MANAGEMENT SCIENCE
Vol. 50, No. 3, March 2004, pp. 299–308
ISSN 0025-1909 | EISSN 1526-5501 | 04 | 5003 | 0299

DOI 10.1287/mnsc.1040.0194

On the Depth and Dynamics of Online Search Behavior

Eric J. Johnson
Department of Marketing, Columbia Business School, Columbia University, New York, New York 10027, ejj3@columbia.edu

Wendy W. Moe
Department of Marketing, McCombs School of Business, University of Texas at Austin, Austin, Texas 78712, wendy.moe@mccombs.utexas.edu

Peter S. Fader
Department of Marketing, The Wharton School, University of Pennsylvania, Philadelphia, Pennsylvania 19104, faderp@wharton.upenn.edu

Steven Bellman
University of Western Australia, Crawley, Western Australia, Australia, sbellman@gsm.uwa.edu.au

Gerald L. Lohse
Accenture, Philadelphia, Pennsylvania 19107, gerald.l.lohse@accenture.com

This paper examines search across competing e-commerce sites. By analyzing panel data from over 10,000 Internet households and three commodity-like products (books, compact discs (CDs), and air travel services), we show that the amount of online search is actually quite limited. On average, households visit only 1.2 book sites, 1.3 CD sites, and 1.8 travel sites during a typical active month in each category. Using probabilistic models, we characterize search behavior at the individual level in terms of (1) depth of search, (2) dynamics of search, and (3) activity of search.

We model an individual's tendency to search as a logarithmic process, finding that shoppers search across very few sites in a given shopping month. We extend the logarithmic model of search to allow for time-varying dynamics that may cause the consumer to evolve and, perhaps, learn to search over time. We find that for two of the three product categories studied, search propensity does not change from month to month. However, in the third product category we find mild evidence of time-varying dynamics, where search *decreases* over time from already low levels. Finally, we model the level of a household's shopping activity and integrate it into our model of search. The results suggest that more-active online shoppers tend also to search across more sites. This consumer characteristic largely drives the dynamics of search that can easily be mistaken as increases from experience at the individual level.

Key words: electronic commerce; dynamic consumer search; stochastic models; consumer behavior
History: Accepted by Chris F. Kemerer, former department editor; received June 2000. This paper was with the authors 10 months for 5 revisions.

1. Introduction

E-commerce has engendered a widely held belief: that because the Internet lowers search costs, people should search more. There is little doubt that search costs, as measured by time, have decreased. Consider the contrast of shopping for a book offline versus online, knowing its title. Traveling from bookstore to bookstore, even in a shopping mall, might take minutes, but a typical search agent (such as Shopping.com) will search dozens of stores and provide prices, including shipping and sales tax, within seconds.

But does decreased search cost on the Internet really lead to increased search? This question is important, both for understanding the nature of consumer behavior and competition in electronic environments, and because increased search is thought to increase efficiency of such markets, in part due to increased price competition.

Previous research has provided some evidence for limited search: Adamic and Huberman (2001) found that the top 1% of sites on the Web capture 50% of all visits to the Web, consistent with the idea that shoppers are limiting their search to a few popular sites. There is an extensive literature describing search for goods in nonelectronic markets, mostly based on self-reports (see Newman 1977 for a review). For example, the number of stores visited for the purchase of a major appliance is reported to be about three (Beatty and Smith 1987). Prior research has often argued that the observed amount of search is surprisingly low (Wilkie and Dickson 1985). However, recent research suggests that consumers' priors may account for some of this departure (Moorthy et al. 1997). The current research allows us to revisit some of these issues, both because of the decreased costs of search provided by the Internet and the use of observational data about search, which although imperfect may be superior to self-reports that have been reported to be poor measures of actual search (Newman 1977).

In this paper, we try to understand these issues by looking directly at search behavior of shoppers as they visit online retailers. Because electronic markets often allow us to observe search (whereas data from offline markets are generally restricted to purchasing transactions only), we think this is an important new source of evidence. In addition, we can examine how the depth of search systematically differs across consumers and can therefore offer models that will allow us to separate the factors that might influence search.

Like several previous studies, we use observable clickstream data to examine consumer search across sites and how search behavior may differ across individuals as well as over time. Other studies in this area have focused on behavior for a given store over time (Moe and Fader 2004) or within a given store session (Bucklin and Sismeiro 2003, Moe 2002). Other papers look at behavior across sites: Johnson et al. (2003), for example, have compared rates of learning and purchasing across sites, and Park and Fader (2003) provide a model that links an individual's visiting behavior at one site to that same individual's visiting behavior at another site. However, our focus is different, concentrating on competition among sites for visits by customers.

Given available data, we face several challenges. Our major data will be site-visit behavior and the tendency for shoppers to visit multiple sites. Because buyers frequently shop in multiple channels and because clickstream data do not identify purchases, it is difficult, if not impossible, to definitively link store visits to a specific purchasing decision. Our strategy, then, is to examine search across a variety of sites within a specified window of time as a close approximation. Because research using Internet clickstream data is a relatively new area, there are some limitations to the data and the type of analysis that can be done. Though this behavior may represent shoppers' overall tendency to search for information and not necessarily their tendency to search for purchase, it does serve as an upper bound. We find that, overall, search for information is low, much lower than might be expected. Because search for purchase may be a subset of information search, our results might be indicative, to some degree, of the amount of prepurchase search in online environments. We find from our analysis that search behavior is actually quite limited, a surprising result given that there are theoretically no physical search or transportation costs online. This seems to suggest that perhaps there are more factors driving search behavior than the physical costs of search (e.g., cognitive search costs or lock-in). We provide a more in-depth discussion of other possibilities at the end of this paper.

