Available online at www.sciencedirect.com
ScienceDirect
European Journal of Operational Research 186 (2008) 748–765
EUROPEAN JOURNAL OF OPERATIONAL RESEARCH
www.elsevier.com/locate/ejor

Decision Support

Enabling e-transactions with multi-attribute preference models

John C. Butler [a], James S. Dyer [b,*], Jianmin Jia [c], Kerem Tomak [d]

[a] *Department of Accounting and MIS, Fisher College of Business, The Ohio State University, Columbus, OH 43210, United States*
[b] *Department of Management Science and Information Systems, The Graduate School of Business, University of Texas at Austin, Austin, TX 78712, United States*
[c] *Department of Marketing, Faculty of Business Administration, The Chinese University of Hong Kong, Shatin, NT, Hong Kong*
[d] *Yahoo! Inc., Search and Marketplace Division, 2811 Mission College, Santa Clara, CA 95054, United States*

Received 10 May 2006; accepted 30 January 2007
Available online 21 March 2007

Abstract

This paper describes potential applications of multi-attribute preference models (MAPM) in e-commerce and offers some guidelines for their implementation. MAPM are methodologies for modeling complex preferences that depend on more than one attribute or criterion, and include multi-attribute utility theory, conjoint analysis, and the Analytic Hierarchy Process. There are numerous examples of applications in e-commerce that would benefit from the acquisition of information regarding the preferences of a consumer, a customer, an advice seeker, or a decision maker. Here, the focus is on applications of MAPM models in B2C and B2B websites, where preferences of consumers are assessed for the purpose of identifying products or services that closely match their needs.

In this paper, we provide an overview of decision aids with the MAPM approach, emphasizing how the MAPM structure of an individual's preferences may be assessed. This discussion is illustrated with examples of the use of alternative MAPM assessment approaches that are incorporated in existing websites. We then discuss how MAPM applications should be tailored for success in these environments.

Keywords: Multi-criteria analysis; Decision analysis

1. Introduction

The purpose of this paper is to describe the potential applications of multi-attribute preference models (MAPM) in e-commerce and to offer some guidelines for their implementation. MAPM are methodologies for modeling complex preferences that depend on more than one attribute or criterion, and include multi-attribute utility theory (Keeney and Raiffa, 1976; Dyer and Sarin, 1979), conjoint analysis (Green et al., 2001), and the Analytic Hierarchy Process (Saaty, 1980; Forman and Gass, 2001). In this paper we focus on the approaches that feature additive, compensatory value functions but many of the ideas apply to outranking approaches like ELECTRE (e.g., Roy, 1996) as discussed in

* Corresponding author.
E-mail address: Jim.Dyer@mccombs.utexas.edu (J.S. Dyer).

0377-2217/$ - see front matter © 2007 Elsevier B.V. All rights reserved.
doi:10.1016/j.ejor.2007.01.051

Stewart and Losa (2003). There are numerous examples of applications in e-commerce that would benefit from the acquisition of information regarding the preferences of a consumer, a customer, an advice seeker, or a decision maker.

A cottage industry has emerged where consumers are given "free" on-line decision aids to support decisions ranging from voting in the last presidential election in the United States to selecting a pet. The main objective of some of these business-to-consumer (B2C) tools is to collect valuable consumer data from real decision makers. Many of these tools attempt to emulate MAPM, but often do so in an ad hoc manner and with varying degrees of success. We believe that the informed use of MAPM would lead to decision support systems that make more valuable contributions and recommendations, leading to more site visits from users, and potentially greater profits. We also believe that these decision support tools can provide additional benefit as add-ons to existing web-based service providers, especially in the business-to-business (B2B) domain.

Like Edwards and Fasolo (2001) we start with the premise that the proper use of MAPM is a normative concept. We further assume that website developers take a prescriptive view and attempt to help consumers use an MAPM to accurately capture their preferences. This assumption is reasonable for websites whose primary value proposition is unbiased advice. However, we recognize that some of these services may be biased towards a particular vendor or product, and also discuss ways in which consumer preferences could be manipulated via a web-based system.

In this paper, we will provide an overview of decision aids that offer MAPM-like support, emphasizing how the MAPM structure of an individual's preferences may be assessed. This discussion will be illustrated with examples of the use of alternative MAPM assessment approaches that are incorporated in existing e-enabled websites. We will then discuss how MAPM applications should be tailored for success in these environments.

2. Web-based decision aids

As the amount of data that needs to be processed through a web interface increases, companies rely on more automated decision aids. One way of automating the decision making task is to add "intelligence" to the decision support tools used for analyzing the information generated in B2C and B2B electronic commerce applications.

2.1. Decision aids, intelligent agents, and shopbots

The topic of computerized decision support systems has been the subject of a great deal of work for several decades now, although the focus of these early systems was mainly on decision aids for managers and executives. An early example of a "man–machine interactive" decision support tool for selecting among alternatives based on trade-offs among multiple attributes is provided by Dyer (1973).

However, the past decade has seen an explosive growth in the use of electronic commerce, primarily through the World Wide Web, and has expanded the number of potential users of these decision support systems to include both personal shoppers and professional purchasing agents. Within this new environment, there is an obvious motivation for marketers to enhance the shopping experience for consumers in order to increase sales and profitability.

As a result, decision aids have been created for a number of websites to support these on-line shoppers and professional buyers. These decision aids may be classified as examples of intelligent agents (West et al., 1999), or more specifically as recommendation systems or shopbots. Typically they assist buyers in searching a website or across several websites for products and services that offer low costs and meet other desired performance objectives.

2.2. Approaches to modeling preferences

There are many different approaches to modeling preferences, and several different strategies have been incorporated into e-commerce sites. One of the most common approaches to obtaining some knowledge about a consumer's preferences is to infer that information based on her past purchase history. The objective of the analysis of this data is to use collaborative filtering to identify subgroups among the population of consumers with similar purchasing patterns (and implicitly, with similar preferences). Based on this classification, the consumer can be presented with alternatives that are commonly purchased by others within her subgroup. Perhaps the most well known example of this type of a decision aid for consumers is the service provided by www.Amazon.com.

The advantage of collaborative filtering is that it is a relatively painless way for the consumer to be provided with alternatives that may be consistent with her preferences. On the other hand, this methodology does not create an explicit model of preferences that can be modified by the consumer. As a result, collaborative filtering may not be useful in the evaluation of new or novel products and services with no purchase histories.

Also, a preference model based on collaborative filtering is not actively developed in cooperation between the consumer and the decision aid. In some cases, this collaboration between the consumer and the agent in the construction of a representation of preferences may increase the consumer's commitment to the process, and her ultimate satisfaction with the choice that is made (Kahn and Huffman, 1998).

In contrast, recommendation agents or shopbots that seek specific information regarding the consumer's preference use content filtering. This information is obtained by identifying the attributes that are relevant for a choice and assessing the importance of these attributes. One possibility is to use collaborative filtering to identify attributes and objectives that the consumer may care about as a first step. This would be particularly useful for an inexperienced consumer. Next, some information may be obtained regarding how varying the levels of performance on individual attributes translates into satisfaction with the product or service. The result is a multi-attribute preference model, or MAPM, that will be the focus of this paper.

2.3. Considerations of context

The efficacy of the use of MAPM as a Web-based decision aid will depend on several aspects of the context of the decision task. These aspects include the consumer's familiarity and level of expertise regarding the product or service, and the significance of the decision in terms of cost and commitment.

In order to construct and utilize a recommendation agent featuring the MAPM approach, the user must be able to express preferences for different levels of performance on attributes both within the same attribute and across attributes. Individuals with low levels of knowledge regarding a product class may have a difficult time making the necessary trade-offs in terms of product attributes, and therefore be very uncomfortable with both the process and the results of the ranking that is offered. At another extreme, an expert may have prior knowledge regarding available products and services, and prefer to search quickly for detailed information on a small number of alternatives without the distraction of responding to preference assessment questions. For similar reasons, consumer purchases that are considered routine, or that require relatively low commitments of resources, may not be appropriate environments for the intrusion of a decision aid.

Therefore, decision aids based on MAPM approaches to approximating consumer preferences may be more appropriate for a relatively inexperienced B2C consumer who at least has sufficient knowledge to describe desired levels of performance on product attributes, and for infrequently purchased durable products such as cameras, laptop computers, and automobiles. MAPM approaches may be even more useful in B2B purchasing environments, where the buyer's specification of preferences in terms of attributes, weights of importance, and desirable performance levels may provide useful information to potential sellers, and vice versa. The participants in a B2B relationship may also be required to justify their choices which is consistent with having a well specified MAPM.

