

Chorus, Caspar G.; Timmerman, Harry J. P.

Article

Revealing consumer preferences by observing information search

Journal of Choice Modelling

Provided in Cooperation with:

Journal of Choice Modelling

Suggested Citation: Chorus, Caspar G.; Timmerman, Harry J. P. (2008) : Revealing consumer preferences by observing information search, Journal of Choice Modelling, ISSN 1755-5345, University of Leeds, Institute for Transport Studies, Leeds, Vol. 1, Iss. 1, pp. 3-25

This Version is available at:

<https://hdl.handle.net/10419/66837>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



<http://creativecommons.org/licenses/by-nc/2.0/uk/>

Revealing consumer preferences by observing information search

Caspar G. Chorus^{1,*} Harry J.P. Timmermans^{2,†}

¹ Section of Transport Policy and Logistics' Organisation,
Delft University of Technology, PO Box 5015, 2600 GA Delft, The Netherlands

² Urban Planning Group, Eindhoven University of Technology, PO Box 513,
5600 MB Eindhoven, The Netherlands

Received 28 August 2007, revised version received 24 October 2007, accepted 9 November 2007

Abstract

Motivated by the notion that consumers' use of the internet creates a wealth of data on information search, we put forward the idea that observed information searches may be used for deriving consumer preferences. First, we derive a theoretical model of consumer behaviour under uncertainty and information availability. In theory, this model provides the opportunity to derive consumer preferences from information search alone. The model is then illustrated, based on an artificial dataset. Estimation results show that information search concerning an uncertain attribute of a good can indeed be used to identify consumer preferences concerning the good. Although the proposed model relies on a number of assumptions (for example the premises behind expected utility maximisation) we show how these can be relaxed without compromising the potential of the approach in general.

Keywords: information search, consumer preferences, choice-modelling

1 Introduction

For decades, the study of consumer preferences for multi-attribute goods has been one of the central topics of marketing research, leading to a variety of methods to derive preferences from revealed or stated behaviour. In the 1970s, so-called compositional models were developed to study consumer preferences, involving respondents to explicitly express their valuations and importance weights for a set of attributes using rating scales (e.g. Wilkie and Pessemier, 1973; Bettman et al., 1975). Later, so-called decompositional or conjoint models (e.g. Louviere

*Corresponding author, T: +31 152 781 144, F: +31 152 782 719, c.g.chorus@tudelft.nl

†T: +31 402 473 315, F: +31 402 438 488, h.j.p.timmermans@bwk.tue.nl



and Woodworth, 1983; Louviere and Hensher, 1983) became popular. Originally, conjoint analysis was based on ranking or rating data, but later also on choices among multi-attribute goods. The use of choice data, instead of rating or ranking data, for the estimation of conjoint models is rooted in microeconomic theories of consumer behaviour (e.g. Lancaster, 1966) and econometric discrete choice theory (e.g. McFadden, 1974; Ben-Akiva and Lerman, 1985).

Notwithstanding the intuitively appealing nature of observed *choices between goods* for revealing consumer preferences, we argue in this paper that it is not only in this act of choosing that consumers reveal their preferences. More specifically, we put forward the idea that consumers, faced with uncertainty, also reveal their preferences by *deciding whether or not to search for information*. Moreover, we theoretically argue how preferences for goods (the relative importance and evaluation of their attributes) can be derived by observing consumers' information search patterns alone.

Our practical motivation for developing this idea stems from the increasing use of the internet as a tool for searching information (e.g. Peterson and Merino, 2003; Ratchford et al., 2003; Moon, 2004; Capra and Pérez-Quinones, 2005). Information search through the internet can be relatively easily traced and this creates a potential wealth of so-called clickstream data. As Bucklin et al. (2002) suggest, such clickstream data presents empirical researchers with a significant opportunity to advance the understanding and prediction of consumer choice behaviour. Or, as Battelle (2005) puts it in his book on search engine Google, people's use of the internet's search engines creates huge and valuable 'databases of intentions'. From the perspective of data-efficiency, it is worthwhile to investigate ways in which such information search data may be put to use for the purpose of deriving consumer preferences.

In this paper, we tailor our model towards the following type of information search: a consumer faces a choice from among two goods, one of them being uncertain in terms of one of its attributes. Information about that attribute is available. Only the consumer's decision whether or not to search for the information is observed. Obviously, this scenario represents only one of many possible information search contexts: for example, consumers may wish to search for information in order to form a consideration set (e.g. Schocker et al., 1991; Roberts and Lattin, 1991, 1997), or with the aim of exploring the potential of a variety of goods (e.g. Marchionini, 2006). Such other contexts should be considered in future work.

Furthermore, it should be noted that this paper complements recent attempts in marketing and transportation to predict the full sequence of possibly multiple information searches followed by a choice among alternatives (e.g. Erdem et al., 2005; Moe, 2006; Chorus et al., 2007b). The models presented in these papers derive estimates of consumer preferences from a combination of observed information search patterns and choices among alternatives. In contrast with these contributions we wish to show here that in order to obtain such estimates, it is not necessary to observe choices among alternatives in the first place: prefer-

ences can be derived from solely observing the decision whether or not to acquire information concerning an uncertain attribute.

As a result, the relative simplicity of the choice context we consider (only one uncertain attribute, no sequential information search process, no choice among alternatives) is motivated theoretically: we attempt to show that consumer preferences may be derived from observations that at first sight contain only little information regarding these preferences. More realistic (and hence more complex) choice situations, involving a combination of different types of information search, can be modelled using the same general principles that we propose and illustrate in this paper using a simple choice context.

The outline of the paper is as follows: the next section provides a formal utilitarian presentation of how consumers reveal their preferences for goods through information search. This is done by modelling the utility of information search in terms of the expected utility of the current choice situation, and the expected utility of the anticipated choice situation after having received the information. We first propose our model at the individual level and discuss its properties. We proceed by providing an econometric specification that is applicable for estimation efforts based on observed information search patterns. Subsequently, we present a fictitious example of consumer choice between computers under conditions of uncertainty and information availability. Artificial datasets of information search behaviour are created using Monte Carlo simulation on a predefined set of preferences for the computers, and we show how the generated information search patterns can be used to derive these true consumer preferences through model estimation. Finally, the general applicability of the proposed approach is discussed, and conclusions are drawn.

2 The model

A prerequisite for being able to use observed information search behaviour to derive consumer preferences is that information search behaviour is modelled in terms of the consumer's preferences for these goods. The most frequently used approach in this regard is to formulate the decision to search for information in terms of a cost-benefit trade-off (e.g. Wilde, 1980; Ratchford, 1982; Punj and Staelin, 1983; Huneke et al., 2004; Chorus et al., 2006); the benefits of information search are generally modelled as the difference between the expected utility of the current choice situation on the one hand, and the expected utility of the anticipated choice situation after having received the information on the other hand (e.g. Raiffa and Schlaifer, 1961; Weibull, 1978; Akerberg, 2003). The costs of information search may involve monetary costs. It is however generally acknowledged that a wide variety of non-monetary costs may also play a role, for example costs of thinking (Shugan, 1980) and opportunity costs of time (e.g. Ratchford, 1982). In addition to theoretical models of information search at the individual level, we need an econometric specification of these models to be

able to actually estimate preferences from observed information search patterns. Recently, random-utility-based specifications of information search and choice processes have been proposed by Erdem et al. (2005) and Chorus et al. (2007b). Our approach to derive consumer preferences from observed information search patterns is rooted in these streams of literature.

