Available online at www.sciencedirect.com
ScienceDirect
Journal of Interactive Marketing 25 (2011) 110 – 122
JOURNAL OF INTERACTIVE MARKETING
www.elsevier.com/locate/intmar

Decision Quality Measures in Recommendation Agents Research

Lerzan Aksoy [a,*], Bruce Cooil [b] & Nicholas H. Lurie [c]

[a] *Fordham University, School of Business, 1790 Broadway, 11th floor, #1129, New York, NY 10023, USA*
[b] *Owen Graduate School of Management, Vanderbilt University, Nashville, TN 37203, USA*
[c] *College of Management, Georgia Institute of Technology, Atlanta, GA 30308-0520, USA*

Available online 3 March 2011

Abstract

Electronic recommendation agents provide a way for online marketers to gather information about consumer preferences and assess the quality of consumer decisions. Much of the literature on recommendation agents, however, employs divergent measures to assess consumer decision quality. Moreover, decision quality measures are dictated by the type of agent employed. This article provides a review of the decision quality measures used in the recommendation agent research to date and proposes novel measures. We classify and examine the assumptions of–and relationships among–preference-dependent, preference-independent, and subjective measures of decision quality. The analysis of data from an experiment that simulates a broad spectrum of recommendation agents shows that the relative utility and the sum of attribute values of the chosen alternative capture the majority of variance in objective decision quality. Although subjective decision quality measures turn out to be poor proxies for objective measures, they provide important incremental information. Managerial implications for deploying electronic recommendation agents to gather information and measure consumer decision quality under different conditions are discussed.

Keywords: Decision quality; Decision making; Measurement; Preferences; Recommendation agents; Online

There is little doubt that the Internet has greatly enhanced the variety of choices available to consumers. Despite its potential to enhance consumer utility, by presenting alternatives that more closely match individual preferences, increased choice may also overwhelm the customer and lead to worse choices (Jacoby 1977; Lurie 2004; Schwartz 2004). In response, several online sites (such as myproductadvisor.com) offer electronic recommendation agents to help consumers make decisions. Electronic recommendation agents have important implications for online service providers (Rust 2001) by influencing consumer perceptions of Web site attributes (Zeithaml, Parasuraman, and Malhotra 2002), changing relevant dimensions of service quality (Fassnatch and Köse 2007; Parasuraman, Zeithaml, and Malhotra 2005; Wolfinbarger and Gilly 2003), affecting the quality of consumer decisions (Aksoy et al. 2006; Häubl and Trifts 2000), changing the relative importance of different product attributes (Diehl, Kornish, and Lynch 2003; Häubl and Murray 2003), helping consumers identify their preferences (Kramer 2007), and offering ways to improve the customer experience (Rayport, Jaworski, and Kyung 2005).

Although decision quality can serve as an important benchmark for assessing service quality (Zeithaml et al. 2006), few mechanisms have been proposed for allowing providers to measure decision quality in any meaningful way. In addition, recommendation agents vary in the type of information they collect and this has implications for the ways in which decision quality can be assessed. Furthermore, there is no clear consensus on which measures provide the most information about different aspects of consumer decision quality and the relationships among these measures. Understanding and measuring customer decision quality is important because a customer who makes a good choice is likely to be happy, satisfied, and more likely to return and repurchase. Assessing decision quality is fundamental to those who study consumer judgment and decision making; yet researchers disagree about how decision quality should be measured (Jacoby 1977; Keller and Staelin, 1987; Keren and de Bruin 2003; Meyer and Johnson 1989). Importantly, there have been few attempts to theoretically and empirically compare alternative measures of

* Corresponding author.
E-mail addresses: aksoy@fordham.edu (L. Aksoy), bruce.cooil@owen.vanderbilt.edu (B. Cooil), lurie@gatech.edu (N.H. Lurie).

1094-9968/$ - see front matter © 2011 Direct Marketing Educational Foundation, Inc. Published by Elsevier Inc. All rights reserved.
doi:10.1016/j.intmar.2011.01.001

decision quality. Consequently, a comparison of the benefits of alternative measures of decision quality is of both practical and theoretical importance.

This paper reviews the different measures that have been used to date in the recommendation agent literature, classifies them based on their characteristics, and empirically examines the relationships among them. We distinguish among preference-dependent measures (PDM), that require knowledge of individual decision maker preferences as well as the attribute values of available alternatives; preference-independent measures (PIM), that do not require knowledge of individual decision maker preferences but do require product attribute values; and subjective measures (SM) of decision quality, that depend neither on knowledge of consumer preferences nor attribute values. To empirically compare these alternative measures, and test the generalizability of these comparisons, we use data from an experiment that simulates recommendations from a broad spectrum of online recommendation agents.

Examining the relationship among different types of measures of decision quality is important from the standpoint of gaining a more complete perspective on the customer; particularly understanding the chain of effects running from marketing strategies and tactics to firm value (Rust et al. 2004). It is also important from a public policy standpoint. For example, in the realm of financial decision making, in which consumers can choose among thousands of funds, and may turn to online recommendation agents such as financialengines.com for advice, it is important that consumers not only feel good about their investments but in fact make investments that allow them to meet their needs in retirement. Thus examining objective as well as subjective decision quality is important. By classifying the measures available to online recommendation agent providers, identifying the best among each type, and empirically assessing the extent to which different measures provide additional information, this article helps researchers, managers, and public policy makers gain a more complete perspective on the quality of consumer decisions.

Because the quality of a decision and how to measure it depends on one's role (e.g., decision maker vs. judge), perspective (e.g., short term vs. long term), and beliefs about whether human decision makers should be judged relative to normative (i.e., perfect information processing) versus prescriptive (i.e., limited information processing) standards (Keren and de Bruin 2003), we avoid claiming that a single definition exists and that a particular measure of decision quality is superior. Instead, we seek to examine the extent to which different measures of decision quality provide new information, and potentially new insights, to online marketers. As part of this analysis we compare different measures relative to weighted additive utility (WADD), one of the most popular measures of decision quality and the basis for most economic models of utility.

In addition, although it is clear from prior research that different recommendation agents influence the degree to which a customer makes a "good" choice (Aksoy et al. 2006; Häubl and Trifts 2000), it is unclear how to best *measure* decision quality given a particular type of recommendation agent. We argue that using and employing different types of electronic agents inherently *determines* the type of information that is gathered and therefore the types of measures that can be used to assess a customer's decision quality. This research therefore provides insights into the implications of deploying different types of recommendation agents and what this means for the online provider's ability to assess the quality of customer decisions.

We address the following managerially-relevant research questions:

1) What is the relationship among preference-dependent (PDM), preference independent (PIM), and subjective (SM) measures of decision quality?
2) What are the best indicators of decision quality relative to weighted-additive utility (WADD; an objective measure of choice quality often used by decision researchers and the basis for most economic models of utility)?
3) How can decision quality, as a multidimensional variable, be summarized in terms of underlying measures, and which are the most important individual measures (overall and within each group of measures)?

