimental manipulations. All manipulations were successful, and no technical problems were encountered with respect to the electronic shopping environment. For the main study, 249 undergraduate business students completed the online shopping task for both product categories.⁸ In addition to partial course credit, subjects in the main study participated in a lottery designed to increase the

⁸The number of respondents per cell in the 2³ between-subjects sub-design ranged from 29 to 33. All subjects had prior experience using the WWW, as did all participants in the pilot study.

validity of the findings by making the shopping task more consequential. Prior to entering the electronic shopping environment, subjects were informed that two randomly selected winners were to receive one of the two products they “purchased” during their shopping trip, plus the difference between \$500 and the price of that product in cash. One tent and one stereo system were dispensed. Because the alternatives used in the study were constructed (see above), the two winners received the real-world model that best matched their chosen alternative (i.e., same brand name and similar attribute levels). Each prize’s cash component was based on the chosen model’s price used in the shopping environment.

Procedure

Data were collected in a university computer lab in sessions of 15 to 20 subjects. Upon arrival, participants were assigned to a personal computer and informed that they would be pilot-testing a new online store by shopping for a product in each of two categories. The experimenter then held a 10-minute practice session aimed at demonstrating the features of the shopping environment.

Before shopping for the first product, subjects were asked to rate their level of product category knowledge and interest (using nine-point rating scales). They then read a description of the task and of the lottery incentive. Subjects in the No-RA conditions were taken to a hierarchically structured website with all six brands listed at the top level and all models for a brand listed at the lower level. Subjects accessed detailed information about a product by first clicking on a brand name and then on a model name. In the conditions in which the RA was available,⁹ subjects started by providing attribute importance weights using a 100-point constant sum scale, minimum-acceptable attribute levels, and the maximum number of alternatives to be included in their personalized recommendation list. From that list, they were able to request detailed information about a particular product. Subjects who used the CM were able to add the attribute information displayed on a product’s page to the CM, from where they eventually made their purchase. Subjects in the

No-CM conditions made their purchase from one of the individual product pages. In all conditions, respondents were asked to confirm the product of their choice before the purchase was finalized.

After subjects made their first purchase, a measure of confidence in their purchase decision (“How confident are you that the product you just purchased is really the best choice for you?”) was obtained using a nine-point rating scale. Next, they were presented with the list of alternatives and asked to report their consideration set (“Please indicate which of the products you considered seriously before making your purchase decision.”). Subjects could then switch from the purchased alternative to each of six (five) nondominated alternatives.¹⁰ This switching task was presented as a series of pairwise comparisons in which complete descriptions of both products were displayed side by side.¹¹ Subjects were encouraged to switch whenever they saw an alternative they preferred over their initial choice. They were informed that the lottery winners would receive whatever product they had “in their basket” after the switching task. The same procedure was repeated for the second product category. After that, subjects completed a questionnaire containing manipulation checks. Finally, the administrator debriefed the participants and concluded the session.

Results

Manipulation Checks

To verify that the experimental manipulations were successful, subjects responded to manipulation-check questions after completion of their second shopping trip. First, they expressed (using a nine-point rating scale) how difficult it was for them to locate the products that best matched their personal preferences. This was used to check the RA manipulation. The mean ratings obtained from the RA and No-RA conditions are 2.95 and 4.67, respectively. This difference in means is highly significant ($p < 0.001$, $\omega^2 = 0.178$) and in the

¹⁰The nondominated products used for this purpose were identical to the ones used during the shopping task. The number of switching opportunities depended upon whether a subject had initially chosen a dominated (6) or nondominated alternative (5).

¹¹This switching task is similar to a method used by Widing and Talarzyk (1993).

⁹In the RA conditions, *all* subjects actually used this decision aid. The same is true for the CM conditions.

Table 1 Model Results: Coefficient Estimates

Predictors	Dependent Variables					
	Amount of Search	Consideration Set Size	Consideration Set Quality	Purchase of Nondominated Alternative	Switching	Confidence in Purchase Decision
Intercept	2.158 ^a (0.032) ***	1.054 (0.027) ***	0.469 (0.098) ***	1.686 (0.182) ***	-0.441 (0.109) ***	6.554 (0.086) ***
Recommendation Agent (RA)	-0.574 (0.063) ***	-0.094 (0.054) *	2.362 (0.199) ***	2.086 (0.359) ***	-1.728 (0.218) ***	0.314 (0.173) *
Comparison Matrix (CM)	0.022 (0.063)	-0.543 (0.054) ***	0.970 (0.195) ***	0.497 (0.358)	-0.452 (0.218) *	-0.068 (0.173)
RA × CM Interaction	-0.191 (0.126)	0.282 (0.109) **	0.455 (0.391)	1.114 (0.717)	-0.094 (0.436)	-0.519 (0.346)
Product Category	0.109 (0.035) **	0.085 (0.042) *	-0.487 (0.106) ***	-0.569 (0.192) xf.**	0.013 (0.185)	0.849 (0.116) ****
Order Position	0.284 (0.035) ***	-0.076 (0.043)	0.383 (0.107) ***	0.212 (0.192)	-0.177 (0.185)	-0.291 (0.116) *