2. Data

We examine consumer search directly, by looking at the shopping patterns of a large panel of Internet users over time. We use data collected by Media Metrix, Inc. (now part of ComScore Networks), a firm that recorded every URL visited by families that are members of its panel by using a small program, the PCMeter, which ran continuously as a background application on the family's home computer.

The U.S. panel we analyzed corresponded to approximately 10,000 households at any given point in time, recruited by Media Metrix to be representative of the U.S. population. We examine the online shopping behavior of each of these households during the 12-month period from July 1997 through June 1998.[1] Specifically, we are interested in the number of unique sites searched by each household within a given product category. A key issue when examining search behavior is defining the period of time that constitutes a search session. For example, the industry practice is to define a user session by "closing" the session after 15 minutes of inactivity. However, that is a fairly narrow definition of a search session, as it is conceivable that a shopper's decision period, and therefore search, may span multiple days, weeks, etc. In fact, search on the Internet can be better characterized as a series of store visits over a span of days, which "builds up" to a purchase (Moe and Fader 2004). Conceptually, these shoppers are accumulating information toward a purchasing threshold before the final purchasing decision is made, perhaps days later, during an entirely different store visit. While similar behavior has been documented offline (Putsis and Srinivasan 1994), this behavior is much more common in the online environment due to the lower costs of "traveling" to the store. One may visit a given site today to examine the product offering, then deliberate the purchase offline, and return to shopping online days later.

An alternative method of defining a search session is to close the session after a purchase is made and assume that all store visits prior to that purchase contributed to that purchase cycle. There are three problems with this method as well. First, not all purchases are made online. A shopper may search

[1] We used the information contained in the URL of the site to determine the product category. The end of this time period corresponds to the entry of one of the online stores, Amazon, into multiple categories, making determination of the product class ambiguous.

Table 1 E-Commerce Sites Included in Books, CDs, and Travel Categories (with Number of Searches per Site)

Books		CDs		Air travel	
Amazon	4,721	MusicBlvd	1,880	City.Net	3,010
Barnes & Noble	2,813	CDNOW	1,655	Preview Travel	2,095
Books.com	230	BestBuy	498	Travelocity	2,048
Superlibrary	229	CDUniverse	442	Expedia	1,988
Borders.com	127	Music Central	271	American Airlines	1,334
Book Zone	67	Tower	227	Southwest Airlines	997
Powells	61	Tunes	189	Delta Air	957
AltBookStore	59	CDWorld	107	ITN	920
BooksNow	30	MassMusic	94	Continental	872
Wordsworth.com	28	Newbury	70	Travel Web	817
Acses	17	Emusic	31	Northwest Airlines	723
Books-a-Million	17	Ktel	20	United Airlines	714
Kingbooks	4	CDConnect	16	US Airways	643
		Music Spot	14	Priceline	512
		CDEurope	13	TWA	418
		CDUSA	11	Best Fares	401
				The Trip	366
				European Travel	279
				Lowest Fare	138
				Cheap Tickets	115
				Alaska Airlines	94
				Travel Zoo	70

Note. The numbers in this table represent the actual number of visits performed by the individuals in the data panel and do not represent the total level of activity at the U.S. population level.

online and make some purchases offline. In that case, search sessions as defined strictly by observed online purchases might in fact include multiple purchase cycles. Second, clickstream data typically includes only the URLs viewed and does not necessarily offer any insight into purchasing activity. While some sites, such as Amazon and CDNOW, may have easily identifiable URLs that indicate a purchase occurred, many sites do not offer any such indication. Third, after a shopper views the product selection at one or more sites, he/she may simply decide not to buy at all for whatever reason, and the next visit may be related to a completely unrelated purchasing decision. In these instances, store visits cannot be linked to any specific future visit that contains a purchasing transaction.

Our objective is to define a session broadly enough to encompass a series of site visits that contribute to the same purchasing cycle, but also narrowly enough not to inadvertently merge multiple cycles. As a result, we chose to examine each household's search behavior at the monthly level, which allows us to avoid making any assumptions (or incorrect inferences) about the link between purchasing and online visit behavior.[2]

We focused on three categories: CDs, books, and air travel, choosing them because: (1) they were relatively frequently purchased online during this time period; (2) they tend to be nondifferentiated goods (a book purchased from Amazon is the same book purchased from a low-cost provider such as Books-a-Million.com), increasing the probability that these categories will be subject to broader search; (3) they vary in price from relatively inexpensive (books and CDs) to more expensive (air travel); and (4) they have been used in similar studies by other researchers. Sites were chosen from each category according to lists of leading online retailers from Media Metrix, BizRate (www.bizrate.com), and Netscape's "What's Related" feature, which uses Alexa's records of consumers' actual surfing behavior to identify related sites. While there may be some sites that are excluded from the analysis, they would probably constitute a very small fraction of the activity in each category. Specifically, our dataset covers consumer search activity across 13 book sites, 16 music sites, and 22 travel sites (Table 1), a more inclusive set of sites within each category than that used by either Brynjolfsson and Smith (2000) or Clemons et al. (2002). Sites included in each category range considerably in terms of visitor traffic (Table 1 also provides the number of searches

[2] We examined the number of these month-long sessions that contained more than one purchasing transaction to confirm that our definition for "session" was not so long as to include multiple purchasing cycles. Because our data only contained purchasing information for a limited number of retailers, we obtained this information for the most popular site in each category for which we had purchasing information. The results revealed that in each case, fewer than 1% of all month-long sessions contained more than one purchasing transaction.

experienced by each site). For example, the smallest CD site in our dataset attracted only 11 unique visitors from the Media Metrix panel, while the largest in the same category attracted over 1,800 unique visitors.