More specifically, we assume i) that a consumer's decision to search for information is based on a trade-off between a cost component and a benefit component; ii) that the benefit of information search can be written in terms of the difference in expected utility of the current and anticipated choice situation; iii) that part of the expected utility of goods and the utility of information cannot be observed by the analyst, leading to additive random error components in the utility functions. In the remainder of this section, we will specify our model at the individual level, after which an econometric formulation is provided.

2.1 The individual level model

Consider a binary choice situation between two multi-attribute goods i and j . Assume that, in the spirit of Lancaster (1966), a consumer perceives both goods as a bundle of attribute values, denoted x_{ik} and x_{jk} , where $k \in K$. A constant β_0 reflects a preference for brand i relative to brand j . Let the consumer's preferences for alternatives i and j be specified through linear additive utility functions $U_i = \beta_0 + \sum_{k \in K} \beta_k x_{ik}$ and $U_j = \sum_{k \in K} \beta_k x_{jk}$. Here, β_k represents a consumer's evaluation of a particular attribute value. Assume now that there is uncertainty attached to one of the attributes of good i , in the eyes of the consumer. Take for example the situation where it is uncertain whether good i possesses some feature l . We conceptualise this uncertainty by assuming that the consumer thinks that x_{il} may take the value 1 (the good does possess the particular feature) with probability $P(x_{il})$. The expected utility the consumer derives from choosing good i is then written as follows:

$$EU_i = \beta_0 + P(x_{il}) \cdot \beta_l + \sum_{k \in K-l} \beta_k x_{ik} \quad (1)$$

The expected utility she may derive from the current choice situation, EU , is then denoted as the maximum of the expected utility of good i and the utility of good j .

Let an information search possibility exist concerning the value of x_{il} (i.e. the consumer may search for information whether or not good i possesses feature l). The expected utility of the anticipated choice situation after having received the information, EU^+ , can then be formalised as follows: the individual knows that, given fully reliable information, she will receive the message "good i possesses feature l " with probability $P(x_{il})$ and the message "good i does not possess feature l " with probability $1 - P(x_{il})$. For example, should our consumer believe that the probability that good i possesses feature l is 90%, we assume that when acquiring fully reliable information she expects to receive a message, saying that

the good indeed possesses feature l , with 90% probability as well. She also knows that, after having received one of these messages, she will maximise her utility by choosing between the good i (with or without feature l) and good j . The following equation gives the notational representation of this argument:

$$\begin{aligned}
 EU^+ = & P(x_{il}) \cdot \max \left\{ \beta_0 + \beta_l + \sum_{k \in K-l} \beta_k x_{ik}, U_j \right\} \\
 & + (1 - P(x_{il})) \cdot \max \left\{ \beta_0 + \sum_{k \in K-l} \beta_k x_{ik}, U_j \right\}
 \end{aligned} \tag{2}$$

As proposed earlier in this Section, we now conceptualise the expected utility of information as the difference between the expected utility of the current choice situation and that of the anticipated choice situation, minus the cost c of information search: $EU_I = EU^+ - EU - c$. Information is searched for when $EU_I > 0$.

2.2 Model properties

Before discussing how the model presented above can be formulated econometrically, let us elaborate on how it captures some basic intuitions concerning the behavioural determinants of consumer information search.

First, let us consider how consumer information search is influenced by a good's performance on other attributes than the uncertain one. We would expect that the expected utility of information about the uncertain attribute x_{il} is relatively high when good i and j are perceived as comparably attractive in terms of the total of their other attributes. In such a situation, whether or not good i possesses feature l will likely determine the consumer's choice. When one of the two goods is superior in terms of the performance of attributes other than l , we expect a consumer to be less interested in knowing the value of x_{il} , since knowing its true value is unlikely to change her preference for the superior good. The following fictitious example illustrates how our model captures this intuition.

We assume the following settings, based on the choice-situation described above: a consumer chooses between two goods i and j . Choice depends on the good's price, a brand preference, and attribute l . She perceives the probability that good i possesses feature l as $P(x_{il}) = 0.5$ and is certain that good j does not possess the feature. The importance of the uncertain attribute, denoted β_l , equals 5. Information I concerning the value of x_{il} can be obtained at no costs (i.e. $c = 0$). We now vary the performance of good i relative to good j in terms of its brand preference (*brandpref*) and the difference in prices (*pricediff*, positive values indicating that i is more expensive than j ; we assume $\beta_{\text{price}} = -1/\text{unit}$). We observe (vertical axis) how this affects the expected utility of searching for information concerning x_{il} . Figure 1 shows how the expected utility of information, given our conceptualisation, is low or non-existing in situations where either i) good i is substantially more expensive than j and the consumer

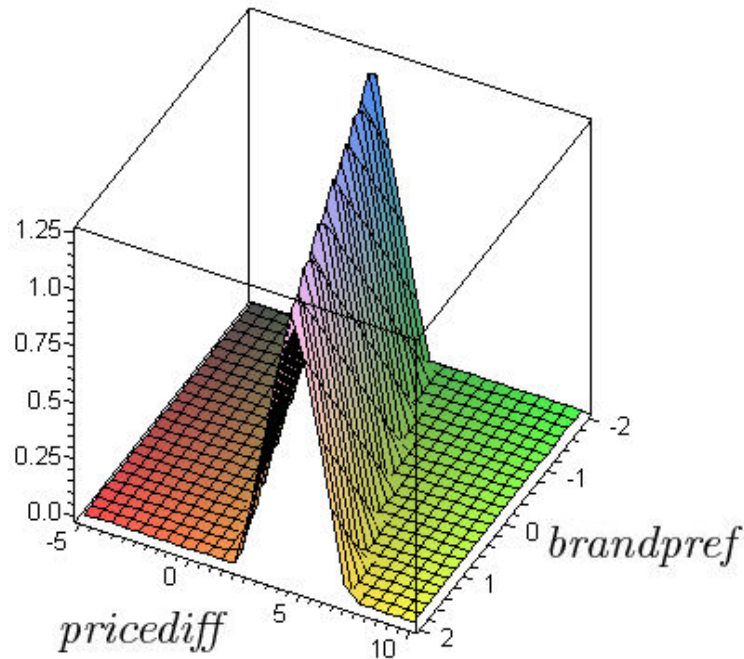


Figure 1: Information utility as a function of the value of certain attributes

holds a strong brand-preference for j , or ii) good i is substantially cheaper than j and the consumer holds a strong brand-preference for i .