^aCell Format:
Coefficient Estimate
(Standard error)
(t value)
Level of Significance

Level of Significance:
* denotes significance at 0.05 level
** denotes significance at 0.01 level
*** denotes significance at 0.001 level

intended direction. The CM manipulation was checked by asking subjects to rate (on a nine-point rating scale) how difficult it was for them to compare different products. The mean ratings obtained from the CM and No-CM conditions are 2.80 and 4.44, respectively. This difference in means is highly significant ($p < 0.001$, $\omega^2 = 0.176$) and in the intended direction. We conclude that both of our manipulations were successful.

Hypothesis Tests

To test the hypotheses regarding the effects of the RA and CM, a GEE model was estimated for each of the six dependent variables.¹² In addition to an intercept, the following predictor variables were included in

these models: main effects for RA, CM, product category, and order position, plus an RA × CM interaction effect. The RA and CM main effects are of primary substantive interest. Product category, order position, and the RA × CM interaction are included for an adequate representation of the data and for exploratory purposes. The two levels of RA and CM were coded -0.5 (not used) and 0.5 (used). Product category was coded -0.5 (tent) and 0.5 (stereo), and order position was coded -0.5 (first product category) and 0.5 (second product category).¹³ The RA × CM interaction

brary function developed by V. J. Carey, Department of Biostatistics, Harvard University.

¹³In models that include interaction terms in the form of products of

¹²All GEE models were estimated using the S-Plus statistical analysis and programming environment (MathSoft, Inc. 1998) and a GEE li-

was treated as the product of the RA and CM predictor variables.

Table 1 provides an overview of the model results. Each column contains the GEE coefficient estimates with respect to one of the six dependent variables. Standard errors are in parentheses. The level of statistical significance of a coefficient is indicated by asterisks.¹⁴ The tests of our hypotheses are based on the coefficients for the RA and CM main effects. Figures 1 through 6 contain the cell means, percentages, or ratios of the dependent variables as a function of whether or not the RA and CM were used. We discuss the effects of the RA and the CM in turn.

The effect of the RA on the amount of search for product information is highly significant ($p < 0.001$) and in the expected direction. As shown in Figure 1, subjects viewed detailed product information for substantially fewer alternatives when the RA was used (6.58 on average) than when it was not (11.78). This provides strong support for H1. As expected, use of the RA also led to smaller consideration sets. The average number of alternatives considered seriously for purchase was 2.78 in the RA conditions and slightly above 3 in the No-RA conditions (See Figure 2). This effect is significant at $p < 0.05$ and supports H2. While reducing consideration set size, use of the RA resulted in a drastic increase in consideration set quality. The average share of the alternatives considered seriously for purchase that were nondominated was 0.85 when the RA was used and 0.42 when it was not (see Figure 3). This effect is highly significant ($p < 0.001$) and provides strong support for H3.

Use of the RA had the following effects on the three

categorical predictor variables, the coefficient estimate and statistical test for a predictor that is included in an interaction term are *not* invariant to the coding of other predictors included in the same interaction term (see Irwin and McClelland 2000). We use standardized mean-centered coding for all main effects. As a result of the mean-centering, all main and interaction effects (i.e., coefficients) are with respect to the average of the two levels of other factors. As a result of standardization, all main effects are expressed in terms of the difference between the two levels of a factor.

¹⁴Because we are testing directional hypotheses regarding the effects of the RA and CM, the level of significance of these effects is based on one-tailed t values. For all other effects, two-tailed t values are used.

