

# Cognitive Lock-In and the Power Law of Practice<sup>1</sup>

February 27, 2002

Eric J. Johnson  
Columbia Business School, Columbia University

Steven Bellman  
University of Western Australia

Gerald L. Lohse  
Accenture

---

<sup>1</sup> We thank Jupiter Jupiter 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 of seminars at Ohio State University, Columbia University, University of Rochester and the University of Texas and from Asim Ansari and John Zhang. Correspondence should be sent to the first author at: Columbia Business School, Columbia University, 3022 Broadway, New York, NY, 10027. Phone 212-854-5068. Email: [ejj3@columbia.edu](mailto:ejj3@columbia.edu).

## **Cognitive Lock-In and the Power Law of Practice**

### **Abstract**

As information technology reduces the role of physical search costs, what are the sources of consumers' loyalty? We suggest that learning is an important factor in electronic environments and that efficiency resulting from learning can be modeled using a strong empirical regularity from cognitive science, the power law of practice. We examine the time spent visiting Web sites by a large panel of Web users and show that most sites can be characterized by decreasing visit times, and that generally those sites with the fastest learning curves show the highest rates of purchasing. [93 words]

## **Cognitive Lock-In and the Power Law of Practice**

### **Introduction**

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—what is labeled in the popular press as ‘perfect competition’ or ‘frictionless capitalism,’ but is more correctly called Bertrand competition (Bakos, 1997, see Lal and Savary, 1999 and Brynjolfsson and Smith 2000 for a discussion).

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, in particular identifying which sites exhibit greater loyalty or ‘stickiness’ and speculating about what causes repeat visits. The most common loyalty metric is the frequency and cumulative duration of visits. For example, eBay is listed in the *New York Times* top ten ‘stickiest’ sites because its users spend a substantial amount of time there, about 90 minutes a month according to Web rating services such as Jupiter Media Metrix<sup>1</sup>, and is consequently thought to be highly successful even though it is visited by less than five percent of the Web audience. Other loyalty metrics relate both visiting loyalty and purchasing loyalty, for example the number of visits per purchase, termed the “browse to

buy” or “book to look ratio,” and the ratio of the number of purchases to the number of repeat purchase visits (Morrisette *et al.*, 1999).

In this paper, we describe a mechanism and model for understanding the development of loyalty in electronic environments and an accompanying metric based upon an empirical generalization from cognitive science, the power law of practice (Newell and Rosenbloom, 1981). To provide an intuitive understanding of the mechanism, imagine a user visiting a Web site for the first time in order to purchase a compact disc. This user must first learn how to use the Web site in order to accomplish this goal. Once the CD has been purchased, we think that having learned to use this site raises its attractiveness relative to competing sites, and, all other things being equal (for example, fulfillment), that 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, increasing the difference in effort between using any another site and simply returning to the site where browsing and buying can be executed at the fastest rate, generating 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 by using various forms of customization, including personalization, recommendations, or easy checkout. Together, learning how to navigate a site and customization can increase the relative attractiveness of the site, generating a type of “cognitive loyalty program.”

A couple of analogies may reinforce this idea. 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, and 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, for example, explains the slow conversion from WordPerfect to Word (Shapiro and Varian, 1999).

In this work, we examine learning in electronic environments by looking at 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 often used in cognitive psychology to study learning—the power law of practice. We then investigate the relationship between this phenomenon of decreasing visit times and repeat-visit loyalty and online purchasing, using data from a panel of consumers using the World Wide Web.

The paper proceeds as follows. We first review the literature describing learning as a power law function and discuss its underlying causes. We then discuss why this type of learning may or may not apply to use of the Web. Using panel data that captures 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 see if such learning is related to subsequent visits and purchases. Finally, we close by discussing the implications of these results for managers of firms competing in electronic environments and for future research in this area.