As a result, some websites offer the consumer the choice of two or more of these types of decision aids. These sites allow her to select the decision support that is most appropriate for the decision context.

3. A framework of MAPM

For the types of problems we will consider in this paper, for example the selection of the best response to a request for quote (RFQ) by a buyer at a B2B site, there are numerous measures of the desirability of the alternatives. Often these measures are conflicting and require trade-offs by the decision maker; the cheapest quotes may be for lower quality products. Under ideal circumstances, a MAPM approach will provide a faithful representation of the decision maker's preferences, and can be used in an automated decision support system to simplify the tasks of identifying and comparing desirable alternatives.

Multi-attribute utility theory provides the theoretical basis for a MAPM, and so we will begin with an informal review of the basic concepts of this theory as it is commonly applied in the e-commerce environment. As we shall discuss, e-commerce applications of conjoint analysis and of the AHP may be viewed as alternate approaches to the assessment of multi-attribute utility models.

A multi-attribute utility model of preference is developed by defining attribute utility or value functions on performance measures of the alternatives, and by aggregating performance on multiple criteria into a single score. As far as we are aware, all of the web-based applications of MAPM use the simplest aggregation rule, the additive multi-attribute model. Further, these applications also explicitly or implicitly treat the consumer choice decision as though it occurs in the context of certain outcomes or attribute values, that is, these applications typically do not provide information on attribute levels defined by probability distributions. Formally, there are distinctions among the assumptions that justify the use of the additive multi-attribute model of preference and among the appropriate assessment methodologies for the domains of certain versus risky choices. This distinction is often emphasized in the literature by using the term value function for the case of certainty and utility function for the case of risk, and we will adopt this convention for emphasis.

Let $v_i(x_i)$ be a single-attribute value function over the attribute measure x_i. This single-attribute value function provides a functional relationship between the observed attribute levels of performance and their perceived values or utilities. The simple additive multi-attribute value function may be written as

$$v(x_1, x_2, \ldots, x_n) = \sum_{i=1}^{n} v_i(x_i). \tag{1}$$

The additive model provides an accurate representation of the preferences of individuals only if their trade-offs among any two of these criteria are not affected by common outcomes on the remaining $n - 2$ criteria (Keeney and Raiffa, 1976). This additive preference independence condition is an important one, and may not be recognized by designers of web-based decision aids.

Further, this additive value function provides a basis for ranking alternatives, but it does not measure preference differences and it should not be used for risky choices. More formally, (2) is an additive conjoint preference structure (Krantz et al., 1971).

To simplify the assessment of the additive preference model, an even stronger independence condition is required, known as difference independence (Dyer and Sarin, 1979). Loosely speaking, this condition requires that common performance levels on the other attributes should not affect the individual's preferences for increases or decreases in the performance levels on any one attribute. This condition allows each individual single-attribute value function $v_i(x_i)$ to be assessed without explicitly considering the values of the other attributes, a common practice in many e-commerce and traditional applications. When this condition is met, the additive value function is also called a measurable value function, and the scores of a measurable value function provide the basis for comparisons of preference differences among alternatives rather than simple rankings of the alternatives.

This assessment is often simplified by scaling each single-attribute value function $v_i(x_i)$ from 0 to 1, and by adjusting for the scaling of these normalized value functions using weights w_i for each attribute. With this choice of scaling, the additive multi-attribute value function can be written as

$$v(x_1, x_2, \ldots, x_n) = \sum_{i=1}^{n} w_i \bar{v}_i(x_i), \tag{2}$$

where $\bar{v}_i(\cdot)$ is a single-attribute value function over attribute i scaled from 0 to 1, $w_i > 0$ is the weight for measure i, and $\sum_{i=1}^{n} w_i = 1$.

This scaling of the single-attribute value functions can be accomplished by assigning a score of 0 to the worst level of the attribute x_i^0 and a score of 1 to the best level x_i^*. Then, other levels of performance between these two extremes might be assigned scores by direct rating, by "part worth scores" from conjoint analysis, or using other assessment approaches. As an alternative, a simple mathematical relationship may be assumed or assessed (Farquhar and Keller, 1989; Kirkwood, 1996).

There are several approaches for assessing weights for the additive multi-attribute value model, but they all should be consistent with the results obtained from trade-off judgments (Keeney and Raiffa, 1976). To illustrate the idea, we consider a trade-off between price and resolution for buying a digital camera, assuming that the performance levels of the other features are the same. Specifically, for two digital camera alternatives $A = $ (€200, 1600×1200 pixels) and $B = $ (____, 1920×1600 pixels), the decision maker is asked to specify a price for B such that the two alternatives are equally attractive.

Suppose the decision maker would be indifferent between A and B if the price of B were set at €300. From (2), we have

$$w_1 \bar{v}_1(€200) + w_2 \bar{v}_2(1600 \times 1200 \text{ pixels}) = w_1 \bar{v}_1(€300) + w_2 \bar{v}_2(1920 \times 1600 \text{ pixels})$$

where $\bar{v}_1$ and $\bar{v}_2$ are value functions for the price and resolution attributes respectively, and w_1 and w_2 are corresponding weights. We may conduct similar trade-offs for other pairs of attributes to obtain several equations like the above. If all the single-attribute value functions have been determined, these indifference equations can be solved for the weights.

As the name implies, conjoint analysis was developed based on the use of statistical techniques to approximate the additive conjoint preference structure (2) (Luce and Tukey, 1964). A typical procedure in many applications of conjoint analysis is to assess the rank-order or overall value for alternatives with different profiles of attribute levels, and then to use the holistic judgment information to estimate discrete levels of single-attribute value functions by regressions, hierarchical Bayes models, or linear programming. When used in this manner, conjoint analysis is a decomposition method, and it does not need to evoke weights through direct assessment. This holistic assessment approach is typically used with hundreds or even thousands of consumers in consumer preference research regarding product designs and service offerings (Green et al., 2001).

However, the conjoint analysis models implemented in the Internet environment to assess buyer preferences use choice-based conjoint models, or hybrid models that require subjects to directly estimate ("self-explicate") both the desirability and the importance of attributes and levels. These methodologies also use direct assessments of preference intensities, or preference differences to more closely mirror the assessment methodologies appropriate for the measurable value function (2).

Finally, the AHP methodology may be given a firm theoretical foundation through an interpretation of its assessment questions as judgments regarding ratios of preference differences. Viewed from this perspective, and with proper scaling of the single-attribute value functions, the AHP is an alternative approach to assessing the measurable value function (2) (Dyer, 1990).

4. Examples of MAPM in e-commerce

There are a number of examples of applications of MAPM in e-commerce. In this section, we highlight some of these examples, and discuss the assessment methods that they employ to approximate models of the users' preferences. This discussion is summarized in Table 1. Although this discussion is written in the present tense for ease of exposition, most of our observations were made in February 2002 and we realize that some of the web sites surveyed in Table 1 may no longer exist in the form that we describe. In some cases, it seems that persons not familiar with MAPM intuitively used related concepts and features in the design of these sites. Often they could benefit from some simple changes in the assessment protocols. In other cases, the site descriptions explicitly indicate that formal MAPM approaches are used in their implementation.

4.1. Examples of MAPM models for B2C decision support

Consumers are often interested in evaluating a number of alternatives prior to finalizing a purchase decision. In recent years, many websites have been created that provide information to consumers regarding specific classes of products. Haubl and Trifts (2000) report that these online shopping assistants may have strong favorable effects on both the quality and efficiency of purchase decisions, allowing shoppers to make better decisions with less effort.

In many cases, the consumers must specify the alternatives they wish to consider, and the site simply provides product specifications along with links to potential dealers and stores. Examples of such sites would include Auto-by-Tel (www.autoby-tel.com) and MySimon.com (www.mysimon.com).

Other sites now offer enhanced decision support to assist the consumer in the initial screening of alternatives, and in the in-depth comparison of this reduced set of alternatives, the consideration set, based on multiple criteria. Dealtime (www.deal-time.com) offers the ability to compare products and services on a variety of attributes. For example, the site provides a comparison of long distance telephone services based on the user's telephone exchange. An initial screen provides a complete list of the alternative service providers in the relevant area, ranked by a single attribute (e.g., annual savings). The user is allowed to choose among several other attributes to re-sort these rankings. The user is then allowed to select a subset of the available alternatives, and a second screen provides a side-by-side comparison of this consideration set on an expanded number of attributes, and with additional details. The process of reducing the number of alternatives considered is often called winnowing (e.g., Edwards and Fasolo, 2001).

