

Cognitive Lock-In and the Power Law of Practice

The authors suggest that learning is an important factor in electronic environments and that efficiency resulting from learning can be modeled with the power law of practice. They show that most Web sites can be characterized by decreasing visit times and that generally those sites with the fastest learning curves show the highest rates of purchasing.

The widespread use of information technology by buyers and sellers is thought to increase competition by lowering search costs. "The competition is only a click away" is a common phrase in the popular press and an oft-cited reason for the failure of Internet ventures to achieve profitability. A potential result of reduced search costs is a decrease in brand loyalty and an increase in price sensitivity. At the extreme, there is the fear of a price-cutting spiral that drives out profits—labeled in the popular press as "perfect competition" or "frictionless capitalism," but more correctly called Bertrand competition (Bakos 1997; for a discussion, see Brynjolfsson and Smith 2000; for dissenting reviews, see Alba et al. 1997; Lal and Sarvary 1999; Lynch and Ariely 2000).

As a result, there has been interest in how to retain customers in electronic environments. The most commonly discussed solution is creating loyalty to Web sites, so research attempts to identify which sites exhibit greater loyalty or "stickiness" and speculates about what causes repeat visits. The most common loyalty metric is the frequency and cumulative duration of visits. For example, eBay is listed on the *New York Times* top ten stickiest sites because, though it has relatively few users, visitors spend approximately 90 minutes a month there, according to Web rating services such as Media Metrix. Consequently, eBay is thought to be highly successful. Other loyalty metrics relate visiting loyalty and purchasing loyalty, such as the number of visits per purchase, which is termed the "browse to buy" or "look to book" ratio (Schonberg et al. 2000), to the percentage of customers who become repeat customers (Win 2001).

In this article, we describe a mechanism and model for understanding the development of loyalty in electronic environments and an accompanying metric based on an empiri-

cal generalization from cognitive science, the power law of practice (Newell and Rosenbloom 1981). For an intuitive understanding of the mechanism, imagine a user visiting a Web site to purchase a compact disc (CD). This user must first learn how to use the Web site to accomplish this goal. We believe that after the CD has been purchased, having learned to use this site raises its attractiveness relative to competing sites for the consumer, and all other things being equal (e.g., fulfillment), the site will be more likely to be used in the future than a competitor. Further use reinforces this difference because practice makes the first site more efficient to use and increases the difference in effort between using any other site and simply returning to the first site, where browsing and buying can be executed at the fastest rate. This reinforcement generates an increasing advantage for the initial site. Sites can actively encourage this learning by implementing a navigation scheme that can be rapidly apprehended by visitors and using various forms of customization, including personalization, recommendations, or easy checkout. Learning how to navigate a site and customization together can increase the relative attractiveness of the site, generating a type of "cognitive loyalty program" that adds another, more cognitive explanation of loyalty to the existing rich set of definitions (Oliver 1999).

Two analogies may reinforce this idea and suggest that our analysis of learning is applicable to nonelectronic environments as well. On a first visit to a new supermarket, some learning takes place. The aisle location of some favorite product classes, the shelf location of some favorite brands, and a preferred shopping pattern through the store may be acquired (Kahn and McAlister 1997). This knowledge of the layout of a physical store, which increases with subsequent visits, makes the store more attractive relative to the competition. We argue that the same process happens with virtual stores. A similar argument has been commonly made about learning software such as word processors. Experience with one system raises the cost of switching to another, which explains, for example, the slow conversion from WordPerfect to Word (Shapiro and Varian 1999).

In this article, we examine learning in electronic environments by studying the time spent visiting individual Web sites. We focus on the cognitive costs of using a site and how they decrease with experience. We argue that this decrease can be modeled with a simple functional form that is used

Eric J. Johnson is Professor of Marketing, Columbia Business School, Columbia University. Steven Bellman is Senior Lecturer, Graduate School of Management, University of Western Australia. Gerald L. Lohse is with the Customer Decision Analytics Practice at Accenture. The authors thank Media Metrix, Inc. for providing the data used here and the supporting firms of the Wharton Forum on Electronic Commerce for their financial support. This research has benefited from the comments of participants in seminars at Ohio State University, Columbia University, University of Rochester, and the University of Texas and helpful comments by Peter Fader, Asim Ansari, and John Zhang.

often in cognitive psychology to study learning—the power law of practice. We then investigate the relationship between the phenomenon of decreasing visit times and repeat visit loyalty and online purchasing using data from a panel of consumers from the World Wide Web.

The article proceeds as follows: We first review the literature that describes learning as a power law function, and discuss its underlying causes. We then discuss why this type of learning might apply to use of the Web. Using panel data that capture the in situ Web surfing of a large consumer panel, we examine the fit of the power law function, and alternatives, to the observed visit times. We then attempt to determine whether such learning is related to purchases. Finally, we discuss the implications of these results for managers of firms competing in electronic environments and for further research in this area.

The Two Components of Search Costs

When information about sellers and their prices is not available completely or free of cost to buyers, sellers are able to charge prices in excess of marginal costs (Bakos 1997; Salop 1979; Stiglitz 1989). Such search costs have two components: *physical search costs*, which represent the time required to find the information needed to make a decision, and *cognitive costs*, which represent the costs of making sense of information sources and thinking about the information that has been gathered (Payne, Bettman, and Johnson 1993; Shugan 1980).

Electronic environments may produce a shift in the relative importance of cognitive and physical search costs. Although the widespread diffusion of information technology markedly lowers physical search costs, it has had less impact on cognitive costs. As West and colleagues (1999) observe, whereas Moore's law has reduced the cost of computing, it has not affected the cost or speed of the human information processor. More important, because the number of stores and products that can be searched online has increased because of low entry costs, electronic commerce potentially increases the relative importance level of cognitive search costs.

Cognitive costs are dynamic and change with experience. With practice, the time required to accomplish a task decreases. For example, it should be much more efficient to search a favorite site—following, we hypothesize, a power relationship with amount of use—than to learn the layout of a novel site. This would imply that perceived switching costs increase the more times a favorite site is visited, which creates a cognitive “lock-in” to that site over time, just as firms can lock in customers with high physical switching costs (Klemperer 1995; Williamson 1975).

The Power Law of Practice

The power law of practice is an empirical generalization of the ubiquitous finding that skill at any task increases rapidly at first, but later, even minor improvements take considerable effort (Newell and Rosenbloom 1981). At the beginning of the twentieth century, task performance was found to

improve exponentially with practice, for example, when using a typewriter (Bair 1902; Swift 1904). The exponential learning curve was one of the first proposed laws of human psychology (Thurstone 1937). Groups, organizations, and people can exhibit learning curves (Argote 1993; Epple, Argote, and Devadas 1991), and since World War II, learning curves have been used to forecast the increasing efficiency over time of industrial manufacturing (Hirsch 1952). Newell and Rosenbloom (1981) review the empirical evidence and show that improvement with practice is not exponential but instead is linear in log–log space; that is, it follows a power function. The power law of practice function and its equivalent log–log form is

$$(1) \quad T = BN^{-\alpha},$$

and

$$(2) \quad \log(T) = \log(B) - \alpha \log(N),$$

where T is the time required to complete the task, the most commonly used dependent measure of performance efficiency, though any dependent measure of efficiency can be used; N is the number of trials; and B is the baseline, an intercept term reflecting the performance time on the first trial ($N = 1$). The rate of improvement, α , is the slope of the learning curve, which forms a straight line when the function is graphed in log–log space.¹

Explanations for the Power Law of Practice

Two explanations have been proposed for the form of the power law of practice, though in most tasks, a combination of both is more likely responsible for log–log improvement over time. According to the method selection explanation (Crossman 1959), when a task is repeated, less efficient methods of accomplishing the task are abandoned in favor of more efficient methods as more efficient methods are discovered. In effect, the person performing the task is learning by trial and error the most efficient combination of methods, which could be revealed more systematically by a time and motion analysis (e.g., Niebel 1972). Over time, it becomes increasingly harder to distinguish minor differences among methods, and this accounts for the gradual slowing down of

¹Systematic deviations from a straight-line power law function have often been observed in previous studies. Improvement in the performance of a task, such as cigar rolling, ultimately reaches an asymptote imposed by the physical limitations of the tools used to perform the task, such as a cigar-rolling machine (Crossman 1959), and the observed data curve upward from a straight line as N increases. When the baseline time is not observed for a person, the empirically estimated power law curve shifts horizontally and appears flatter than curves estimated from subjects for whom the first observed trial is the baseline. Newell and Rosenbloom (1981) augment the simple power function form to derive a general power law of practice:

$$T = A + B(N + E)^{-\alpha},$$

where A is the asymptote, the minimum possible time in which the task can be performed, and E , prior experience, is the number of trials in which the person learned to perform the task before observation.

improvement. Card, Moran, and Newell (1983) demonstrate that improvement in the task of text editing could be modeled by the selection of the most efficient combination of task components.