3. Analysis, Models, and Results

Using these data, we first examine the number of stores visited in a typical month. Figure 1 shows the average number of websites visited by households each month in which they were actively shopping in the product category. For example, the average number of CD sites searched in a household's first month of shopping is 1.23 and increases to 1.62 in the fourth shopping month for those households that shopped in four (or more) different months in the dataset.

Two patterns are striking in the data. First, the overall level of search is low, initially ranging from 1.1 stores for books to 1.8 for travel. In fact, 70% of the CD shoppers, 70% of the book shoppers, and 42% of the travel shoppers were observed as being loyal to just one site throughout the duration of our data. Second, there seems to be an increase in search from month to month. This seems to suggest that Internet search, while currently fairly low, may be increasing over time. One possible explanation for the increasing trend seen in Figure 1 is that consumers may be evolving and searching more as they gain experience, consistent with the idea that time will lead to lower prices and reduced dispersion. We will more closely examine this dynamic of search in the next section.

However, aggregate patterns like those in Figure 1 can be misleading, because they can, in principle, mask different underlying trends that may exist at the individual level. To model these trends, and to provide a more accurate portrait of shopping behavior, we examine search using a model that allows us to decompose this data into three components: (1) Depth of search: the decision of the household to visit more than one store in a given month, (2) dynamics of search: the evolution of the number of stores visited over time, and (3) activity of search: the overall amount of category-level shopping activity for each household in a given product class.

Figure 1 **Average Number of Online Stores Visited During Each Observed Shopping Month**

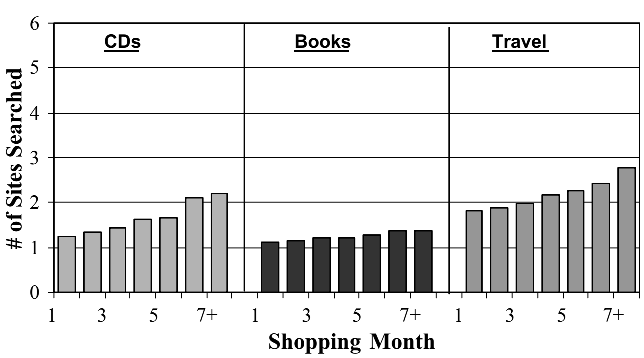

3.1. Depth of Search

Consumer search has been theoretically modeled as a process in which the consumer's decision to seek out additional information is a function of the expected benefit of that added information (Diamond 1987). As a consumer obtains more information from visiting additional store sites, the expected benefit of seeking new information decreases, resulting in a lower probability of soliciting information from an additional source. To capture this process, we model the probability that individual i searches an xth site as a decrement of the probability of visiting the $(x_i - 1)$st site:

$$\Pr[X_i = x_i] = \frac{(x_i - 1)\theta_i}{x_i} \Pr[X_i = x_i - 1], \quad x_i = 2, 3, \ldots, \tag{1}$$

where θ_i is an individual-specific search propensity parameter $(0 < \theta_i < 1)$. Lower values of this parameter indicate an increasing likelihood of stopping with a small search set, but even at its maximum value $(\theta_i \to 1)$, the probability of searching an additional site will always be a decreasing function of x_i.

The model presented in Equation (1) is well-suited for the type of search behavior described above. Working backwards through this recursive relationship, we can obtain the logarithmic distribution (Johnson et al. 1993)

$$\Pr[X_i = x_i] = \frac{a_i \theta_i^{x_i}}{x_i}, \qquad x_i = 1, 2, \ldots, \tag{2}$$

where $a_i = -[\ln(1 - \theta_i)]^{-1}$ and $0 < \theta_i < 1$. To illustrate the process, Figure 2 plots the shape of the logarithmic distribution for a variety of θs.

Figure 2 **Logarithmic Distribution**

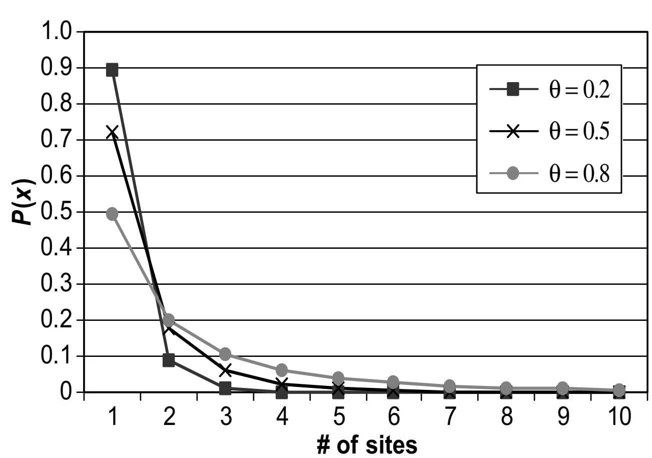

3.2. Dynamics of Search

The logarithmic distribution given in Equations (1) and (2) models the number of sites searched by an individual in a given session. However, in our data, we observe each individual across multiple search sessions. This allows us to more closely examine potential dynamics that may exist in search propensity as shoppers gain more experience in a particular category.

Recall that the data show that search propensity is surprisingly low given the ease of visiting multiple stores online. One explanation is that people are still learning to search, and over time, as consumers familiarize themselves with the Internet environment, search propensity will increase. This argument would be consistent with the apparent trends seen in Figure 1. We explicitly test this argument by including a model component that captures and measures the change in search propensity from session to session.