In situations with no strong brand preferences nor substantial price-differences, or where they cancel out, expected utility of information is relatively high. Note that in fact, since we assumed that good j does not possess feature l , the expected utility of information is highest when j is slightly more attractive than i in terms of price and brand preference. Completely in line with intuition, our model predicts that consumers derive most utility from information search concerning an uncertain attribute when the true value of the attribute determines choice between the goods.

Second, let us consider how consumer information search concerning an uncertain attribute is influenced by the degree of uncertainty and the importance a consumer attaches to the attribute. We would expect that the more uncertain a consumer is concerning the true value of good i 's attribute l , the more she will be inclined to search for information, *ceteris paribus*. We also expect that the more important the uncertain attribute is in the eyes of the consumer, the more she will be inclined to search for information concerning the attribute, *ceteris paribus*. In notation, we would expect that higher values for β_l induce higher expected utility of information and that when $P(x_{il})$ approaches either 0 or 1, the expected utility of information decreases. Figure 2 visualises how our model captures these intuitions, by plotting the expected utility of information against

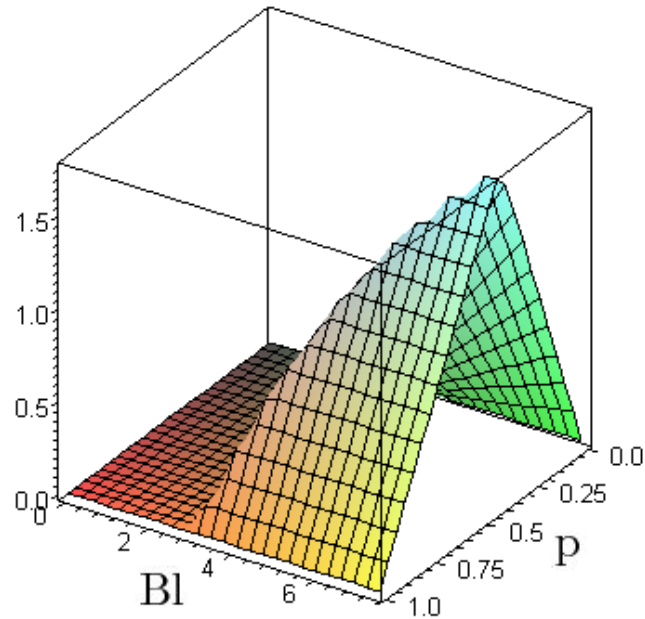


Figure 2: Information utility as a function of attribute importance and uncertainty

attribute-importance (β_l) and the level of uncertainty (p). We adopt the following settings: there is a small brand-preference for good j over i ($\beta_0 = -3$), prices of the two goods are equal. We vary $P(x_{il})$ within $[0,1]$, and β_l between 0 and 7.5.

As anticipated, expected information utility is low or non-existing when the individual is relatively certain that good i does (not) possess feature l , and for low values of β_l . In these situations, there is little uncertainty, and the uncertainty that does exist does not influence consumer choice due to the relative irrelevance of attribute l . However, when $P(x_{il})$ approaches 0.5, and β_l increases, the expected utility of information increases as well: the attribute becomes more relevant, and its value highly uncertain, so the information becomes valuable. Again, the proposed model of expected information utility appears to behave in line with our intuitions regarding the behavioural determinants of information search.

Summarising, the proposed model of information search appears to be sensitive to consumer preferences in ways that are in line with intuition. This provides first confidence that the model may be applied to derive consumer preferences from observed information search patterns. However, in order to become applicable for data-analysis, we need an econometric specification of our individual-level model. Such a specification is presented below.

2.3 Econometric model specification

Assume that an analyst observes information search from N consumers, one choice per consumer n , and wishes to derive (average) preferences for goods from these observations. Adopting a random-utility framework, we assume that the analyst is only able to observe part of the (expected) utility that a consumer derives from the available goods, or from information search. That is, these utilities are –from the analyst’s point of view– composed out of an observed and an unobserved (random) part. We assume the following distributions for the random utility components: δ_i^n represents the part of the utility of good i that is unobservable by the analyst and reflects, for example, variation across consumers concerning intrinsic preferences for good i relative to good j (the average intrinsic or brand preference is given by β_0). We assume that δ_i^n is normally distributed across individuals: $\delta_i^n \sim N(0, \sigma_i)$. Furthermore, the part of the utility of (not) searching for information that is unobservable by the analyst (ε_l^n) is assumed to be distributed *iid* Extreme Value Type I with standard deviation $\pi/\sqrt{6}$. For ease of presentation, we assume also that every certain attribute is perceived and evaluated equally by each consumer. However, different consumers may attach different probabilities $P^n(x_{il})$ concerning the availability of feature l . This leads to the following formulation of the observed expected utility, given δ_i^n , derived by individual n from the anticipated choice situation, after having received information:

$$EU^{n+}(\delta_i^n) = P^n(x_{il}) \cdot \max \left\{ \beta_0 + \beta_l + \sum_{k \in K-l} \beta_k x_{ik} + \delta_i^n, U_j \right\} \\ + (1 - P^n(x_{il})) \cdot \max \left\{ \beta_0 + \sum_{k \in K-l} \beta_k x_{ik} + \delta_i^n, U_j \right\} \quad (3)$$

The observed expected utility of information search, given δ_i^n , becomes $EU_I^n(\delta_i^n) = EU^{n+}(\delta_i^n) - EU^n(\delta_i^n) - c^n$, where

$$EU^n(\delta_i^n) = \max \left\{ \beta_0 + \beta_l P^n(x_{il}) + \sum_{k \in K-l} \beta_k x_{ik} + \delta_i^n, U_j \right\} \quad (4)$$

Given these formulations, computing the probability that an observed consumer n will search for information now involves the evaluation of two integrals: one to integrate out δ_i^n relating to the consumer’s intrinsic preference for good i , another one to integrate out the *iid* errors relating to her inclination (not) to search for information. The distributional assumptions of the random errors result in a mixed binary logit model (e.g. McFadden and Train, 2000); the closed form solution of the integral associated with the *iid* errors (resulting in binary logit) is mixed over the probability density function of δ_i^n . Acknowledging that the non-random part of the utility of not acquiring information equals zero by

definition, the choice probability for information search by consumer n may thus be denoted as:

$$P(I^n) = \int_{\delta_i^n} \left(\frac{\exp [EU_I^n (\delta_i^n)]}{\exp [EU_I^n (\delta_i^n)] + 1} \right) \cdot f(\delta_i^n) d\delta_i^n \quad (5)$$

Let us denote observed information search as y_I^n , assuming it takes the value 1 when the individual chooses to acquire information, and 0 otherwise. We can now estimate the parameters of the choice model, i.e. consumer preferences, by maximising the (log of) the sample likelihood L for the observed information search patterns:

$$L = \prod_{n=1}^N \left([P(I^n)]^{y_I^n} \cdot [1 - P(I^n)]^{1-y_I^n} \right) \quad (6)$$

The above model *theoretically* provides us with a tool to derive consumer preferences from observed information search patterns. However, it is yet unclear whether this theoretical notion holds in practice. It seems that the assumed relation between consumers' preferences and the resulting information search patterns is rather subtle and indirect when compared to the relation, assumed in conventional choice models, between preferences and choices among goods. In other words, it is not clear whether the parameters (preferences) that underlie a data-generating process (information search behaviour) that is characterised by the above equations can be identified by maximum likelihood estimation based on the observed data. Another question that we cannot answer just by inspecting the model itself is how many parameters are identifiable based on observed information search patterns concerning one of the attributes. In the notation of our model: given that good i 's attribute l is uncertain and information search concerning the attribute is observed, does this observation enable us to estimate not only β_l , but also β_k for some other attribute k ? In the next Section, we will address these questions concerning identification of consumer preferences by estimating our model of information search on artificial datasets.