The other explanation of practice law effects focuses on the cognitive processing of the input and output of the task rather than on the methods used in its performance. Rosenbloom and Newell (1987) explain log-log improvement as due to the "chunking" of patterns in the task environment, in much the same way that complex patterns can be memorized as a limited number of higher-order chunks (Miller 1958; Servan-Schreiber and Anderson 1990). Input-output patterns that occur often are readily learned in the first few trials, but rarer input patterns that occur maybe once in a thousand times require thousands of trials to chunk.

Applying the Power Law to Electronic Markets

Although the power law of practice has been found to operate in such diverse areas as perceptual motor skills (Snoddy 1926), perception (Kolers 1975; Neisser, Novick, and Lazar 1963), motor behavior (Card, English, and Burr 1978), elementary decisions (Seibel 1963), memory retrieval (Anderson 1983), and human-computer interaction (Card, Moran, and Newell 1983), there are many reasons to be skeptical of its applicability to consumer behavior on the Web and in other electronic environments.

First, there are theoretical reasons that the power law may not apply. Time spent at a site is routinely used as a measure of interest in the site (Novak and Hoffman 1997), which would seem to predict increasing, not decreasing, visit duration. Similarly, consumers spend more time looking at stimuli describing the alternatives they eventually choose (Payne 1976). In addition, purchasing usually requires at least one more page view than browsing (to enter data on the purchase form page), so any correlation between visit time and purchasing should work against the power law.

Second, there are several pragmatic concerns. If the content of a Web site changes regularly or, as is the case with dynamically generated Web pages, is different for every visit or when new navigation features are introduced to the site, each visit will involve a mixture of old (practiced) tasks and new (unpracticed) tasks, which attenuates any learning process. Thus, visits potentially consist of many aggregated tasks. Some tasks, such as site registration, are only performed on the first visit. Similarly, many classic power law studies observe hundreds or thousands of repetitions of a task. In contrast, the subjects in our Web data set have made many fewer visits to individual sites. The time between visits, which may be seconds in laboratory studies, is much greater in our data and varies significantly. The median time between visits to the same site is more than four days.

Third, if there are unobserved visits to Web sites, before panel membership or at another location such as at work, we will have underestimated the number of visits, which leads to underestimates of both learning parameters and reduces our ability to observe a power law.

Fourth, our data are likely to be much noisier than those from a typical power law study. Our data come from pan-

elists surfing in their living rooms, not in tightly controlled lab conditions. Their goals for visiting sites and the tasks they perform probably vary widely across visits.²

These reasons suggest that though the power law might be, in theory, a useful metric for understanding real-world learning, it is not obvious that it is either applicable or detectable in data collected from real-life Internet users.

Modeling the Learning of Web Sites

Data

The data we used came from the Media Metrix panel database, which records all the Web pages seen by a sample of personal computer (PC)-owning households in the 48 contiguous United States (Media Metrix is now a division of comScore Networks; www.comscore.com). During the period of analysis, Media Metrix maintained an average of 10,000 households in its panel every month. During the 12 months, from July 1997 to June 1998, examined in this study, the number of participants in the panel averaged 19,466 per month, roughly 2 per household. On each PC in the household, Media Metrix installs a software application that monitors all Web-browsing activity. Members of the household must log in to this monitoring software when they start the computer or take over the computer from another member of the household, as well as at half-hour intervals. This ensures that PC activity is assigned to the unique user who performed it. Media Metrix surveys more than 150 variables for each panelist, detailing among other things each person's age, sex, income, and education. The URL of each Web page viewed by members of the household, the date and time at which it was accessed, and the number of seconds for which the Web page was the active window on the computer screen are routinely logged by the software. Media Metrix records all the page views made by a household, even if these page files have come from a cache on the local computer. Although the Media Metrix panel contains participants of all ages, we restricted our analysis to a database of page views from panelists between 18 and 70 years of age, thus eliminating younger users who were unlikely to be purchasing on the Web.

²We could not take advantage of the general form of the power law function to model any systematic deviations that might be present in the data because of the low number of visits made by the majority of panelists. Very few would have made enough trials to hit up against their personal asymptotic performance. It is unlikely that a constant asymptote exists for physical performance of the site navigation task, because of typical variance in network delays across Web sessions experienced by most Web visitors. Because we have data from in-home Web surfing only, we may be missing many observations that occurred when the panelists visited these sites from other locations. In addition, many of our subjects may have visited these sites before they joined the panel, so the number of trials is underestimated. The number of prior trials, E , can be estimated by means of a grid search for an $E \geq 0$ that minimizes a loss function (Newell and Rosenbloom 1981). However, stable estimates of the number of prior visits require solid estimates of the power law function itself based on a large number of observed visits, and that is precisely what we do not have for most of our subjects.

Site Selection

We selected the books, music, and travel categories because they register the highest numbers for repeat visits and repeat online purchasing among online merchants (see also Brynjolfsson and Smith 2000; Clemons, Hann, and Hitt 2002; Johnson et al. 2002). Sites in each category were chosen from lists of leading online retailers from Media Metrix, BizRate (www.bizrate.com), and Netscape's "What's Related" feature, a service provided by Alexa (www.alexa.com) that defines related sites by observing which sites are visited by users. Table 1 shows the sites considered from each of the three categories.³ Although there are certainly more sites on the Web in each category, the number of users from the Media Metrix panel who visited other sites was too low for us to conduct meaningful analyses.

During the period we examined, July 1997–February 1999, the two largest online booksellers, Amazon.com and Barnes and Noble, also started to sell music and other categories. Although we could identify the category being browsed on these sites from the URL, we could not easily assign the time spent on the site to the different categories. We ended our analysis of data from the books and music categories after June 1998, when Amazon opened its music store (Amazon.com 1998).

For a subset of the sites in each category, noted by an asterisk in Table 1, we were able to determine whether a purchase had been made from the site with a reasonably high degree of certainty. These were sites that confirmed purchases with a "thank you" page that has the same text in the URL for every purchase made on the site. We used this subset of sites to examine the relationship between the parameters of the power law and whether a purchase had been made. Although this measure confirms a purchase, it does not provide the size of the purchase.

Defining Visits

Each row of the Media Metrix data contains a URL, a household identifier, the date and time the page became active (became the window on the desktop with "focus" attached to it), and the number of seconds it remained active.⁴ For our purposes, we defined a visit to a site as an unbroken sequence of URLs related to the same storefront. Our goal was to (1) eliminate visits that were accidental (e.g., typing

the wrong URL, clicking on the wrong link, being misdirected from a search engine); (2) identify a series of page views of a site that should be considered one visit, despite a brief side trip to another site; and (3) eliminate visits that were artificially lengthened because the user walked away from the computer, minimized the browser and did something else on the machine, and so forth.