Figure 1 Effects of RA and CM on Amount of Search Within Online Store

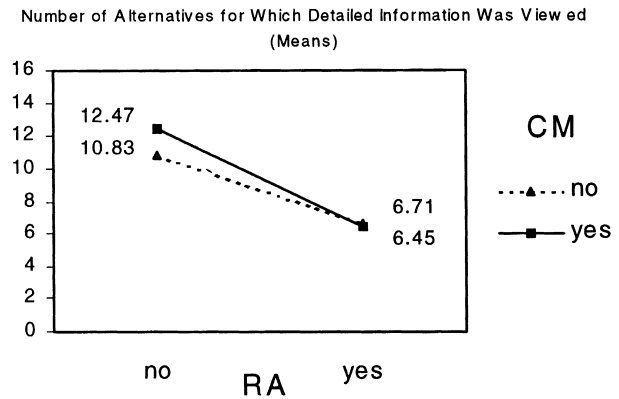


Figure 2 Effects of RA and CM on Consideration Set Size

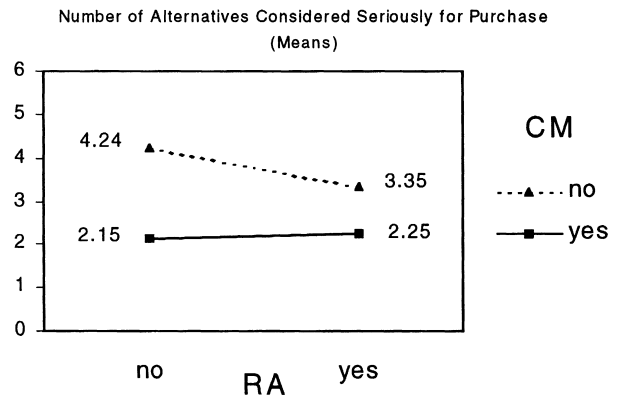


Figure 3 Effects of RA and CM on Consideration Set Quality

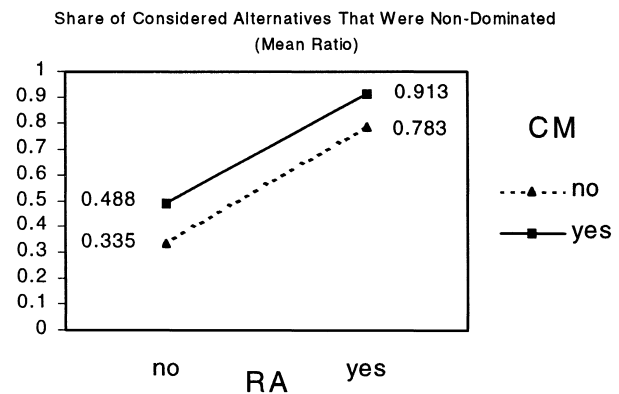


Figure 4 Effects of RA and CM on Purchase of Nondominated Alternative

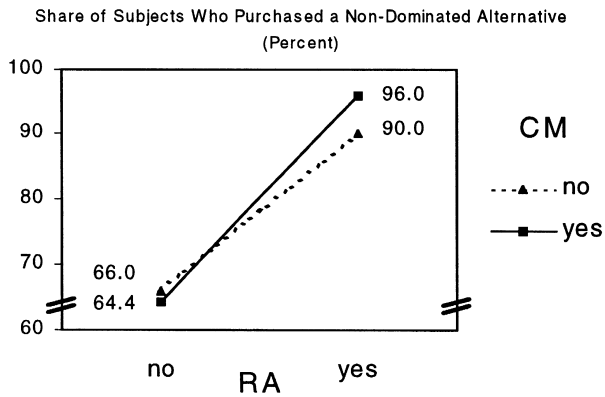
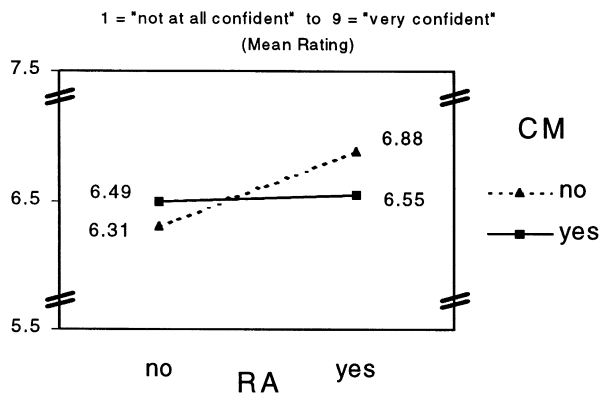


Figure 5 Effects of RA and CM on Switching During Post-Purchase Switching Task



Figure 6 Effects of RA and CM on Confidence in Purchase Decision



measures of decision quality.¹⁵ The share of subjects who purchased a nondominated alternative was about 93% when the RA was used, and about 65% when it was not (see Figure 4). This highly significant result ($p < 0.001$) supports H4. The proportion of subjects who switched to another alternative when given an opportunity to do so during the post-purchase switching task was only slightly above 20% in the RA conditions but about 60% among those subjects who did not use the RA (see Figure 5). This effect is also highly significant ($p < 0.001$) and provides strong support for H5. Finally, use of the RA resulted in a significant ($p < 0.05$) improvement in consumers' confidence in their choice (means of 6.71 vs. 6.41) as predicted by H6 (see Figure 6). This effect of the RA on the subjective measure of decision quality is noticeably weaker than its effect on the two objective measures. In sum, use of the RA reduces search effort for product information, decreases the size but increases the quality of consideration sets, and improves the quality of purchase decisions.