Specifically, we allow the search propensity parameter, θ_i, to vary from visit to visit for each panelist. In other words, the probability that individual i will visit x_{ij} sites in the jth month is a function of a search propensity parameter, θ_{ij}:

$$\Pr[X_{ij} = x_{ij}] = \frac{a_{ij}\theta_{ij}^{x_{ij}}}{x_{ij}}, \qquad x_{ij} = 1, 2, \ldots. \tag{3}$$

Because θ_{ij} is constrained to be between zero and one, we apply a logistic transformation $\theta_{ij} = \exp\{\theta^*_{ij}\}/[1+\exp\{\theta^*_{ij}\}]$ and assume that $\theta^*_{ij} \sim N(\mu_j, \phi)$. We operationalize changes in individual search propensity over time through the mean of the θ^*_{ij} distribution

$$\mu_j = \beta_0 + \beta_1 \ln(j). \tag{4}$$

The parameter β_0 represents the baseline propensity to search, while the parameter β_1 indicates how search propensity changes as a function of experience. If β_1 is positive, then search increases with experience. However, if β_1 is negative, then search decreases as a function of experience. To accommodate heterogeneity, we assume that β_0 and β_1 are normally distributed, with mean and standard deviation to be estimated.

3.3. Activity of Search

An alternative explanation can also generate the patterns shown in Figure 1 and offers a very different interpretation of these data. Specifically, the pattern that we observe in Figure 1 could stem from a selection effect. Because we are able to observe more activity from the frequent shoppers in our dataset, we observe an increasingly greater proportion of these more-active households as we move from left to right in the figure. These relatively active shoppers (i.e., those who visit a particular category more frequently) may be inherently different from those less active and may tend to search across more sites. As a result, the increase we see may not represent more stores being visited over time by a typical panelist, but rather a change in the mix of shoppers as we move to a greater number of active months. Simply put, heavier users may be represented more in later months. There may be no household-level dynamics whatsoever, but only an apparent pattern resulting from heterogeneity alone.

To explore the relationship between the level of shopping activity levels and depth of search, we extend Equation (4) to include a second individual-specific covariate—a measure of category-level search activity, q_i:

$$\mu_{ij} = \beta_0 + \beta_1 \ln(j) + \beta_2 \ln(q_i). \tag{5}$$

Again, the coefficient β_2 is assumed to be normally distributed with mean and variance to be estimated.

One measure we could use to characterize a household's activity level is the proportion of months for which the household was actively shopping. However, there is a critical drawback of using such a measure—not all households are present in the panel for equal amounts of time. For example, a household that was in the panel for only a month and was also active in that month would be represented by this measure as being active 100% of the time (or every month), even if this individual's true activity level is significantly less frequent (e.g., every six months). However, because we are only able to observe a small sample of some shopper's lifetime behavior, any measure based on this limited history is likely to misrepresent that individual's latent and true tendencies in behavior. Therefore, instead of directly measuring shopping activity from the data, a better measure would incorporate a Bayesian shrinkage estimate that allows for some regression to the mean, derived from a separate category-level model of activity.[3]

We examine the number of months each individual is actively shopping as a zero-truncated binomial process (because only those who have conducted at least one search are included in our dataset). Each individual household in the panel has a probability, q_i, of visiting a product category in any given month they are in the panel (regardless of which store(s) they choose to visit). Given that household i was in our dataset

[3] We estimated the search depth models using both a Bayesian estimate of search activity (q_i) as a covariate, as well as using a direct measure of search activity (# of months active/# of months in panel) for each of the three product categories. We found that, in each case, the model using the Bayesian estimate of search activity fits significantly better than the model with a direct measure of activity.

(but not necessarily active in a given category) for T_i months, the probability of shopping in J_i of those T_i months is

$$P_i(J_i \mid T_i, q_i) = \frac{\binom{T_i}{J_i} q_i^{J_i}(1-q_i)^{T_i-J_i}}{1-(1-q_i)^{T_i}}. \quad (6)$$

Furthermore, we assume q_i to be distributed across the population according to a beta distribution with parameters k and m to allow for heterogeneity in these activity levels. This mixture model will provide a general measure and method of assessing each individual's shopping activity. We use the expectation of q_i conditional on each individual's observed behavior. This completes the description of our complete model specification.

We estimate this hierarchical Bayesian model using winBUGS, which uses Markov chain Monte Carlo methods to obtain a distribution for each model parameter. To estimate the model, we followed standard practice, simulating multiple chains of 10,000 iterations each, and discarding the first 5,000 as an initial burn-in. To confirm the adequacy of convergence, we monitored each chain and plotted their traces. All plots show convergence, and the Gelman-Rubin convergence statistic (Gelman and Rubin 1992) also demonstrates evidence of sufficient convergence. Appropriate and diffuse priors and hyperpriors were specified.[4] Specifically, the hyperpriors in the search depth model were specified as

$$\mu_\beta \sim \text{normal}(0, 10) \quad \text{and} \quad \sigma_\beta \sim \text{gamma}(1, 1).$$

The priors in the search activity model were specified as

$$k \sim \text{gamma}(1, 1) \quad \text{and} \quad m \sim \text{gamma}(1, 1).$$

We also estimated a model that allowed for correlation between depth of search and activity of search. We did this by specifying the coefficients in Equation (5) and the parameters of the activity model in Equation (6) as coming from a multivariate normal distribution with a correlated error structure

$$\begin{bmatrix} \beta_{0i} \\ \beta_{1i} \\ \beta_{2i} \\ q_i \end{bmatrix} \sim \text{MVN}(\mu, \Sigma), \quad (7)$$

where μ is normally distributed and Σ is inverse Wishart.

[4] We estimated the model using a number of different priors, some more diffuse and some less diffuse than those specified in this paper, and find that the results are not sensitive to prior specifications.