To define visits, we first examined the distribution of the time between page views for individual panelists visiting the same site. These gap times, or interpage times, were the number of seconds between the time when the panelist stopped actively viewing one page from the site and the time when another page from the same site became active. Most gaps between page views were instantaneous (0 seconds duration), as is expected if pages are viewed consecutively. Approximately two-thirds of interpage gaps were less than a minute in duration, and beyond one minute, the distribution flattened out rapidly, with 95% of all gaps less than 15 minutes long. We therefore used 15 minutes between page views as the cut-off to distinguish one visit to the same Web site from a repeat visit. With this definition of a repeat visit, the median time between repeat visits across all three product classes is 4.5 days (books 6.2 days, music and travel both 4.2 days). In addition, we eliminated any visits that had a total duration of less than 5 seconds (a typical page load time) or exceeded 3 hours (which we assumed reflected an unattended browser). These numbers are similar to the definitions used by Media Metrix and other firms to define visits, and a sensitivity analysis showed that our conclusions were robust to these assumptions. To provide enough data points to allow at least one degree of freedom for testing a power law relationship, only the panelists who made three or more visits to a site in one of the three categories were retained in the data set ($N = 7034$). To provide stable estimates, we examined all sites that had at least 30 visitors (providing at least ten observations per parameter).

Analysis

From the 20-month database of page views, we extracted a separate data set for each site, sorting these data sets by date and time for each panelist. The active viewing time for each page during a visit was summed to yield total visit duration in seconds. After using the natural log function to transform visit number and visit duration, we estimated the power law

³Many of these Web companies have several different Web sites or pseudonyms that Media Metrix identifies with a single domain name. For example, Barnes and Noble has seven Web addresses for its site, six of which are hosted on America Online servers. Because it is important for our analysis that we identify all the related sites at which a visitor could learn a particular interface, we independently checked Media Metrix's roll-up definitions of domain names for the sites we considered. We searched for sites that had similar words in their URLs for one month, June 1998, and checked whether these sites belonged to companies on our list and were pseudonyms for identical storefronts. We verified the number of page views for our roll-up definitions with the Media Metrix counts for the same domain names.

⁴Our data are superior to typical Web server log file data in this respect. Web server log files record only the date and time a

requested file was sent to the requesting Internet provider address. If the file was sent successfully, it can be assumed that the receiver at least began to read the file. The time spent reading the file is unobserved, but it can be assumed to equal the time between the first request and a second request for a page from the same site. If no further request is made, it is typically assumed that the page was read for 30 minutes and then the session with that site ended. There can be many problems with these assumptions. The Media Metrix data show that active viewing often ceases before the next request is made from a site; for example, a visitor may focus on another application (e.g., sending an e-mail), which makes this second application the active window instead of the browser. See Novak and Hoffman (1997) and Drèze and Zufryden (1998) for further discussion of these issues.

TABLE 1
Retail Sites Used in the Analysis

Travel Sites (July 1997–February 1999)		
AAA.com	ETN.nl	PreviewTravel.com*
AlaskaAir.com	Expedia.com	Priceline.com*
AA.com	HotelDiscount.com	Southwest.com*
Amtrak.com	1096HOTEL.com	TheTrip.com
Avis.com*	ITN.net	TravelWeb.com
BestFares.com	LVRS.com	TravelZoo.com
CheapTickets.com	LowestFare.com	Travelocity.com
City.Net	MapBlast.com	TWA.com
Continental.com	MapQuest.com	UAL.com
Delta-Air.com	NWA.com	USAirways.com
Book Sites (July 1997–June 1998)		
Acses.com	Books.com	Kingbooks.com*
AltBookStore.com	BooksaMillion.com	Powells.com*
Amazon.com*	BooksNow.com*	Superlibrary.com
BarnesandNoble.com	Borders.com*	Wordsworth.com*
BookZone.com*		
Music Sites (July 1997–June 1998)		
BestBuy.com*	CDWorld.com*	MusicCentral.com
CDConnection.com	eMusic.com*	MusicSpot.com
CDEurope.com	Ktel.com	Newbury.com*
CDNow.com*	MassMusic.com	TowerRecords.com*
CDUnlverse.com*	MusicBoulevard.com	Tunes.com*
CdUSA.com		

*Purchases can be identified from Media Metrix data (URL) with a high level of confidence.

using two approaches. The first is an individual-level linear regression,

$$(3) \quad \log(T) = \beta + \alpha \log(N),$$

where T is the visit duration, N is the number of that visit, β is the intercept (which can be interpreted as an estimate of the log of B , the initial visit baseline time), and α is the learning rate. This approach makes no assumptions about the sign of α , though the power law posits a negative estimate. These individual linear regressions avoid many of the problems associated with the analysis of aggregate practice law data (Delaney et al. 1998). The mean of the individual-level estimates of α for each site provides an unbiased indicator of the mean power law slope for that site (Lorch and Myers 1990), and we conducted a series of one-tailed t -tests to compare the value of α with 0.⁵

Although these individual-level estimates are unbiased, they are a conservative measure and limit the number of predictor variables, which provides limited flexibility in testing alternative models. Our second estimation approach therefore was to use a hierarchical (random effects) linear model that allows heterogeneity in β and α and provides empirical Bayes estimates for each panelist:

$$(4) \quad \text{Log}(T)_{ij} = (\beta_j + \lambda_{1j}) + (\alpha_j + \lambda_{2j})\log(N_{ij}) + \varepsilon_{ij},$$

⁵We also examined aggregate patterns for the power law, a method that is inferior because of heterogeneity across consumers. The power law results are qualitatively similar. For example, an analysis of Amazon.com shows an α of -0.31 with an R^2 of $.45$, a result that does not change much if we alter the number of visits used in estimation from 3 to 5 to 20.

where β_j is the intercept for site j , and α_j is the slope of the learning curve for site j . In addition, we estimated λ_{1j} and λ_{2j} , which represent individual-level heterogeneity in estimates of β and α , respectively. We assumed that λ_1 and λ_2 were distributed normally and independently and that ε_{ij} had mean 0 and was independent.

Results

The Power Law and Repeat Visits to Web Sites

Table 2 shows the mean individual-level estimates for β (the intercept) and α (the learning rate), as well as the mean of the empirical Bayes estimates including heterogeneity, for the 36 sites. The sample-weighted average learning rate for the individual-level estimates, is -0.19 (95% confidence interval = -0.21 to -0.18 ; Hunter and Schmidt 1990). With two exceptions, Delta-Air.com and HotelDiscount.com, the individual-level means are negative, so visit duration declines as more visits are made, as we would expect if the power law of practice applied to Web site visits. Of the 36 sites, 28 (78%) had significantly more negative than positive individual-level estimates of α , and the number of negative estimates was significantly more than would be expected by chance (50%). There were no significant positive slopes.

The empirical Bayes estimates generally agreed with the individual-level regression estimates. All but 3 sites had negative empirical Bayes mean slopes, and 30 of the 36 sites (83%) had negative slopes and a mean α that was significantly negative, $p < .05$. The empirical Bayes model enabled us to test the estimates for the fixed components of the slope