CallVision (www.callvision.com) provides a natural extension of the capabilities of the Dealtime site by providing technologies to assist a consumer

Table 1
Summary of selected websites

Applications	Provide/request	Steps
B2C		
Auto-By-Tel www.autobytel.com MySimon.com www.mysimon.com	Product information	Select from a set of alternatives, receive matching products and sellers
Dealtime www.dealtime.com	Product information and comparisons	Sort the alternatives by attributes, and receive detailed information on consideration set
CallVision (I-Optima SST) www.callvision.com	Product categories and attributes	Input information about current product characteristics, receive ranking of the plans
Best Place to Live (Money Magazine) www.money.cnn.com	Recommendation of best city to live in	Provide weights for seven pre-specified attributes and receive ranking of cities
Active Decisions www.activebuyersguide.com (no longer a standalone site) My Product Advisor www.myproductadvisor.com	Product categories and attributes	Choose the important attributes, indicate desirability of levels of performance, determine importance weights, receive product rankings
Online Insight www.onlineinsight.com	Match consumers with products based on product features	Provide importance ratings of the attributes, make choices revealing trade-offs, receive product list
B2B		
Perfect Source www.perfect.com	Match suppliers to buyers based on vendor and product features	Provide importance weighting of requirements, receive candidate list
SAS www.sas.com	Perform weighted supplier ranking to identify the best suppliers	Choose important criteria and use a dynamic weighting approach to rank the suppliers
Asparity Decision Solutions www.plansmartchoice.com	Match employees to health care product plans	Choose the important attributes, indicate desirability of levels of performance, determine importance weights, receive product rankings
Colmart www.colmart.com	Match buyers and sellers of paper products	Choose the important attributes, indicate desirability of levels of performance, determine importance weights, receive product rankings

considering multiple services or products based on a MAPM evaluation of alternatives. For example, a consumer using Call Vision's I-Optima SST to evaluate wireless telephone services would be presented with the following product attributes (or objectives): price, quality, coverage, and data services. The "weights of importance" for the MAPM model are estimated through comparisons of the importance of the other three attributes with the importance of price, using percentages.

The site uses a novel approach based on comparisons to the consumer's current service provider to estimate the scores of the alternatives on the criteria. The consumer is asked to identify her current service provider, and to rate its performance on the four attributes above. In addition, the consumer is asked to input her current monthly bill, and information regarding such usage issues as the percentage of long distance calls, roaming calls, and any required features such as Internet access. The result is a ranking of the plans offered in the consumer's area, based on these inputs. As before, the consumer can choose one or more alternatives from this ranked list for a side-by-side comparison that offers additional details.

Active decisions (www.activebuyersguide.com) provides a more sophisticated example of a multidimensional search based on MAPM concepts. As shown in Fig. 1, a user considering the purchase of a camera is able to search for a product based on how she anticipates using the product or based on the product attributes. Novices may prefer to think in terms of usage – e.g., I like 8×10 prints – while experts may prefer to think in terms of the attributes – e.g., I need at least 4 megapixels (Mitchell and Dacin, 1996). The current implementation of Actives Buyers Guide does not allow the user to provide weights for the attributes, but the user can sort on any attribute. Previous implementations of Active Buyers Guide and other products offered by Active decisions do allow the specification of weights as discussed in Section 5.

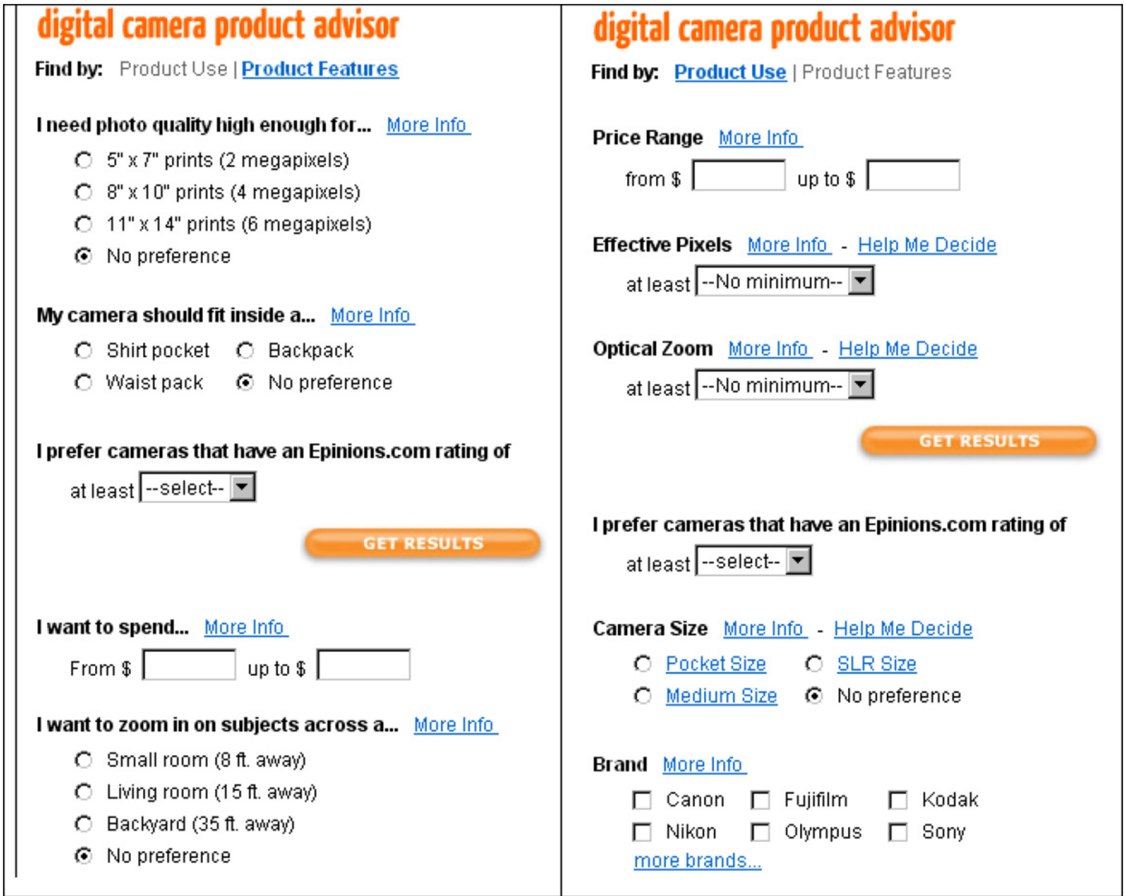

Active Buyers Guide Focus on Usage	Active Buyers Guide Focus on Attributes

Fig. 1. Usage or attribute focus at Active Buyers Guide.

Another website that does address the importance of attributes is My Product Advisor (www.myproductadvisor.com). As shown in Fig. 2, My Product Advisor initially asks the user to think about how she will use the camera, in addition to how important each aspect of usage is relative to the others. The user is given a list of "other" attributes that she might include in the analysis, and has the option of specifying value function scores for different levels of some of the attributes to approximate the value functions of these attributes. For example, it may be reasonable to assume that the value function for cost is linear, particularly after a consumer has limited the range of costs to consider. However, the value function for resolution may need to be assessed as it is likely to be concave for all but the most expert users. Finally, she can specify the attribute weights for her selected attributes, as shown in Fig. 3.

These responses provide an approximation to the weights and single-attribute value function scores of the additive MAPM model (2). Given this information, a list of cameras is generated and ranked using the underlying MAPM model. Once again, additional information regarding each camera is available by clicking on the appropriate icon.

Online Insight's Solution Suite (www.onlineinsight.com) provides another example of a consumer-facing software product that allows consumers to express their personal preferences and trade-offs for products and services. Based on the elicitation of these preferences, the software then matches the consumer with the appropriate products. Each recommended product is supported with a graphical representation showing how it ranks against the needs and preferences expressed by the consumer. The stated objective of this feedback is to increase consumer confidence and trust, and to increase consumer loyalty.