In two of the three product categories we tested (books and music), the correlation between search activity, q_i, and the baseline propensity to search, β_0, was not significantly different from zero. In the air travel category, the mean correlation was significant; however, the overall fit of the correlated model was poorer in comparison to the uncorrelated model in terms of the Bayesian information criterion (198,386 versus 197,980). Therefore, for the remainder of the paper, we will discuss the results from the uncorrelated model with the depth of search component specified by Equation (5) and the activity of search component specified by Equation (6).

3.4. Results

The resulting parameters estimated are presented in Table 2, along with the 90% and 95% confidence bounds calculated from the MCMC simulations. For all three product categories, more-active households

Table 2a Parameter Estimates for Book Category

Param.	Mean	Std. dev.	2.50%	5.00%	Median	95.00%	97.50%
k	1.067	0.062	0.948	0.959	1.069	1.160	1.173
m	5.007	0.279	4.468	4.535	5.013	5.438	5.509
β_0	−0.078	0.071	−0.197	−0.189	−0.067	0.020	0.036
β_1	0.059	0.049	−0.032	−0.017	0.056	0.146	0.163
β_2	0.857	0.042	0.772	0.782	0.853	0.917	0.925
ϕ	12.760	2.295	8.745	9.294	12.610	16.620	17.140
$\sigma_{\beta 0}$	8.481	1.384	6.247	6.564	8.261	10.900	11.450
$\sigma_{\beta 1}$	7.738	1.796	4.753	5.045	7.561	11.100	11.650
$\sigma_{\beta 2}$	12.370	1.833	8.757	9.117	12.420	15.470	15.860

Table 2b Parameter Estimates for Music Category

Param.	Mean	Std. dev.	2.50%	5.00%	Median	95.00%	97.50%
k	0.786	0.050	0.693	0.707	0.783	0.871	0.889
m	5.706	0.366	5.046	5.136	5.695	6.326	6.476
β_0	0.880	0.107	0.691	0.716	0.881	1.041	1.065
β_1	−0.007	0.065	−0.143	−0.117	−0.004	0.093	0.109
β_2	0.839	0.062	0.739	0.748	0.838	0.936	0.949
ϕ	9.063	2.494	5.026	5.658	8.564	13.830	14.880
$\sigma_{\beta 0}$	6.408	1.677	3.735	4.024	6.126	9.601	10.140
$\sigma_{\beta 1}$	5.585	1.583	3.174	3.483	5.305	8.320	9.426
$\sigma_{\beta 2}$	13.110	2.626	9.057	9.344	12.900	18.030	19.170

Table 2c Parameter Estimates for Travel Category

Param.	Mean	Std. dev.	2.50%	5.00%	Median	95.00%	97.50%
k	1.029	0.042	0.951	0.963	1.027	1.098	1.111
m	3.444	0.138	3.193	3.226	3.438	3.680	3.724
β_0	1.378	0.065	1.258	1.269	1.378	1.479	1.514
β_1	−0.148	0.032	−0.214	−0.199	−0.148	−0.093	−0.085
β_2	0.697	0.044	0.629	0.634	0.693	0.772	0.791
ϕ	14.740	2.085	11.020	11.330	15.090	17.910	18.460
$\sigma_{\beta 0}$	8.053	1.280	5.824	6.164	7.939	10.360	11.040
$\sigma_{\beta 1}$	8.749	1.286	6.578	6.874	8.560	11.070	11.440
$\sigma_{\beta 2}$	11.120	1.892	8.020	8.404	10.860	14.350	15.450

seem to search more store sites than the less active households (i.e., β_2 is significantly positive). However, evidence of any month-to-month dynamics at the household level is limited. We first address the issue of time dynamics before discussing the relationship between search depth and search activity.

For two of the three product categories we examined (books and music), we find no significant effect of time dynamics, β_1, on propensity to search ($p > 0.10$). For air travel, however, we do find significant, but negative, time dynamics ($E[\beta_1] = -0.148$). In other words, as we observe subsequent shopping sessions for an individual, the mean propensity to search actually declines. This is contrary to arguments suggesting that consumers are still learning how to search online, and the currently low levels of search will increase as they gain experience and search becomes less difficult. Instead, our results indicate that the effect on search propensity from increased consumer experience is reversed. We find that the level of search being conducted now, however low it may be, is only decreasing as consumers gain experience in that product category (as indicated by $\beta_1 < 0$). Rather than learning how to search, consumers seem to be gravitating toward a preferred site over time, at least in the case of air travel.

In contrast, the effect of category-level search activity, β_2, is significant and positive across all three product categories ($p < 0.05$). This strongly suggests that the downward trend seen in Figure 1 is the result of a selection effect where the more frequent searchers tend to search more sites in any given active month. Figure 3 plots the relationship between search activity and the depth of search. In the upper part of the figure, the solid line indicates the median expected number of sites searched at each activity level, q, and the dotted lines indicate the 90% confidence bounds. The lower portion of the figure is a histogram representing the distribution of panelists based on their category-level search activity. Overall, we see very low levels of search that increase for more-active shoppers, but these more-active shoppers are less abundant in the sample. Taken together, the two portions of Figure 3 demonstrate how the selection effects arise from the raw data.

In general, there seems to be very little search in the book and CD categories. The travel category reflects a bit more search, but still far less than one might expect. The fact that search in the travel category is actually more than the amount of search seen in the books or CD categories may seem logical given that the size of purchase is large, and prices are dynamic. But on the other hand, because many travel sites feature comparison-shopping tools within the site itself, it may seem a bit surprising that individuals search more across travel sites than they do in the books and CD categories. These latter categories consist of sites that are pure e-commerce vendors and not price-search agents, or "bots." This observation provides an initial indication that search 'bots are not the reason behind the limited search that we find in these categories; we discuss this issue in more detail later in the paper.