We now turn to the effects of the CM. As predicted, use of the CM led to an increase in the average number of alternatives for which subjects viewed detailed product information. However, this effect is not significant ($p > 0.1$), and thus H7 is not supported. By contrast, the CM's impact on consideration set size is very substantial. The average number of alternatives considered seriously for purchase was 2.19 when respondents used the CM and 3.77 when they did not (see Figure 2). This effect is highly significant ($p < 0.001$) and provides strong support for H8. In addition, the average share of the alternatives considered for purchase that were nondominated was about 0.68 when the CM was used and about 0.57 otherwise (see

¹⁵The pairwise Spearman correlation coefficients among these three measures are $\rho = -0.44$ ($p < 0.001$) between purchase of nondominated alternative and switching, $\rho = 0.09$ ($p < 0.05$) between purchase of nondominated alternative and confidence, and $\rho = -0.12$ ($p < 0.01$) between switching and confidence. All three correlations are in the expected direction. The fact that the pairwise correlations between each of the two objective measures and the subjective one are of small magnitude shows that the two types of indicators are not merely redundant with each other and that obtaining objective measures of decision quality provides valuable insight above and beyond what subjective indicators may reveal.

Figure 3). This favorable effect of the CM on consideration set quality is also highly significant ($p < 0.001$), thus supporting H9.

Consistent with H10, use of the CM led to an increase in the share of subjects who purchased a non-dominated alternative. This effect is only marginally significant ($p < 0.1$). As predicted by H11, the proportion of subjects who switched to another alternative when given an opportunity to do so during the post-purchase switching task was lower when the CM was used (38%) than when it was not (44%) (see Figure 5). This positive effect of the CM on decision quality is significant at $p < 0.05$. Finally, we do not find any support for H12—the CM did not reliably affect subjects' confidence in their choice. In sum, use of the CM leads to a decrease in the size but an increase in the quality of consideration sets, and it tends to have a favorable effect on objective decision quality.

Other Results

In addition to the tests of our hypotheses, several other results are of interest. First, interaction effects between RA and CM were included in the models discussed above. For one of the six dependent variables, consideration set size, this interaction is statistically significant ($p < 0.01$). Follow-up tests suggest that use of the RA leads to a reduction in consideration set size when the CM is not available ($p < 0.01$) but has no such effect when the CM is available ($p > 0.7$).

For an adequate representation of the data, main effects for product category and order position were included in the models. While we have no substantive interest in these two factors, we note that each affects some of the dependent variables. Product category appears to have an impact on five of the outcomes. However, inclusion of subjects' knowledge and interest with respect to the category as predictors renders the effects of product category on all but two dependent variables (consideration set quality, purchase of a non-dominated alternative) statistically insignificant. The effects of order position may reflect increasing familiarity with the shopping environment over time.

To examine whether the effects of the RA and CM were moderated by product category, order position, knowledge, or interest, a number of additional GEE models were estimated for each of the dependent variables. The moderating relationships were expressed

and tested in the form of interaction terms between RA, CM, or the $RA \times CM$ interaction on one hand and a potential moderating variable on the other. Each of these terms was included in a separate model that also contained main effects for RA, CM, product category, and order position, as well as an $RA \times CM$ interaction effect. None of these moderating effects were statistically significant (at $p = 0.05$) with respect to any of the six dependent variables. This suggests that the generalizability of the substantive findings across product categories is satisfactory.

Discussion

A characteristic feature of electronic shopping environments is the lack of physical constraints with respect to product display. The virtually infinite "shelfspace" available in online stores allows vendors to offer an extremely large number of alternatives within a product category. From a consumer perspective, having access to a very large number of products is highly desirable. At the same time, however, consumers have limited cognitive resources and may simply be unable to process the potentially vast amounts of information about these alternatives. A potential solution to this dilemma is to provide consumers with sophisticated interactive decision aids designed to help them effectively manage and capitalize on the enormous amounts of product information that may be available in electronic shopping environments.