Our results show that PIM measures of decision quality capture about 63% of the variance of PDM measures whereas subjective SM measures capture about 9% of the variance of preference-dependent and preference independent measures. This suggests that the information collected by recommendation agents can provide insights into consumer welfare that are not provided by traditional surveys of customer satisfaction. In addition, a principal component analysis shows that decision quality is a multi-dimensional construct best assessed by combining select measures of decision quality.

Using recommendation agents as a context to study decision quality is particularly appropriate for the following reasons:

1) The study of recommendation agents is a popular topic and hence understanding the connections among different types of recommendation agents and different measures of decision quality is relevant to a large group of researchers.
2) The use of recommendation agents in practice has become much more mainstream and therefore this study is relevant to a large group of service providers.
3) There is a strong connection between the type of recommendation agent used and the type of measures that can be calculated. In particular, certain types of agents make it impossible to calculate preference dependent measures because no information about consumer preferences is collected.
4) There is currently no consensus in research conducted on recommendation agents as to which decision quality measures provide the most information and under which circumstances.
5) To investigate the relationships among decision quality measures, an environment in which there is variance in decision quality is required. Recommendation agents provide a good context for manipulating decision quality.

In the next section we provide an overview of different types of electronic recommendation agents and their implications for measuring decision quality. Next, we describe alternative measures of decision quality and compare these measures in terms of the extent to which they require knowledge of consumer preferences and knowledge of attribute values. In outlining alternative measures, our goal is to illustrate the assumptions, informational and otherwise, of different measures of decision quality.

How Recommendation Agents Affect the Measurement of Decision Quality

Recommendation agents vary in the extent to which they elicit consumer preferences and use product attributes to make recommendations. Some agents, like about.com, provide alphabetically ordered lists of choice alternatives that neither account for consumer preferences nor attribute values; while agents that rely on collaborative filtering techniques, such as Amazon.com, base their recommendations on customers' purchase or browsing histories, rather than product attribute values or elicited consumer preferences (Ansari, Essegaier, and Kohli 2000; Iaobucci, Arabie, and Bodapati 2000). Many of the recommendation agents studied in academic research (e.g., Aksoy et al. 2006; Diehl, Kornish, and Lynch 2003; Häubl and Trifts 2000) ask consumers to provide importance weights and use these weights along with attribute values to rank alternatives; others elicit preferences through conjoint tasks (De Bruyn et al. 2008); while others generate product recommendations using only attribute values by weighting each attribute equally (Olson and Widing 2002). Notably, because different types of agents use different types of information to make recommendations, providing consumers with different types of agents has important implications for measuring decision quality. For example, some measures of decision quality require knowledge of consumer preferences while others can be calculated independently from preferences. Thus the availability of information, and types of measures, that can be used to assess decision quality is determined by the particular online recommendation agent employed.

Table 1 defines the measures identified and tested in this article, their requirements in terms of service provider knowledge of consumer preferences and attribute weights, and provides examples of service providers who could apply these measures.

Preference-Dependent Measures (PDM)

We identify preference dependent measures as those that require knowledge of the decision maker's preferences, often determined by eliciting decision makers' preference weights. The literature has established that recommendations from agents that take into account consumer attribute weights and preferred decision strategies lead to better decisions compared to randomly ordered alternatives (Aksoy et al. 2006; Ariely 2000; Diehl, Kornish and Lynch 2003; Häubl and Trifts 2000). Further, Punj and Moore (2007) found that "smart" agents, that recommend products that meet consumer preferences, are more effective in helping consumers make less effortful decisions compared to "knowledgeable" agents that identify only whether a particular product is carried by a retailer.

Weighted-Additive Utility

Consistent with multi-attribute utility theory (McFadden 1986), one of the most common approaches to measuring decision quality is to calculate the utility of chosen alternatives using a weighted-additive, or linear, model (Payne, Bettman, and Johnson 1993) in which the products of decision maker preferences (attribute weights) and attribute values are summed across attributes. Weighted-additive utility (WADD) is one of the most popular measures and the basis for most economic models of utility (Payne, Bettman, and Johnson 1993), where the higher the utility, the better the choice is assumed to be (Barron and Barrett 1996). To calculate this measure, preference weights and attribute values must be known. This approach also assumes that attributes can be expressed on common interval scales in which unit changes are equivalent for different attributes and for which every unit increase in a given attribute has an equal impact on utility.

Relative Utility

This is a variation of weighted-additive utility that accounts for utility of the chosen alternative relative to the best and worst alternatives in the choice set. Choice of the best alternative leads to a relative utility of one whereas choice of the worst alternative leads to a relative utility of zero (see Table 1). It is thus particularly appropriate for making comparisons across choice sets that vary in absolute utility. If there is no variation in choice set quality, weighted-additive and relative utility are functionally equivalent. Relative utility involves the same informational requirements and assumptions as weighted-additive utility.

Euclidean Distance in Utility (Euclidean Distance)

Another approach assesses quality of choice in terms of how close a consumer comes to the best choice available to her (Sproles 1983). The weighted-additive utility of the best alternative in the choice set for a given individual and that of the chosen alternative are determined and a distance index is calculated. Some have used a measure of absolute difference in utility (e.g., Aksoy et al. 2006; Hahn, Lawson, and Lee 1992) while others have used a relative index (e.g., Keller and Staelin 1987).

Preference-Independent Measures (PIM)

Unlike preference-dependent measures, which rely on knowing the decision maker's attribute importance weights, preference-independent measures require no such knowledge. Instead, preference-independent measures assess choice quality by comparing the attribute values of chosen and non-chosen alternatives.

Non Dominated Choice (Non Dominated)

Choice of non dominated alternatives (i.e., alternatives that are superior to other alternatives in the choice set on at least one attribute) is one way to objectively measure choice quality independently from preferences (Häubl and Trifts 2000; Swaminathan 2003). In addition to not requiring knowledge of consumer preferences, an advantage of this approach is it makes no assumptions about decision processes. That is, the choice of a dominant alternative is a good choice regardless of individual preferences or decision rules employed (Payne, Bettman, and Johnson 1993). A disadvantage is that its usage depends on the existence of non-dominated and dominated alternatives in the choice set—which may be rare in real-world settings (Häubl and Trifts 2000; Lurie 2004).

Number of Dominant Alternatives (Dominant Alternatives)

This is a related measure employed by Aksoy et al. (2006) that counts the number of alternatives that dominate the chosen alternative. The advantage of this measure over the non-dominated choice measure is that it offers an ordinal, rather than binary, assessment of decision quality. Like the non-dominated choice measure, this measure can be determined independent of preferences and external ratings and makes no assumptions about the decision process but requires that dominant and dominated alternatives exist in the choice set.

Attribute Sum

Another preference-independent measure of choice quality is the sum of the attribute values of the chosen alternative. Attribute sum is equivalent to weighted additive utility if all attributes have equal weight (Payne, Bettman, and Johnson 1993). In addition to assuming that all attributes are equally important and that consumers engage in compensatory decision processes, this approach also assumes that attributes can be expressed on common interval scales in which unit changes are equivalent for different attributes and for which every unit increase in a given attribute has an equal impact on utility.