Figure 3 **Effect of Search Activity: Estimated Number of Store Visits (Top) and Number of Panelists as a Function of q, Propensity to Search**

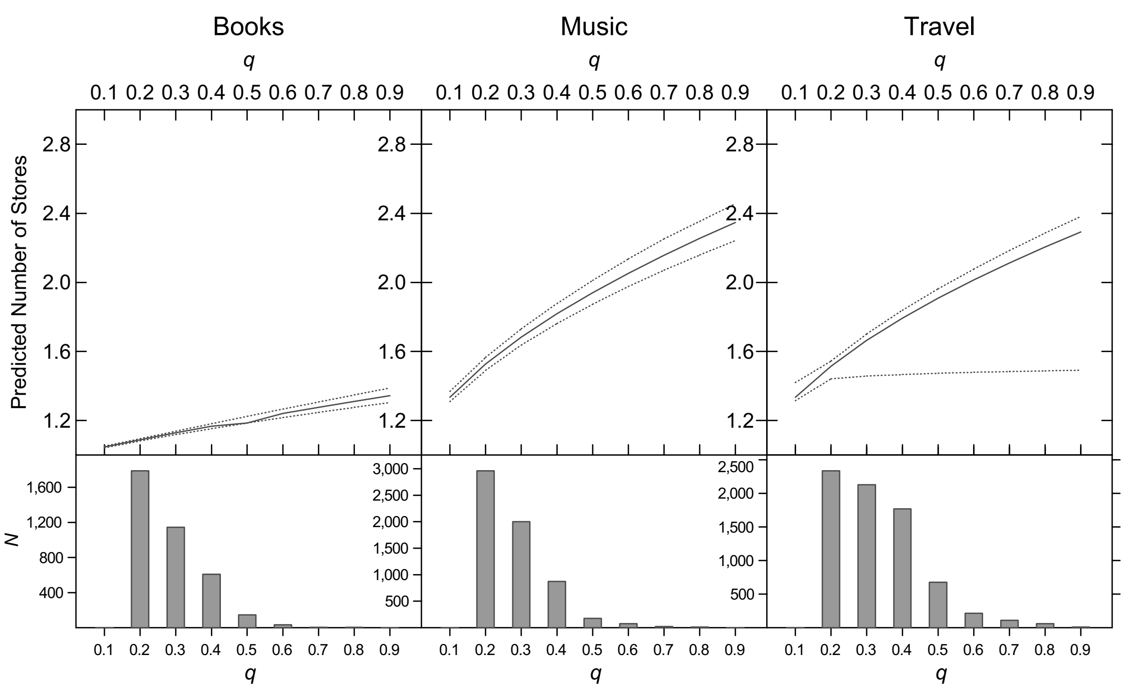

4. Discussion

4.1. Summary and Implications for Research

In summary, these three categories show fairly low levels of search overall. While more-active shoppers tend to visit more sites in any given month, there is no evidence that experience increases the number of sites visited. We might expect the greatest returns to search for travel services, both because prices can change over time and because this is a more expensive purchase. However, we find that experience leads to a slight decrease in the number of visited sites.

Our results suggest that people visit few stores online despite the fact that consumers are "just a mouse click away" from other stores. Browsing behavior varied by product category and level of activity, but showed no increase with experience.

These findings may provide an important explanatory mechanism for Internet researchers who have been puzzled by the existence of substantial price dispersion. In understanding our results, it is important to realize that we examined product classes and time periods that are very similar to those used by researchers who have examined changes in price and price dispersion on the Web. The attraction of books, CDs, and travel for examining these questions is their apparent commodity-like status, which should provoke higher levels of search. Also, our data have been collected during time frames that substantially overlap the periods used in prior studies of prices for CDs and books (Brynjolfsson and Smith 2000) and airline travel (Clemons et al. 2002). This minimizes the potential for historical differences and suggests that the lack of search is consistent with the observed price dispersion. The methods we develop here may be useful in future work that attempts to connect search with pricing in a more direct manner, if and when appropriate data become available.

Our results have implications concerning the use of self-reports of search, because they stand in such sharp contrast to survey-based measures of online search. Self-reported data are subject to the fallibility of people's memories, idiosyncratic scale use, and even deliberate alteration through social desirability biases, and have been found to have very little correlation with actual search. Industry analyses based on self-reports suggest higher levels of search than we observe (e.g., a reported 2.4 visits per purchase noted by McQuivey 1999). We believe additional research is needed to compare self-reported behavior to actual Web usage, but the type of panel data examined here reflects actual behavior and should serve as the appropriate baseline for such efforts. It is interesting to note that the one reported example of decreases in prices and dispersion (Brown and Goolsbee 2001) employs self-reports based on Forrester's research. Brown and Goolsbee suggest that recent declines in life insurance prices are the result of increased use of Internet shopping sites for insurance. An interesting next step would be to examine actual (as opposed to self-reported) usage data to see if observed search is correlated with lower prices paid. Similarly, it would be interesting to examine if observed search helps explain which shoppers will be most helped by buying online (Scott Morton et al. 2003).

4.2. Caveats and Limitations

4.2.1. Data. In this study, the analysis of depth of search and activity level of search was conducted at the household level, rather than at the level of the individual consumer. By aggregating consumers to the household level, we overstate both the activity level and depth of search. Households with multiple members shopping in different months would have an activity level that is higher than that of any one individual in the household. Similarly, households with multiple members who are each loyal to a different site will demonstrate greater search propensities at the aggregate household levels. Therefore, even though we observe low levels of search at the household level, search at the individual consumer level could be even lower.