3 Illustration on artificial data

This section consists of two parts: first, we will generate artificial data concerning information search patterns. This is done by feeding a set of true preferences (and Monte-Carlo-generated random error components) in the model presented above, in order to simulate information search. Second, we will use the artificially generated data on information search for maximum likelihood estimation and investigate whether the obtained parameter estimates correspond to the true preferences used to generate the data.

3.1 Data generation

We assume the following hypothetical situation: consumers are faced with a choice between two computers A and B, and perceive them as a bundle of the attribute *price*, another attribute denoted *feature* (say, a software package), and a set of remaining attributes regarding which there is no difference between computer A and B in the eyes of consumers. Computer A is more expensive than computer B and consumers know the price difference. They also know that computer B does not possess the software package and they attach a probability $P(\text{feature}=1)$ to the event that computer A possesses the software package; this probability differs between consumers, and we assume that consumers perceive it as independent from the attribute *price* (see further below for a discussion concerning the implications of this assumption). Consumer preferences are as follows: -1 util per dollar regarding the attribute *price* and 50 utils for the attribute *feature*. That is, consumers are willing to pay no more than 50 dollar for the software package. In case consumers believe that computer A does not possess the software package ($P(\text{feature}=1) = 0$), then computer B is preferred on average, as it is cheaper and perceived to be equivalent in terms of every attribute. In case consumers believe that A does possess the software package ($P(\text{feature}=1) = 1$), consumer preference depends on the price difference between the two computers. Under conditions of uncertainty ($0 < P(\text{feature}=1) < 1$) computer A may or may not be preferred to computer B, depending on the price difference between the two and the magnitude of $P(\text{feature}=1)$. By replacing the attribute *price* by *pricediff* (reflecting the difference in price between A and B), we normalise the utility of computer B to zero. Adding δ_A^n to reflect intrinsic preferences for brand A, we arrive at the following expected utility function for computer A:

$$EU_A^n(\delta_A^n) = \beta_{\text{pricediff}} \cdot \text{pricediff} + \beta_{\text{feature}} \cdot P(\text{feature} = 1) + \delta_A^n \quad (7)$$

Consumers now face a choice whether or not to acquire fully reliable information concerning the value of *feature*, denoted *infosearch*. The information comes at a certain information cost, denoted *cost* and perceived by consumers to equal -4 utils. Note that *cost* may consist of monetary as well as non-monetary components. Based on the described choice situation, we generate 10 datasets of 2000 cases (consumers) each. Each case presents a choice whether or not to acquire the available information. The data generation process consists of three steps:

1. Each case is systematically assigned *pricediff* from the set {10, 20, 30, 40, 50}, reflecting that A is always more expensive than B. Also, each case is randomly assigned a probability $P(\text{feature}=1)$ drawn from the 0-1 interval, representing the individual's belief strength that computer A contains the software package.
2. We then apply Monte Carlo simulation to draw random error components for each case. One error component, reflecting the consumer's preference for

computer brand A over B, is drawn from a standard normal distribution and added to the (expected) utility of computer A. Two other errors, reflecting the individual's inclination (not) to search for information, are drawn from an *iid* Extreme Value Type I distribution with standard deviation $\pi/\sqrt{6}$, and added to the expected utility of (not) acquiring information.

3. Finally, we compute for each case the expected utility of information search, and based on this utility simulate a consumer's choice whether or not to acquire the available information (see the Appendix for an example of this computation). For the given model settings, roughly half of the 20,000 hypothetical consumers choose to search for information.

3.2 Model estimation and comparison of estimates with true values

The generated datasets contain four columns each: three independent variables (*pricediff*, $P(\text{feature}=1)$, *cost*) and one dependent variable (*infosearch*). We now present an attempt to estimate parameters for *feature* as well as *pricediff* and *cost*, by relating these independents to *infosearch*, using the proposed model of information search. The likelihood function and estimation process were coded in GAUSS 7.0, using the MaxLik-module. Note that calculation of the likelihood function involves integration of a binary logit function over a standard normal probability density function, for which there is no closed form solution. We compute the integral through simulation, using 1,000 intelligent Halton draws (e.g. Train, 2003) for each case. We gratefully use Kenneth Train's GAUSS code for making these draws. This high number of intelligent draws results in a very high level of coverage of the probability density function and a high level of stability of parameter estimates.

It turns out that, notwithstanding the subtlety and indirect nature of the assumed relation between consumer preferences and information search behaviour, we are able to empirically identify from the generated data all three parameters that were used for data generation (see Table 1). That is, not only a parameter for *feature* (true value = 50) itself is estimated, indicating the relative importance of *feature*, but also parameters for the other two independents, *pricediff* (true value = -1) and *costs* (true value = -4)¹. The ρ^2 measure averages 0.58

¹ Note that when the cost of information search can safely be assumed to be more or less constant for different choice situations, as is the case in our example, its value can in principle be identified by means of systematically varying the (uncertain) attributes of choice-alternatives i and j . This systematic variation will lead to variation in information benefits ($EU^+ - EU$), which in turn allows one to infer from observed information search, in theory, whether these benefits outweigh information costs. When information costs systematically vary across situations due to variation in observed costs such as monetary costs, identification is also possible: a separate parameter is estimated for the observed component, and the remaining information cost is identified in the way discussed above. When information costs vary across observations due to variation in unobserved costs such as mental effort, identification is intrinsically difficult. In these cases, the latent costs may be modelled as a function of observable variables such as

Table 1: Parameters estimated from observed information search

Set	<i>pricediff</i>	SE	t^\dagger	<i>feature</i>	SE	t^\dagger	<i>costs</i>	SE	t^\dagger
<i>true</i>	-1	-	-	50	-	-	-4	-	-
1	-1.02	0.05	-0.37	49.41	2.16	-0.27	-3.80	0.18	1.12
2	-1.01	0.05	-0.22	49.04	2.16	-0.44	-3.87	0.18	0.71
3	-1.06	0.05	-1.24	51.14	2.23	0.51	-4.05	0.19	-0.25
4	-0.99	0.04	0.21	49.45	2.21	-0.25	-3.96	0.19	0.23
5	-0.99	0.04	0.32	47.84	2.09	-1.03	-3.82	0.18	1.04
6	-1.07	0.05	-1.43	51.75	2.31	0.76	-3.94	0.19	0.29
7	-1.03	0.05	-0.67	49.86	2.20	-0.06	-3.93	0.19	0.37
8	-1.05	0.05	-1.00	50.28	2.22	0.13	-3.87	0.18	0.70
9	-1.06	0.05	-1.32	50.69	2.24	0.31	-3.86	0.18	0.78
10	-1.01	0.04	-0.17	48.69	2.14	-0.62	-3.89	0.18	0.63

[†]Note that the t -values refer to the differences between the estimated parameters and the true values (not the differences with respect to zero): $t = (\text{estimated value} - \text{true value}) / \text{SE}$.

over the ten models. The estimated parameters seem to be very close to the true values and all of the estimates have small standard errors².