Difference in Attribute Sum (Attribute Difference)

A related measure proposed here for the first time is to assess the difference between the attribute sum of the chosen alternative and that of the best alternative in the choice set. According to this measure, a chosen alternative with a lower attribute difference score is assumed to be a better choice than one that has a larger difference score. As with the Euclidean Distance in Utility measure, an advantage of this measure over the attribute sum measure is that it accounts for differences across choice sets by offering a relative assessment of decision quality. The assumptions of this difference measure are identical to those of the attribute sum measure.

Subjective Measures (SM)

Although objective measures of decision quality are of high relevance, subjective evaluations (i.e., how happy the decision maker is with the choice) need not be correlated with the objective measures and therefore can provide additional valuable insights into decision effectiveness in the context of recommendation agents (Häubl and Trifts 2000; Lilien et al. 2004). Some research suggests that personalized recommendations, while enhancing objective decision quality, can lower choice satisfaction (Häubl, Dellaert, and Usta 2010). Subjective measures are particularly useful for assessing consumer evaluations of the choice process and their feelings post purchase.

Perceived Fit with Preferences (Fit)

One subjective measure that may be particularly relevant in the context of agent-assisted choice is the extent to which consumers believe that they chose an alternative that fit their preferences. This measure offers a nice comparison to the objective Euclidean distance measure.

Choice Confidence (confidence)

Several authors (e.g., Häubl and Trifts 2000; Keller and Staelin 1987) have used choice confidence as a measure of subjective choice quality in addition to objective decision quality measures. In addition to depending on individual characteristics, such as product experience (Bearden, Hardesty, and Rose 2001), confidence is likely to depend on the choice set size and difficulty in making tradeoffs among alternatives (Payne, Bettman, and Johnson 1993).

Choice Satisfaction (Satisfaction)

Unlike perceived fit, in which consumers are asked to reflect on their preferences, choice satisfaction provides an assessment of overall happiness with choice. This suggests it is more akin to distance in individual utility. At the same time, previous research suggests that measuring satisfaction may shift consumer preferences by changing the mental processes through which alternatives are compared (Shiv and Huber 2000).

Choice Liking (Liking)

Choice liking serves as a perceptual measure of the utility associated with a product alternative. Unlike the perceived fit with preferences measure, which asks consumers to compare their choice with their ideal on each attribute, choice liking asks the consumer for a Gestalt assessment. Importantly, research suggests that measuring liking does not lead to significant shifts in preferences (Shiv and Huber 2000).

Interest in Choice (Interest)

Interest in choice reflects how excited the decision maker is with her choice. This may have important implications for her efforts to fully use the chosen product and whether she will experience post-purchase regret if the product does not meet her needs. As with the liking measure, interest in choice offers a perceptual alternative to utility-based measures by asking consumers to provide an overall evaluation of their choice.

Method

To better understand the relationships among alternative measures of decision quality, and how these relationships depend on the type of electronic recommendation agent used, an

Table 1
Alternative measures of decision quality and implications of electronic recommendation agents.

Measure	Description/operationalization	Requirements and assumptions	Examples of service providers that could use measure
Preference-dependent (PDM)			
Weighted-additive utility (WADD)[a]	Strict linear utility of the chosen alternative calculated as the sum of attribute values multiplied by the decision maker's relative preference weights $U_i = \sum_{k=1}^{m} w_k x_{k(choice)}$, where $x_{k(choice)}$ is the observed value of attribute k for the chosen alternative (where all attributes are expressed on common interval scales and higher values are preferred on all attributes), wk is the relative importance of that attribute to the decision maker, on a 0 to 1 scale, $w_k = a_k / \sum_{k=1}^{m} a_k$ and a_k is the stated preference for attribute k (1 to 10 scale).	Decision maker preferences known Attribute values known Decision maker engages in comprehensive review of all important information and makes decision in compensatory fashion.	myproductadvisor.com shopping.yahoo.com/smartsort
Relative utility[a]	Utility of the chosen alternative relative to the best and worst options in the choice set. $U_i = \frac{U_{choice} - U_{worst}}{U_{best} - U_{worst}}$ U_i ranges from 0 to 1 with 0 being the worst choice and 1 being the best choice.	Decision maker preferences known Attribute values known Decision maker engages in comprehensive review of all important information and makes decision in compensatory fashion.	myproductadvisor.com shopping.yahoo.com/smartsort myproductadvisor.com
Euclidean distance in utility[a]	Difference in utility between the chosen and best alternative (reverse scaled). $U_i = U_{best} - U_{choice}$	Decision maker preferences known Attribute values known Decision maker engages in comprehensive review of all important information and makes decision in compensatory fashion.	shopping.yahoo.com/smartsort
Preference-independent (PIM)			
Non dominated choice[a]	Choice of an alternative that is not simply dominated by any other alternative in the choice set (i.e., the chosen alternative is better than or equal to other alternatives on every attribute). 1 if non-dominated; 0 otherwise.	Attribute values known Simple dominance exists in choice set No assumptions about choice processes	myproductadvisor.com shopping.yahoo.com/smartsort cdw.com

Number of dominant alternatives[a]	Number of alternatives that simply dominate the chosen alternative.	Attribute values known Simple dominance exists multiple times in choice set No assumptions about choice processes	myproductadvisor.com shopping.yahoo.com/smartsort cdw.com
Attribute sum[a]	Sum of attribute values of chosen alternative where higher values are preferred on all attributes. $U_i = \sum_{k=1}^{m} x_{k(choice)}$	Attribute values known All attributes are equally important Attributes can be expressed on common interval scales in which unit changes are equivalent for different attributes and for which every unit increase in a given attribute has an equal impact on utility. Decisions made in compensatory fashion.	myproductadvisor.com shopping.yahoo.com/smartsort cdw.com
Difference in attribute sum[a]	Difference between the sum of attribute values of the alternative with the highest sum and that of the chosen alternative. $U_i = \max_i \sum_{k=1}^{m} x_{ik} - \sum_{k=1}^{m} x_{k(choice)}$ where $\max_i$ is the choice with the highest sum of attribute values.	Attribute values known All attributes are equally important Attributes can be expressed on common interval scales in which unit changes are equivalent for different attributes and for which every unit increase in a given attribute has an equal impact on utility Decisions made in compensatory fashion.	myproductadvisor.com shopping.yahoo.com/smartsort cdw.com
Subjective (SM)			
Perceived fit with preferences[a]	How well do you think that the alternative you chose fits your preferences (1 = does not fit my preferences well; 7 = fits my preferences well)?	No assumptions about choice processes	Any service provider
Choice confidence[a]	How confident are you with the choice that you have made (1 = not at all confident; 7 = very confident)?	No assumptions about choice processes	Any service provider
Choice satisfaction[a]	How satisfied are you with the choice that you have made (1 = not at all satisfied; 7 = very satisfied)?	No assumptions about choice processes	Any service provider
Choice liking[a]	How much do you think you would like the alternative that you chose (1 = not at all; 7 = very much)?	No assumptions about choice processes	Any service provider
Interest in choice[a]	Please indicate your interest in the alternative that you chose (1 = not at all interested; 7 = very interested).	No assumptions about choice processes	Any service provider

[a] Tested in this article.

empirical analysis of decision quality measures from a study that manipulated the presence and type of electronic agent used to make recommendations. We were particularly interested in examining: 1) the relationships among PIM, PDM, and SM measures of decision quality; 2) the best indicators of decision quality relative to weighted-additive utility (WADD; one of the most popular measures and the basis for most economic utility models) and 3) how decision quality, as a multidimensional variable, can be summarized in terms of the underlying measures, and which of these are the most important (overall and within each group of measures).

