Title
A Bayesian Model of How People Search Online Consumer Reviews

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 30(30)

ISSN
1069-7977

Authors
Lelis, Stelios
Howes, Andrew

Publication Date
2008

Peer reviewed

A Bayesian Model of How People Search Online Consumer Reviews

Stelios Lelis (stelios.lelis@gmail.com)
Manchester Business School, Booth Street West
Manchester, M15 6PB UK

Andrew Howes (andrew.howes@mbs.ac.uk)
Manchester Business School, Booth Street West
Manchester, M15 6PB UK

Abstract

In this paper we describe a model of how people search online consumer reviews in service of purchasing decisions. The model is similar to other recent models of information seeking in that it updates estimates of products' utilities using Bayesian inference. It is different, in that it stops seeking further information when the confidence that one of the alternatives is the best exceeds a threshold. Findings from a controlled experiment support the model by suggesting that high variance in review ratings causes people to seek more information.

Keywords: Information search; online consumer reviews; user modeling; threshold models.

Introduction

Online opinion forums like Epinions and c|net, and review sections of retailer websites are used by people eager to both share and gather views on interesting products. Using these sites, those who have made purchases can contribute reviews, and those who are planning to make a purchase can find information that may assist decision making.

In recent years, consumer opinions have become an important component of the product related information that is available to potential buyers (Bei, Chen, & Widdows, 2004). However, despite their apparent importance, there is very little work exploring their impact on how people search for information in service of purchasing decisions. Yet, developing an understanding of how potential buyers consult online consumer reviews could provide insights into how consumer search processes can be facilitated through interface design (Miles, Howes, & Davies, 2000).

In this paper we report work towards an understanding of why and how people consult online consumer reviews. We propose a criterion-dependent Bayesian choice model of online consumer reviews search and inspection. From this model, we derive two predictions regarding the depth of consumers' opinion search and subsequently we report a test of the predictions that used a controlled experiment. To foreshadow the results, the experiment provides some evidence to support the model but some of the findings were inconclusive. Possible explanations are outlined. Finally, recommendations for further investigation are discussed.

Literature Review

Online Opinions and Decision Making

Most research in the effects of online opinions on decision making concentrate on their influence on attitudes towards products and product choices, and on the mediators of opinions' persuasive power (Huang & Chen, 2006; Park & Han, 2008; Xue & Phelps, 2004). Further, limited evidence suggests that consumers consult online opinions to reduce buying-related risks and decrease decision effort (Hennig-Thurau & Walsh, 2004), and that the mere availability of one opinion for a single option decreases search time (Smith, Menon, & Sivakumar, 2005). However, these studies do not allow us to ascertain the process of how people consult online consumer reviews.

One possibility is that online decision-making processes are similar to offline decision-making processes. Many studies of traditional offline word-of-mouth suggest that advice from family and friends facilitate consumers in reducing buying related risks and informational uncertainties; for example see (Lutz & Reilly, 1974). Taken together, these studies suggest that a potential explanation of consumers' opinion seeking should incorporate choice related uncertainties.

Information Search and Stopping Rules

In cognitive science the study of information search tasks has been strongly influenced by Pirolli and Card's information foraging theory (IFT) (1999), which predicts user behavior in general search tasks. More specifically, it predicts that a patch of information should be left, so as to exploit another, when the rate of within patch gains diminishes below the expected average rate of gain. Similarly, Fu and Pirolli's SNIF-ACT model (2007), predicts users' link selection on a Web page and when the current web page will be left.

Traditional information economics puts the search for information in a decision making context. Likewise we set opinion seeking in a product choice situation. Furthermore, information economics posits that people search for information until the costs of searching for more outweigh the potential gains of acquiring it (Stigler, 1961). Similarly Fu and Gray (2006) assume that information seeking in a map-navigation task stops when the estimated utility of the information is lower than the information seeking cost. Both information utility and gain are opertionalised in terms of time. In our model, we view information collection as a means to reduce choice related uncertainties.