TABLE 2
Estimated Power Law Functions

Site	N	Individual-Level Ordinary Least Squares Power Law Estimates		Empirical Bayes Power Law Estimates		Empirical Bayes Linear Model Estimates	
		β	α	β	α	β	α
Travel Sites	6146						
Map Quest.com	1482	5.37	-.118***	5.39	-.053**	5.31	-.007*
Travelocity.com	1394	5.52	-.176***	5.59	-.081***	5.45	-.009***
Expedia.com	1227	5.41	-.102***	5.42	-.032*	5.37	-.003
PreviewTravel.com	1167	5.13	-.164***	5.13	-.053**	5.04	-.002
City.net	1005	4.87	-.215***	5.00	-.149***	4.81	-.029***
Southwest.com	620	5.56	-.279***	5.69	-.138***	5.46	-.015***
AA.com	595	5.34	-.167***	5.36	-.073**	5.22	-.002
Delta-Air.com	425	5.03	.009	5.06	.010	5.05	.005
NWA.com	402	5.67	-.321***	5.78	-.228***	5.37	-.020***
Continental.com	331	5.27	-.236***	5.36	-.143***	5.12	-.016**
UAL.com	326	5.19	-.141*	5.32	-.127**	5.14	-.022**
ITN.net	326	5.03	-.298***	5.58	-.090**	5.42	-.007
Priceline.com	292	5.35	-.230**	5.89	-.344***	5.53	-.098***
USAirways.com	284	5.05	-.423***	5.33	-.335***	4.83	-.049***
TravelWeb.com	261	5.14	-.359***	5.16	-.035	5.12	-.009
TheTrip.com	213	5.21	-.287***	5.30	-.116**	5.08	-.006
BestFares.com	203	5.53	-.379***	5.55	-.158***	5.27	-.014**
Amtrak.com	198	5.38	-.602***	5.68	-.414***	5.04	-.048***
MapBlast.com	181	5.35	-.083	5.35	-.023	5.32	-.004
TWA.com	151	5.43	-.388***	5.39	-.098*	5.22	-.008
TravelZoo.com	150	5.12	-.301***	5.27	-.197**	4.97	-.025*
AAA.com	104	5.13	-.159	5.53	-.302***	5.07	-.039**
LowestFare.com	99	4.26	-.082	4.77	.042	4.79	.019
CheapTickets.com	95	5.29	-.509***	5.81	-.513***	5.21	-.122***
Avis.com	79	5.44	-.167	5.50	-.076	5.36	.002
1096HOTEL.com	77	5.06	-.243*	5.30	-.210*	4.96	-.021
AlaskaAir.com	49	5.13	-.286*	5.30	-.175*	5.02	-.022
ETN.nl	43	5.03	-.534**	5.41	-.427**	5.01	-.142*
LVRs.com	43	5.14	-.329	5.42	-.242*	4.99	-.015
HotelDiscount.com	39	4.56	.028	4.96	-.031	4.88	.009
BIC				257,471		257,708	
Book Sites	1282						
Amazon.com	1044	5.17	-.175***	5.27	-.077***	5.13	-.006*
BarnesandNoble.com	370	4.78	-.044	4.76	.013	4.76	.007
BIC				30,796		30,816	
Music Sites	534						
CDNow.com	256	5.29	-.169**	5.24	-.022	5.18	.004
MusicBoulevard.com	206	5.11	-.189**	5.15	-.078*	5.00	-.003
BestBuy.com	75	4.92	-.286*	5.15	-.230**	4.76	-.019*
CDUniverse.com	42	4.89	-.343*	5.00	-.197*	4.70	-.027
BIC				11,706		11,730	

* $p < .05$ (one-tailed).

** $p < .01$ (one-tailed).

*** $p < .001$ (one-tailed).

Notes: All β significantly > 0 , $p < .001$. BIC = Bayesian information criterion.


and the intercept across all the sites in a product category. In all three categories, the negative slope (α) and positive intercept (β) were significant, $p < .001$, and the majority (77.8%) of the learning coefficients (α) for specific sites were both significant and negative.

Figure 1 illustrates the estimated learning functions for both book sites, the four music sites, and some of the most frequently visited travel sites. As can be seen in Figure 1, there are significant differences in the learning rates across

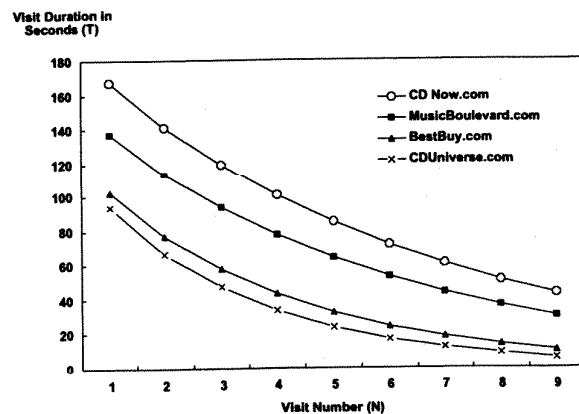
the sites in all three categories. In the case of books, the learning rate for Amazon is much faster than that for Barnes and Noble. These learning curves conform to the conventional wisdom that, initially at least, Barnes and Noble's online store lagged Amazon in the quality of its interface design. Nielsen (1999), for example, said "the best major site was probably amazon.com as of late 1998," and many commentators accused Barnes and Noble of playing "catch-up" in its approach to online design.

FIGURE 1
Power Law Learning Curves for Sites from the
Travel, Music, and Books Categories

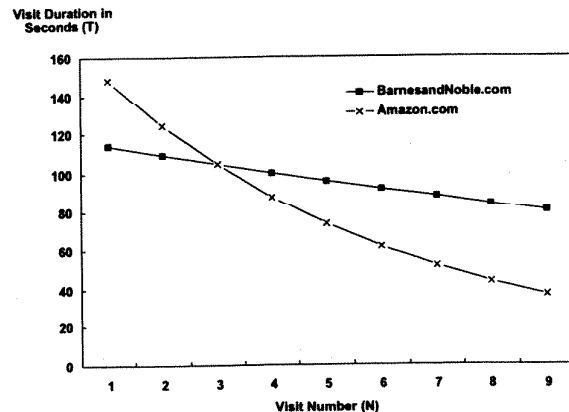
A: Travel



B: Music



C: Books



We should note, however, that there are several reasons that differences in slopes and intercepts must be interpreted with some caution. Across categories, the nature of the task may change. Finding books may involve different decisions than finding an appropriate airline ticket. Across sites, the set of users attracted to the site, their online experience, network connection speed, and other variables may also differ. The major point to be drawn from Figure 1 and Table 2, therefore, is that for most sites, the power law of practice provides a good account of visit times. The dynamic nature of Web content makes it difficult to relate specific characteristics of these particular Web sites to their power law parameters. Without an archive of server images for these Web sites collected at regular intervals, it is practically impossible to ascertain all the changes in content and design made on these sites during the time of observation. However, such research is possible to conduct prospectively, as are studies that explore these issues in experimental contexts.

Alternative Models and Tests

Although theory and evidence from other studies of practice suggest that a decrease in task duration is best modeled by a power law, we compared the results from the power law regression analysis with a likely alternative, a simple linear model, similar to the one used in Equation 3 but with a simple linear representation of the number of visits. The natural log of visit time T remains the dependent variable, because this transformation normalizes the distribution of visit times.⁶ To compare models, we used the Bayesian information criterion (BIC). All models had the same number of parameters. As can be seen in Table 2, the power law model was a superior model to the linear model of learning in all three product classes.⁷

In addition to comparing the two functional forms, we can construct an ordinal test of the differences in visit duration (untransformed) for the first three visits made by each panelist. If the data follow an exponential trend, the difference in duration between Trial 1 (t_1) and Trial 2 (t_2) will be greater than the difference in duration between Trial 2 (t_2) and Trial 3 (t_3). That is,

$$(5) \quad (t_1 - t_2) > (t_2 - t_3).$$

If, however, these differences follow a linear trend, the probability of observing a first difference greater than the second difference will not differ from chance ($p = .5$). In other words, with a linear slope, only approximately 50% of subjects will have a first difference ($t_1 - t_2$) greater than the second difference ($t_2 - t_3$), whereas for an exponentially decreasing slope, this number should exceed 50%. Table 3 shows the results of a series of binomial tests for each site

⁶Similar analyses with an untransformed dependent measure show a weaker pattern of results than the log-transformed visit times.

⁷We performed similar tests using individual-level regressions with similar results: The fit of the linear model is worse, overall, than the fit of the power law model, and only five sites (13.9%) have more significant estimates of α from the linear model than from the power law model.

with more than 30 visitors. At each of these sites, more than 50% of users had a first difference ($t_1 - t_2$) greater than the second ($t_2 - t_3$), and for 30 of the sites (83.3%), this difference was significant. We also examined the differences in duration of the second, third, and fourth visits, though fewer panelists recorded this many visits. Again, for the majority of the sites (63.9%), the percentage of visitors with a second difference ($t_2 - t_3$) greater than the third ($t_3 - t_4$) was significantly greater than 50%. If the signs of these differences are considered independent trials, the overall percentage for $(t_1 - t_2) > (t_2 - t_3)$ is 57.7% and for $(t_2 - t_3) > (t_3 - t_4)$ is 56.8%. Both are significantly different from the 50% that would result if a linear model was the best description of the

data. These results strengthen our claim that the decline in visit duration with successive visits is exponential and better modeled with a power function than a simple linear function.

A major difference between laboratory applications of the power law and the real-world task that we observe is the variability in the periods between trials. In laboratory studies, one task occurs right after another with little intervening time. However, in our naturalistic application, trials may occur on the same day or months apart.⁸ We examined

⁸We thank an anonymous reviewer for this insightful suggestion.