A potential disadvantage of the type of clickstream data used here is that we identify all browsing activities, not just search behavior associated with purchases. This is a limitation of the data, but provides an initial view of online behavior in this very new and developing area of research. While such browsing may be an important part of consumer search, we are unable to partition our observations into those associated with purchase-oriented behavior and those that may be associated with other activities (e.g., using Amazon.com to find biographical information about a certain author). While generally this may lead to an overestimation of the amount of search in our dataset, it is possible that purchase-related browsing might feature more extensive search than other "look up" activities.

Finally, the time period studied in this paper represents a fairly early stage in e-commerce. As such, the subset of the consumer population shopping online may be very different from the consumer population as a whole. Studies conducted around the same period of time have found that online consumers tend to be more time constrained than the average consumer (Bellman et al. 1999). As such, this subset of the consumer population that shops online may consist largely of those consumers who are unable to spare the time to search across multiple sites. Projecting the search behavior of this subset of shoppers to the consumer population as a whole may lead to biased conclusions.

4.2.2. The Role of Shopbots. While the dataset used here contains a rich and realistic portrayal of online shopping, it is an early snapshot of such behavior. While it could be argued that the low levels of search documented here will change as the online market matures, our analysis of time dynamics does not suggest this was happening in the time frame observed. All three classes showed no increase in the number of sites visited. One major trend that may well modify these results is the widespread adoption of price search agents, or price robots, termed 'bots for short. We examined our data for the use of such 'bots, but found that their use at this time was at very low levels. For example, the most popular was Acses, a price search agent for books (which later evolved into today's DealTime) that had a total of 17 users in the Media Metrix panel. As consumers may become more sophisticated in their search over time and as search agent technology evolves, we may see automated search further lowering search costs and minimizing prices and price dispersion. However, recent research indicates that even those who use shopping robots seem to display loyalty to sites previously used. Brynjolfsson and Smith (2001) examined users of EvenBetter.com, a popular search engine for books (also an early version of DealTime). It is interesting to note that even though this subset of Internet consumers was highly price sensitive and tended to patronize those stores with cheaper prices, over 51% of these customers did not choose the retailer with the cheapest price. Even with this self-selected group of consumers, the retailer brand had a significant advantage. Amazon, for example, commanded a price premium of $1.85 over unbranded retailers, according to their model.

4.3. Why Is Search So Limited?

One possible explanation for why we see so little search is that the current market for these goods is efficient. We chose these products, however, because prior evidence from contemporary studies suggest that this is not the case. Both Bailey (1998) and Brynjolfsson and Smith (2000) find significant dispersion for CDs and books. Prices for identical items differed, on average, by 33% for books and 25% for CDs. Clay et al. (2001) find significant dispersion and differentiated retail strategies in the online book industry, with the largest retailer, Amazon, presenting a price premium 10%–25% over the lowest-cost provider. Similarly, Clemons et al. (2002) find that dispersion exists even when they control statistically for the quality of the airline tickets in their study. Thus, the dispersion found in these studies suggests that these markets have not yet produced the kind of highly efficient markets thought to be a natural hallmark of e-commerce.

Another possibility is that normative models of search are not complete. Clearly, while any modifications of standard search models exceed the scope of this paper, it is worth noting that the heart of most search models is the trade-off between the cost of search, usually measured by time, and the benefit of that search to the consumer. Other domains involving the allocation of time have benefited from richer descriptive frameworks incorporating the effects of context and mental accounts. Two examples that seem relevant are descriptive theories of time-money trade-offs (Loewenstein and Prelec 1992, Loewenstein and Thaler 1997) and the observation that out-of-pocket costs are overweighted relative to opportunity costs (Thaler 1999).

A third possibility is the realization that search costs are not constant over time and that they change as consumers gain experience shopping with a particular online store. For example, by visiting a site, one learns its navigational scheme, which reduces the cost of using that site in the future. Similarly, the site can make changes, through customization, user-based recommendations, and memorization of names, addresses, and payment details that lower the cost of that site relative to others. This idea that search costs are dynamic is analogous to the concept of lock-in, which has been discussed in markets for technology goods, and is a topic we are exploring elsewhere (Johnson et al. 2003).

4.4. Managerial Implications

From a managerial perspective, this paper suggests that, despite claims to the contrary, the Internet does not produce enormous amounts of search, even for products that are difficult to differentiate. However, this conclusion, while a fair characterization of the aggregate pattern of search, does depend upon the amount of category-level activity: More-active households search more. From the perspective of the firm, these results can inform market selection and help identify customer segments and tactics: Infrequent customers who tend to search less may not need as much in the way of incentives for loyalty. At the same time, the most active, and perhaps valuable, customers are those most likely to shop around.

This research also reinforces the value of looking at household, as opposed to aggregate, data. The aggregate pattern, as portrayed in Figure 1, seems to support an explanation due to dynamic growth. However, careful modeling of the household-level data indicates that this is not occurring, and that the pattern results from differential activity. In fact, for one product class, air travel, we actually see a slight decrease in search over time.

The bottom portion of Figure 3 makes this quite clear in terms of potential marketing segmentation and strategy: The very active shoppers are few in numbers in all categories, and shopping around seems limited to a few high-activity (high q) shoppers. At the same time, these may be customers of high value, so paying attention to their shopping habits and developing tools to increase their loyalty may be prudent.

Beyond the specifics of these results, the methodology we introduced can serve as a continuing tool for e-tailers interested in the dynamic development of their customer base. Thus, we think that a tracking study producing a time-varying equivalent of Figure 3 would help managers identify changes in shopping activity and loyalty. In addition to this type of longitudinal extension to our work, it might also be worthwhile to profile households by looking at their search behavior across multiple categories. This might help managers anticipate and identify the kinds of behavioral patterns described above.