Experimental Context

To examine the relationships among these decision quality measures, it was necessary to create *variance* in decision quality such that different experimental conditions led consumers to make a range of decisions—from better, to average, to poor. This was done by randomly assigning participants to conditions that simulated a variety of online recommendation agents. In particular, 313 participants (half male and half female with an average age of 21) were randomly assigned to 10 conditions that varied in the degree to which the consumers' preferences (attribute weights and preferred decision strategy) were incorporated in product recommendations. This allowed us to examine potential differences in decision quality measures among agents that do not use consumer preferences in their recommendation process (such as http://www.amazon.com) or fully integrate consumer preferences in the recommendation process (such as myproductadvisor.com). Through this experimental design, we were able to simulate and mimic a broad spectrum of recommendation agent types currently offered by online providers to their customers.

Basic Method

Stimulus

In the study (programmed using Authorware software with stimuli used by Aksoy et al. 2006), the electronic agent made recommendations about a database of 32 cellular phones. Each cell phone was defined by four attributes–low price, light weight, long talk time, and long stand-by time–that could take on values ranging from 1 to 10 where higher values were always preferred. Among the 32 alternatives, 10 were non-dominated (dominant) and 22 were dominated. Appendix A shows the attribute values for the 32 alternatives used in the study. Participants were asked to search and then make a choice. Incentives were not used. However, as Aksoy et al. (2006) point out, the differences across experimental groups in decision quality are not consistent with an incentive argument in that those who spent the most time on the task did *not* make better decisions.

Experimental Conditions

Data were collected from undergraduate business students who provided attribute preference weights and made selections from among the 32 cell phones under 10 distinct experimental conditions.[1] The participants went through several distinct steps before arriving at their choice. They accessed the simulated web site that included cell phones in its product portfolio and were instructed to choose a cell phone for themselves using this website and its recommendation agent. They were then asked to rate their preferred attribute importance weights for the 4 cell phone features. The 10 conditions differed in terms of whether (1) the attribute weights used by the agent were the same as, similar, or dissimilar to those of the participant and (2) the participant perceived the decision process used by the agent to be similar or dissimilar from the participant's preferred decision process. Attribute weight similarity was manipulated by adding a small or large random value to the participant's stated preference weights. Participants were then presented product lists that were generated based on preferences that were similar or dissimilar from their own which meant those in the similar condition viewed a rank ordered alternative list that more closely matched their stated preferences and those in the dissimilar condition viewed a rank ordered alternative list that was discrepant with their stated preferences. Participants then searched through the 32 rank ordered alternatives by clicking on each alternative on the recommended list and were ultimately asked to choose a cell phone at the end of the task. A control condition was also included, where alternatives were presented in random order. A filler task was added to the control conditions to eliminate differences in cognitive load between these and the other conditions.

Preference-dependent, preference-independent, and subjective measures of decision quality were collected and/or calculated for all participants. Preference dependent measures were *weighted-additive utility, relative utility*, and *Euclidean distance*. Preference independent measures were *non-dominated choice, dominant alternatives, attribute sum,* and *attribute difference*. Subjective measures were *fit, confidence, satisfaction, liking*, and *interest*. (See Table 1 for details.) In addition, the covariates of age, gender, online searches per month, perceived purchase online relative to classmates, purchases online per year, purchases online per month, self-assessed information seeking tendency (Bearden, Netemeyer, and Teel 1990), and self-assessed subjective knowledge about cell phones (Flynn and Goldsmith 1999) were measured.

Results

To make the different measures of decision quality comparable, each measure, with the exception of the non-dominated choice, was transformed into a ranked variable (i.e., instead of the original raw decision quality, z_i, we use the rank $y_i = n - k + 1$, where z_i is the kth largest quantity measured, so that if n is the number of observations, $100 * [yi/(n+1)]$ is the percentile corresponding to z_i), where larger ranks correspond to higher utility. Transforming to ranks provides a common

[1] The authors would like to acknowledge data overlap with Aksoy et al. (2006). The aforementioned research, however, utilizes only a subset of the extensive decision quality data measured, calculated, and analyzed in the current research.

scale that also allows us to interpret the correlations as the standardized change in rank (or percentile) on one measure relative to a change in rank (or percentile) on the other. Although this use of ranks does, in a sense, degrade the scales used by the PDM measures, it clearly facilitates comparisons among all measures. Also, the correlations among the ranks are smaller in absolute terms, than among the original raw measures, and therefore this transformation provides a more conservative assessment of the relationships among measures.

To gain a sense of the overall differences among preference-dependent, preference-independent, and subjective measures of decision quality, partial correlations among each of the ranked measures were calculated. In computing partial correlations (see Table 2) we controlled for all subject and experimental variables.

Comparing partial and regular correlations shows that they generally differ from the regular correlations by less than 5%, except between the objective and subjective measures, where partial correlations are sometimes smaller by 15% (but often the reduction is less than 10%). We also formally tested whether the correlation matrix for the measures estimated from subjects receiving "no help" (38% of the sample, consisting of those cases where alternatives were not ordered, or where the agent's attribute weights and decision strategies were both dissimilar) was significantly different from the correlation matrix estimated from the rest of the sample. This test (Jenrich 1970) shows no significant difference ($p=.46$). In summary, the relationships among the measures are relatively uniform across experimental conditions.

Table 2 organizes the decision quality measures by type (preference-dependent, preference-independent, and subjective), and orders them according to their correlation with two primary preference-dependent measures, *weighted additive utility* and *relative utility*. All correlations among PDM and PIM measures are significant at the .001 level (1-sided). This is also true among all subjective measures, and among PDM, PIM, and the first three subjective measures (*fit*, *confidence*, and *satisfaction*). Correlations are highest for: (1) the PIM measures *dominant alternatives*, *attribute difference*, and *non dominated*, where *dominant alternatives* and *non dominated* are nearly equivalent, and otherwise all pairwise R-square values exceed 75%; (2) the PDM and PIM measures: *weighted additive utility*, *relative utility* and *attribute sum*, where pairwise R-squares all exceed 70%; and (3) the subjective pairs, *confidence* and *satisfaction* (R-square=45%), and *satisfaction* and *liking* (R-square=47%).