The objective of the present study was to examine the effects of such interactive decision aids on various aspects of consumer decision making in an online shopping context. In particular, we focused on two tools that represent obvious choices given the well-established notion that consumers often reach purchase decisions via a two-stage process. The RA assists consumers in the initial screening of alternatives, and the CM facilitates in-depth comparisons of selected alternatives that are considered seriously for purchase. The results of our study indicate that use of these tools has a substantial impact on the amount of search for product information, the size and quality of shoppers' consideration sets, and the quality of their purchase decisions.

In a nutshell, the two interactive decision aids allow

consumers to make much *better decisions* while expending substantially *less effort*. Given the well-established notion that a trade-off between effort and accuracy is inherent to human decision making in traditional environments (Payne et al. 1993), it is interesting that tools like the RA and the CM can simultaneously increase decision quality *and* reduce effort. The findings of our study show how drastically interactive decision aids implemented in online shopping environments may transform the way in which consumers search for product information and make purchase decisions.

The present study examines the effects of interactive decision aids on consumer decision making in a particular electronic shopping setting. How well do our empirical results generalize to other types of online retail environments? While both the RA and the CM were operationalized as decision aids available within an online retailer's site, the relevance of our findings is not limited to interactive tools that are exclusive to an individual merchant. In particular, the results apply to all within- and cross-store decision aids that allow online shoppers direct access to a common database of products and that do not discriminate between products on the basis of which vendor they are associated with. To the extent that such decision aids are implemented in the context of multi-retailer online malls, cross-merchant comparison schemes, or groups of stores that allow for unrestricted cross-store comparisons, our results are of relevance.

Three important boundary conditions with respect to the present study's findings should be made explicit. First, the focus of this research is on consumers' goal-directed shopping behavior in an online environment, rather than on exploratory navigation behavior. Thus, no conclusions about the latter may be drawn based on the results reported here. Second, our findings do not pertain to situations in which a consumer has not yet selected an electronic store, online mall, or other entity with common product offerings. Hierarchical decision processes, such as first selecting one of a set of competing online merchants and subsequently selecting a product from that merchant's offerings, should be examined in future research. Finally, the RA used in the present study is a high-quality decision aid. Real-world recommendation systems may suffer from

substantial imperfections. For example, they may neglect relevant attributes, overlook some attractive alternatives entirely, or even be systematically biased in favor of a subset of products (e.g., those of a certain brand). Therefore, we do not conclude that recommendation agents will always and unconditionally lead to desirable outcomes for consumers. Rather, the effects of the RA found in the present study should be viewed as a demonstration of the potential effects of typical, well-functioning recommendation tools.

While the present study provides valuable insights into the effects of interactive decision aids on consumer decision making in online shopping environments, further research will be needed to obtain a deeper understanding of these effects. In particular, an examination of potential moderators would be valuable. Factors that might moderate the effects reported here include the number of available alternatives, the amount of risk associated with a purchase, and consumers' confidence in the integrity of the interactive decision aids.

In conclusion, the findings of the present study suggest that interactive decision aids designed to assist consumers in the initial screening of available products and to facilitate in-depth comparisons among selected alternatives may have highly desirable properties in terms of consumer decision making. Such tools allow shoppers to more easily detect products that are overpriced or otherwise dominated by competing alternatives, thus increasing market efficiency. More generally, the availability of interactive decision aids in online shopping environments should enhance the ability of individuals to identify products that match their personal preferences and, therefore, lead to substantial positive welfare effects for consumers.¹⁶

¹⁶This research was supported by a Nova Corporation Faculty Fellowship awarded to the first author by the University of Alberta, a contribution from Macromedia Inc., as well as grants from Telus Communications Inc. and the University of Alberta through its Social Science Research Program. The authors thank Terry Elrod for his valuable assistance, V. J. Carey for making available his S-Plus function for fitting GEE models, Barry Ard for his programming assistance, and Adam Finn, Julie Irwin, as well as the Editor, the Area Editor, and two anonymous *Marketing Science* reviewers for their insightful comments. Correspondence should be addressed to the first author.