## The Two Components of Search Costs

When information about sellers and their prices is not available completely and free-of-cost to buyers, sellers are able to charge prices in excess of marginal costs (Bakos, 1997; Diamond, 1985; Salop, 1979; Stiglitz, 1989). Such search costs have two components: *physical search costs* representing the time required to find the information required to make a decision, and *cognitive costs*, the cost of making sense of information sources, and the costs of 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. While the widespread diffusion of information technology markedly lowers physical search costs, it has had less impact upon cognitive costs. As West *et al.* (1999) observe, while Moore's law has reduced the cost of computing, it has not affected the costs or speed of the human information processor. More importantly, perhaps, because the number of stores and number of products that could be searched online has increased due to low entry costs, electronic commerce potentially increases the absolute as well as the relative 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 become 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, creating over time a cognitive “lock-in” to that site, just as, by analogy, 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 improves rapidly at first but later on even minor improvements take considerable effort (Newell and Rosenbloom, 1981). At the beginning of this century it was noticed that task performance improved exponentially with practice, for example, when using a typewriter (Bair, 1903; Swift, 1904). The exponential “learning curve” was one of the first proposed “laws” of human psychology (Thurstone, 1937). Groups and organizations, as well as individuals, 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) reviewed the empirical evidence and showed 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:

$$T = BN^\alpha \quad (1)$$

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

where  $T$  is the *time* required to complete the task, the most commonly used dependent measure of performance efficiency, although any dependent measure of efficiency can be used;  $N$  is the *number* of trials and  $B$  is the *baseline*, the performance time on the first trial ( $N = 1$ ). The rate of improvement,  $\alpha$ , is the slope of the learning curve, which forms a straight line when the function is graphed in log-log space.<sup>2</sup>

## **Explanations for the Power Law of Practice**

Two explanations have been proposed for the form of the power law of practice, although in most tasks a combination of both will most likely be 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 these 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, 1962). Over time, it becomes increasingly harder to distinguish minor differences between methods, and this accounts for the gradual slowing-down of improvement. Card, Moran, and Newell (1983) demonstrated that improvement in the task of text editing could be modeled by selection of the most efficient combination of task components. The second explanation of practice law effects focuses on the cognitive processing of the inputs and outputs of the task, rather than 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 “seven, plus or minus two” (Miller, 1956), 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 patterns that occur maybe once in a thousand times require thousands of trials to chunk.

## **Applying the Power Law to Electronic Markets.**

While 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), problem solving (Neves and Anderson, 1981), 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 why 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), and that would seem to predict increasing, and not decreasing visit duration. Similarly, it is well known that consumers spend more time looking at stimuli describing alternatives they eventually choose (e.g., 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 a number of pragmatic concerns:

If the content of a Web site changes regularly, or, as will be the case with dynamically generated Web pages, is different for every visit, each visit will involve a mixture of old (practiced) tasks and new (unpracticed) tasks, attenuating any learning process.

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.

- Likewise, if there are unobserved visits to Web sites, either prior to panel membership, or at another location, such as at work, we will have underestimated the number of visits leading to underestimates of learning reducing our ability to observe a power law.
- Finally, our data are likely to be even noisier than those from a typical power law study. Our data comes from panelists surfing in their living rooms, not in tightly controlled lab conditions. Their goals for visiting sites, and the tasks they performed, probably varied widely across visits.<sup>3</sup>

Together, these reasons suggest that while the power law might be, in theory, a useful metric for Web 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 Jupiter Media Metrix panel database, which records all the Web pages seen by a sample of PC-owning households in the 48 contiguous United States. During the time we analyze, Jupiter Media Metrix maintained an average of 10,000 households in their panel every month. Over the twelve months from July 1997 to June 1998 examined in this study, the number of individuals in the panel averaged 19,466 per month, roughly two per household. On each PC in the household, Jupiter Media Metrix installs a software application that monitors all Web-browsing activity. Members of the household must login to this monitoring software when they start the computer, or take over the computer from another member of the

household, and at half-hour intervals. This ensures that PC activity is assigned to the unique individual who performed it. Jupiter Media Metrix surveys over 150 variables for each individual, detailing, among other things, each individual's age, gender, income, and education. The URL of each Web page viewed by individual 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 an individual's computer screen, are routinely logged by the software. Jupiter 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 Jupiter Media Metrix panel contains individuals of all ages, we restricted our analysis to a database of page views from panelists aged between 18 and 70.