Research Question 1: What is the relationship among preference-dependent (PDM), preference independent (PIM), and subjective (SM) measures of decision quality?

To address this question, we first conducted a canonical correlation analysis between PDM and PIM measures controlling for all subject and experimental variables. Canonical correlation analysis has been shown to be relatively robust to the non-normality of the underlying variables, and in many cases rank correlations tend to provide the greatest protection from false positives when using Barlett's test for residual correlation (Branco et al. 2005; Mantalos and Shukur 2007). The redundancy index shows that 63% of the standardized variance in each group of objective measures is explained by the other and only the first canonical variate pair is significant ($p<.001$ based on Rao's F-test; Bartlett's test for significant residual correlation beyond one canonical variate is not significant, $p=.52$). The preference dependent variate is virtually indistinguishable from *relative utility* (correlation=.996) and has a canonical correlation of .89 with the preference independent variate, which is essentially equivalent to the *attribute sum* variable (correlation=.999). Therefore, the two objective

Table 2
Partial correlations among decision quality measures, conditional on all subject and experimental variables (n=313).

Measure	Preference-dependent			Preference-independent				Subjective				
	1	2	3	4	5	6	7	8	9	10	11	12
Preference-dependent measures (PDM)												
1. Weighted additive utility	–	0.94	.73	.84	.47	.46	.45	.31	.26	.28	.21	.15
2. Relative utility		–	.76	.88	.48	.47	.46	.30	.29	.28	.20	.18
3. Euclidean distance			–	.62	.35	.36	.33	.28	.24	.21	.10 [a]	.08 [b]
Preference-independent measures (PIM)												
4. Attribute sum				–	.52	.51	.49	.24	.27	.26	.16	.17
5. Dominant alternatives					–	.89	.99	.17	.20	.21	.14	.22
6. Attribute difference						–	.87	.15	.15	.16	.10 [a]	.17
7. Non dominated							–	.16	.19	.20	.12 [a]	.22
Subjective measures (SM)												
8. Fit								–	.47	.60	.61	.37
9. Confidence									–	.67	.53	.33
10. Satisfaction										–	.68	.44
11. Liking											–	.43
12. Interest												–

Note. Except for *non dominated*, which is an indicator for whether the chosen alternative was not dominated, all correlations are with respect to ranked variables. Correlations among ranked variables would remain unchanged if ranks were replaced by empirical percentiles, so that the 100th percentile represents the highest possible value on each quality measure. That is, ranks were assigned so that larger ranks correspond to higher utility. Unless otherwise noted, all correlations are significant ($p<.001$).

[a] $p<.05$.

[b] $p>.05$.

measures of *relative utility* and *attribute sum* provide the most information about PDM and PIM respectively.

Next, we conducted a separate canonical correlation analysis to examine the relationship between all of the objective (PDM and PIM) and all of the subjective measures (SM) controlling for all subject and experimental design variables (see Table 3). Overall there is a significant relationship between the objective and subjective measures ($p<.001$). However, there are three significant pairs of canonical variates showing that different subjective measures are correlated with different objective measures (Bartlett's test for significant residual correlation indicates significant structure beyond one pair, $p=.020$; marginal significance beyond 2 pairs, $p=.096$; but no significant structure beyond three, $p=.584$). The first three canonical correlations are .38, .25, and .23 (R-squared values of 14, 6 and 5%, respectively). The first pair of canonical variates represents the correlation between objective (PDM and PIM) versus subjective measures as a whole. The second pair of canonical variates shows a correlation between the objective variate consisting primarily of *Euclidean distance* and *dominant alternatives*, and the subjective variate consisting primarily of a contrast between *fit* with the other subjective measures (especially *interest*). In the third pair, the objective measure is a contrast between *weighted additive utility* and PDM and PIM measures other than *relative utility*; the subjective variate contrasts *confidence* with *liking*, *fit* and *satisfaction*. Despite the significant correlations between canonical variate pairs, the Stewart–Love canonical redundancy index indicates that only 9% of the standardized variance in each group of measures (objective versus subjective) is explained by the other. Clearly, the subjective measures provide substantial additional information that is not provided by the objective measures.

Results comparing PDM and PIM decision quality show there are substantial gains from knowing consumer preferences—such as those gathered by recommendation agents. At the same time, PIM measures offer a reasonable way to assess decision quality when preferences are unknown. This suggests that setting up mechanisms to store and compare product attribute values may be beneficial even if recommendation agents are not employed. These results also suggest that, among the PDM and PIM groups, there are certain measures (i.e., *relative utility* and *attribute sum*) that are more effective than other measures at capturing the variance associated with these groups.

Results comparing SM to PIM and PDM measures of decision quality suggest that SM measures are poor proxies for the objective measures. This is in line with the findings of Häubl and Trifts (2000); Häubl, Dellaert, and Usta (2010) who found that the positive effect of recommendation agents on SM measures of decision quality (confidence) was much weaker than that for the objective measures, and that the correlations between SM and the

Table 3
Canonical objective and subjective variate pairs, conditional on all subject and experimental variables (n=313).

Variate pair	1	2	3
Canonical correlations	.380	.245	.232
Canonical loadings: correlation with each measure			
Objective variates			
	1	2	3
	Overall PDM and PIM	Attribute proximity	Utility beyond attribute proximity
Reification			
Preference-dependent (PDM)			
Weighted additive utility	.88	.17	.25
Relative utility	.92	.03	.08
Euclidean distance	.82	.51	−.16
Preference-independent (PIM)			
Attribute sum	.83	−.03	−.12
Dominant alternatives	.67	−.40	−.11
Attribute difference	.54	−.22	−.07
Non dominated	.64	−.41	−.17
Subjective variates			
	1	2	3
	Overall subjective	Disinterested fit	Desired fit
Reification			
Subjective (SM)			
Fit	.83	.27	.42
Confidence	.81	−.15	−.19
Satisfaction	.79	−.15	.26
Liking	.54	−.33	.70
Interest	.55	−.72	.04

objective measures of decision quality were weak. Therefore, if service providers wish to have a complete view of consumer decision quality they should not rely on consumers' self-assessments; rather, the objective quality of consumers' choices should also be assessed. This implies that attribute values and, ideally, consumer preferences be collected (perhaps by recommendation agents). These results however also suggest that, although they are not substitutes for PDM and PIM measures, the SM measures of decision quality provide substantial additional information beyond the objective measures.

Research Question 2: What are the best indicators of decision quality relative to weighted-additive utility (WADD; an objective measure of choice quality often used by decision researchers and the basis for most economic models of utility)?