Threshold models take a different view on when people stop acquiring information (For an excellent discussion see (Hausmann & Lage, 2008)). The core idea is that people stop searching for information when their confidence that one of the decision alternatives outperforms the rest reaches or exceeds a threshold. In our model we utilize a similar 'desired level of confidence' criterion of stopping opinion seeking. We assume that prospective purchasers stop reading opinions, i.e. they stop searching for information, when they have decided which product is best given some desired level of confidence.

The Model

We view and model opinion search and acquisition as a means to reduce choice related uncertainties. In line with proposed models of information acquisition (Hagerty & Aaker, 1984; Moorthy, Ratchford, & Talukdar, 1997), we assume that the consumer faces uncertainty about how the products under consideration perform, and that opinion acquisition reduces this uncertainty. That is, the consumer is not sure of products' true value. Rather he holds beliefs about each product's true mean value, which in the model are represented by product-specific distributions g_i, where i stands for product. The consumer evaluates the products using a utility function $U(g_i)$. Therefore, consumers' uncertainty about how the products perform is represented in utility terms with product-specific utility distributions denoted by $f_i = U(g_i)$. We assume that the utility function is the identity one, consequently $f_i = g_i$. The consumer updates his beliefs about product's true mean value (and thus utility) by acquiring reviews.

At any point during the search process the consumer faces uncertainty about which product out-performs the rest. We denote this uncertainty as $p(t)$, where t stands for time step. It is the probability that the product with the current highest mean estimated utility will turn out not to be so. Figure 1 pictures consumer's product-specific utility distributions in a binary choice; $p(t)$ is analogous to the degree of overlap between the two distributions. As long as the uncertainty, namely $p(t)$, is high, the decision maker keeps acquiring reviews, and updates his beliefs about products' true mean value. However, as soon as $p(t)$ falls below a threshold value, opinion acquisition stops and the product with the highest mean estimated utility is selected. The stopping rule is a variation of satisficing (Simon, 1955) in which the aspiration level is not product value, but rather the confidence that one of the products is better than the others[1].

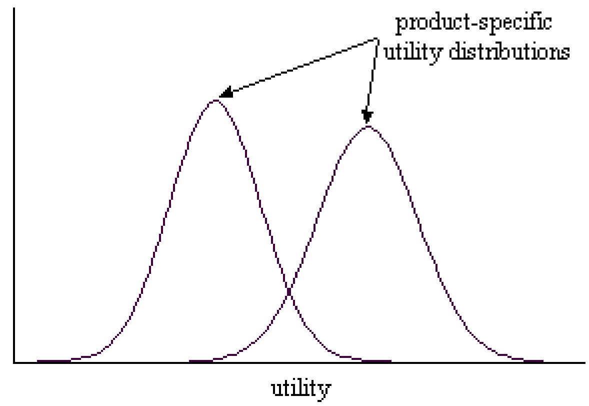

Figure 1: Graphical representation of consumer's product-specific utility distributions.

Model Description

The model describes information search in a binary choice. We assume that the true value of product i follows a Normal distribution. The decision maker does not know the mean of that distribution, i.e. product's mean value, but he does know its variance[2]. Each individual review is a 'signal' of product value and follows the above mentioned normal distribution[3]. The decision maker holds beliefs about the mean of the true product value, g_i, and product's utility, f_i, and as more opinions are acquired the beliefs are being updated. We assume decision maker's prior beliefs of product's i true mean value at time step 0 to follow a Normal distribution with mean $\mu_i(0)$ and variance $\sigma_i^2(0)$.

At each time step t the consumer (i) selects a product for which to read a review, (ii) inspects the selected product's next review and updates his beliefs about product's true mean value and utility, (iii) calculates the probability $p(t)$, and (iv) decides whether to stop the search process and make a choice or to obtain more reviews. Below we describe each action in detail.

(i) *Selection of product for which to read review*: At time step $t = 1$ the product is randomly selected from consumer's consideration set. At time step $t > 1$ the product for which to read a review is that of previous time step $t - 1$, unless all of the product's reviews have been read, or a switch to another product threshold has been reached. The switch threshold is defined in terms of p's difference between the two last time steps, $p(t$-

[1] The stochastic nature of utility in the model and the definition of $p(t)$ resembles Random Utility Models (RUM) (Baltas & Doyle, 2001). In this family of models it is assumed that products' utilities follow random distributions, and that the decision maker is a rational utility maximizer. Therefore, in a binary choice between products i and j, the decision maker will choose the product i with probability $P(i) = P(U_i > U_j)$. Apparently, our model's $p(t)$, and RUM's $P(i)$ are related. If $E(U_i) < E(U_j)$ then $p(t) = P(i)$, whereas, if $E(U_i) > E(U_j)$, $p(t) = P(U_i < U_j) = 1 - P(i)$.