An example of the dialogue used by the Online Insight solution may be useful. Suppose that a com-

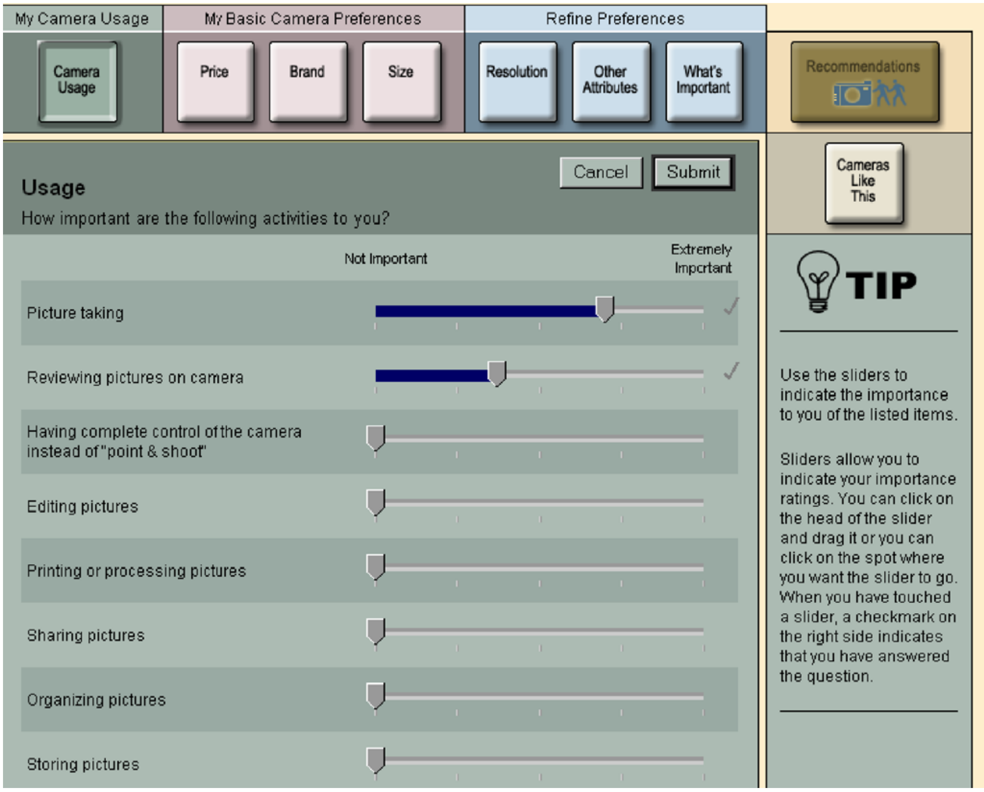

Fig. 2. Relative importance of camera usage features at My Product Advisor.

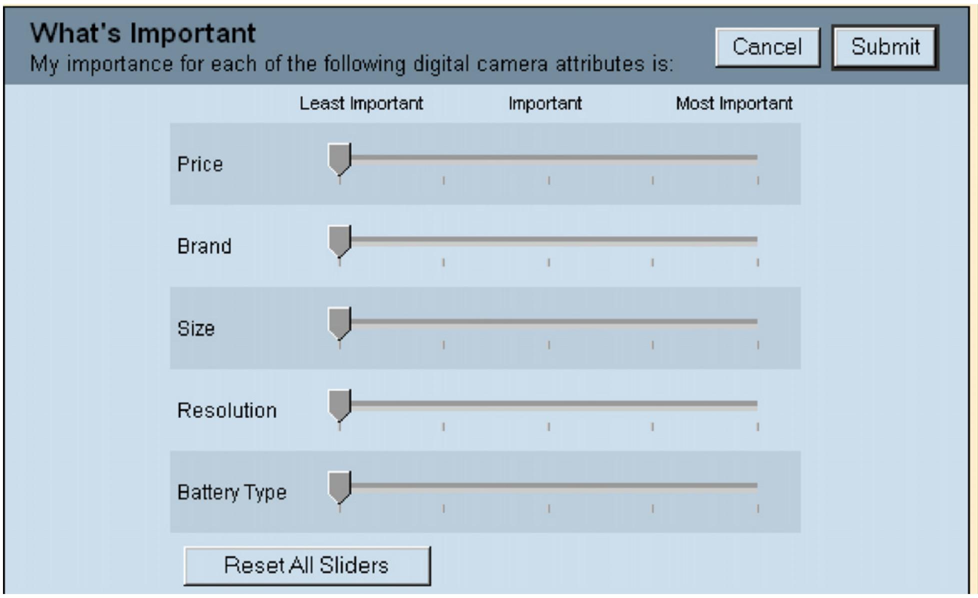

Fig. 3. Specifying attribute weights at My Product Advisor.

pany has several options in a particular product line that may be differentiated by the features: brand, color, top speed, warranty, and price. The collaborative conversation with a consumer begins by obtaining a rudimentary importance rating of these features on a 1–5 scale, with 5 being the most important. In some cases, the system will assume ordinal preference information concerning such features as price (cheaper is better), speed (faster is better), and warranty (longer is better). For categorical features, such as brand and color, the consumer is also asked to make preference ratings using the same 5-point scale.

Next, the consumer is asked to help refine the underlying model of her preferences by asking such questions as the following: Which of the following products would you purchase?

	Product *A*	Product *B*
Brand	Brand X	Brand Y
Price	2000	1800
Color	Green	Blue

Based on a series of questions such as this, the underlying preference model is "tweaked" using conjoint analysis to assess the single-attribute value functions.

There are other examples of recommendation systems that are not intended to induce a purchase, per se. For example, Money Magazine offers a support tool to help users choose a city to reside. The user can weight seven pre-specified criteria, including the possibility of allocating no weight, and these weights are used to rank cities. The criteria include aspects of weather, crime, culture and job growth.

4.2. Examples of MAPM models for B2B decision support

There are numerous examples of commercial B2B software and websites that offer users the opportunity to match products and services based on multiple criteria. In many cases, the details of how these products function are proprietary, but some of them appear to be using the basic concepts of MAPM. We will first provide examples of software products that are available to implement such ideas, and then discuss a few representative examples of websites that offer MAPM matching capabilities. More examples of the use of MAPM in B2B applications can be found in Teich et al. (2004). We focus on situations where the goal is to find a stand-alone alternative but note that MAPM can easily be extended to decision contexts involving portfolios of several alternatives (Kleinmuntz, 2007).

4.2.1. MAPM in B2B software

Perfect Source (www.perfect.com) is a modular set of software products that provide an end-to-end strategic sourcing system. The module to support buyer purchasing includes the ability to introduce several "requirements" or criteria (e.g., shipping, return policy, and restocking fee), and to determine the "importance weighting" of these requirements using intuitively appealing slider bars. Suppliers use a complementary module allowing them to respond in terms of multiple requirements as well. The resulting system is called a weighted scorecard, which scores and ranks bids and then optimizes them into suggested multi-source award packages line-item by line-item.

SAS® (www.sas.com) offers a Supplier Relationship Management software suite that is designed to assist procurement professionals in managing their strategic sourcing efforts. The suite is composed of four related modules, with the first being a data base that supports the other decision support tools. The Spend Analysis module allows the user to rank suppliers according to the criteria most important to the company. The criteria that are often used in such applications include supplier type, quality of supply, lead times, purchase price, and financial viability. The software uses a "dynamic weighting" of the chosen criteria to add "balance and flexibility" to the process of ranking the suppliers.

In addition, the Procurement Scorecard module provides the ability to monitor the selected suppliers and the entire procurement organization against a specific set of "key performance indicators". This scorecard will alert the appropriate personnel when procurement objectives are in danger of not being met. Finally, the Sourcing Strategy module includes a wizard-driven application that guides business users through the process of identifying measurable objectives and defining business rules to constrain the model, which then determines the suppliers to use and calculates the amount to be spent with each of them to achieve the stated objectives.

As Milgrom (2004) notes, electronic markets will be utilized only when they offer value relative to traditional supply channels. These value-added services could include reduced transaction costs or enhanced information. Milgrom (2004) asserts that supporting the evaluation of "true multi-dimensional custom offers in response to complex, weighted, concurrent, multi-dimensional RFQs" allows a "better match between buyers and sellers... helping tailor the non-price aspects of the transaction...".

4.2.2. MAPM in B2B applications

Asparity Decision Solutions (www.plansmart-choice.com) has developed a decision support application to assist end users in sorting through the options and choosing a health plan that best fits their needs. The preference module asks the user to perform three steps. First, she clicks on the attributes or features of a health plan that are important to her, and, second, she rates the level of importance of the chosen attributes. Third, she must respond to a series of trade-off questions. The preference module then creates a unique preference function for each user, matches it with the available health plans, and presents the results in rank order. We use the term ''preference function'' because it is not clear if a utility or value function or some other preference representation is being utilized by the web site. Summaries of this preference information are available to employers to aid them in negotiating new options with health care providers.

The objective of the Colmart.com Internet trading site (www.colmart.com) is to provide a forum for buyers and suppliers to buy and sell paper products. Much like the Asparity Decision Solution, the buyer is allowed to set ''priorities'' or weights on product attributes from a dropdown list, using a scale from 1 to 5. The buyer then specifies ''ideal values'' for each of the attributes, and an approximation to a single attribute value function is used to score each alternative based on its distance from the specified value. The alternatives are then selected based on a weighted sum of these distance measures, which are a proxy for an additive multi-attribute value function. As an alternative, the user can enter the range of acceptable levels on each attribute, and the software will find all of the supplier alternatives that match the profile. Keskinocak et al. (2001) provides additional details regarding these matching processes.