Previous research has proposed weighted-additive utility (WADD) as the normative standard for choice quality (Payne, Bettman and Johnson 1993). Given this, the best scientific linear model was explored for *weighted-additive utility*, by considering all possible regressions on the other measures and selecting the models that minimized the Bayesian Information Criterion (BIC; Schwarz 1978; Bozdogan 1987). Again, with the exception of the measure *non dominated*, each measure was expressed in rank form, so that the coefficients represent the marginal increase in *weighted additive utility*, in percentiles, per one unit percentile increase in the predictor measure. When all measures are considered, the best scientific model is simply

$$\textit{weighted-additive utility} = .99\,[\textit{relative utility}]. \tag{1}$$

This means that, on average, *weighted-additive utility* increases .99 percentile per percentile increase in *relative utility* (the 100th percentile represents the maximum level of each measure), where the coefficient's standard error is .009. This regression accounts for 89% of the variance in *weighted additive utility*. The only other measures that have significant ($p<.05$) incremental predictive value when added to this model are the SM of *fit* ($p=.005$) and *satisfaction* ($p=.021$). No other measures, including *satisfaction*, have significant incremental value once *fit* is added to the model. In other words, *relative utility* is the best predictor of *weighted-additive utility* and including *fit* in the model leads to a small but significant improvement in the model fit.

The best scientific model based on PIM and SM is:

$$\textit{weighted-additive utility} = \underset{(.026)}{.82}\,[\textit{attribute sum}] + \underset{(.027)}{.17}[\textit{fit}] \tag{2}$$

(coefficient standard errors in parenthesis), and this regression accounts for 71% of the variance in *weighted additive utility*. Only the subjective measure *satisfaction* ($p=.021$) has significant marginal incremental predictive value when added to this model.

Finally, the best scientific model based on only SM is:

$$\textit{weighted-additive utility} = \underset{(.035)}{28.4} + \underset{(.071)}{.23}[\textit{fit}] + \underset{(.071)}{.21}[\textit{satisfaction}] \tag{3}$$

This model accounts for only 13% of the variance in *weighted-additive utility* and none of the other subjective measures have significant incremental predictive value. The best scientific models in Equation (2) and (3) indicate that among the SM, *satisfaction* and *fit* are the measures most closely related to *weighted-additive utility*, and other SM have no significant incremental value in predicting *weighted-additive utility*. This is not to say that the SM measures are not capturing something important, but they are generally assessing another aspect of decision quality.

Research Question 3: How can decision quality, as a multidimensional variable, be summarized in terms of underlying measures, and which are the most important individual measures (overall and within each group of measures)?

To address this question, two principal component analyses were conducted using the rank transformed variables. The first was based on the regular correlations and the second on the partial correlations from Table 2 (shown in parentheses in Table 4). There is very little difference between the two sets of analyses and their corresponding eigenvalues. These analyses reveal three latent dimensions of decision quality summarized by three components that capture nearly 80% (46%+20%+13%) of the total standardized variance (see Table 4). These three dimensions are:

1) *Overall decision quality*, which is positively correlated with all measures, accounts for 46% of the total standardized variance, and has the highest nominal correlations with two PDM measures (*weighted additive utility* and *relative utility*). Among the PIM measures, *overall decision quality* is most strongly correlated with *attribute sum*. Among subjective measures, its largest correlation is with *satisfaction.*
2) *Subjective (SM) measures* versus *objective (PDM and PIM) measures of decision quality*. Among these, *liking* is the most correlated with the objective measures (.71) followed by *satisfaction, fit, confidence* and *interest.* The SM measures are less correlated with PIM measures and least correlated with PDM measures alone.
3) *Utility* versus *relative optimality and interest*, where utility is captured by all PDM measures and the PIM measure *attribute sum*, and relative optimality and interest is captured by all PIM measures other than the *attribute sum* and SM *interest*.

In other words, decision quality can indeed be described as an overall dimension that includes PDM, PIM and SM measures lending support to the conceptualization of decision quality proposed in this paper, that the objective measures (PDM and PIM) are distinct from SM measures. In particular, SM measures can provide distinct information about decision quality. In addition, the combination of *weighted additive utility* (PDM) and *attribute sum* (PIM) are distinct from the remaining PIM measures such as *dominant alternatives, attribute difference* and *non dominated.*

Conclusions and Managerial Implications

Electronic recommendation agents offer a way to help customers sort through and make good choices among ever

Table 4
Principal components of sample correlation matrix (and partial correlation matrix; n=313).

Component	1	2	3
Reification	Overall decision quality	Subjective measures versus PDM/PIM	Utility versus relative optimality and interest
Proportion of variance (%)	46 (45)	20 (21)	13 (12)
Correlation with each measure			
Preference-dependent Measures (PDM)			
Weighted additive utility	.83 (.82)	−.20 (−.16)	.42 (.45)
Relative utility	.84 (.84)	−.22 (−.17)	.43 (.46)
Euclidean distance	.68 (.67)	−.22 (−.15)	.44 (.52)
Preference-independent measures (PIM)			
Attribute sum	.80 (.88)	−.27 (−.29)	.34 (.19)
Dominant alternatives	.76 (.76)	−.35 (−.31)	−.52 (−.53)
Attribute difference	.71 (.80)	−.39 (−.40)	−.48 (−.40)
Non dominated	.74 (.73)	−.35 (−.31)	−.54 (−.55)
Subjective measures (SM)			
Fit	.53 (.50)	.58 (.59)	.03 (.05)
Confidence	.54 (.49)	.54 (.58)	−.01 (−.01)
Satisfaction	.59 (.53)	.66 (.68)	−.05 (−.07)
Liking	.47 (.43)	.71 (.71)	−.08 (−.09)
Interest	.42 (.38)	.45 (.46)	−.24 (−.27)

Note. Parenthetical values refer to the principal component analysis of the partial correlations from Table 2.
The correlation loadings of these components are the partial correlations with each measure.

increasing numbers of alternatives. Although decision quality has long been a focus for marketing researchers, marketing practitioners have had few ways to use decision quality as a metric for assessing the impact of marketing activities. Electronic recommendation agents may serve to bridge this gap by collecting information that can be used to assess consumer decisions while providing consumers with an incentive to provide this information. Like real salespeople, electronic agents can serve dual roles as information providers and information collectors (Liu and Comer 2007).

This research contributes to theory and practice in several ways. First, it reviews and categorizes the various decision quality measures available to online providers into PDM, PIM, and SM measures and identifies the role of recommendation agents in collecting information used in these measures. Results indicate that overall decision quality can indeed be described in terms of PDM, PIM and SM measures. Second, it empirically assesses the relationship among the measures and shows the extent to which PIM measures capture the information in PDM measures and the extent to which SM measures provide additional information beyond that captured by the objective measures. Third, for each of the three types of measures, it identifies which individual measures are the best exemplars and which are the best predictors of weighted-additive utility—a popular measure of decision quality and the basis for economic utility models. Finally, it empirically identifies the factors and underlying measures that best represent decision quality as a multidimensional construct.