[2] Off-course the assumption of known variance is a simplification. The model can be extended to incorporate unknown variance but on the current state of development the simplified model proved capable of yielding testable predictions.

[3] Hu et al. (2007) recently demonstrated that the distribution of online consumer reviews is not Normal but J-shaped. However, at the current state of model development reviews were assumed to follow a Normal distribution.

2) – $p(t$-1), and the product is switched if the absolute of the difference is below a threshold value θ.

(ii) *Inspect review and update beliefs about product's true mean value and utility*: The distribution of beliefs about product's true mean value (and thus utility) is updated by incorporating review's rating. Given that both prior beliefs and reviews follow a Normal distribution, the posterior distribution is still Normal with mean and variance given by (DeGroot, 1970; Roberts & Urban, 1988):

$$\mu_i(t) = \frac{\tau_t}{\tau_t + 1} \cdot \mu_i(t-1) + \frac{1}{\tau_t + 1} \cdot x$$

$$\sigma^2_{\mu i}(t) = \frac{\tau_t}{\tau_t + 1} \cdot \sigma^2_{\mu i}(t-1)$$

Where x is review's rating, and τ_t is the strength in prior beliefs at time step t. τ changes with time as $\tau_t = \tau_{t-1} + 1$.

(iii) *Calculate p(t)*: The distribution of p(t), the probability that the product with the current highest mean estimated utility will turn out not to be so is Normally distributed as is the difference of two Normal distributions. Calculating $p(t)$ is then straightforward[4].

(iv) *Decision whether to stop the search process and make a choice or not*: The search process stops either if $p(t)$ is bellow a threshold value k, or if all the reviews of all of the products have been obtained. In any of the two situations the product with the highest mean estimated utility is chosen. If it has been decided not to stop the search process the model moves on to the next time step.

Predictions

The model predicts the relative number of reviews that people will read given prior beliefs and three parameter values, k, θ, and τ_0. To verify its behavior we created different scenarios of prior beliefs of products' true mean value and we ran numerical simulations across a large parameter space. The scenarios involved choices between two products. They were built to examine whether a small gap between the means of products' prior beliefs will result in more search compared to a large one, and whether high prior beliefs' variance will entail more reviews to be read compared to low variance. The gap between prior beliefs' mean values was manipulated in five levels, namely 0.4, 0.6, 0.8, 1 and 1.2. The variance was also manipulated in five levels, specifically 0.2, 0.4, 0.6, 0.8 and 1, and was the same for each product within each scenario. The product ratings with which beliefs were updated were held constant across all scenarios and were of mean 4.4 for the product with the higher mean of prior beliefs and 3.6 for the other. The actual ratings were (5, 5, 5, 5, 5, 5, 4, 4, 3, 3) and (5, 5, 4, 4, 3, 3, 3, 3, 3, 3), and of identical variance equal to 0.71.

Each scenario was run separately and for five times across the following parameter space: τ = 1:1:10, k = 0.01:0.01:0.1, θ = 0.005:0.005:0.05. By firstly averaging the number of reviews read across the parameter space for each run, and then across each scenario, we obtained the average number of reviews read for each scenario and confidence intervals. As Figure 2 indicates, the model predicts that more reviews will be acquired when the gap between prior-belief means is small and/or when the variance is high.