4.5. Conclusion

We see this research as an initial demonstration that data from the Web are changing our view of search from an unobservable explanatory variable in the analysis of markets to one that can and should be observed and included in empirical analyses of market behavior. The kind of analysis we do here cannot easily be done in offline environments, and the kind of data used here can track search at a fine level. Much of the challenge in the evolving analysis of clickstream data will be to develop efficient methods for theory testing, given the large amount of data provided by shoppers on the Internet. While these challenges are significant, we believe the analysis of search data will be extraordinarily useful in expanding our understanding of the role of search in explaining marketplace behavior, and the analysis offered here represents a first step.

Acknowledgments

The authors thank ComScore Media Metrix, Inc., for providing the data used here, and the supporting firms of the Wharton Forum on Electronic Commerce for their financial support.

References

Adamic, Lada A., Bernardo A. Huberman. 2001. The Web's hidden order. *Comm. ACM* **44**(9) 55–59.

Bailey, Joseph. 1998. Electronic commerce: Prices and consumer issues for three products: Books, compact discs, and software. *Organ. Econom. Co-Operation Development* **98**(4).

Beatty, Sharon E., Scott M. Smith. 1987. External search effort: An investigation across several product categories. *J. Consumer Res.* **14**(1) 83–95.

Bellman, Steven, Gerald L. Lohse, Eric J. Johnson. 1999. Predictors of online buying behavior. *Comm. ACM* **42** 32–38.

Brown, Jeffrey R., Austan Goolsbee. 2001. Does the Internet make markets more competitive? Evidence from the life insurance industry. NBER working paper, Cambridge, MA.

Brynjolfsson, E., M. D. Smith. 2000. Frictionless commerce? A comparison of Internet and conventional retailers. *Management Sci.* **46**(4) 563–585.

Bucklin, Randolf E., Catarina Sismeiro. 2003. A model of Web site browsing behavior estimated on clickstream data. *J. Marketing Res.* **40**(3) 249–267.

Clay, Karen, Ramayya Krishnan, Eric Wolff. 2001. Prices and price dispersion on the Web: Evidence from the online book industry. *J. Indust. Econom.* **49** 521–539.

Clemons, Eric, Il-Horn Hann, Loren Hitt. 2002. Price dispersion and differentiation in online travel: An empirical investigation. *Management Sci.* **48** 534–549.

Diamond, Peter. A. 1989. Search theory. John Eatwell, Murray Milgate, Peter Newman, eds. *The New Palgrave: Allocation, Information, and Markets*. Norton, New York, 271–286.

Gelman, A., D. B. Rubin. 1992. Inference from iterative simulation using multiple sequences. *Statist. Sci.* **7** 457–511.

Johnson, Eric J., Steven Bellman, Gerald L. Lohse. 2003. Cognitive lock-in and the power law of practice. *J. Marketing* **67**(2) 62–75.

Johnson, Normal L., Samuel Kotz, Adrienne W. Kemp. 1993. *Univariate Discrete Distributions*, 2nd ed. John Wiley and Sons, New York.

Loewenstein, George, Drazen Prelec. 1992. Anomalies in interpersonal choice: Evidence and an interpretation. Jon Elster, George Loewenstein, eds. *Choice Over Time*. Russell Sage Foundation, New York.

Loewenstein, George, Richard H. Thaler. 1997. Intertemporal choice. William Mark Hogarth, Robin M. Goldstein, eds. *Research on Judgment and Decision Making: Currents, Connections, and Controversies. Cambridge Series on Judgment and Decision Making*. Cambridge University Press, New York.

McQuivey, James. 1999. Web buyers shop around. Forrester brief. Forrester Research, www.forrester.com.

Moe, Wendy W. 2003. Buying, searching, or browsing: Differentiating between online shoppers using in-store navigational clickstream. *J. Consumer Psych.* **13**(1, 2) 29–40.

Moe, Wendy W., Peter F. Fader. 2004. Dynamic conversion behavior at e-commerce sites. *Management Sci.* **50**(3) 326–335.

Moorthy, Sridhar, Brian T. Ratchford, Debabrata Talukdar. 1997. Consumer information search revisited: Theory and empirical analysis. *J. Consumer Res.* **23**(4) 263–277.

Newman, J. R. 1977. Consumer external search: Amount and determinants. Arch Woodside, Jagdish Sheth, Peter Bennet, eds. *Consumer and Industrial Buying Behavior*. North-Holland, New York.

Park, Young-Hoon, Peter S. Fader. 2003. Modeling browsing behavior at multiple websites. *Marketing Sci.* Forthcoming.

Putsis, W. P., N. Srinivasan. 1994. Buying or just browsing—The duration of purchase deliberation. *J. Marketing Res.* **31**(3) 393–402.

Scott Morton, Fiona, Florian Zettelmeyer, Jorge Silva-Risso. 2003. Consumer information and price discrimination: Does the Internet affect the pricing of new cars to women and minorities? *Quant. Marketing Econom.* **1**(1) 65–92.

Smith, M. D., E. Brynjolfsson. 2001. Consumer decision-making at an Internet shopbot: Brand still matters. *J. Indust. Econom.* **49**(4) 541–558.

Thaler, Richard H. 1999. Mental accounting matters. *J. Behavioral Decision Making* **12**(3) 183–206.

Wilkie, William, Peter Dickson. 1985. Shipping for appliances—Consumers' strategies and patterns of information search. Marketing Science Institute, Cambridge, MA.