The empirical analysis leads to five specific conclusions and recommendations:

1) Whenever the service provider has information on consumer preferences, including preference information–such as that collected by recommendation agents–it provides significant improvements in the assessment of decision quality. The *relative utility* measure provides the most information in this context. Furthermore, this single measure captures most of the variance in objective decision quality. Hence, it is not necessary to calculate multiple measures of preference dependent measures. This of course implies that either the firm is already collecting consumer preference information or that it should whenever it has the opportunity. For service providers who already employ recommendation agents, our results suggest an opportunity to use information that is already being collected to assess the extent to which electronic agents are helping consumers make better choices. For service providers that are not collecting this information, the costs of doing so, from both a budget and customer effort perspective, need to be evaluated and ROI calculated.
2) Whenever the service provider has no information on consumer preferences such as attribute weights, our research suggests that calculating decision quality independent of this information can provide important insights. Even if the firm does not collect customer preference information, PIM measures can be utilized to calculate decision quality. In such circumstances, the PIM measure *attribute sum* is the best approximation to preference-based measures. Using this single metric can capture the majority of the variance and eliminate the need to calculate multiple measures.
3) If the popular WADD is the preferred decision quality measure *but* customer preference information is *not* available, a combination of the PIM measure *attribute sum* and the SM measure *fit* can provide a great deal of information (e.g., together they explain 71% of the variance in WADD utility). If only SM measures are considered, then *fit* and *satisfaction* provide 13% of the information that

would be captured if preference information exists and WADD is calculated. This demonstrates that, even in the absence of preference information, a combination of PIM and SM measures can capture a great deal of the information that WADD would have provided. Although SM measures alone provide some insight, they capture a small proportion of the information that WADD provides.

4) Moreover, SM measures are generally poor predictors of objective decision quality and therefore should not be used interchangeably or as “proxies” for objective decision quality. Although one can argue that objective measures should not matter as long as the customer is happy with her/his choice, it is quite possible to have a change in happiness levels over time. Particularly for experience-type products, even though a customer may indicate high levels of happiness at the time of choice, he/she may discover through product usage that the product was not what they had expected. This is an important issue from a managerial perspective as customers that decide post-use that they are not happy with a product are likely to return it, in turn generating cost to the firm.
5) Nonetheless, SM measures provide important insights beyond the objective measures of decision quality. Although subjective measures provide limited insights into normative indicators of decision quality they do offer important complementary information about individual decisions. This suggests the need to measure subjective decision quality measures in addition to objective decision quality metrics.

As an initial effort to empirically compare a wide variety of decision quality measures, this paper provides a way for managers and researchers to determine which measures are likely to be most informative of consumer decision quality and to make informed judgments about the benefits of collecting subjective measures, identifying attribute values, and deploying mechanisms, such as recommendation agents, to collect consumer preferences. Understanding the degree to which customers make the right choices for themselves is managerially relevant as it provides the basis for customer satisfaction, customer retention and other loyalty outcomes such as word of mouth. Future research should examine the extent to which alternative measures of decision quality predict customer loyalty, word of mouth, and customer lifetime value. Furthermore, gaging decision quality can help managers manage the process by which they create happy customers. By understanding how consumers feel about their choices, and how close consumers come to making choices that match their preferences, managers can 1) gage how consumers are doing, 2) create benchmarks for comparison purposes, and 3) identify areas for improvement or opportunities to cultivate stronger customer relationships.

Finally, although the focus and context of this paper is on online decision quality, the measurement of decision quality is important to a variety of disciplines, including organizational behavior, decision support systems, and marketing. To some extent, these results should generalize and provide fruitful guidance in many areas to researchers who use decision quality as a variable of interest in their studies.

Acknowledgment

This work was supported, in part, by the Dean’s Fund for Faculty Research of the Owen Graduate School of Management, Vanderbilt University.

Appendix A. Alternatives in the Database

Alternative	Low price	Light weight	Long talk time	Long standby time	Number of dominating alternatives[a]
AAB	5	10	2	5	3
ABB	**3**	**6**	**10**	**3**	**0**
ABC	**10**	**6**	**5**	**6**	**0**
BBC	**3**	**10**	**8**	**9**	**0**
BCC	3	10	8	6	1
BCD	**10**	**10**	**4**	**5**	**0**
CCD	3	7	5	4	5
CDD	1	10	2	5	6
DCE	1	7	6	10	2
DDE	1	6	4	2	20
EFG	9	7	4	2	2
FGH	9	3	3	3	5
GHI	**5**	**3**	**8**	**3**	**0**
HIJ	1	4	1	1	27
IJK	**9**	**7**	**6**	**10**	**0**
JKL	9	10	4	5	1
KLM	1	9	4	7	2
LMN	**1**	**9**	**6**	**10**	**0**
MNO	7	7	4	3	7
NOP	7	10	4	5	2
OPQ	**7**	**7**	**7**	**6**	**0**
PQR	3	1	8	4	4
QRS	1	4	5	3	10
RST	1	5	1	1	26
STU	**7**	**8**	**4**	**7**	**0**
TUV	1	10	4	5	5
UVW	**5**	**8**	**5**	**2**	**0**
VWX	5	5	3	8	2
WXY	1	3	4	2	22
XYZ	1	8	5	1	4
YZA	1	8	3	3	9
ZAB	5	7	6	10	1

[a]The number of alternatives that dominates this alternative.
Note. Ten is the best rating on each attribute. Boldface indicates non-dominated alternatives.

References

Aksoy, Lerzan, Paul Bloom, Nicholas H. Lurie, and Bruce Cooil (2006), “Should Recommendation Agents Think Like People?,” *Journal of Service Research*, 8, 4, 297–315.

Ansari, Asim, Skander Essegaier, and Rajeev Kohli (2000), “Internet Recommendation Systems,” *Journal of Marketing Research*, 37, August, 363–75.

Ariely, Dan (2000), “Controlling the Information Flow: Effects on Consumers’ Decision Making and Preferences,” *Journal of Consumer Research*, 27, September, 233–48.

Barron, F. Hutton and Bruce E. Barrett (1996), "Decision Quality Using Ranked Attribute Weights," *Management Science*, 42, 11, 1515–23.

Bearden, William O., Richard G. Netemeyer, and Jesse E. Teel (1990), "Further Validation of the Consumer Susceptibility to Interpersonal Influence Scale," in *Advances in Consumer Research, Vol. 17,* Marvin E. Goldberg, Gerald Gorn, and Richard Pollay, editors. Provo, UT: Association for Consumer Research, 770–6.

———, David M. Hardesty, and Randall L. Rose (2001), "Consumer Self-Confidence: Refinements in Conceptualization and Measurement," *Journal of Consumer Research*, 28, June, 121–34.

Bozdogan, Hamparsum (1987), "Model Selection and Akaike's Information Criterion (AIC): The General Theory and Its Analytical Extensions," *Psychometrika*, 52, 3, 345–70.

Branco, Joao Antonio, Christophe Croux, Peter Filzmoser, and Maria de Rosario Oliveira (2005), "Robust Canonical Correlations: A Comparative Study," *Computational Statistics*, 20, June, 203–29.