### **Site Selection**

We selected the books, music, and travel categories because these categories register the highest numbers for repeat visits and repeat online purchasing among online merchants (see also Brynjolfsson and Smith, 1999; Johnson *et al.*, 2001; Clemons *et al.*, 1999). Sites in each category were chosen from lists of leading online retailers from Jupiter Media Metrix ([www.mediametrix.com](http://www.mediametrix.com)), BizRate ([www.bizrate.com](http://www.bizrate.com)), and the Netscape's "What's Related" feature, a service, provided by Alexa ([www.alexa.com](http://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<sup>4</sup>. Although there are certainly more sites on the Web in each category than are in these, the number of individuals from the Jupiter Media Metrix panel who visited these other sites was too low for meaningful analyses to be conducted.

---

Table 1 about here

---

Over the period we examined, July 1997 to 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 couldn't easily assign 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, 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 purchasing.

### ***Defining Visits.***

Each line of the Jupiter Media Metrix data contains a URL, a household identifier, the date and time when that page became active (became the window on the desktop with 'focus' attached to it), and the number of seconds it remained active. 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 which were accidental (typing the wrong URL, clicking on the wrong link, or being 'misdirected' from a search engine); (2) identifying series of page views of a site that should be considered as one visit, with a brief side-trip to another site; and (3) eliminating visits which were artificially lengthened because the user walked away from the computer, minimized the browser and did something else on the machine, etc.

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 inter-page times, were the number of seconds between the time when an individual 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 expected if pages were viewed consecutively. About two thirds of inter-page gaps were less than a minute in duration, and beyond one minute the distribution flattened out rapidly, with 95% of all gaps being less than 15 minutes long. We therefore used 15 minutes between page views as the cutoff distinguishing one visit to the same Web site from a repeat visit. Using 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 three hours (which we assumed reflected an unattended browser). These numbers are similar to the definitions used by Jupiter Media Metrix and other firms to define visits, and a sensitivity analysis showed that our conclusions were robust to these definitional assumptions. To provide enough data points to allow at least one degree of freedom for testing a power law relationship, only those panelists who made three or more visits to a site in one of the three categories were retained in the data set ( $N = 7,034$ ), and to provide stable estimates, we examined all sites that had at least 30 visitors (providing at least 10 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 individual. The active viewing time for each of the pages seen 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 using two approaches, the first, individual-level linear regressions:

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

where  $T$  is the visit duration,  $N$  is the number of that visit (again, only visits longer than 5 seconds and shorter than 3 hours were counted),  $\beta$  is the intercept (which can be interpreted as an estimate of the log of  $B$ , the initial visit baseline time), and  $\alpha$  is the learning rate. This approach makes no assumptions about the sign of  $\alpha$ , although 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  $\alpha$  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 comparing the value of  $\alpha$  to 0.<sup>5</sup>

While these individual-level estimates are unbiased, they are a conservative measure, and limit the number of predictor variables, allowing us limited flexibility in testing alternative models. Our second estimation approach was, therefore, to use a hierarchical (random effects) linear model that allows heterogeneity in  $\beta$  and  $\alpha$ , providing empirical Bayes estimates for each panelist:

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

where  $\beta_j$  is the intercept for site  $j$ , and  $\alpha_j$  is the slope of the learning curve for site  $j$ . In addition, we estimated  $\lambda_{1i}$  and  $\lambda_{2i}$  that represent individual-level heterogeneity in

estimates of  $\beta$  and  $\alpha$  respectively. We assumed that  $\lambda_1$  and  $\lambda_2$  were distributed normally and independent and also that  $\varepsilon_{ij}$  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  $\beta$  (the intercept) and  $\alpha$  (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 was  $-.19$  (95% confidence interval:  $-.21, -.18$ ; Hunter and Schmidt, 1990). With two exceptions, Delta-Air.com and CDUniverse.com, the individual-level means were negative, indicating that visit duration declines as more visits are made, as we would expect if the power law of practice applied to Web site visits. For 29 (85%) of the 34 sites with negative slopes the mean  $\alpha$  was significantly negative,  $p < .05$ , and for 28 (80%) of the 35 sites with more negative than positive individual-level estimates of  $\alpha$ , the number of negative estimates was significantly more than would be expected by chance (50%). The Delta-Air.com and CDUniverse.com means were positive, but not significantly different from zero, indicating that visits to these sites may have fluctuated around a constant mean duration.