Perhaps more importantly than assessing whether the estimated preferences seemingly correspond to the true preferences is to assess whether they are statistically indistinguishable from each other. This assessment is performed by testing, for each of the 30 estimated parameters, the following hypothesis through a one-sample t -test: the estimated parameter is equal to the true value. See Table 1 as well for the results of this test. Acknowledging that a significance level of 5% (10%) implies that t -values lower than 1.96 (1.65) signal that the hypothesis cannot be rejected, it is clear that at any conventional level of significance the estimated parameters are statistically indistinguishable from their true values.

4 Model applicability

From the analyses presented above it may be concluded that, in principle, consumer information search patterns concerning an uncertain attribute of a multi-attribute good can be used to identify the true consumer preferences for that good (i.e. the preferences underlying the data generating process). More specifically, it appears that information search patterns concerning an uncertain attribute (*feature*) not only serve to reveal consumers' evaluation of that particular attribute,

the number of alternatives and attributes involved. This way, observable variables can act as proxies for information cost, in principle enabling identification. It is expected that identification of information costs from observed search behaviour may prove difficult in most real life settings. This is an issue that should be addressed in future research.

² Note here that one reason for these small standard errors lies in the fairly large number of cases that are constructed.

but also of another attribute of the good (*pricediff*), as well as their evaluation of the costs of information search (*costs*). We feel that at this point some reflection is in place, and we therefore in the remainder of this section explore the general applicability of our approach.

4.1 Extension towards partially reliable information

Until here, we have assumed that available information is (perceived as) fully reliable. In practice, this assumption will not always hold. Information may often be (perceived as) only partially unreliable, which is likely to result in lower inclination to search for information than our model predicts. Should the model that the analyst uses for deriving preferences from observed information search be – erroneously – based on the premise of fully reliable information, this would lead to biased estimates for consumer preferences. In the following, we show that there is a straightforward way to incorporate the notion of perceived information unreliability in our model of information search, based on the concept of Bayesian perception updating.

Assume the exact same choice situation that is described in the model-section above – we here focus on the individual-level model for simplicity of notation. Assume that the information I concerning the value of the uncertain attribute x_{il} is not (perceived as) fully reliable. That is, there is a perceived non-zero probability $P(I : x_{il} = 1 | x_{il} = 0)$ that the information I received by the consumer states that the good possesses feature l when in fact this is not the case. Similarly, there may be a perceived non-zero probability $P(I : x_{il} = 0 | x_{il} = 1)$. What are the effects of these non-zero probabilities on the expected utility of information within an expected utility maximisation-framework of consumer choice?

First, the probabilities that govern a consumer's perception of what message she thinks she will receive when acquiring information will change. In case the information is perceived as completely unreliable, the consumer will perceive received messages to be random draws, irrespective of her initial perceptions concerning the uncertain attribute $P(x_{il} = 1)$ and $P(x_{il} = 0)$. In case of fully reliable information, her initial perceptions concerning the uncertain attribute fully determine her perceived probability of receiving particular messages as discussed above: $P(I : x_{il} = 1) = P(x_{il} = 1)$, $P(I : x_{il} = 0) = P(x_{il} = 0)$; In general, we can write the perceived probability of receiving particular messages as follows:

$$\begin{aligned} P(I : x_{il} = 1) &= P(x_{il} = 1) \cdot P(I : x_{il} = 1 | x_{il} = 1) \\ &\quad + P(x_{il} = 0) \cdot P(I : x_{il} = 1 | x_{il} = 0) \\ P(I : x_{il} = 0) &= 1 - P(I : x_{il} = 1) \end{aligned} \tag{8}$$

Note that the following must hold:

$$P(I : x_{il} = 1 | x_{il} = 0) + P(I : x_{il} = 0 | x_{il} = 0) = 1 \tag{9}$$

and

$$P(I : x_{il} = 0 | x_{il} = 1) + P(I : x_{il} = 1 | x_{il} = 1) = 1. \tag{10}$$

It can be seen that the perceived probabilities of receiving partially reliable messages reduce to the probabilities associated with the case of fully reliable information when $P(I : x_{il} = 1 | x_{il} = 0) = P(I : x_{il} = 0 | x_{il} = 1) = 0$.

Besides that information unreliability affects consumer beliefs about what message may be received when searching for information, it is expected that the unreliability also affects the impact of a received message on consumer perceptions concerning the uncertain attribute. The more reliable the information is perceived to be, the more a consumer will take messages seriously, up to the point that she is willing to replace here initial perceptions completely by received fully reliable information. Information that is believed to be fully unreliable will be disregarded by consumers. Using Bayes' law of conditional probabilities, we arrive at the following general formulation for a consumer's perception updating process (e.g. Edwards et al., 1963)³:

$$\begin{aligned} P(x_{il} = 1 | I : x_{il} = 1) &= \frac{P(I : x_{il} = 1 | x_{il} = 1) \cdot P(x_{il} = 1)}{P(I : x_{il} = 1)} \\ P(x_{il} = 1 | I : x_{il} = 0) &= \frac{P(I : x_{il} = 0 | x_{il} = 1) \cdot P(x_{il} = 1)}{P(I : x_{il} = 0)} \end{aligned} \quad (11)$$

Based on the above two equations, it is seen that in case of partial unreliability of information, consumer choice after having received the information remains a choice under uncertainty. This is in contrast to the case of fully reliable information, where all uncertainty is removed upon receiving a message. The updated expected utility of good i , denoted EU_i^* , is based on the updated perceptions concerning the uncertain attribute. Combining our derivations concerning i) what message is expected to be received and ii) the effect of received messages on consumer perceptions of the uncertain attribute, we can now write the expected utility of the anticipated choice situation after having searched for partially reliable information as follows:

$$\begin{aligned} EU^+ &= P(I : x_{il} = 1) \cdot \max \{EU_i^*(I : x_{il} = 1), U_j\} \\ &+ (1 - P(I : x_{il} = 1)) \cdot \max \{EU_i^*(I : x_{il} = 0), U_j\} \end{aligned} \quad (12)$$

Subtracting from Equation 12 the expected utility of the current choice situation and the costs of information acquisition yields the expected utility of searching partially reliable information.