Appendix A Product Descriptions: Backpacking Tents

Brand	Model	Pole Material (aluminum type)	Fly Fabric	Weight (kilograms)	Durability Rating (0 to 100 points)	Vestibule	Warranty (years)	Price (\$)
Eureka	Adventurer	regular-strength	2.3 oz. nylon	3.3	84	yes	2	329.99
Eureka	Challenger	regular-strength	2.3 oz. nylon	3.5	86	yes	2	324.99
Eureka	Drifter	high-strength	2.3 oz. nylon	3.3	89	yes	4	359.99
Eureka	Huntsman	regular-strength	2.3 oz. nylon	3.7	88	yes	3	324.99
Eureka	Mountaineer	regular-strength	2.3 oz. nylon	3.3	86	yes	3	324.99
Eureka	Naturalist	ultrahigh-strength	2.3 oz. nylon	3.3	88	yes	4	369.99
Eureka	Outfitter	high-strength	2.3 oz. nylon	3.6	89	yes	2	329.99
Eureka	Traveler	ultrahigh-strength	2.3 oz. nylon	3.3	92	yes	4	314.99
Eureka	Wanderer	high-strength	2.3 oz. nylon	3.5	86	yes	2	329.99
Kelty	Glacier Lake	ultrahigh-strength	2.3 oz. nylon	3.4	84	yes	2	324.99
Kelty	Lakeside	ultrahigh-strength	2.3 oz. nylon	4.0	89	yes	3	324.99
Kelty	Mountain Springs	high-strength	2.3 oz. nylon	4.1	91	yes	3	324.99
Kelty	Oasis	ultrahigh-strength	2.3 oz. nylon	3.5	91	yes	4	369.99
Kelty	Raging Tide	ultrahigh-strength	2.3 oz. nylon	4.2	91	yes	3	324.99
Kelty	River Rapid	ultrahigh-strength	2.3 oz. nylon	3.9	85	yes	3	324.99
Kelty	Seabreeze	ultrahigh-strength	2.3 oz. nylon	3.1	91	yes	4	319.99
Kelty	Swift Current	ultrahigh-strength	2.3 oz. nylon	3.5	89	yes	4	354.99
Kelty	Waterfall	ultrahigh-strength	2.3 oz. nylon	3.5	82	yes	3	334.99
Outbound	Galaxy	ultrahigh-strength	2.3 oz. nylon	3.3	91	yes	2	359.99
Outbound	Lunar Eclipse	ultrahigh-strength	2.3 oz. nylon	3.9	83	yes	4	324.99
Outbound	Moonscape	ultrahigh-strength	2.3 oz. nylon	3.1	92	yes	4	324.99
Outbound	Neptune	ultrahigh-strength	2.3 oz. nylon	3.6	86	yes	3	329.99
Outbound	North Star	ultrahigh-strength	2.3 oz. nylon	3.3	86	yes	3	334.99
Outbound	Skyline	ultrahigh-strength	2.3 oz. nylon	3.3	84	yes	3	329.99
Outbound	Stargazer	ultrahigh-strength	2.3 oz. nylon	3.3	90	yes	3	354.99
Outbound	Sunlight	ultrahigh-strength	2.3 oz. nylon	3.3	89	yes	3	359.99
Outbound	Westwind	high-strength	2.3 oz. nylon	4.1	89	yes	2	324.99
REI	Bear Paw	regular-strength	2.3 oz. nylon	3.6	84	yes	2	324.99
REI	Coyote	high-strength	2.3 oz. nylon	4.2	87	yes	4	324.99
REI	Eagle	high-strength	2.3 oz. nylon	3.3	82	yes	4	334.99
REI	Grizzly	ultrahigh-strength	2.3 oz. nylon	4.2	86	yes	4	324.99
REI	Mountain Lion	ultrahigh-strength	2.3 oz. nylon	3.2	91	yes	4	314.99
REI	Night Owl	ultrahigh-strength	2.3 oz. nylon	3.3	91	yes	4	369.99
REI	Raven	ultrahigh-strength	2.3 oz. nylon	3.8	89	yes	4	344.99
REI	Red Fox	ultrahigh-strength	2.3 oz. nylon	3.3	91	yes	3	359.99
REI	Timberwolf	ultrahigh-strength	2.3 oz. nylon	3.6	91	yes	4	349.99
Sierra Designs	Backtrail	regular-strength	2.3 oz. nylon	3.3	84	yes	3	324.99
Sierra Designs	Badlands	regular-strength	2.3 oz. nylon	3.8	88	yes	3	324.99
Sierra Designs	Big Country	high-strength	2.3 oz. nylon	3.4	91	yes	4	359.99
Sierra Designs	Forest Mist	high-strength	2.3 oz. nylon	3.5	82	yes	2	329.99
Sierra Designs	Landscape	high-strength	2.3 oz. nylon	3.8	89	yes	2	329.99
Sierra Designs	Mountain Range	ultrahigh-strength	2.3 oz. nylon	3.3	93	yes	4	319.99
Sierra Designs	Rocky Ridge	ultrahigh-strength	2.3 oz. nylon	3.4	90	yes	4	369.99
Sierra Designs	South Ridge	regular-strength	2.3 oz. nylon	3.6	83	yes	3	324.99
Sierra Designs	Summit	regular-strength	2.3 oz. nylon	3.3	82	yes	2	324.99
Quest	Daydream	high-strength	2.3 oz. nylon	3.9	91	yes	2	324.99
Quest	Freestyle	regular-strength	2.3 oz. nylon	4.0	88	yes	4	329.99
Quest	Great Escape	high-strength	2.3 oz. nylon	3.3	88	yes	4	369.99
Quest	Journey	ultrahigh-strength	2.3 oz. nylon	3.2	93	yes	4	324.99
Quest	Peacemaker	ultrahigh-strength	2.3 oz. nylon	4.2	91	yes	4	324.99
Quest	Serenity	high-strength	2.3 oz. nylon	3.4	89	yes	4	349.99
Quest	Solitude	regular-strength	2.3 oz. nylon	3.3	84	yes	3	324.99
Quest	Spirit	ultrahigh-strength	2.3 oz. nylon	4.0	84	yes	4	324.99
Quest	Tranquility	ultrahigh-strength	2.3 oz. nylon	3.8	91	yes	4	359.99