This process is an example of a multi-attribute reverse auction (or procurement auction) (e.g., Bichler and Kalagnanam, 2005; Parkes and Kalagnanam, 2005; Teich et al., 2006). These approaches allow potential suppliers to bid on multiple attributes, in some cases in an iterative fashion, and use the potential buyers' MAPM to score the bids. There must be some trust on the part of the buyers that they will not be exploited by providing their MAPM and typically the potential sellers are given limited feedback in the form of changes in buyer bid evaluations. Early experimental work suggests that buyers are better off using these systems and seller profits are not significantly degraded (Chen-Ritzo et al., 2005; Bichler, 2000).

Online negotiation is another possible application of MAPM in online settings. As outlined in Kersten and Noronha (1999) these negotiation support systems (NSS) can be used to improve the outcomes of both parties by directing the process towards pareto-dominating alternatives based on the specified preferences. NSS also require the negotiators to trust the system enough to provide detailed preference information.

5. Some guidelines for building MAPM applications

The examples surveyed above provide some indication of the widespread adoption of MAPM as a basis for modeling consumer preferences in the B2C and B2B environments. This discussion also features a variety of assessment protocols, and degrees of sophistication in the MAPM modeling efforts. In some cases, these differences may reflect conscious decisions by the website designer in trading off accuracy of the MAPM representation versus the cognitive burden imposed on the user. In other cases, these differences may reflect a lack of awareness of some basic concepts of MAPM that would enhance the users' experience with the website.

In this section, we begin with a discussion of some recommendations that should be given serious consideration in most website implementations of MAPM-like decision aids. Next, we acknowledge that a clever implementation of a MAPM assessment technique might actually be used to manipulate preferences, and identify some ways in which this might be done. On the one hand, a web designer might wish to take advantage of some of these insights to select how information is presented to a consumer to influence her choices, within ethically appropriate standards. On the other hand, the user of such websites may also want to be aware of the potential for the manipulation of her expressed preferences, so that she may avoid being guided to a misstatement of her preferences, and led to alternatives that she would otherwise not prefer.

5.1. Avoiding common mistakes in MAPM website design

Our review of the applications of MAPM approaches to support consumer decisions in B2C and B2B websites indicates that several mistakes are commonly made in their designs. Further, these mistakes are relatively easy to overcome in this environment, and should be considered in the design of any decision aid based on preference modeling. These mistakes are not unique to web-based implementations of MAPM, and Keeney (2002) discusses some of these issues in a more general setting.

5.1.1. Using attributes rather than objectives

The identification of the appropriate objectives for an application of a MAPM is an important step, and requires an understanding of a subtle but important notion. That is, it is important to separate *fundamental objectives* from *means* to accomplish objectives. A simple example, borrowed from Clemen and Reilly (2001) may help to illustrate this concept. For an individual, working fewer hours may seem to be an important objective. However, it may be the case that this is only a means of accomplishing something else that is really important to the individual, such as spending more time with family. In this case, spending more time with family would be a fundamental objective, while working fewer hours would be a means objective.

This distinction is an important one that relates to the design of e-commerce B2B and B2C website interfaces. Most sources of information regarding products and services provide information related to the *attributes* of the product or service. These are typically things that can be objectively measured, and relate in some manner to aesthetic or operational features. However, there is no *a priori* reason to believe that these product attributes will be the appropriate measures for a consumer. In many cases, these product attributes may be relevant means objectives, but may need to be combined in some unique and personal manner by the consumer to provide information related to a fundamental objective. These product attributes may also relate to more than one fundamental objective.

A focus on product or service attributes in a web design is a natural mistake, and is also common in traditional "hard copy" reviews of products and services. However, when a MAPM model is developed based on these attributes, there are at least two issues that may lead to poor approximations of the user's preferences.

First, the dialogue with the consumer that is required to develop a model of her preferences requires several value trade-offs, where comparisons may be required between hypothetical (or real) alternatives that differ on some attributes. These trade-offs may be very difficult for the individual to make when the attributes do not map directly into her fundamental objectives. For example, a consumer considering the purchase of an automobile may be given attribute values for horsepower and torque, which may be viewed as proxy attributes for performance. Unless the consumer is an automotive engineer, it may be very difficult for her to thoughtfully make value trade-offs between, say, an additional 20 lb-ft of torque versus a reduction of 10 hp.

Perhaps her fundamental objective is more closely related to acceleration, and some measure of performance on this product attribute may be much better suited to her concerns and values. Naturally, we must also recognize that the expertise of the user of a web system may have an impact on the selection of the appropriate measures of performance for the product or service. This suggests that a web-based system that offers the consumer the choice among several different product and service attributes would be preferable to one that implicitly assumes a universal set of fundamental objectives among all users. It should also provide sufficient information regarding the available attributes so that the user, with a click of the mouse, can obtain useful information to clarify the meaning and potential significance of the attribute.

However, this flexibility regarding the attributes has its own potential for misuse, since a consumer interested in automobile performance might be tempted to select horsepower, torque, and 0–100 kmh acceleration as relevant attributes for her MAPM. Some careful thought might lead to the recognition that horsepower and torque (combined with weight), may be means objectives related to the fundamental objective of acceleration, so including all three could lead to "double counting" regarding this objective, and place too much weight on performance relative to the consumer's true preferences.

This latter point emphasizes the second concern regarding the use of attribute values in a MAPM model. As we discussed, most website applications of MAPM or MAPM-like approaches to recommending alternatives use an additive weighted model. The independence assumptions that are required to ensure that the additive form is appropriate are typically not verified, and the attributes that are used in the decision aid may not be compatible with these assumptions. In such a case, these attributes may be relevant for the decision as means objectives, but they may need to be combined in a non-linear fashion to reflect the preferences of the consumer. Montibeller et al. (2005) discuss the complications of transforming a mental model of a decision situation into a value hierarchy.

A related issue is that many websites specify the attributes, and implicitly the objectives, for the decision maker. For example, Fig. 4 shows the seven attributes used to rank cities to live in by Money Magazine. While it is clear that "Weather" is likely to be an important consideration for most consumers, the website requires that these concerns be captured in "Sunny Weather". A previous implementation of the same website was used by one of the authors several years ago. The author was surprised by the recommendations as he preferred warm, winterless, weather and was told he should live in New England. After some investigating he realized that this was due to the high weight he specified for the attribute "Weather". The website defined "Weather" as having four seasons which was not consistent with the author's preferences for a single, warm season. After lowering the weight on the weather attribute – although the author felt that *his* definition of weather was very important – the recommendations were improved.

Butler et al. (2006) develop an alternative MAPM procedure where the attributes are combined to form mental models that predict the performance of alternatives on the fundamental objectives of the decision maker. In the first step the decision maker specifies weights that reflect each attribute's "impact" on each of the objectives, and therefore creates a predictive model for that objective. In the second step the decision maker specifies weights on each major objective, and therefore creates a preference model for the overall objective. Then the weights assessed in steps one and two are combined to determine the appropriate weights on the attributes to be used as a basis for ranking alternatives and for choice. The predictive model weights are statements of belief regarding the achievement of each objective while the objective model weights are statements of preference, and we believe this distinction should be explicitly recognized. Ignoring this distinction and confounding judgments of belief and judgments of preferences violates a fundamental principle of classical decision theory.

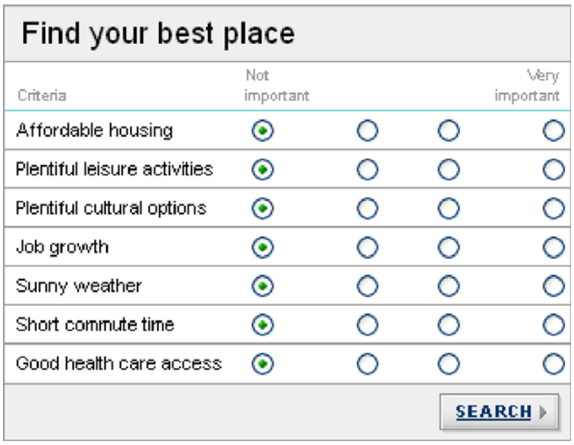

Fig. 4. Selecting a place to live at Money Magazine.

The designer of a MAPM website should be aware of these issues, and carefully select the product attributes that will be made available to the user. In some cases, an interactive dialogue focused on identifying the user's fundamental objectives might be considered to guide the selection of the appropriate attributes. For example, Active Buyers Guide allows the user to specify either desired camera usage or desired attribute levels as shown in Fig. 1. The first step at www.myproductadvisor.com is to specify the consumer's objectives in terms of how a digital camera is to be used before asking about attributes (Fig. 2). Later in the process the consumer is asked about camera resolution as measured by megapixels but this quantity is related to what consumers actually value: photo clarity and photo quality. Finally, as illustrated in Fig. 5, MySimon shows actual images at a variety of megapixel choices so the user can see the implications of choosing a particular camera resolution. Whenever possible, the web designer should enable the user to understand how a product attribute relates to her objectives.