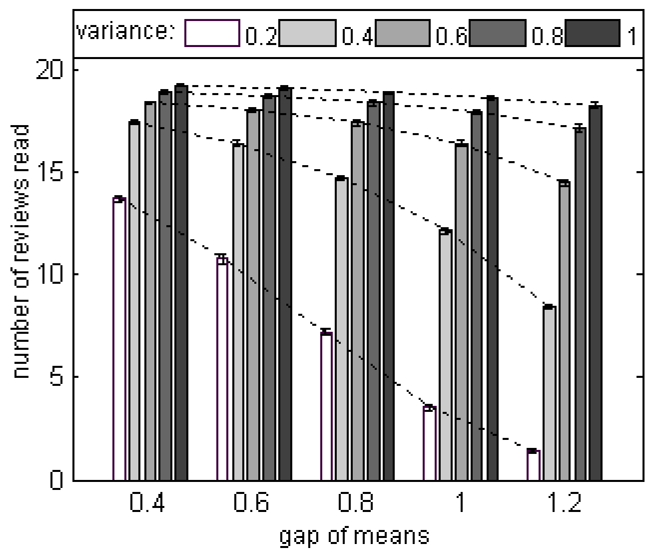

Figure 2: Model Predictions: The average number of reviews acquired for different values of prior belief's gap of means and variance.

Experiment

To test the predictions of the model we conducted a controlled experiment in which we manipulated participants' prior beliefs about products and measured search activity.

Design

A 2 × 2 full factorial experimental design was used. The manipulated factors were prior beliefs' gap of means (high, low) and prior beliefs' variance (high, low). The gap of means was manipulated between subjects and the variance within. The main task consisted of selecting a digital camera among three alternatives in each of two categories (5 and 7

[4] $p(t)$ is exactly the same probability of selecting a product in the binary probit choice model. For the purpose of illustration let product i have a higher mean estimated utility value than product j. Then $p(t)$ is the probability of selecting product j. In the binary probit choice model this probability is $\Phi[(\mu_j - \mu_i) / \sigma]$, where $\sigma^2 = \sigma_{\mu i}^2 + \sigma_{\mu j}^2$ and the covariance of the two Nomral distributions is assumed to be zero (Louviere, Hensher, & Swait, 2000, p. 362). Exactly the same formula is obtained by standardizing the Normal distribution of $p(t)$.

mega-pixels resolution)[5]. The manipulation of prior beliefs was accomplished by splitting the task into two phases. In the first phase participants were exposed to product features and five reviewers' ratings for each alternative. The manipulation was introduced through these ratings. In the second phase participants were free to read consumer reviews for each camera. Second phase reviewer ratings were similar across all conditions to ensure that any effect on information search was due to the manipulation on the first phase and only.

First Phase

Product information presented in the first phase was carefully constructed to ensure that any effects in the 2nd phase were due to the differentiated ratings alone. Choice alternatives were very similar, non-dominating, each had a fictitious name, and was accompanied by a camera picture. Differentiated product features, cameras' names and images, and order of presentation were randomized for each choice. Along with this information, each camera was described with 5 consumer ratings. The ratings were from 1 to 5 and their mean ranked the three alternatives as best, 2nd-best and worst. However, average ratings were not presented. Best camera's ratings were different between the variance conditions, but the same within the gap conditions. The ratings of the 2nd-best camera were manipulated across all conditions, while worst camera's ratings were not manipulated. Table 1 summarizes products' ratings across all conditions. The difference between best and 2nd-best products average ratings was 0.4 in the low gap condition, and 0.8 in the high. Best product's variance was 0.2 in the low variance condition and 0.8 in the high variance condition, while 2nd-best product's was 0.3 and 1.2 respectively. The order of the ratings was randomized for each product and choice. Finally, all products information was presented in a single web page, with layout very similar to the 'Compare Products' pages of real web sites like Epinions for example.

Table 1: Products' ratings across conditions.

Product		Low Variance	High Variance
Best		5, 5, 5, 4, 4	5, 5, 5, 5, 3
2nd-best			
	Low Gap	5, 4, 4, 4, 4	5, 5, 5, 3, 3
	High Gap	4, 4, 4, 4, 3	5, 5, 3, 3, 3
Worst		4, 4, 3, 3, 2	4, 4, 3, 3, 2

Second Phase

In the second phase participants could read consumer reviews for each alternative. Firstly, they were forwarded to a web page with alternatives' names, pictures, and one link per camera leading to its first consumer review. Alternatives order of presentation was the same as in the 1st phase. Then, participants could navigate through product reviews by following 'Next' and 'Previous' links and also could at anytime return to the first page by clicking another link. There were 10 reviews per alternative. Each review included the star rating, 1 to 5, the title, and the opinion. Each camera's review ratings were exactly the same across all conditions and were also of average variance. The reviews were presented in one of a small number of semi-random sequences (sequences, for example, that had all low ratings at the beginning were avoided). Table 2 summarizes all products' ratings.