The empirical Bayes estimates generally agreed with the individual-level regression estimates. The empirical Bayes model allowed us to test the estimates for the fixed components of the slope and the intercept across all the sites in a product category. In all three categories, the negative slope ( $\alpha$ ) and positive intercept ( $\beta$ ) were significant,  $p < .0001$ , as were the majority (78%) of the learning coefficients ( $\alpha$ ) for specific sites.

Figure 1 illustrates the estimated learning functions for both book sites, and for some of the most frequently visited travel sites. As can be seen in Figure 1, there is significant variance in the learning rates across the sites in both categories. In the case of books, at least, it is interesting to note that 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 behind Amazon in the quality of its interface design. Nielsen (1999) for example, cited Amazon "as the best major Website as of late 1998" and many commentators accused Barnes and Noble of playing "catch-up" in its approach to online design.

We should also 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 well involve different decisions than finding an appropriate airline ticket. And across sites, the set of individuals attracted to the site, as well as 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 provided 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, and would be particularly

powerful if conducted in conjunction with random assignment to experimental conditions.

---

Figure 1 about here

---

### ***Alternative Models and Tests***

Although theory and evidence from other studies of practice suggest that a decrease in task duration will be best modeled by a power law, we compared the results from the power law regression analysis with a likely alternative, a simple linear model, exactly like the one used in Equation 3, but with a simple linear representation of the number of visits. The natural log of visit time  $T$  was still used as the dependent variable, as this transformation normalized the distribution of visit times.<sup>6</sup> To compare models we used the Bayesian information criterion (BIC) for comparison and 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 for all three product classes.<sup>7</sup>

In addition to comparing the two functional forms, we can construct an ordinal test on 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:

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

If, on the other hand, these differences follow a linear trend, the probability of observing a first difference greater than the second difference will not differ from chance ( $p = 0.5$ ). In other words, with a linear slope, only about 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 of the sites with more than 30 visitors. At every one of these sites, more than 50% of individuals had a first difference ( $t_1 - t_2$ ) greater than the second ( $t_2 - t_3$ ), and for 28 of the sites (78%), this difference was significant. We also examined the differences in duration of the second, third, and fourth visits, although fewer panelists recorded this many visits. Again, for the majority of the sites (61%), more than 50% of visitors had a second difference ( $t_2 - t_3$ ) greater than the third ( $t_3 - t_4$ ). If the signs of these differences are considered independent trials, the overall percentages for  $(t_1 - t_2) > (t_2 - t_3)$ , 57%, and for  $(t_2 - t_3) > (t_3 - t_4)$ , 56%, were both significantly different from the 50% that would obtain 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 best modeled with a power function rather than a simple linear function.

---

Table 3 about here

---

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.<sup>8</sup> We examined whether we could improve the fit of the power law by including the interval between repeat visits as a covariate in the empirical Bayes estimation:

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

where  $G_{ijN}$  is the interval time (or gap) preceding the  $N$ th visit ( $N > 1$ ) by individual  $i$  to site  $j$  (log transformed to normalize the distribution of  $G$ ),  $\gamma_j$  is the fixed effect of the gap in time between visits to site  $j$ , and  $\lambda_3$  is a normally-distributed random variable

accounting for individual-level heterogeneity in  $\gamma$ . These intervals were significant and improved the fit of the model. Yet the power law still described the data:  $\alpha$  remained significant in all three categories,  $p < .0001$ . This alternative model represents an important modification of the power law when applied to non-experimental Web data. While traditional applications of the power law emphasize the amount of practice, and ignore its timing, this modified power law suggests that in these data the density of practice matters.