5.1.2. Making value trade-offs independent of attribute range information

Many of the website applications present the consumer with a list of attributes, and ask her to rate their relative importance, perhaps on a scale from 1 to 5, without any information regarding the ranges over which these attributes might actually vary among the alternatives to be compared. The responses to these types of questions may include significant errors, at best, or may be very misleading when used in a MAPM model of preference. This issue is prevalent in self-explicated and hybrid implementations of conjoint analysis (Gibson, 2001) and also in applications of the AHP (Dyer, 1990).

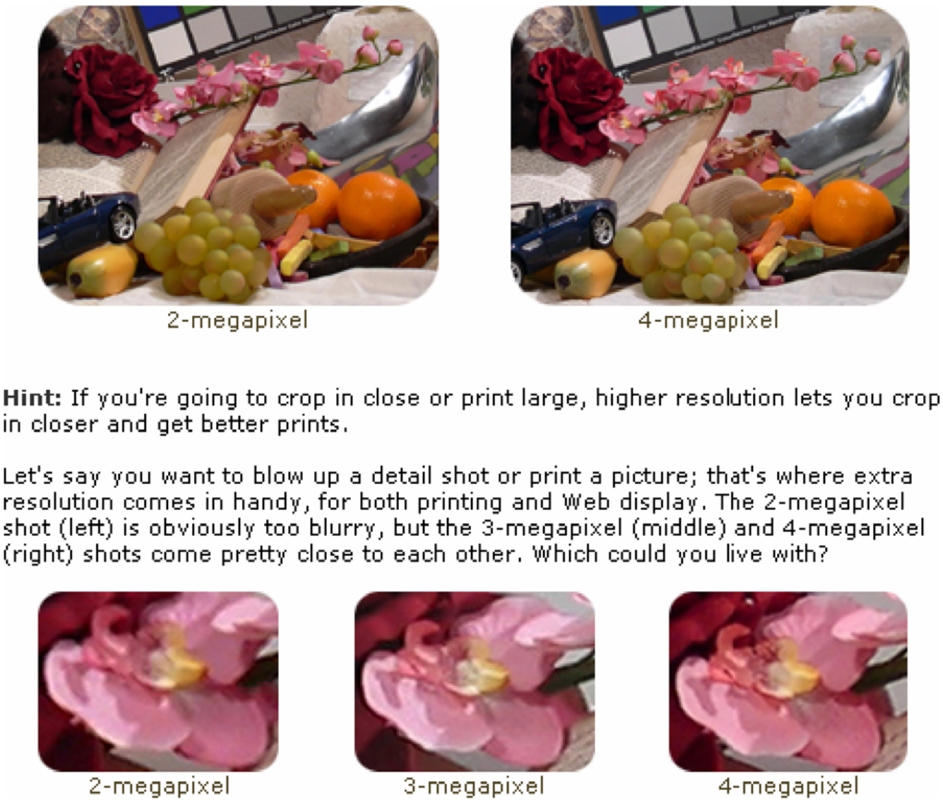

Fig. 5. Mapping megapixels to photo quality at MySimon.

The weights in the additive value model (2) reflect the relative importance of the ranges of attribute outcomes in determining the overall value. This is different from the notion of the intrinsic importance of the attribute per se. Returning to the example of the automobile purchaser, most consumers would probably agree that the most important objective in purchasing an automobile would be to minimize the cost of the vehicle that is chosen. However, once the consumer restricts herself to a consideration set of vehicles within the same general price class, the importance of cost may diminish.

For example, if the consumer says that she has decided to purchase a mid-size, four-door sedan, and would like to have the assistance of an automated decision aid in selecting among the alternatives, she may recognize that the cost of these alternatives varies from €24 000 to €28 000, but they differ significantly in terms of such features as horsepower, safety ratings, and quality ratings. Within this subset of all automobiles that are available for purchase, she may consider it much more important to have an increase from the lowest to the highest safety rating among this group, than to decrease the cost of the purchase from €28 000 to €24 000. In such a case, the weight in the MAPM model on the attribute of safety should be larger than the weight on cost (see Keeney, 2002 or Kirkwood, 1992 for an elaboration regarding this important point).

The swing weight method (von Winterfeldt and Edwards, 1986) provides an alternative to the trade-off method discussed earlier, and requires the decision-maker to consider the relative importance of changes from the worst to the best levels on each of the objectives. Fischer (1995) has investigated this phenomenon, and offers suggestions for ensuring that consumers are sensitive to the attribute ranges. Otherwise, there may be significant errors in the underlying model.

We are not aware of a website that intentionally focuses the user on the range of attribute performance when assessing attribute weights. For example, there is no indication of range information at My Product Advisor (Figs. 2 and 3) or Money Magazine (Fig. 4).

5.1.3. Making the process an unpleasant experience

The standard protocol for the elicitation of a MAPM model can require a number of tedious trade-off questions within and across product attributes that may be hard to compare, poorly understood, and not obviously representative of any real alternatives. The process can be time-consuming, and frustrating for the user.

Ideally, a consumer will be offered a process that is educational, informative, and fun. It is well known that decision makers will choose decision processes or heuristics based on recognition of the cognitive cost of the processes versus the perceived benefits of greater decision accuracy (Payne et al., 1993). A web-based MAPM decision aid should be implemented with the recognition that the number of questions to be asked, and the cognitive burden that is placed on the user should be commensurate with the potential gains or losses associated with some errors in the resulting recommendations.

Milgrom (2004) argues that for e-markets to succeed, they must be scalable. As an example, one necessary condition for scalability in the B2B environment is that suppliers are able to automate the screening of RFQs so that they only receive those requests that they are most likely to be able to meet. Similarly, buyers need help sifting through a potentially huge number of offers.

An increase in the quality of information and the reduction of the transaction costs should be goals of any recommendation system, search engine or intelligent agent. Many recommendation websites ask the user how many results to return. While this certainly reduces the cognitive burden, there is no mechanism to insure that the user will choose to look at the "right" number of alternatives. Other authors have incorporated a measure of cognitive burden directly in the MAPM model (e.g., Montgomery et al., 2004) forcing the decision maker to make a potentially difficult trade-off between some measure of mental processing load and the size of the consideration set. A MAPM allows a more straightforward way to make the comparison of alternatives less difficult: (1) by scoring the alternatives in an "intuitive" unit to ease comparison across the alternatives presented, and (2) by allowing the user to specify a minimum "significant" difference to reduce the number of alternatives presented. Both are related to a rescaling of the MAPM scores of alternatives, as discussed in Butler et al. (2001).

Sarin (1979) and Kirkwood and Sarin (1985) have investigated how MAPM models can be approximated with relatively little cognitive burden on the user, and how errors in the specification of these models may impact the rankings of alternatives that they generate. Ravinder and Kleinmuntz (1991), Butler et al. (1997) and Jia et al. (1998) also provide results that explore the implications of errors in assessments on the actual quality of decisions. This research may provide useful guidelines for determining the numbers of trade-off questions that should be asked, or perhaps for identifying questions that will reduce the consideration set most efficiently.

The task of making the process pleasant seems to contradict the other recommendations; it may be difficult and unpleasant to think hard about preferences and carefully identify objectives and attributes, and to make weight judgments. Here again the contrast between a domain expert and novice is important. An expert may even enjoy thinking hard about her trade-offs and preferences. The cost of the item and the importance of the decision may also influence the willingness of the consumer to think hard about these issues. Perhaps the best solution is to provide multiple mechanisms for users with different abilities and desires. For example, Fig. 1 from Active Buyers Guide shows two interfaces to identify evaluation attributes.

5.2. The construction of preferences

A great deal of research has confirmed what a little introspection by anyone would recognize: preferences do not "exist" in the minds of individuals, but are constructed as the need for them arises (for discussions, see Slovic, 1995 and Bettman et al., 1998). Further, the preferences that are constructed by an individual may depend in a significant manner on the context of the choice situation, even when the alternatives are identical (Loewenstein, 2001). This insight offers two contrasting points of view that relate to the implementation of web-based MAPM decision aids.

First, to assess an accurate model of an individual's preferences, the process of questioning should be designed to minimize the impacts of well-known biases. For example, Tversky et al. (1998) documented the existence of the prominence effect: The more prominent (important) attribute will weigh more heavily in a choice between two alternatives than in a trade-off question between the same alternatives (Fischer et al., 1999).