TABLE 3
Binomial Test of Differences in Visit Duration

Site	N_1^a	$\%(t_1 - t_2) > (t_2 - t_3)$	N_2^b	$\%(t_2 - t_3) > (t_3 - t_4)$
Travel Sites				
MapQuest.com	1482	55.7***	970	57.5***
Travelocity.com	1354	55.8***	932	53.2*
Expedia.com	1204	56.0***	837	58.1***
PreviewTravel.com	1156	56.1***	712	55.8***
City.net	1003	55.9***	602	58.3***
AA.com	583	56.1**	371	58.0***
Southwest.com	575	53.6*	394	57.4**
Delta-Air.com	417	60.4***	272	59.2***
NWA.com	337	58.5***	237	62.9***
Continental.com	331	58.9***	225	55.6*
UAL.com	325	62.2***	190	53.2
USAirways.com	284	60.9***	175	59.4**
Priceline.com	255	57.6**	128	59.4*
BestFares.com	201	62.7***	143	62.2**
Amtrak.com	198	57.6*	109	66.1***
TravelWeb.com	190	61.1***	107	57.0
MapBlast.com	181	58.0*	107	59.8*
TheTrip.com	166	60.2**	127	62.2**
ITN.net	153	66.0***	91	58.2*
TWA.com	150	64.0***	94	48.9
TravelZoo.com	150	67.3***	89	58.4*
AAA.com	101	60.4*	53	56.6
CheapTickets.com	95	64.2**	61	55.7
1096HOTEL.com	73	53.4	35	57.1
Avis.com	65	61.5*	35	45.7
AlaskaAir.com	46	58.7	28	50.0
ETN.nl	43	58.1	24	70.8*
LVRG.com	35	62.9*	21	33.3
LowestFare.com	33	75.8***	17	52.9
HotelDiscount.com	25	80.0***	12	58.3
Book Sites				
Amazon.com	962	61.0***	603	53.9*
BarnesandNoble.com	360	56.9**	204	50.0
Music Sites				
CDNow.com	250	58.8**	152	57.2*
MusicBoulevard.com	176	53.4	99	61.6**
BestBuy.com	75	56.0	50	52.0
CDUniverse.com	42	52.4	23	65.2*
Overall	13076	57.7***	8329	56.8***

* $p < .05$ (one-tailed).

** $p < .01$ (one-tailed).

*** $p < .001$ (one-tailed).

^aNumber of visitors with three valid trials.

^bNumber of visitors with four valid trials.

whether we could improve the fit of the power law by including the interval between repeat visits as a covariate in the following empirical Bayes estimation:

$$(6) \quad \text{Log}(T)_{ij} = (\beta_j + \lambda_{1i}) + (\alpha_j + \lambda_{2i})\log(N_{ij}) \\ + (\gamma_j + \lambda_{3i})\log(G_{ijN}) + \varepsilon_{ij}$$

where G_{ijN} is the interval time (or gap) preceding the N th visit ($N > 1$) by user i to site j (log transformed to normalize the distribution of G), γ_j is the fixed effect of the gap in time between visits to site j , and λ_3 is a normally distributed random variable accounting for individual-level heterogeneity in γ . These intervals were significant and positive (travel = .045, $p < .0001$; books = .056, $p < .0001$; music = .031, $p < .001$), and the inclusion of a gap parameter improved the fit of the model, which indicates that the longer the time between visits, the longer the visit takes. Yet the power law still described the data; α remained significantly negative in two of the three categories (travel $p < .0001$; books $p = .273$, not significant [n.s.]; music $p = .006$). This alternative model represents an important modification of the power law when applied to nonexperimental Web data. Whereas traditional applications of the power law emphasize the amount of practice and ignore its timing, this modified power law suggests that the density of practice matters in these data.

Alternative Explanations

An alternative explanation for this power law function is that it does not reflect learning on the part of the user but rather adaptation on the part of the network to the user's needs. Specifically, many Internet service providers and browsers cache copies of popular pages, that is, keep local copies of Web pages so they can be retrieved faster after the initial access.

To control for caching, we reran the power law model and added a variable that distinguished the first (and presumably uncached) visit to the site from all subsequent visits. If the decrease in visit times we observed was due to caching, we would expect this variable to be significant and the power law relationship to disappear or be greatly diminished. Although the inclusion of this control variable diminished the size of the slope coefficient, α , most remained negative and significant. The first trial dummy variable was significant for travel sites ($F(1, 65000) = 61.69$, $p < .0001$) and book sites ($F(1, 7504) = 4.32$, $p = .038$) but not for music sites ($F(1, 2962) = 1.29$, n.s.). However, for all three categories—travel ($F(30, 65000) = 7.40$, $p < .0001$), books ($F(2, 7504) = 2.97$, $p = .026$), and music ($F(4, 2962) = 2.42$, $p = .023$)— α remained significantly negative. Similar results were found at the individual-site level. For example, 16 (53.3%) of the 30 travel sites possessed a significant, negative slope coefficient, and 23 (76.7%) of 30 remained negative. In addition, we compared the power law and linear models with the cache term included in both models. This enabled us to test whether the apparent increase in fit of the power law compared with a linear learning function was due to lengthy first visits followed by subsequent caching. However, for all three categories, the power law model had a lower BIC than the linear model.

We also examined the possibility that the slope coefficient, α , might reflect not learning, but rather a decrease in

interest in the site. We examined the correlation between a panelist's individual-level α for a site and the number of observations (visits) used to estimate that α . These correlations showed no systematic pattern across product classes ($r = -.07$, $-.002$, and $.04$ for books, music, and travel, respectively) but are statistically significant given the large sample sizes. This analysis, along with our subsequent demonstration that faster learning leads to increased probability of buying, suggests that a decrease in interest does not account for our observed results.

Another reasonable alternative explanation for the observed decrease in visit duration is that people allocated a certain amount of time to Web surfing per session, but with the number of Web sites increasing over the period spanned by our data set from 646,000 in January 1997 to 4.06 million in January 1999 (www.iconocast.com), less time could be devoted to any one site. If this hypothesis is correct, the number of sites in any product class visited per month by a household should constantly increase, and each should receive a decreasing share of session time. However, the number of sites visited per month appears to be constant within a product class over time (Johnson et al. 2002).

Although our results and the power law model were consistent with a learning account, our results also parallel survey evidence that new Internet users navigate the Web in a more exploratory, experiential mode compared with experienced users (Novak, Hoffman, and Yung 2000). This transition from initial exploration to more efficient, goal-directed navigation may be another factor in diminishing visit times at specific sites, may apply to overall Web surfing behavior, and may explain connections with purchasing, but it does not rule out the underlying operation of the power law of practice.

Does Learning Lead to Buying?

Although we have found strong evidence at the individual level for the power law of practice in Web browsing behavior, is the power law consistently related to the buying behavior of Web site visitors? Are visitors more likely to buy from the sites they know best and can navigate more efficiently? If this is true, we should find a relationship between the two learning parameters, α and β , and the probability of making a purchase on any particular visit. We expect a negative relationship with purchasing for both parameters, in that faster initial visits (lower β) and faster slopes (lower α) may produce a greater likelihood of buying. To test this, we included the individual empirical Bayes estimates of α and β as predictors, as well as a variable $N - 1$, where N is the number of visits to the product category. Prior analysis suggests that buying probability increases as more visits to a category are made (Moe and Fader 2001). In addition, although we had no a priori theory of how the effect of the power law parameters might change over time, we included the interactions between N and the power law coefficients. We use $N - 1$ rather than N because this enables meaningful interpretations of these interaction terms (Irwin and McClelland 2001). When $N - 1 = 0$, the model predicts purchase probability for the first visit ($N = 1$) using only the intercept and the two learning parameters.