### ***Other Alternative Explanations.***

One 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 adding 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 expected this variable to be significant, and the power law relationship to disappear or be greatly diminished. While the inclusion of this control variable diminished the size of the slope coefficient,  $\alpha$ , 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 ( $F(2,7504) = 2.97, p = .052$ ), and music ( $F(4,2962) = 2.42, p = .046$ ),  $\alpha$  remained significant. Similar results were

found at the individual-site level. For example, 17 of the 30 travel sites still possessed a significant negative slope coefficient and 25 out of 30 remained negative. In addition, we compared the power law and linear models discussed above with the cache term included in both models. This allowed us to test whether the apparent increase in fit of the power law relative to a linear learning function is due to lengthy first visits followed by subsequent caching. However, for all three categories the power law model still had a lower BIC than a linear model.

We also examined the possibility that the slope coefficient  $\alpha$  might reflect not learning, but rather a decrease in interest for the site. We examined the correlation between a panelist's individual-level  $\alpha$  for a site and the number of observations (visits) used to estimate that  $\alpha$ . 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 dataset from 646,000 in January 1997 to 4.06 million in January 1999 (*www.iconocast.com*), less and less time could be devoted to any one site. If this hypothesis was correct, then the number of sites in any product class visited per month by a household should be constantly increasing, with each receiving a decreasing share of session time. In fact, the

number of sites visited per month appears to be constant within a product class over time (Johnson *et al.*, 2001).

While our results and the power 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 to 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 as well as applying to overall Web surfing behavior, 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, does the power law effect influence the buying behavior of Web site visitors? Are visitors more likely to buy from the sites they know best and can navigate more efficiently? To explore a possible relationship, we estimated the following logit model:

$$Buy_N = \gamma_0 + \gamma_1 \alpha + \gamma_2 \beta + \gamma_3 N + \gamma_4 \alpha N + \gamma_5 \beta N + \varepsilon_{ij} . \quad (8)$$

where  $Buy_N$  is 1 if visit number  $N$  by an individual to a site results in a purchase, 0 otherwise;  $\alpha$  is that individual's learning rate for this site;  $N$  is the visit number;  $\alpha N$  is the interaction of  $\alpha$  and the visit number  $N$ ; similarly,  $\beta$  is the individual's power law function intercept (i.e., the estimated log of first visit duration) and  $\beta N$  is the corresponding interaction; and  $\gamma_0$ , the intercept, and  $\gamma_1$ ,  $\gamma_2$ ,  $\gamma_3$ ,  $\gamma_4$ , and  $\gamma_5$  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 a specific visit. For all three product classes, the main effect of  $\alpha$ , the learning rate, was, as predicted, negative and significant, as was the effect of  $\beta$  for two of the three product classes, music and travel. In addition, there was a significant tendency in both these product classes for the probability of a purchase to increase with an increase in visits. The next two columns of Table 4 show that the number of visits to the site moderated some of these effects. Figure 2 plots the significant interaction effect of  $\alpha$  and  $N$  on purchase probability for music sites, over the range  $\pm 1.5$  standard deviations from the mean  $\alpha$ , and over the first ten trials at a site, holding  $\beta$  constant at the sample mean (Jaccard, Turrisi, and Wan 1990). Visitors with the fastest learning rates ( $\alpha$ ) had the highest probability of purchase at all trials. The plots for the significant interaction effects of first-visit duration ( $\beta$ ) and  $N$  for music and travel sites were very similar to Figure 2. Visitors with faster first-visit times had a higher probability of purchasing at all trials, although there was a slight tendency in travel for this effect decrease with more visits.