We presented an example of a weight assessment question when considering the price and resolution of a digital camera in Section 2: $A = (€200, 1600 \times 1200$ pixels) and $B = (\_\_\_, 1920 \times 1600$ pixels); the decision maker is asked to specify a price for camera B so that the two alternatives are equally attractive. Fischer et al. (1999) and Tversky et al. (1998) would call this a matching task because the decision maker is asked to provide a response in terms of one of the attributes. Alternatively, a decision maker might be asked to choose between two fully specified alternatives in a choice task. For example, would you prefer camera $A' = (€200, 1600 \times 1200$ pixels) or $B' = (€300, 1920 \times 1600$ pixels).

Previous research suggests that in a matching task a decision maker is much more likely to place *less* weight on the prominent attribute when compared to a choice task. The experimental evidence is consistent with the conjecture that in the more qualitative choice task the prominent attribute is not formally priced out in the mind of the decision maker, while the stimulus in the matching task forces the decision maker to explicitly make the trade-off (Fischer et al., 1999). Since choice and matching questions may be used to assess the weights in a MAPM model, Keeney (2002) suggests using both and pointing out any implied inconsistencies in the responses to the decision maker to resolve this bias appropriately. Of course, this may increase the cognitive burden on the decision maker, especially if the approach is implemented in an automated environment, and may diminish her willingness to use the decision aid. This is again an issue of effort versus accuracy that must be considered in the design of a useful recommendation system.

More of the recommendation systems we surveyed have assessment procedures featuring choice tasks as opposed to matching tasks. However there have been some attempts to combine the two as illustrated in Fig. 6. A previous version of Active Buyers Guide asked consumers not only to state a preference (choice) for one of two options, but also to specify the strength of preference for one option over the other. This type of question may help balance the two types of assessment tasks.

Second, the observation that preferences may be constructed suggests that a website may be designed to manipulate preferences to increase the likelihood that a particular product or service will be recommended for consideration. A simple example should make this clear. As we have noted, a MAPM-based website will typically provide a list of objectives or attributes as the basis for the evaluation of a product or services. It is likely that many subjects may accept the set of objectives or attributes suggested for an evaluation, without expending the cognitive effort necessary to identify other attributes that might be more relevant to reflect their fundamental objectives. To the extent that preferences are actually constructed during the process of preference elicitation, the set of attributes that are provided may even influence the fundamental objectives that the user considers to be relevant for the evaluation of the product or service. Relatively inexperienced consumers lacking expertise regarding a product may be particularly susceptible to manipulation regarding the appropriate attributes to use in the evaluation and selection of an alternative, as suggested in the discussion of the use of a website to recommend a city in which to reside (Fig. 4).

Efforts to suggest product attributes that are favorable to a particular alternative are certainly

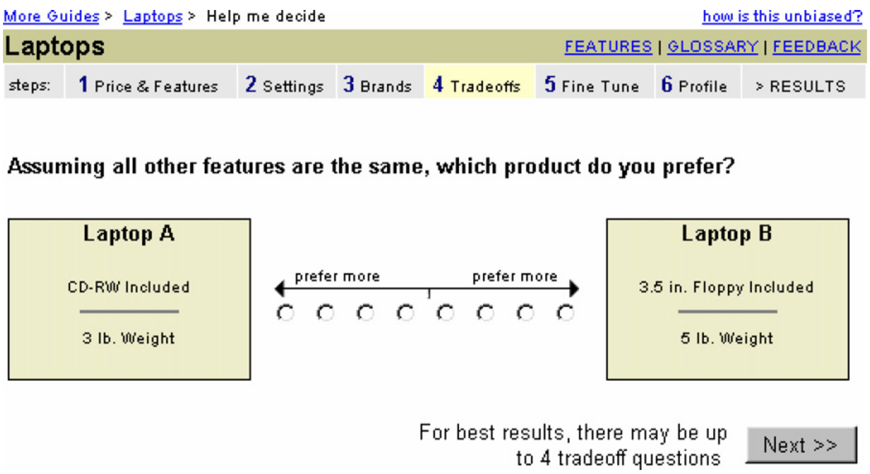

Fig. 6. Example of a combined choice/matching task from an old implementation of Active Buyers Guide.

not new. Product advertisements in the traditional media of newspapers and magazines will often offer side-by-side comparison matrices with rival brands, but with selected product criteria that clearly favor the sponsoring product. Even the choice of which alternative will serve as the basis of comparison can impact the assessed preferences. This suggests that "independent" shopbots and electronic agents in the B2C environment should be required to reveal any ties to particular products or services that they may be used to evaluate.

Preferences may also be biased, or constructed, based on even more subtle manipulation of MAPM assessment methodologies. Stillwell et al. (1987), Weber et al. (1988) and Borcherding and von Winterfeldt (1988) have shown that providing additional attributes related to some fundamental objective will cause a user of the MAPM methodology to assign too much "weight" to the objective, as implied by the numerical weights assessed for the individual attributes. Returning to the automobile example, an evaluation system that included a large number of attributes related to performance, such as torque, horsepower, power-to-weight ratio, 0–100 kmh acceleration, 60–80 kmh acceleration, and so forth, would encourage users to construct preference models that would give relatively more "weight" to the performance of the automobile. And, a consumer who used such a system could be influenced to conclude that performance is actually more important as a fundamental objective than she may otherwise have believed.

These observations suggest that a cleverly designed website could be used to construct the preferences of consumers in such a way that they would be influenced to choose more expensive vacations, say, or more sophisticated cameras or computers, simply through the dialogue used to elicit their preferences. This insight should come as no surprise, of course, since skillful salespersons have been doing the same thing for years. There is no reason to naively assume that the results of an interaction with an automated decision aid produce more "objective" product evaluations than one carried out with the aid of a salesperson.

An example of research that relates to this phenomenon is provided by Mandel and Johnson (1998), who show that the background of a website can suggest the importance of some of the product attributes, and alter user choices based on a single exposure to the decision aid. This is clearly a ripe area for further research, as noted by Geoffrion and Krishnan (2001), and also for a recognition that the age-old adage "buyer beware" will apply in the e-commerce environment as well.

6. Conclusions

It seems clear that MAPM will play an important role in the implementation of websites that provide shopping assistance to consumers in B2C environments and to consumers in B2B environments. We have reviewed a number of current implementations. When viewed from the perspective of MAPM and the lessons that have been learned based on research related to its application, many of the existing websites fall short of offering the quality of decision support that is feasible. There are some exceptions of course, and we have highlighted some sites that do represent insightful applications of MAPM concepts. We emphasize again that our observations are based on the implementations active at these sites in 2002, and that some of them may have been modified or deleted in the interim.

Many of the recommendations we have made are likely to increase the cognitive burden of the consumer. It is important to verify that the additional effort is rewarded with additional benefit (Payne et al., 1993). For example, Butler et al. (2006) investigate the trade-off between the number of required assessments associated with explicitly linking attributes to objectives and the resulting decision quality. They find that for all but very small decision problems, it is worthwhile for the decision maker to expend the effort associated with thinking hard about how the attributes impact her objectives. Similar work by Jia et al. (1998) suggests that while weight approximations do perform relatively well, it is usually worth the effort to perform a formal weight assessment. Nonetheless, we do not advocate adding more and more assessment burden without an investigation of the associated benefits.

Some scholars have argued that individuals do not make decisions based on a thoughtful consideration of objectives and their relative importance, but rather use a variety of decision heuristics that are strongly affected by the decision domain (e.g., Loewenstein, 2001). We would not advocate that every consumer would (or should) seek the assistance of a website based on MAPM principles when making routine purchases. Nevertheless, it is clear that many consumers are using the web as an aid to shopping decisions, whether the actual purchase is then transacted electronically or in a traditional retail environment.

As we have noted, many B2C websites, such as Active Buyers Guide, offer the consumer multiple approaches to identifying a most preferred alternative, in recognition of the influence of different factors on the consumer's choice of a decision aid. As we have discussed, MAPM has a role to play in this environment, but it may certainly be the case that some web-oriented shoppers will prefer product recommendations based on collaborative filtering or alternate approaches. These alternative approaches are most likely to be preferred by inexperienced consumers who need help specifying the attributes, the relationships between the attributes, and their relative importance.

The B2B environment should be an even more fertile ground, however, for MAPM applications, and some evidence in support of this observation is provided by the number of sophisticated software products and websites that we have reviewed above. The same scholars who have been skeptical of the potential for rational decision support systems for consumers offer an insight into why these applications may be successful for B2B consumer support. Loewenstein (2001) argues that individuals do not actually make decisions based on rational theories of decision making, but may "make sense" of those decisions in retrospect by describing the process in those terms.