Note: Boldface indicates nondominated alternatives.

References

- Alba, Joseph, John Lynch, Barton Weitz, Chris Janiszewski, Richard Lutz, Alan Sawyer, Stacey Wood. 1997. Interactive home shopping: consumer, retailer, and manufacturer incentives to participate in electronic marketplaces. *J. Marketing* 61 (July) 38–53.
- Ariely, Dan. 2000. Controlling the information flow: effects on consumers' decision making and preferences. *J. Consumer Res.* Forthcoming.
- Bakos, J. Yannis. 1997. Reducing buyer search costs: implications for electronic marketplaces. *Management Sci.* 43 (12) 1676–1692.
- Beach, Lee R. 1993. Broadening the definition of decision making: the role of prechoice screening of options. *Psych. Sci.* 4 (4) 215–220.
- Benbasat, Izak, Barrie R. Nault. 1990. An evaluation of empirical research in managerial support systems. *Decision Support Systems* 6 (2) 203–226.
- Bettman, James R., Eric J. Johnson, John W. Payne. 1990. A componential analysis of cognitive effort in choice. *Organ. Behavior Human Decision Processes* 45 111–139.
- Diggle, Peter J., Kung-Yee Liang, Scott L. Zeger. 1995. *Analysis of Longitudinal Data*. Clarendon Press, Oxford, UK.
- Einhorn, H., R. Hogarth. 1978. Confidence in judgment: persistence of the illusion of validity. *Psych. Rev.* 85 395–416.
- Feinberg, Fred M., Joel Huber. 1996. A theory of cutoff formation under imperfect information. *Management Sci.* 42 (1) 65–84.
- Goslar, Martin D., Gary I. Green, Terry H. Hughes. 1986. Decision support systems: an empirical assessment for decision making. *Decision Sci.* 17 79–91.
- Guttman, Robert H., Alexandros G. Moukas, Pattie Maes. 1998. Agent-mediated electronic commerce: a survey. *Knowledge Engrg. Rev.* 13 (2) 147–160.
- GVU. 1999. *GVU's WWW User Survey*. Georgia Institute of Technology: Graphic, Visualization, and Usability Center [http://www.gvu.gatech.edu/user-surveys].
- Hauser, John R., Birger Wernerfelt. 1990. An evaluation cost model of consideration sets. *J. Consumer Res.* 16 (March) 393–408.
- Hoch, Stephen J., David A. Schkade. 1996. A psychological approach to decision support systems. *Management Sci.* 42 (1) 51–64.
- Hoffman, Donna L., Thomas P. Novak. 1996. Marketing in hypermedia computer-mediated environments: conceptual foundations. *J. Marketing* 60 (July) 50–68.
- Irwin, Julie R., Gary H. McClelland. 2000. Heuristics for moderated regression models. *J. Marketing Res.* forthcoming.
- Johnson, Eric J., John W. Payne. 1985. Effort and accuracy in choice. *Management Sci.* 31 394–414.
- , ———, James R. Bettman. 1988. Information displays and preference reversals. *Organ. Behavior Human Decision Process.* 42 1–21.
- Kasper, George M. 1996. A theory of decision support system design for user calibration. *Inform. Systems Res.* 7 (2) 215–232.
- Keller, Kevin L., Richard Staelin. 1987. Effects of quality and quantity of information on decision effectiveness. *J. Consumer Res.* 14 200–213.
- , ———. 1989. Assessing biases in measuring decision effectiveness and information overload. *J. Consumer Res.* 15 504–508.
- Kleinmuntz, Don N., David A. Schkade. 1993. Information displays and decision processes. *Psych. Sci.* 4 221–227.
- Lehmann, Donald R., Yigang Pan. 1994. Context effects, new brand entry, and consideration sets. *J. Marketing Res.* 31 (3) 364–374.
- Liang, Kung-Yee, Scott L. Zeger. 1986. Longitudinal data analysis using generalized linear models. *Biometrika* 73 (1) 13–22.