Table 2: Products' review ratings in 2nd phase.

Product	Ratings
Best	5, 5, 5, 5, 5, 5, 5, 4, 4, 3
2nd best	5, 5, 5, 4, 4, 4, 4, 3, 3, 3
Worst	4, 4, 4, 4, 3, 3, 3, 3, 2, 2

The actual reviews utilized were real ones downloaded from amazon.co.uk, selected according to specific criteria and slightly edited if needed[6]. Two sets of reviews were constructed according to the ratings depicted in Table 2. Consensus, in the sense that there were not two reviews in the same set offering apparently contradicting comments for the same product was intentionally build.

Finally, the order of reviews set, as well as choice (5 or 7 mega-pixels), and variance condition were counterbalanced across participants.

Procedure

Eighteen participants completed the task for £5 reward, all of them students and native English speakers. Participants first performed a practice task to get use to the environment and were informed that after each task they would be asked to justify their choice and that, a £20 prize would be awarded to the participant who would offer the best justifications[7]. After that, they proceeded to the experimental tasks.

[5] Choices between three alternatives were utilized in the experiment although the reported model describes opinion search in binary choices. However, one of the products was clearly inferior and we expected participants to concentrate on the best two products.

[6] In order for a review to be selected it should satisfy specific criteria, either as it was presented in amazon.co.uk, or after very slight editing usually deleting a few words. The criteria were (i) be of length 100 to 150 words, (ii) not to mention attributes different to the ones of our fictitious products, (iii) not to mention specific characteristics of other reviews, (iv) not to mention amazon, (v) not to mention that the camera comes with extras like case, camera dock etc., (vi) not to compare the reviewed camera to competitors, or describe the choice between two cameras, and (vii) not to largely review the video capability. After the selection, the reviews were slightly edited if needed. Spelling mistakes were corrected, mentioned attribute values were changed according to the camera category, and brand and model names were replaced by the fictitious ones. For a full list of the reviews utilized please contact the authors.

[7] Participants were asked to justify their choice as a motive to perform the task properly. Of course, it is reasonable to expect people to behave differently depending on whether a justification

Results

The dependent variables were the number of reviews read and the time participants took to make their choice in the 2nd phase. Both measure the extent of opinion seeking. Mixed ANOVA tests of the dependent variables indicated a significant effect of the variance manipulation at the alpha level of .05. The effect was observed both on the number of reviews read and the time taken to make the choice. Participants read more reviews in the high variance condition ($M = 22.78$, $SD = 6.4$) than the low variance one ($M = 19.94$, $SD = 8.47$), $F(1, 16) = 5.40$, $p = .034$, $\eta^2 = 0.252$. Similarly, participants took more time to complete the task in the high variance condition ($M = 372.4$, $SD = 196.7$) than the low ($M = 309.7$, $SD = 219$), $F(1, 16) = 6.23$, $p = .024$, $\eta^2 = 0.28$. However, there was no effect of the gap manipulation either on the number of reviews read or on the time taken to indicate choice ($Fs < 1$). To test for practice effects we introduced the order of the variance conditions as a between subjects factor. Neither order, nor variance by order effects were obtained, $p > .5$, ruling out any likely practice ones.

Discussion

In this paper we presented a model of opinion seeking that combined Bayesian update with a decision making criterion based on a desired level of confidence. From the model we derived two predictions regarding the amount of information that people would seek in service of a purchasing decision. The prediction that increased opinion variance would lead to more information gathering was supported and suggests that the model might explain the rational basis for information gathering in consumer decision making. That is, people acquire opinions to discriminate between alternatives, and increase the certainty that one alternative outperforms the rest until a confidence threshold has been reached and whereupon opinion seeking ceases.