In a B2B environment, the purchasing agent may feel the need to document the rationale for the decisions that are made, and to leave a "decision audit trail" in terms of the rational models that supported major purchasing commitments. The attributes and their weights may reflect an explicit interpretation of company policies, and perhaps reflect the consensus of several group discussions that are then delegated to one person to implement. Within these environments, the logic and appeal of MAPM would seem to make it a strong candidate for the documentation of purchasing actions.

References

Bettman, J.R., Luce, M.F., Payne, J.W., 1998. Constructive consumer choice processes. Journal of Consumer Research 25, 187–217.

Bichler, M., 2000. An experimental analysis of multi-attribute auctions. Decision Support Systems 29, 249–268.

Bichler, M., Kalagnanam, J., 2005. Configurable offers and winner determination in multi-attribute auctions. European Journal of Operational Research 160, 380–394.

Borcherding, K., von Winterfeldt, D., 1988. The effect of varying value trees on multi-attribute evaluations. Acta Psychologica 68, 153–170.

Butler, J.C., Jia, J., Dyer, J.S., 1997. Simulation techniques for the sensitivity analysis of multi-criteria decision models. European Journal of Operational Research 103, 531–545.

Butler, J.C., Morrice, D., Mullarkey, P., 2001. A multiple attribute utility theory approach to ranking and selection. Management Science 47, 800–816.

Butler, J.C., Dyer, J.S., Jia, J., 2006. Using attributes to predict objectives in preference models. Decision Analysis 3, 100–116.

Chen-Ritzo, C., Harrison, T.P., Kwasnica, A.M., Thomas, D.J., 2005. Better, faster, cheaper: An experimental analysis of a multi-attribute reverse auction mechanism with restricted information feedback. Management Science 51, 1753–1762.

Clemen, R.T., Reilly, T., 2001. Making Hard Decisions with DecisionTools. Duxbury Press.

Dyer, J.S., 1973. A time-sharing computer program for the solution of the multiple criteria problem. Management Science 19, 1379–1383.

Dyer, J.S., 1990. Remarks on the analytic hierarchy process. Management Science 36, 249–258.

Dyer, J.S., Sarin, R.A., 1979. Measurable multi-attribute value functions. Operations Research 22, 810–822.

Edwards, W., Fasolo, B., 2001. Decision Technology. Annual Review of Psychology 52, 581–606.

Farquhar, P.H., Keller, L.R., 1989. Preference intensity measurement. Annals of Operations Research 19, 205–217.

Fischer, G.W., 1995. Range sensitivity of attribute weights in multi-attribute value models. Organizational Behavior and Human Decision Processes 62, 252–266.

Fischer, G.W., Carmon, Z., Ariely, D., Zauberman, G., 1999. Goal-based construction of preferences: Task goals and the prominence effect. Management Science 45, 1057–1075.

Forman, E., Gass, S.I., 2001. The analytic hierarchy process: An exposition. Operations Research 49, 469–486.

Geoffrion, A.M., Krishnan, R., 2001. Prospects for operations research in the e-business era. Interfaces 31, 6–36.

Gibson, L., 2001. What's wrong with conjoint analysis. Marketing Research 13, 16–19.

Green, P., Krieger, A.M., Wind, Y., 2001. Thirty years of conjoint analysis: Reflections and prospects. Interfaces 31, S56–S73.

Haubl, G., Trifts, V., 2000. Consumer decision making in online shopping environments. Marketing Science 19, 4–21.

Jia, J., Fischer, G.W., Dyer, J.S., 1998. Attribute weighting methods and decision quality in the presence of response error: A simulation study. Journal of Behavioral Decision Making 11, 85–105.

Kahn, B.E., Huffman, C., 1998. Mass customization. In: Dorf, R.C. (Ed.), The Handbook of Technology Management. CRC Press, Inc.

Keeney, R.L., 2002. Common mistakes in making value trade-offs. Operations Research 50, 935–945.

Keeney, R.L., Raiffa, H., 1976. Decision with Multiple Objectives: Preference and Value Trade-offs. Wiley, New York.

Kersten, G.E., Noronha, S.J., 1999. WWW-based negotiation support: Design, implementation, and use. Decision Support Systems 25, 135–154.

Keskinocak, P., Goodwin, R., Rama, F.W., Akkiraju, R., Murthy, R.S., 2001. Decision support for managing an electronic supply chain. Electronic Commerce Research 1, 15–31.

Kirkwood, C.W., 1992. Estimating the impact of uncertainty on a deterministic multi-attribute evaluation. Management Science 38, 819–826.

Kirkwood, C.W., 1996. Strategic Decision Making. Duxbury Press, Belmont, CA.

Kirkwood, C., Sarin, R., 1985. Ranking with partial information: A method and an application. Operations Research 33, 38–48.

Kleinmuntz, D.N., 2007. Resource Allocation Decisions. In: Edwards, W., Miles, R.F. Jr.Jr., von Winterfeldt, D. (Eds.), Advances in Decision Analysis. Cambridge University Press.

Krantz, D.H., Luce, R.D., Suppes, P., Tversky, A., 1971. Foundations of MeasurementAdditive and Polynomial Representations, Vol. I. Academic Press.

Loewenstein, G., 2001. The creative destruction of decision research. Journal of Consumer Research 28, 499–505.

Luce, R.D., Tukey, J.W., 1964. Simultaneous conjoint measurement: A new type of fundamental measurement. Journal of Mathematical Psychology 1, 1–27.

Mandel, N., Johnson, E., 1998. When web pages influence choice: Effects of visual primes on experts and novices. Journal of Consumer Research 29, 235–245.

Milgrom, P., 2004. An economist's vision of the B-to-B marketplace. Executive Whitepaper, available from: www.perfect.com.

Mitchell, A.A., Dacin, P.A., 1996. The assessment of alternative measures of consumer expertise. Journal of Consumer Research 23, 219–239.

Montgomery, A.L., Hosanagar, K., Krishnan, R., Clay, K.B., 2004. Designing a better shopbot. Management Science 50, 189–206.

Montibeller, G., Ackermann, F., Belton, V., Ensslin, L., 2005. Reasoning Maps for Decision Aid: An Integrated Approach for Problem Structuring and Multi-Criteria Evaluation. Working Paper 2005/11, Department of Management Science, University of Strathclyde, (available from: http://www.managementscience.org/research/ab0511.asp).

Parkes, D.C., Kalagnanam, J., 2005. Models for iterative multiattribute procurement auctions. Management Science 51, 435–451.

Payne, J.W., Bettman, J.R., Johnson, E.J., 1993. The Adaptive Decision Maker. University Press, Cambridge.

Ravinder, H.V., Kleinmuntz, D.N., 1991. Random error in additive decompositions of multi-attribute utility. Journal of Behavioral Decision Making 4, 83–100.

Roy, B., 1996. Multicriteria Methodology for Decision Aiding. Kluwer Academic Publishers, Dordrecht.

Saaty, T.L., 1980. The Analytic Hierarchy Process. McGraw-Hill, New York.

Sarin, R., 1979. Ranking of multi-attribute alternatives with an application to coal power plant siting. Working Paper, Purdue University.

Slovic, P., 1995. The construction of preferences. American Psychologist 50, 364–371.

Stewart, T.J., Losa, F.B., 2003. Towards reconciling outranking and value measurement practice. European Journal of Operational Research 145, 645–659.

Stillwell, W.G., von Winterfeldt, D., John, R.S., 1987. Comparing hierarchical and nonhierarchical weighting methods for eliciting multi-attribute value models. Management Science 33, 442–450.

Teich, J.E., Wallenius, H., Wallenius, J., Koppius, O., 2004. Emerging Multi-Issue E-Auctions. European Journal of Operational Research 159, 1–16.

Teich, J.E., Wallenius, H., Wallenius, J., Zaitsev, A., 2006. A multi-attribute e-auction mechanism for procurement: Theoretical foundations. European Journal of Operational Research 175, 90–100.

Tversky, A., Sattath, S., Slovic, P., 1998. Contingent weighting in judgment and choice. Psychological Review 95, 371–384.

von Winterfeldt, D., Edwards, W., 1986. Decision Analysis and Behavioral Research. Cambridge University Press.

Weber, M., Eisenfuhr, F., von Winterfeldt, D., 1988. The effects of splitting attributes on weights in multi-attribute utility measurement. Management Science 34, 431–445.

West, P., Ariely, D., Bellman, S., Bradlow, E., Huber, J., Johnson, E., Kahn, B., Little, J., Schkade, D., 1999. Agents to the rescue. Marketing Letters 10, 285–300.