As an illustration of the workings of this model extension towards including partially reliable information, consider the following choice situation: a consumer chooses between two goods i and j . She perceives the probability that good i possesses feature l as $P(x_{il}) = 0.8$ and is certain that good j does not possess the attribute. The importance of the uncertain attribute, denoted β_l , equals 5. There is a brand-preference for good j over i : $\beta_0 = -3$. Prices of the

³ We give the probabilities for $x_{il} = 1$, subsequent derivation of those for $x_{il} = 0$ is straightforward.

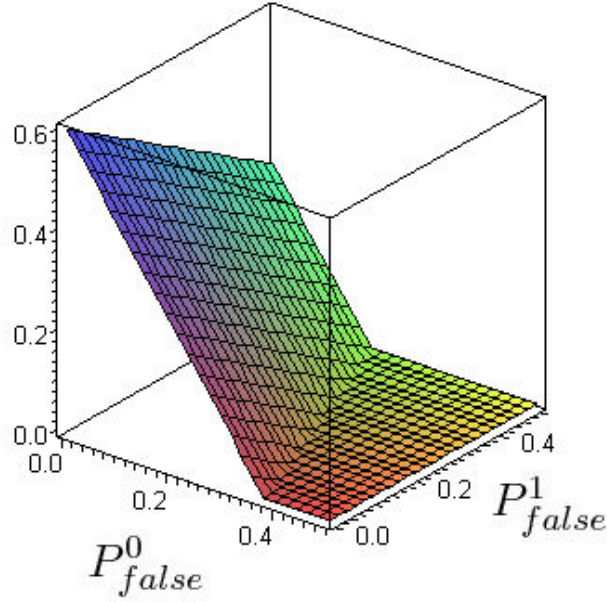


Figure 3: Information utility as a function of the probability of receiving incorrect messages

two goods are equal. Thus, the expected utility of the current choice situation equals $\max(-3 + 0.8 \cdot 5, 0) = 1$. Information I concerning the value of x_{il} can be obtained at no costs (i.e. $c = 0$). However, information is perceived as only partially reliable in the sense that there is a perceived probability that a received message claims that $x_{il} = 1$ where in fact $x_{il} = 0$, and vice versa. Figure 3 shows the expected information utility for varying levels of magnitude of the probabilities of receiving incorrect messages $P(I : x_{il} = 1 | x_{il} = 0)$, denoted P_{false}^1 , and $P(I : x_{il} = 0 | x_{il} = 1)$, denoted P_{false}^0 . It is directly seen that the expected utility of information increases as the perceived probability of receiving incorrect messages approaches zero. As the perceived probability of receiving incorrect messages increases, expected information utility drops, as would be expected. Note that should the perceived probability of receiving incorrect messages approach 1, information value would rise again. This is logical, as consumers then know that when receiving a message, they may derive the true value of x_{il} with certainty: when a message $I : x_{il} = 0$ is received, the updated perception consists of knowing with certainty that $x_{il} = 1$. Information utility is lowest when the probability of receiving incorrect messages ≈ 0.5 .

In summary, it appears that extending the proposed model with Bayes' law of conditional probabilities makes it possible to provide a meaningful account of the way in which consumers deal with partially reliable information.

4.2 Applicability beyond expected-utility maximisation premises

It should again be emphasised here that our approach is tested in this paper on a dataset generated by a process that is congruent with the model subsequently used for estimation. That is, both data generation and model estimation are based on the notions that information search is the result of a cost-benefit trade-off, and that the benefit consists of the difference in expected utility between the anticipated and current choice situation.

The more the behavioural dynamics underlying observed information search patterns in real life differ from these model premises, the less this particular model will be able to provide meaningful estimates for consumer preferences. This may be seen as a disadvantage, but one should realise that similar assumptions regarding the match between modelled and observed behaviour of course apply to any model that is used to derive preferences from choices based on a set of behavioural assumptions. For example, conventional stated choice models are predominantly based on utility-maximisation premises. Notwithstanding this, it should be noted that in order to extract preferences from observed information search, stronger behavioural assumptions are needed than is generally the case in conventional models with a choice task.

However, it should be noted that the general argument made in this paper – preferences for goods can be derived from observed information search patterns - remains valid as long as consumers' decisions to search for information can to a reasonable extent be described in terms of the preferences for the goods the information relates to (i.e. as long as decisions to search for information are not made in a completely random way). This latter assumption seems fairly reasonable, indicating the potential applicability of the ideas developed here for a wide range of perspectives on consumer choice.

4.3 The case of multidimensional uncertainty

Where in this paper, we focus on the situation where only one of a good's attributes is uncertain, it seems plausible to assume that in many real life situations, consumers will perceive multiple attributes as uncertain and may acquire multiple bits of information to reduce the uncertainty. In cases where it may be safely assumed that the true values of the uncertain attributes are perceived among consumers as being mutually independent, the model proposed earlier in this paper can be extended towards the situation of multidimensional uncertainty as follows.

Firstly, the model extension must take into account how the individual combines his or her initial knowledge with a received bit information. In case of fully reliable information, this updating process is trivial as it amounts to replacing the initial perceptions with received information. When information is not (assumed to be) perceived as fully reliable, updating mechanisms such as Bayesian learning may be appropriate.

Secondly, the expected utility of the anticipated choice situation after having received information concerning one of the uncertain attributes must incorporate the possibility that further information is acquired after having received the message. That is, the expected utility of acquiring information concerning one of the uncertain goods is a function of how the forward looking consumer evaluates the utility of future decisions whether or not to search for information, based on the set of messages already received. Naturally, one might wonder if consumers may be safely assumed to fully anticipate such elaborate sequences of potential information search. The assumption of more or less myopic behaviour, modelling individuals' anticipation of only a limited number of future information searches (e.g. no or only one additional search after having received a message), may be more intuitive (e.g. Hauser et al., 1993; Chorus et al., 2006, 2007b).

Econometrically, dealing with situations where uncertainty and information search are multidimensional implies that information search and subsequent product choice are modelled as a series of conditional probabilities (e.g. Lerman and Mahmassani, 1985). That is, the probability of observing a particular sequence of information acquisition and subsequent product choice can be written as the product of probabilities associated with observing each separate decision to acquire a bit of information (or choose a product), conditional on all information already searched for at the time of that decision. Chorus et al. (2007b) develop such a sequence of conditional probabilities to model multidimensional information search, in the context of travellers' route- and mode-choices under conditions of uncertain travel times and costs.

Things become more complicated when we do not want to assume that the true values of the uncertain attributes are perceived among consumers as being mutually independent: in many real-world situations, uncertain attributes like price may for example be (perceived to be) positively correlated with uncertain quality-related attributes, as a higher price may signal higher quality of the good (e.g. Wolinsky, 1983; Milgrom and Roberts, 1986). In such cases of mutual dependency among uncertain attributes, the benefit of searching for information concerning on one of these attributes will generally differ from what models assuming independence across attributes would predict. Take for example the situation where price is perceived among consumers to positively correlate with quality. Information about price will then also be perceived to contain additional (indirect) information concerning quality. As a result, its value for consumers will be higher than if price were perceived to be uncorrelated with quality. Also, the value of subsequent information concerning uncertain quality-related attributes will be lower than in situations where price-information does not provide indirect quality-related information. The value of information in situations where several uncertain attributes are perceived to be correlated can be modelled using Bayesian Belief Networks (Pearl, 1988), where consumer beliefs are modelled as a coherent set of conditional probabilities. However, incorporating the notion of mutual dependencies across uncertain attributes will generally lead to more complex and less tractable models.