De Bruyn, Arnaud, John C. Liechty, Eelko K.R.E. Huizingh, and Gary L. Lilien (2008), "Offering Online Recommendations with Minimum Customer Input through Conjoint-Based Decision Aids," *Marketing Science*, 27, 3, 443–60.

Diehl, Kristin, Laura J. Kornish, and John G. Lynch Jr. (2003), "Smart Agents: When Lower Search Costs for Quality Information Increase Price Sensitivity," *Journal of Consumer Research*, 30, June, 56–71.

Fassnatch, Martin and Ibrahim Köse (2007), "Consequences of Web-Based Service Quality: Uncovering a Multi-Faceted Chain of Effects," *Journal of Interactive Marketing*, 21, Summer, 35–54.

Flynn, Leisa R. and Ronald E. Goldsmith (1999), "A Short, Reliable Measure of Subjective Knowledge," *Journal of Business Research*, 46, September, 57–66.

Hahn, Minhi, Robert Lawson, and Young Gyu Lee (1992), "The Effect of Time Pressure and Information Load on Decision Quality," *Psychology and Marketing*, 9, September–October, 365–78.

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

——— and Kyle B. Murray (2003), "Preference Construction and Persistence in Digital Marketplaces: The Role of Electronic Recommendation Agents," *Journal of Consumer Psychology*, 13, 1, 75–91.

———, Benedict G.C. Dellaert, and Murat Usta (2010), "Ironic Effects of Personalized Product Recommendations on Subjective Decision Outcomes," in *Proceedings of the Society for Consumer Psychology Winter Conference*, Margaret G. Meloy and Adam Duhachek, editors. Elsevier, 68.

Iaobucci, Dawn, Phipps Arabie, and Anand Bodapati (2000), "Recommendation Agents on the Internet," *Journal of Interactive Marketing*, 14, Summer, 2–11.

Jacoby, Jacob (1977), "Information Load and Decision Quality: Some Contested Issues," *Journal of Marketing Research*, 14, November, 569–73.

Jenrich, Robert (1970), "An Asymptotic χ^2 test for the Equality of Two Correlation Matrices," *Journal of the American Statistical Association*, 65, June, 904–12.

Keller, Kevin Lane and Richard Staelin (1987), "Effects of Quality and Quantity of Information on Decision Effectiveness," *Journal of Consumer Research*, 14, September, 200–13.

Keren, Gideon and Wändi Bruine de Bruin (2003), "On the Assessment of Decision Quality: Considerations Regarding Utility, Conflict and Accountability," in *Thinking: Psychological Perspectives on Reasoning, Judgment and Decision Making,* David Hardman and Laura Macchi, editors. New York: Wiley, 347–63.

Kramer, Thomas (2007), "The Effect of Measurement Task Transparency on Preference Construction and Evaluations of Personalized Recommendations," *Journal of Marketing Research*, 44, 2, 224–33.

Lilien, Gary L., Arvind Rangaswamy, Gerrit H. Van Bruggen, and Katrin Starke (2004), "DSS Effectiveness in Marketing Resource Allocation Decisions: Reality vs. Perception," *Information Systems Research*, 15, September, 216–35.

Liu, Sandra S. and Lucette B. Comer (2007), "Salespeople as Information Gatherers: Associated Success Factors," *Industrial Marketing Management*, 36, 565–74.

Lurie, Nicholas H. (2004), "Decision Making in Information-Rich Environments: The Role of Information Structure," *Journal of Consumer Research*, 30, March, 473–86.

Mantalos, Panagiotis and Ghazi Shukur (2007), "The Robustness of the Reset Test to Non-Normal Error Terms," *Computational Economics*, 30, November, 393–408.

McFadden, Daniel (1986), "The Choice Theory Approach to Marketing Research," *Marketing Science*, 5, Fall, 275–97.

Meyer, Robert J. and Eric C. Johnson (1989), "Information Overload and the Non-Robustness of Linear Models: A Comment on Keller and Staelin," *Journal of Consumer Research*, 15, March, 498–503.

Olson, Erik L. and Rorbert E. Widing II (2002), "Are Interactive Decision Aids Better than Passive Decision Aids? A Comparison with Implications for Information Providers on the Internet," *Journal of Interactive Marketing*, 16, 2, 22–33.

Parasuraman, Parsu, Valarie A. Zeithaml, and Arvind Malhotra (2005), "E-S-Qual: A Multiple-Item Scale for Assessing Electronic Service Quality," *Journal of Service Research*, 7, February, 213–33.

Payne, John W., James R. Bettman, and Eric Johnson (1993), *The Adaptive Decision Maker*. Cambridge, UK: Cambridge.

Punj, Girish N. and Robert Moore (2007), "Smart Versus Knowledgeable Online Recommendation Agents," *Journal of Interactive Marketing*, 21, 4, 46–60.

Rayport, Jeffrey F., Bernard J. Jaworski, and Ellie J. Kyung (2005), "Best Face Forward: Improving Companies' Service Interfaces with Customers," *Journal of Interactive Marketing*, 19, Fall, 67–80.

Rust, Roland T. (2001), "The Rise of E-Service," *Journal of Service Research*, 3, May, 283–4.

———, Tim Ambler, Gregory S. Carpenter, V. Kumar, and Rajendra K. Srivastava (2004), "Measuring Marketing Productivity: Current Knowledge and Future Directions," *Journal of Marketing*, 68, 4, 76–89.

Schwartz, Barry (2004), *The Paradox of Choice: Why More is Less*. New York: Ecco.

Schwarz, Gideon (1978), "Estimating the Dimension of a Model," *Annals of Statistics*, 6, March, 461–4.

Shiv, Baba and Joel Huber (2000), "The Impact of Anticipating Satisfaction on Consumer Choice," *Journal of Consumer Research*, 27, September, 202–16.

Sproles, George B. (1983), "Conceptualization and Measurement of Optimal Consumer Decision-Making," *Journal of Consumer Affairs*, 17, Winter, 421–83.

Swaminathan, Vanitha (2003), "The Impact of Recommendation Agents on Consumer Evaluation and Choice: The Moderating Role of Category Risk, Product Complexity, and Consumer Knowledge," *Journal of Consumer Psychology*, 13, 1&2, 93–101.

Wolfinbarger, Mary and Mary Gilly (2003), "eTailQ: Dimensionalizing, Measuring and Predicting Etail Quality," *Journal of Retailing*, 79, Fall, 183–98.

Zeithaml, Valarie A., Parsu Parasuraman, and Arvind Malhotra (2002), "Service Quality Delivery through Web Sites: A Critical Review of Extant Knowledge," *Journal of the Academy of Marketing Science*, 30, 4, 362–75.

———, Ruth N. Bolton, John Deighton, Timothy L. Keiningham, Katherine N. Lemon, and J. Andrew Petersen (2006), "Forward-Looking Focus: Can Firms Have Adaptive Foresight?," *Journal of Service Research*, 9, November, 168–83.