We estimated the following logit model for each product class:

$$(7) \quad \text{Buy}_N = \gamma_0 + \gamma_1\alpha + \gamma_2\beta + \gamma_3(N - 1) + \gamma_4\alpha(N - 1) + \gamma_5\beta(N - 1) + \varepsilon_{ij}$$

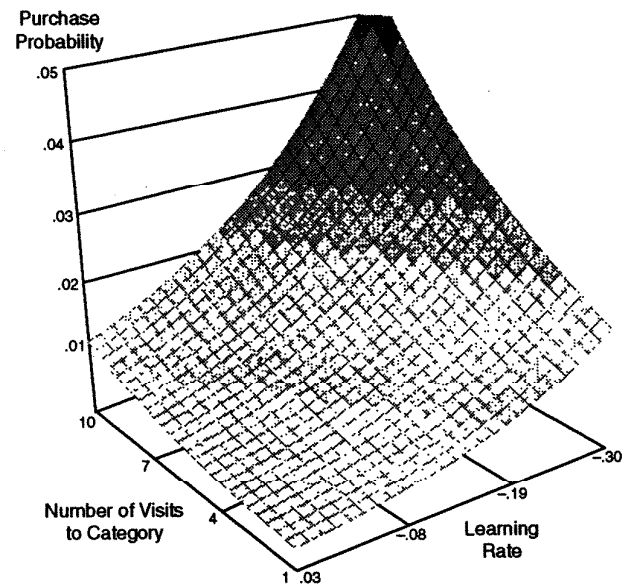
where Buy_N is 1 if category visit number N by a visitor to a site results in a purchase, 0 otherwise; α is that visitor's learning rate for this site; β is the visitor's power law function intercept (i.e., the estimated log of first visit time); and N is the category visit number. In addition, $\alpha(N - 1)$ is the interaction of α and the category visit number N ; similarly, $\beta(N - 1)$ is the corresponding interaction, and γ_0 , the intercept, and $\gamma_1, \gamma_2, \gamma_3, \gamma_4$, and γ_5 are all parameters to be estimated. The results are shown in Table 4. The logit model explained a significant amount of variance in buying (versus not buying) during specific visits.

For all three product classes, the main effect of α was negative and significant, as we predicted. The effect of β for two of the three product classes, music and travel, was significant in the predicted direction. As we expected, there was a significant tendency in both product classes for the probability of a purchase to increase with an increase in category visits.

The next two columns of Table 4 show that the number of visits to the site moderated some of these effects. For music, both interactions suggest that the effect of learning, α , decreases over time, whereas for travel, the effect of β seems to increase over time. However, they are very small effects compared with the simple effect of α and, within the range of learning we observed, slightly attenuate but do not reverse the beneficial effects of learning.

To illustrate the entire pattern, Figure 2 plots the variation in purchase probability for music sites over a range of α and N that is observed in the data we used to estimate the model. In Figure 2, α ranges ± 1.5 standard deviations from its mean (Jaccard, Turrisi, and Wan 1990), and the number of visits to the category N increases from one to ten, when β is held constant at the sample mean. Visitors with the fastest learning rates (α) had the highest probability of purchase at all trials. For example, changing the learning rate from -1 to -2 doubles the probability of purchase from .01 to .02 on

FIGURE 2
Probability of Purchase: Variation Over the Observed Range for Learning Rate (α) and Number of Visits to the Product Category (N) for Music Sites



the fourth visit. The effect of N can be seen as the entire plane tilts upward, but this effect is small compared with the increase due to the learning rate. Finally, the interaction between α and N produces a slight flattening of the slope of the effect of learning as N increases, but this effect is obscured by the effect of the logit transform and only noticeable outside the range of visits we typically observed.

The plots for the significant interaction effects of first visit duration, β , and N for the music and travel sites were similar to Figure 2. Visitors with faster first visit times had a higher probability of purchasing at all trials, though there

TABLE 4
Logistic Regression Predicting Buying on Visit N from Learning Rate (α), Baseline Time (β), and Category Visit Number (N)

Parameter	n ^a	Intercept	Predictors					Likelihood (5 Degrees of Freedom)
			α	β	$N - 1$	$\alpha \times (N - 1)$	$\beta \times (N - 1)$	
		γ_0	γ_1	γ_2	γ_3	γ_4	γ_5	
Books	2824	-3.42***	-5.54*	.035	-.004	.192	.005	14.23*
Music	1526	-2.50***	-5.52**	-.80***	.010***	.068*	.009***	13.02*
Travel	57639	-3.03***	-2.19***	-.45***	.001***	.000	-.002***	516.08***

* $p < .05$ (one-tailed).

** $p < .01$ (one-tailed).

*** $p < .001$ (one-tailed).

^aNumber of observations (visits). Number of valid visits per individual ranges from 3 to 678.

was a slight tendency in travel for this effect to decrease with more visits.

Limitations

The data from the time period we examined are rather sparse, because the frequency of online buying was relatively limited compared with subsequent periods. Similarly, the number of stores visited is limited, which makes the analysis of visit patterns difficult. Analysis of more recent data may not only replicate our current results but also be able to test new hypotheses in data sets that observe more frequent purchase visits. Another significant limitation is that these data lack several covariates that would increase our ability to predict visit times. We lack information about connection speed; details about the contents of each Web page and product offering; and details about caching, network delays, and so forth. Finally, unobserved prior visits not only flatten the learning curve, making it harder to detect learning improvement, but also introduce selection effects that may be alternative explanations for the relationships we observe. For example, the increasing number of visits may in reality have no effect on purchase probability but appears to do so because our data omit early visit purchases made before panel membership. Such information is becoming increasingly available, especially in controlled lab studies, and we believe our current work is a first step toward more sophisticated models that will provide excellent accounts of viewing time and purchase behavior.

Discussion

Implications for Web Competition

We have shown that visit duration declines the more often a site is visited. This decrease in visit time follows the same power law that describes learning rates in other domains of individual, group, and organizational behavior. Just as practice improves proficiency with other tasks, visitors to a Web site appear to learn to be more efficient at using that Web site the more often they use it. This is consistent with the small amount of competitive search observed in similar analyses of the Media Metrix data set, in which most panelists are loyal to just one store in the books, music, and travel categories (Johnson et al. 2002).⁹ Consistent with this view, we find a relationship between the ease of learning a Web site and the probability of purchasing.

The major implication of the power law of practice is that a navigation design that can be learned rapidly is one of a Web site's strongest assets. Although it is inconceivable that a Web site would be designed to be difficult to use, our results show considerable variation in ease of learning across sites and, perhaps most important, indicate that easier learning of a Web site leads to an increased probability of

⁹A possible explanation for the low level of comparison shopping is that people use one site to comparison shop, that is, a pricebot. We found little usage of pricebots (e.g., Acscs) in the Media Metrix data.

purchase. This suggests that the layout of a site can be an important strategic tool for online stores. Our advice for managers of Web sites with rapid learning rates is to maintain the current navigation design if possible. Altering the navigation design of a site reduces the cognitive lock-in effect of practiced efficiency and reduces an important competitive strength. If customers must learn a site design all over again, they might decide to learn someone else's instead. Customers come back on repeat visits to find new content, and the more varied the content, the more they will be encouraged to return. Whereas content should be refreshed often, changes in site design should be reviewed carefully.

Interface design can be exploited by both incumbents and competitors. An existing firm with a large customer base can extend to new product categories by using its familiar navigation design to encourage purchasing. This seems to be the heart of what might be termed Amazon's "tabbing" strategy, which introduces additional product classes (e.g., CDs) using the same navigational structure as previous categories (e.g., books) use and adds these new product classes as tabs along the top of a page. Such techniques lower the cost of using the site for new categories; some tasks will be new, but others are already completed (such as registering) or more easily accomplished. Site designers can take advantage of the power law to sequentially space the introduction of new features, allowing sufficient time between changes for the previous feature to be fully mastered so that cognitive resources can be devoted to mastering the new feature.¹⁰

Within legal limits, competitors can copy many design features of a familiar user interface. Most Web sites have already recognized the value of intuitive navigation design, and sites that have made successful innovations in site design have had many imitators. Some elements of site navigation, such as the ubiquitous use of tabs, quick search boxes, cookie-set preferences, and sometimes the whole look and feel of a competitor, are easily copied. Other navigation elements are harder to reproduce; for example, Amazon.com applied for a process patent for its 1-Click feature and, since an out-of-court settlement with major rival Barnes and Noble (Cox 2002), has licensed its use. An additional competitive advantage can be elements that customize the site in ways that make it easier to use. For example, the accuracy of purchase recommendations based on previous purchases at one store cannot easily be duplicated by that store's competitors, so it represents a difficult-to-imitate source of learning.