4.4 When a consumer's perceptions are unknown

An assumption that has been made throughout the analyses presented here is that the analyst is aware of the probabilities associated with the uncertain attribute. Although this assumption is in a sense equivalent to the assumption -made in many conventional choice based models- that the analyst knows the consumer's perceptions of the values of attributes of goods, it is somewhat less intuitive. Especially when the uncertain attribute is described by non-binary probability distributions, for example in the situation where a good's price is uncertain in the eyes of the consumer, analysts may have difficulties with making valid assumptions concerning consumer perceptions of the uncertain attribute.

Besides the obvious approach of simply asking participants about their perception of (uncertain) attributes, analysts may need to partly base their assumptions concerning consumer perceptions regarding the (uncertain) attributes on intelligent guesses and analogies. Another possible way to deal with unknown consumer perceptions is to formulate a probability distribution describing the analyst's knowledge concerning consumers' perceptions of an uncertain attribute. Conditional on a particular set of perceptions drawn from this distribution, expected information utility may be computed using the equations presented here. The unconditional expected information utility may subsequently be computed by averaging (integrating) over the distribution representing the analyst's knowledge. Finally, it is well known that it is theoretically possible to infer perceived uncertainty, in terms of subjective probabilities, directly from observed choices (e.g. Savage, 1954). However, the theoretical ability to identify subjective probabilities and preferences simultaneously from information search patterns may be difficult to apply in practice, particularly when uncertainty is multidimensional.

4.5 Final remarks

Summarising, it may be said that the result presented earlier in this paper (information search patterns may be used to accurately estimate consumer preferences) should be seen in light of a number of simplifying assumptions made in the process of constructing the model and generating the data. However, it also seems that these assumptions may be relaxed without compromising the applicability of the proposed model in general. It should be kept in mind that such relaxations will generally lead to increases in model complexity.

5 Conclusions

Motivated by the notion that consumers' use of the internet creates a wealth of data on information search processes, this paper puts forward the idea that observed information searches may be used for deriving consumer preferences. In addition to a theoretical underpinning of the notion of extracting preferences from information search within an expected utility framework, we provide an example

of how this process may work in practice. The analyses, based on an artificial dataset of information search behaviour, show that observed information search concerning an uncertain attribute of a good is sufficient to estimate consumer preferences concerning the good. Estimated preferences are statistically indistinguishable from true preferences (i.e. the preferences used to generate the data). Although the model presented here relies on a number of simplifying assumptions, take for example the premises behind expected utility maximisation, it appears that these can generally be relaxed without compromising the potential of the approach in general.

It should be noted here again that in this paper, we have considered a specific type of choice situation, involving the choice between two goods where only one attribute is uncertain, whereas many other types of situations may be relevant in a consumer choice context. As noted in the introduction, we feel that the work presented here contains a first attempt to use information search data in general, by focusing on this context in particular. Furthermore, the relative simplicity of the choice context we consider is motivated by our attempt to show that consumer preferences may be derived from observed information search alone, which at first sight contains only little information on these preferences. It is to be expected that the analysis of most real life (internet-based) choice situations involves more complicated types of information search than the ones discussed in this paper. The argumentations, formalisations and illustrations presented here are designed to provide a fruitful starting point for such more involved analyses.

We feel that further research efforts that use information search observations for preference estimation in other types of choice contexts (e.g. consideration set formation, exploratory search, multi-attribute uncertainty) are needed to gain further insight into the potential of the approach presented here.

Most importantly of course, model estimation on non-artificial data is certainly needed to show the applicability of the approach in practice. Several data-collection methods may prove useful in this regard. From an experimental economics perspective, our model may be tested in a highly controlled laboratory setting where preferences and information costs are imposed by means of a carefully designed induced value scheme (Smith, 1976). Given this input, our model predicts under what circumstances (e.g. levels of uncertainty, importance of uncertain attributes) information will be acquired. Analysis of information search in this experimental environment can tell us how well the model matches observed behaviour.

A more direct approach would consist of estimating the developed information search model on observed stated choice data. The difference with the experimental economics perspective lies in the fact that stated choice-experiments do not induce preferences; instead, they are designed to infer preferences from choices. In order for the experiment to ‘match’ the developed model, participants must face a number of alternatives with uncertain attributes, and be given the opportunity to either acquire information or directly execute an uncertain alternative. Chorus et al. (2007a) provide an example of such an experiment in the context

of travellers' route- and mode-choices. Having collected data on choices among alternatives as well as on information search patterns, our model can be tested by estimating separate parameters (e.g. a parameter reflecting travel time importance in a travel context) using the model of choice among alternatives on the one hand, and the information search model on the other hand. Statistical analysis of the two estimated parameter sets may then determine whether or not the estimated parameters differ between the two models. If the two parameters sets do not differ in a statistical sense, this implies that for the given dataset, inferring preferences from information search yields the same results as conventional choice model estimation. This would be evidence of the validity and applicability of the proposed model.

A third data-collection perspective involves using revealed choices. Although this type of data obviously yields higher levels of external validity than the experimental data discussed above, it is clear that revealed data-collection methods provide much less scope for control. As such, data is likely to be relatively noisy, which creates difficulties when trying to infer whether preferences derived from information search statistically match those derived from conventional choice models. The highest levels of control in a real world setting are likely to be found in an internet-based shopping environment. In the recent tradition of clickstream-data analysis (Bucklin et al., 2002), both information search patterns as well as choices among goods may be adequately observed in such an e-shopping environment. A crucial difference with the experimental approaches outlined above is that the analyst cannot control uncertainty levels. In contrast, there is a need to infer buyers' perceptions of uncertainty concerning the attributes of goods. Section 4 discusses ways how this may be done.

In sum: empirical testing of the model developed in this paper, although not necessarily an easy thing to do, is certainly possible and can be achieved using a number of stated and revealed data-collection methods.

Acknowledgements

We thank three anonymous referees for making many useful remarks on an earlier version of this paper.

A An example of a simulated choice for information search

Remember the following parameter settings:

- -1 util per dollar regarding the attribute *pricediff*,
- 50 utils for the attribute *feature* and
- -4 for *cost*.

Now take for example the following case:

- $pricediff = 10$,
- $P(feature=1) = 0.8$,
- brand preference for A over B = -0.99 (drawn from $N(0, 1)$),
- the drawn *iid* error for information utility = 1.14,
- the one drawn for the utility of no information acquisition = -0.59.

This gives the following results:

- The utility of computer A with *feature* equals $10 \cdot (-1) + 50 - 0.99 = 39.01$.
- The utility of A without *feature* equals $10 \cdot (-1) - 0.99 = -10.99$.
- The expected utility of A equals $0.8 \cdot 39.01 + 0.2 \cdot (-10.99) = 29.01$.
- The utility of B equals 0 by definition.