However, no evidence was found in support of the prediction that a smaller gap of means would result in more information seeking. There are three likely reasons behind the failure to find evidence which upon refinement might reveal an effect. First, the small difference of the gap of means between the two gap conditions might be ineffective. The difference between the two best products' average ratings was only 0.4. A larger difference might induce an effect. Second, there were significant individual differences resulting in large between subjects variance. Manipulating gap within subjects will make an effect more likely. Finally, the experimental environment favored within-alternatives processing and eventually most participants processed the information accordingly. However, note that for the gap manipulation to have en effect on information search, comparisons between products' estimated utility should be performed. An environment equally favoring within- and between-alternatives processing might reveal an effect.

Our results further support the existing literature that views information search as a means to discriminate between alternatives. Harvey and Bolger (2001) describe a study in which they examined whether people collect information either according to a compensatory choice process, or in order to screen out options, or to facilitate discrimination between alternatives. The experiments suggested that people collect information to discriminate between alternatives. Our results provide further support. When it was harder to discriminate between alternatives, that is when ratings' variance was high, participants acquired more information.

The reported model falls in the general category of threshold models (Bockenholt, Albert, & Aschenbrenner, 1991; Hausmann & Lage, 2008). It extents the latter ones by applying the 'desired level of confidence' stopping rule on opinion seeking in service of purchasing decisions, and by integrating Bayesian update of products' utility estimations.

Our model also differs from Fu and Gray's Bayesian satisficing model (BSM) (2006). BSM is defined in terms of two processes; the estimation of the utility of information, and the decision on when to stop seeking information. In the first process, the model updates its estimation of the utility of information through a global Bayesian learning mechanism that combines new observations with prior knowledge of task performance. That is, Bayesian learning occurs across consecutive choices and not within each choice. In the second process, the model stops seeking information when the estimated utility of the information is lower than the information seeking cost. In contrast to BSM, our model utilizes a Bayesian mechanism to update the distribution of options' utility, and does so during the course of each choice. Further, our stopping rule is not based on the interplay between the utility of information and cost, but on the confidence that one of the options outperforms the others.

In contrast to Information Foraging Theory (IFT), our model predicts when opinion gathering ceases. Although IFT may well explain how prospective buyers locate products' attribute information and consumer opinions on the World Wide Web, it does not offer an explanation of the effect of opinion variance on information seeking. There are two different processes at work: (i) locate information, and (ii) selectively acquire and integrate the information in the internal choice representation. IFT focuses on the former, whereas our model on the latter.

The described model has certain simplification assumptions which need to be elaborated and refined in the future. First, the unrealistic assumption of known true product value variance by the decision maker (nevertheless, we do not expect model's predictions to change by dropping it). Second, the assumption of constant instead of dynamic switch product and stop opinion seeking thresholds. The thresholds may be dynamically influenced by many factors like for example alternatives' attractiveness, costs of information access and the development of relatively limited confidence that one of

is asked or not. However, which procedure has greater validity is an open question, and asking to or not to justify the choice is not expected to change the results.

the decision alternatives outperforms the rest (Bockenholt et al., 1991; Hausmann & Lage, 2008, p. 237; Saad & Russo, 1996). Likely context and environmental effects on thresholds should be experimentally tested and the model updated accordingly. Third, information utility, gain and access costs are not considered in our model. However, they could be incorporated by viewing the utility of information as the amount by which it increases, or decreases, the confidence that one of the alternatives outperforms the rest. Under this assumption the gap between information economics and threshold models can be bridged, and normative predictions of optimal stopping can be generated and tested against actual search behavior. Fourth, the model needs to be quantitative compared to alternatives. Fifth, distributions of reviews ratings are assumed to follow a Normal distribution. How the model would behave with non-Normal distributions, and particularly with realistic J-shaped ones (Hu et al., 2007), should be examined.

To conclude, we have demonstrated that a model of opinion seeking that combined Bayesian update with a criterion of desired level of confidence can be applied to predict users' depth of opinion search. Undoubtedly, the model needs to be further tested and elaborated, yet it already serves as a basis for the rational explanation of opinion seeking in consumer decision making.

References

Baltas, G., & Doyle, P. (2001). Random utility models in marketing research: a survey. *Journal of Business Research, 51*, 115-125.