Another example from the short history of Web retail competition of how information can provide lock-in is eBay's seller ratings, which lock sellers into the service and diminish the risk for buyers. This feature enabled eBay to maintain an 80% market share when well-known competitors such as Yahoo! were offering similar auction services.

¹⁰We thank an anonymous reviewer for this recommendation.

Managers of Web sites with customers locked in by the ease of using the site may be able to take advantage of cognitive switching costs and charge price premiums. Smith and Brynjolfsson (2001) provide evidence that Amazon and Barnes and Noble charge a price premium over less well-known and, therefore, more risky sites. Sites that have easy-to-learn but difficult-to-imitate interfaces may also realize premiums in valuation. In the absence of other switching costs or loyalty schemes, cognitive lock-in implies an installed base of loyal customers whose lifetime value will provide a steady stream of earnings in the future (Shapiro and Varian 1999).

Further Research and Extensions

We used the analogy of the familiarity of a supermarket's layout as a form of cognitive lock-in, and we believe our results may be applicable far beyond the Web. For a broad range of products, ranging from video cassette recorders and personal digital assistants to services such as electronic organizers or voice mail menus, ease of learning relative to the competition is a relevant competitive attribute, not just because ease of use is itself good but also because it increases switching costs. Although this observation is not new, our work proposes a framework for modeling and metrics for assessing ease of learning that might be helpful. This framework could be used to study learning and loyalty in many environments in which cognitive costs are a newly important factor because technological advances have minimized physical costs.

Focusing on the Web, many new metrics have been proposed for measuring the attractiveness of Web sites, such as stickiness and interactivity (Novak and Hoffman 1997). Many of these measures assume a positive correlation between a visitor's involvement with the site and the duration of his or her visit or the number of pages viewed. We suggest that this relationship between visit length and interest is typical of a visitor's initial online behavior after adoption of the Web, but it is important, especially with regard to experienced Web surfers, to distinguish between utilitarian transactional and informational sites and hedonic media and entertainment sites (for a similar classification, see Hoffman and Novak 1996; Zeff and Aronson 1999). When a site's primary purpose is to encourage transactions, a decreasing pattern of visit times may be a good outcome. However, for a media site likely supported by advertising revenues, we expect the opposite pattern, or perhaps a constant mean duration, to characterize a successful site.¹¹

¹¹The drop in visit times we observe and model as exponential learning could represent the failure of Web shopping sites to generate a "compelling online experience" (Novak, Hoffman, and Yung 2000, p. 23) that keeps visitors browsing for long periods on successive visits. Instead, visitors are returning only to buy, which takes much less time than browsing. The more rapidly visitors switch from an experiential, exploratory mode of visiting to a goal-directed purchasing mode, the more purchasing occurs. Again, we thank an anonymous reviewer for these observations.

An area of further research of much interest for online retailers is identifying what makes a site easy to learn. What are the determinants of low initial visit times? What features of a Web site determine subsequent learning? Additional research could characterize the attributes of various Web sites, in terms of both infrastructure (servers, caching) and page design (limited graphics, useful search capabilities), and relate them to observed visit times. Such empirical research would help the development of a better cognitive science of online shopping (Nielsen 2000). Experimental work in this direction recently has been reported by Zauberma (2002) and Murray and Häubl (2002).

Economic theory suggests that the low physical costs of information search on the Web should encourage extensive search (e.g., Bakos 1997). However, when the data are examined, Web information search is fairly limited (Johnson et al. 2002), and this, coupled with our finding about cognitive switching costs, argues for the development of a behavioral search theory that extends economic theory beyond its concentration on physical costs. Cognitive switching costs are difficult to value in monetary terms, at least for the consumer evaluating the decision to search multiple sites versus staying with one familiar site. It would be worthwhile to examine whether this observed loyalty is a rational adaptation to search costs or if there are systematic deviations that can be predicted from an alternative theoretical framework.

The reaction of markets to cognitive lock-in is another interesting topic for additional research. Just as they consider other sources of switching costs, customers who anticipate that adopting a site as a favorite will lock them in should adopt the standard strategies for minimizing the effects of lock-in (Shapiro and Varian 1999). First, they should sell their loyalty dearly, choosing the site that pays the most for their lifetime value or offers the most support for relearning another site's navigation. Second, they should always have an escape strategy. For example, consumers should choose sites or tools that minimize switching costs. One example that has not been widely adopted is a nonproprietary shopping wallet that can be used for quick buying from multiple sites.

Conclusion

We suggest that the power law of practice, an empirical generalization from cognitive science, applies to visits to Web sites. Our results show that visits to Web sites are best characterized by decreasing visit times and that this rate of learning is related to the probability of purchasing.

We suggest that cognitive rather than physical costs are important in online competition and that this has several implications for Web site managers. Cognitive lock-in also has welfare implications for consumers, and we suggest some strategies they can adopt to reduce its effects. The phenomenon of cognitive lock-in due to the power law of practice is an important area for further research. Although we have empirically examined the applicability of this idea using Web sites, we believe such cognitive lock-in is an increasingly important factor for a broad range of products.