Now, the utility of information equals

$$0.8 \cdot \max(39.01, 0) + 0.2 \cdot \max(-10.99, 0) - \max(29.01, 0) - 4 + 1.14 = -0.66.$$

The utility of not searching for information equals -0.59. Since $-0.59 > -0.66$, the consumer does not search for information concerning *feature*.

References

- Ackerberg, D. A., 2003. Advertising, learning, and consumer choice in experience good markets: An empirical examination. *International Economic Review* 44 (3), 1007–1040.
- Battelle, J., 2005. *The Search: How Google and Its Rivals Rewrote the Rules of Business and Transformed Our Culture*. Portfolio, New York.
- Ben-Akiva, M., Lerman, S. R., 1985. *Discrete Choice Analysis: Theory and Application to Travel Demand*. MIT Press, Cambridge, MA.
- Bettman, J. R., Capon, N., Lutz, R. J., 1975. Cognitive algebra in multi-attribute attitude models. *Journal of Marketing Research* 12, 151–164.

- Bucklin, R. E., Lattin, J. M., Ansari, A., Gupta, S., Bell, D., Coupey, E., Little, J. D. C., Mela, C., Montgomery, A., Steckel, J., 2002. Choice and the internet: From clickstream to research stream. *Marketing Letters* 13 (3), 245–258.
- Capra, R. G., Pérez-Quinones, M. A., 2005. Using web search engines to find and refine information. *Computer* 38 (10), 36–42.
- Chorus, C. G., Arentze, T. A., Molin, E. J. E., Timmermans, H. J. P., van Wee, B., 2006. The value of travel information: Decision strategy-specific conceptualizations and numerical examples. *Transportation Research Part B: Methodological* 40 (6), 504–519.
- Chorus, C. G., Molin, E. J. E., Arentze, T. A., Hoogendoorn, S. P., Timmermans, H. J. P., van Wee, B., 2007a. Validation of a multimodal travel simulator with travel information provision. *Transportation Research Part C: Emerging Technologies* 15 (3), 191–207.
- Chorus, C. G., Walker, J., Ben-Akiva, M., 2007b. Traveler decision making under conditions of knowledge limitations and information provision. paper presented at the 86th annual meeting of the Transportation Research Board, Washington, DC.
- Edwards, W., Lindman, H. R., Savage, L. J., 1963. Bayesian statistical inference for psychological research. *Psychological Review* 70, 193–242.
- Erdem, T., Keane, M. P., Öncü, T. S., Strebel, J., 2005. Learning about computers: An analysis of information search and technology choice. *Quantitative Marketing and Economics* 3 (3), 207–246.
- Hauser, J. R., Urban, G. L., Weinberg, B. D., 1993. How consumers allocate their time when searching for information. *Journal of Marketing Research* 30, 452–466.
- Huneke, M. E., Cole, C., Levin, I. P., 2004. How varying levels of knowledge and motivation affect search and confidence during consideration and choice. *Marketing Letters* 15 (2-3), 67–79.
- Lancaster, K. J., 1966. A new approach to consumer theory. *Journal of Political Economy* 74, 132–157.
- Lerman, S. R., Mahmassani, H., 1985. The econometrics of search. *Environment and Planning A* 17, 1009–1024.
- Louviere, J. J., Hensher, D. A., 1983. Using discrete choice models with experimental design data to forecast consumer demand for a unique cultural event. *Journal of Consumer Research* 10, 348–361.
- Louviere, J. J., Woodworth, G., 1983. Design and analysis of simulated consumer choice and allocation experiments: A method based on aggregate data. *Journal of Marketing Research* 20, 350–367.
- Marchionini, G., 2006. Exploratory search: From finding to understanding. *Communications of the ACM* 49 (4), 41–46.
- McFadden, D., 1974. Conditional logit analysis of qualitative choice behaviour. In: Zarembka, P. (Ed.), *Frontiers in Econometrics*. Academic Press, New York, pp. 105–142.
- McFadden, D., Train, K., 2000. Mixed MNL Models for discrete response. *Journal of Applied Econometrics* 15, 447–470.
- Milgrom, P., Roberts, J., 1986. Price and advertising signals of product quality. *Journal of Political Economy* 94, 796–821.
- Moe, W. W., 2006. An empirical two-stage choice model with varying decision rules applied to internet clickstream data. *Journal of Marketing Research* 43 (4), 680–692.
- Moon, B.-J., 2004. Consumer adoption of the internet as an information search and product purchase channel: Some research hypotheses. *International Journal of Internet*

- Marketing and Advertising 1, 104–118.
- Pearl, J., 1988. Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference. Morgan Kaufmann Publishers, San Mateo, CA.
- Peterson, R. A., Merino, M. C., 2003. Consumer information search behavior and the internet. *Psychology and Marketing* 20 (2), 99–121.
- Punj, G. N., Staelin, R., 1983. A model of consumer information search behavior for new automobiles. *Journal of Consumer Research* 9, 366–380.
- Raiffa, H., Schlaifer, R., 1961. Applied Statistical Decision Theory. Harvard University Press, Boston, MA.
- Ratchford, B. T., 1982. Cost-benefit models for explaining consumer choice and information seeking behavior. *Management Science* 28, 197–212.
- Ratchford, B. T., Lee, M. S., Talukdar, D., 2003. The impact of the internet on information search for automobiles. *Journal of Marketing Research* 40 (2), 193–209.
- Roberts, J. H., Lattin, J. M., 1991. Development and testing of a model of consideration set composition. *Journal of Marketing Research* 28, 429–440.
- Roberts, J. H., Lattin, J. M., 1997. Consideration: Review of research and prospects for future insights. *Journal of Marketing Research* 34 (3), 406–410.
- Savage, L. J., 1954. The Foundations of Statistics. Wiley, New York.
- Schocker, A. D., Ben-Akiva, M., Boccara, B., Nedungadi, P., 1991. Consideration set influences on consumer decision-making and choice: Issues, models, and suggestions. *Marketing Letters* 2, 181–197.
- Shugan, S. M., 1980. The cost of thinking. *Journal of Consumer Research* 7, 99–111.
- Smith, V. L., 1976. Experimental economics: Induced value theory. *The American Economic Review* 66, 274–279.
- Train, K., 2003. Discrete Choice Methods with Simulation. Cambridge University Press, Cambridge, MA.
- Weibull, J. W., 1978. A search model for microeconomic analysis with spatial applications. In: Karlqvist, A., Lundqvist, L., Snickars, F., Weibull, J. W. (Eds.), *Spatial Interaction Theory and Planning Models*. North-Holland Publishing Company, Amsterdam, pp. 47–73.
- Wilde, L. J., 1980. The economics of consumer information acquisition. *Journal of Business* 53, 143–158.
- Wilkie, W. L., Pessemier, E. A., 1973. Issues in marketing's use of multi-attribute attitude models. *Journal of Marketing Research* 10, 428–441.
- Wolinsky, A., 1983. Prices as signals of product quality. *The Review of Economic Studies* 50, 647–658.