- MathSoft, Inc. 1998. *S-Plus 4.5*. MathSoft, Seattle, WA.
- McCullagh, P., J. A. Nelder. 1989. *Generalized Linear Models*. Second Edition, Chapman and Hall, London, UK.
- Moorthy, Sridhar, Brian T. Ratchford, Debabrata Talukdar. 1997. Consumer information search revisited: theory and empirical analysis. *J. Consumer Res.* 23 (March) 263–277.
- Muthukrishnan, A. V. 1995. Decision ambiguity and incumbent brand advantage. *J. Consumer Res.* 22 (1) 98–109.
- Nedungadi, Prakash. 1990. Recall and consumer consideration sets: influencing choice without altering brand evaluations. *J. Consumer Res.* 17 (December) 263–276.
- Payne, John W. 1982. Contingent decision behavior. *Psych. Bull.* 92 382–402.
- , James R. Bettman, Eric J. Johnson. 1988. Adaptive strategy selection in decision making. *J. Experiment. Psych.: Learning, Memory, and Cognition* 14 534–552.
- , ———, ———. 1993. *The Adaptive Decision Maker*. Cambridge University Press, Cambridge, UK.
- Pearson, J. Michael, J. P. Shim. 1994. An empirical investigation into decision support systems capabilities: a proposed taxonomy. *Inform. Management* 27 45–57.
- Ratchford, Brian T., Narasimhan Srinivasan. 1993. An empirical investigation of returns to search. *Marketing Sci.* 12 (Winter) 73–87.
- Roberts, John H., James M. Lattin. 1991. Development and testing of a model of consideration set composition. *J. Marketing Res.* 28 (November) 429–40.
- Russo, J. Edward. 1977. The value of unit price information. *J. Marketing Res.* 14 (May) 193–201.
- Shugan, Steven M. 1980. The cost of thinking. *J. Consumer Res.* 7 (September) 99–111.
- Simon, Herbert A. 1955. A behavioral model of rational choice. *Quart. J. Econom.* 69 99–118.
- Singh, Danièle T., Michael J. Ginzberg. 1996. An empirical investigation of the impact of process monitoring on computer-mediated decision making performance. *Organ. Behavior Human Decision Process.* 67 (2) 156–169.
- Todd, Peter, Izak Benbasat. 1992. The use of information in decision making: an experimental investigation of the impact of computer-based decision aids. *MIS Quart.* 16 (September) 373–393.
- , ———. 1994. The influence of decision aids on choice strategies: an experimental analysis of the role of cognitive effort. *Organ. Behavior Human Decision Process.* 60 36–74.
- Widing, Robert E. II, W. Wayne Talarzyk. 1993. Electronic information systems for consumers: an evaluation of computer-assisted formats in multiple decision environments. *J. Marketing Res.* 30 (May) 125–141.

HÄUBL AND TRIFTS

Consumer Decision Making in Online Shopping Environments

- Winer, Russell S., John Deighton, Sunil Gupta, Eric J. Johnson, Barbara Mellers, Vicki G. Morwitz, Thomas O'Guinn, Arvind Rangaswamy, Alan G. Sawyer. 1997. Choice in computer-mediated environments. *Marketing Lett.* **8** (3) 287-296.
- Wright, Peter. 1975. Consumer choice strategies: simplifying vs. optimizing. *J. Marketing Res.* **12** (February) 60-67.
- Zeger, Scott L., Kung-Yee Liang. 1986. Longitudinal data analysis for discrete and continuous outcomes. *Biometrics* **42** 121-130.
- Zack, Michael H. 1993. Interactivity and communication mode choice in ongoing management groups. *Inform. Systems Res.* **4** (3) 207-239.

This paper was received May 13, 1998, and has been with the authors 9 months for 3 revisions; Processed by William Boulding.