REFERENCES

- Alba, Joseph W., John Lynch, Barton Weitz, Chris Janiszewski, Richard Lutz, Alan Sawyer, and Stacy Wood (1997), "Interactive Home Shopping: Consumer, Retailer, and Manufacturer Incentives to Participate in Electronic Marketplaces," *Journal of Marketing*, 61 (July), 38-53.
- Amazon.com (1998), "Amazon.com Opens Music Store, Provides A Whole New Way To Discover Music," press release, Seattle, WA (June 11), (accessed January 8, 2003), [available at http://www.corporate-ir.net/ireye/ir_site.zhtml?ticker=AMZN&script=410&layout=8&item_id=233837].
- Anderson, John R. (1983), "Retrieval of Information from Long-Term Memory," *Science*, 220 (April), 25-30.
- Argote, Linda (1993), "Group and Organizational Learning Curves: Individual, System and Environmental Components," *British Journal of Social Psychology*, 32 (March), 31-51.
- Bair, J.H. (1902), "The Practice Curve: A Study in the Formation of Habits," *Psychological Monographs*, 5 (2), 1-70.
- Bakos, J. Yannis (1997), "Reducing Buyer Search Costs: Implications for Electronic Marketplaces," *Management Science*, 43 (December), 1676-92.
- Brynjolfsson, Erik and Michael D. Smith (2000), "Frictionless Commerce? A Comparison of Internet and Conventional Retailers," *Management Science*, 46 (April), 563-85.
- Card, Stuart K., William K. English, and Betty J. Burr (1978), "Evaluation of Mouse, Rate-Controlled Isometric Joystick, Step Keys, and Text Keys for Text Selection on a CRT," *Ergonomics*, 21 (August), 601-13.
- , Thomas P. Moran, and Allen Newell (1983), *The Psychology of Human-Computer Interaction*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Clemons, Eric K., Il-Horn Hann, and Lorin M. Hitt (2002), "Price Dispersion and Differentiation in Online Travel: An Empirical Investigation," *Management Science*, 48 (April), 534-49.
- Cox, Beth (2002), "One Settlement Over One-Click," *AtNewYork.com*, (March 7), (accessed January 8, 2003), [available at <http://www.atnewyork.com/news/article.php/987061>].
- Crossman, E.R.F.W. (1959), "A Theory of the Acquisition of Speed-Skill," *Ergonomics*, 2 (February), 153-66.
- Delaney, Peter F., Lynne M. Reder, James J. Staszewski, and Frank E. Ritter (1998), "The Strategy-Specific Nature of Improvement: The Power Law Applies by Strategy Within Task," *Psychological Science*, 9 (January), 1-7.
- Drèze, Xavier and Fred Zufryden (1998), "Is Internet Advertising Ready for Prime Time?" *Journal of Advertising Research*, 38 (May/June), 7-18.
- Epple, Dennis, Linda Argote, and Rukmini Devadas (1991), "Organizational Learning Curves: A Method for Investigating Intra-Plant Transfer of Knowledge Acquired Through Learning by Doing," *Organization Science*, 2 (February), 58-70.
- Hirsch, Werner Z. (1952), "Manufacturing Progress Functions," *Review of Economics and Statistics*, 34 (May), 143-55.
- Hoffman, Donna L. and Thomas P. Novak (1996), "Marketing in Hypermedia Computer-Mediated Environments: Conceptual Foundations," *Journal of Marketing*, 60 (July), 50-68.
- Hunter, John E. and Frank L. Schmidt (1990), *Methods of Meta-Analysis: Correcting Error and Bias in Research Findings*. Thousand Oaks, CA: Sage Publications.
- Irwin, Julie R. and Gary H. McClelland (2001), "Misleading Heuristics and Moderated Multiple Regression Models," *Journal of Marketing Research*, 38 (February), 100-109.
- Jaccard, James, Robert Turrisi, and Choi K. Wan (1990), *Interaction Effects in Multiple Regression*. Newbury Park, CA: Sage Publications.
- Johnson, Eric J., Wendy Moe, Pete Fader, Steve Bellman, and Jerry Lohse (2002), "On the Depth and Dynamics of Online Search Behavior," working paper, Columbia Business School, New York, (accessed January 8, 2003), [available at http://www.cebiz.org/downloads/search_depthdynamics.pdf].
- Kahn, Barbara E. and Leigh McAlister (1997), *Grocery Revolution: The New Focus on the Consumer*. Reading, MA: Addison-Wesley.
- Klemperer, Paul (1995), "Competition When Consumers Have Switching Costs: An Overview with Applications to Industrial Organization, Macroeconomics, and International Trade," *Review of Economic Studies*, 62 (4), 515-39.
- Kolers, Paul A. (1975), "Memorial Consequences of Automatized Encoding," *Journal of Experimental Psychology: Human Learning and Memory*, 1 (November), 689-701.
- Lal, Rajiv and Miklos Sarvary (1999), "When and How Is the Internet Likely to Decrease Price Competition?" *Marketing Science*, 18 (4), 485-503.
- Lorch, Robert F. and Jerome L. Myers (1990), "Regression Analysis of Repeated Measures Data in Cognitive Research," *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 16 (January), 149-57.
- Lynch, John G., Jr., and Dan Ariely (2000), "Wine Online: Search Costs Affect Competition on Price, Quality, and Distribution," *Marketing Science*, 19 (Winter), 83-103.
- Miller, George A. (1958), "Free Recall of Redundant Strings of Letters," *Journal of Experimental Psychology*, 56 (December), 485-91.
- Moe, Wendy W. and Peter S. Fader (2001), "Uncovering Patterns in Cybershopping," *California Management Review*, 43 (Summer), 106-117.
- Murray, Kyle B. and Gerald Häubl (2002), "The Fiction of No Friction: A User Skills Approach to Cognitive Lock-In," in *Advances in Consumer Research*, Vol. 29, Susan M. Broniarczyk and Kent Nakamoto, eds. Valdosta, GA: Association for Consumer Research, 11-18.
- Neisser, Ulric, Robert Novick, and Robert Lazar (1963), "Searching for Ten Targets Simultaneously," *Perceptual and Motor Skills*, 17 (3), 955-61.
- Newell, Allen and Paul S. Rosenbloom (1981), "Mechanisms of Skill Acquisition and the Law of Practice," in *Cognitive Skills and Their Acquisition*, John R. Anderson, ed. Hillsdale, NJ: Lawrence Erlbaum Associates, 1-55.
- Niebel, Benjamin W. (1972), *Motion and Time Study*, 5th ed. Homewood, IL: Richard D. Irwin.
- Nielsen, Jakob (1999), "Ten Good Deeds in Web Design," *Alertbox*, (October 3), (accessed January 8, 2003), [available at <http://www.useit.com/alertbox/991003.html>].
- (2000), "Why Doc Searls Doesn't Sell Any Books," *Alertbox*, (August 6), (accessed January 8, 2003), [available at <http://www.useit.com/alertbox/20000806.html>].
- Novak, Thomas P. and Donna L. Hoffman (1997), "New Metrics for New Media: Toward the Development of Web Measurement Standards," *World Wide Web Journal*, 2 (Winter), 213-46.
- , ———, and Yiu-Fai Yung (2000), "Measuring the Customer Experience in Online Environments: A Structural Modeling Approach," *Marketing Science*, 19 (Winter), 22-42.
- Oliver, Richard L. (1999), "Whence Consumer Loyalty?" *Journal of Marketing*, 63 (Special Issue), 33-44.
- Payne, John W. (1976), "Task Complexity and Contingent Processing in Decision-Making: An Information Search and Protocol Analysis," *Organizational Behavior and Human Performance*, 16 (August), 366-87.
- , James R. Bettman, and Eric J. Johnson (1993), *The Adaptive Decision Maker*. New York: Cambridge University Press.
- Rosenbloom, Paul and Allen Newell (1987), "Learning by Chunking: A Production System Model of Practice," in *Production System Models of Learning and Development*, David Klahr, Pat Langley, and Robert Neches, eds. Cambridge, MA: MIT Press, 221-86.
- Salop, Steven C. (1979), "Monopolistic Competition with Outside Goods," *Bell Journal of Economics*, 10 (Spring), 141-56.

- Schonberg, Edith, Thomas Cofino, Robert Hoch, Mark Podlaseck, and Susan L. Spraragen (2000), "Measuring Success," *Communications of the ACM*, 43 (August), 53-57.
- Seibel, Robert (1963), "Discrimination Reaction Time for a 1,023 Alternative Task," *Journal of Experimental Psychology*, 66 (3), 215-26.
- Servan-Schreiber, Emile and John R. Anderson (1990), "Learning Artificial Grammars with Competitive Chunking," *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 16 (July), 592-608.
- Shapiro, Carl and Hal R. Varian (1999), *Information Rules: A Strategic Guide to the Network Economy*. Boston, MA: Harvard Business School Press.
- Shugan, Stephen M. (1980), "The Cost of Thinking," *Journal of Consumer Research*, 7 (September), 99-111.
- Smith, Michael D. and Erik Brynjolfsson (2001), "Consumer Decision-Making at an Internet Shopbot: Brand Still Matters," *Journal of Industrial Economics*, 49 (4), 541-58.
- Snoddy, G.S. (1926), "Learning and Stability," *Journal of Applied Psychology*, 10 (February), 1-36.
- Stiglitz, Joseph E. (1989), "Imperfect Information in the Product Market," in *Handbook of Industrial Organization*, Richard Schmalensee and Robert Willig, eds. New York: North Holland, 769-847.
- Swift, Edgar James (1904), "The Acquisition of Skill in Type-Writing; A Contribution to the Psychology of Learning," *Psychological Bulletin*, 1 (August), 295-305.
- Thurstone, L.L. (1937), "Psychology as a Quantitative Rational Science," *Science*, 85 (March 5), 228-32.
- West, Patricia M., Dan Ariely, Steve Bellman, Eric Bradlow, Joel Huber, Eric Johnson, Barbara Kahn, John Little, and David Schkade (1999), "Agents to the Rescue?" *Marketing Letters*, 10 (August), 285-300.
- Williamson, Oliver E. (1975), *Markets and Hierarchies: Analysis and Antitrust Implications*. New York: The Free Press.
- Winer, Russell S. (2001), "A Framework for Customer Relationship Management," *California Management Review*, 43 (Summer), 89-105.
- Zauberman, Gal (2002), "Lock-In Over Time: Time Preferences, Prediction Accuracy, and the Information Cost Structure," in *Advances in Consumer Research*, Vol. 29, Susan M. Broniarczyk and Kent Nakamoto, eds. Valdosta, GA: Association for Consumer Research, 9-10.
- Zeff, Robbin and Brad Aronson (1999), *Advertising on the Internet*, 2d ed. New York: John Wiley & Sons.