Bei, L.-T., Chen, E., & Widdows, R. (2004). Consumers' Online Information Search Behavior and the Phenomenon of Search vs. Experience Products. *Journal of Family and Economic Issues, 25*(4), 449-467.

Bockenholt, U., Albert, D., & Aschenbrenner, M. (1991). The Effects of Attractiveness, Dominance, and Attribute Differences on Information Acquisition in Multiattribute Binary Choice. *Organizational Behavior and Human Decision Processes, 49*, 258-281.

DeGroot, M. (1970). *Optimal Statistical Decisions*: McGrow-Hill.

Fu, W.-T., & Gray, W. (2006). Suboptimal tradeoffs in information seeking. *Cognitive Psychology, 52*, 195-242.

Fu, W.-T., & Pirolli, P. (2007). A Cognitive Model of User Navigation on the World Wide Web. *Human-Computer Interaction, 22*(4).

Hagerty, M., & Aaker, D. (1984). A Normative Model of Consumer Information Processing. *Marketing Science, 3*(3), 227-246.

Harvey, N., & Bolger, F. (2001). Collecting information: Optimizing outcomes, screening options, or fascilitating discrimination? *The Quarterly Journal of Experimental Psychology, 54A*(1), 269-301.

Hausmann, D., & Lage, D. (2008). Sequential evidence accumulation in decision making: The individual desired level of confidence can explain the extent of information search. *Judgment and Decision Making, 3*(3), 229-243.

Hennig-Thurau, T., & Walsh, G. (2004). Electronic Word-of-Mouth: Motives for and Consequences of Reading Customer Articulations on the Internet. *International Journal of Electronic Commerce, 8*(2), 51-74.

Hu, N., Pavlou, P., & Zhang, J. (2007). Why do Online Product Reviews have a J-shaped Distribution? Overcoming Biases in Online Word-of-Mouth Communication. (Working Paper)

Huang, J.-H., & Chen, Y.-F. (2006). Herding in Online Product Choice. *Psychology & Marketing, 23*(5), 413-428.

Louviere, J., Hensher, D., & Swait, J. (2000). *Stated Choice Methods: Analysis and Applicaitons*: Cambridge University Press.

Lutz, R., & Reilly, P. (1974). An Exploration of the Effects of Perceived Social and Performance Risk on Consumer Information Acquisition. *Advances in Consumer Research, 1*, 393-405.

Miles, G. E., Howes, A., & Davies, A. (2000). A framework for understanding human factors in web-based electronic commerce. *International Journal of Human-Computer Studies, 52*(1), 131-163.

Moorthy, S., Ratchford, B., & Talukdar, D. (1997). Consumer Information Search Revisited: Theory and Empirical Analysis. *Journal of Consumer Research, 23*(March 1997), 263-277.

Park, D.-H., & Han, I. (2008). *Integrating Conflicting Reviews: Attributional Hypotheses of Consumer Response to Information Uncertainty Depending on Prior Brand Attitude.* Paper presented at the Hawaii International Conference on System Sciences.

Pirolli, P., & Card, S. (1999). Information Foraging. *Psychological Review, 106*(4), 643-675.

Roberts, J., & Urban, G. (1988). Modeling Multiattribute Utility, Risk, and Belief Dynamics for new Consumer Durable Brand Choice. *Management Science, 34*(2), 167-185.

Saad, G., & Russo, J. E. (1996). Stopping Criteria in Sequential Choice. *Organizational Behavior and Human Decision Processes, 67*(3), 258-270.

Simon, H. (1955). A Behavioral Model of Rational Choice. *Quarterly Journal of Economics, 69*(February), 99-118.

Smith, D., Menon, S., & Sivakumar, K. (2005). Online Peer and Editorial Recommendations, Trust, and Choice in Visrual Markets. *Journal of Interactive Marketing, 19*(3), 15-37.

Stigler, G. (1961). The Economics of Information. *Journal of Political Economy, 69*(January/February), 213-225.

Xue, F., & Phelps, J. (2004). Internet-facilitated consumer-to-consumer communication: the moderating role of receiver characteristics. *Int. J. Internet Marketing and Advertising, 1*(2), 121-136.