---

Table 4 about here  
Figure 2 about here

---

### ***Limitations***

The data from the time period we examined is rather sparse, since the frequency of online buying was relatively limited compared to subsequent periods. Similarly, even the number of stores visited is limited, making the analysis of visit patterns difficult. Analysis of more recent data may not only replicate our current results but may 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 offerings, and details about caching, network delays, etc. Such information is becoming increasingly available, and we think that our current work is a first step towards 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 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 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 Jupiter Media Metrix data set, with most panelists being loyal to just one store in any of the books, music, and travel categories (Johnson *et al.*, 2001).<sup>9</sup> Most importantly, perhaps, we found 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 be one of a Web site's strongest assets. While 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 most importantly, perhaps, our results indicate that learning a Web site leads to an increased probability of 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 not change your navigation design if you don't have to. Altering the navigation design of a site reduces

the cognitive lock-in effect of practiced efficiency with the old navigation design, and reduces an important competitive strength. If your customers have to learn your site design all over again, they might decide to learn someone else's instead. Of course, customers come back on repeat visits to find new content, and the more varied the content the more they will be encouraged to return. But while content should be refreshed often, changes in site design should be reviewed very carefully.

Interface design can be exploited both by 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 (for example, CD's) using the same navigational structure as previous categories (such as books), adding these new product classes as 'tabs' along the top of a page.

Of course, within legal limits, competitors can copy many design features of a 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's 1-Click® feature is now protected by a process patent (Amazon.com, 1999) which is currently the subject of litigation. 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 upon previous purchases at one store

cannot easily be duplicated by that store's competitors, and thus represent a difficult to imitate source of learning.

Other examples from the short history of Web retail competition of how information can provide lockin include eBay's seller ratings that lock sellers into the service, and diminish the risk for buyers, features that allowed eBay to maintain an 80% market share when well known competitors such as Yahoo were offering similar auction services for free.

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 (Brynjolfsson and Smith, 2000, provide evidence that Amazon and Barnes and Noble already charge a price premium over less well-known, and therefore more risky sites). Sites which 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).

### ***Future Research and Extensions***

We earlier used the analogy of the familiarity of a supermarket's layout as a form of cognitive lock-in, and we think that our results may be applicable far beyond the Web. For a broad range of products, ranging from VCR's and Personal Digital Assistants to services such as touch-phone brokerages or voicemail menus, ease of learning relative to the competition is a relevant competitive attribute, not just because ease of use is itself good, but because it increases switching costs. While this observation is not new, our work proposes a framework for modeling and metrics for assessing ease of learning that

we think might be helpful. This framework could be applied to study learning and loyalty in many environments where 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 their visit or the number of pages viewed. We suggest that this relationship between visit length and interest may be typical of an individual's initial online behavior after adoption of the Web, but suggest that it is important, especially for experienced Web surfers, to distinguish between *utilitarian* transactional and informational sites and *hedonic* media and entertainment sites (for a similar classification, see Zeff and Aronson, 1999). When a site's primary purpose is to encourage transactions, it may be that a decreasing pattern of visit times is a good outcome. However, for a media site, most likely supported by advertising revenues, we might expect the opposite pattern, or perhaps a constant mean duration, to characterize a successful site.

An area of future research with much interest for Web site managers is the investigation of 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? Future research could characterize the attributes of various Web sites, both in terms of infrastructure (servers, caching, etc.) and page design (limited graphics, useful search capabilities) and relate them to observed visit times. Such empirical research would help develop a better cognitive science of online shopping (Nielsen, 2000)

Economic theory suggests that the low physical costs of information search on the Web should be encouraging extensive search (e.g., Bakos, 1997). When the data are examined, Web information search is, in fact, fairly limited (Johnson *et al.*, 2001), and this, coupled with our finding of cognitive switching costs argues for the development of a behavioral search theory that could extend 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 interesting 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 future research. Much like 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 lower switching costs. One example, which has not been widely adopted, is a non-proprietary shopping wallet that can be used for quick buying from multiple sites.

### **Conclusion**

We suggested 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 a number of implications for Web site managers. Cognitive lock-in also has welfare implications for consumers, and we suggested some strategies they could adopt to reduce its effects. The phenomenon of cognitive lock-in due to the power law of practice, will, we believe, be an important area for future research. While we have empirically examined the applicability of this idea using Web sites, we believe such cognitive lock-in is an increasingly important factor in a broad range of products.

## References

- Amazon.com (1998), "Amazon.com Opens Music Store, Provides A Whole New Way To Discover Music," Press Release, Seattle, WA, June 11, 1998; see [www.amazon.com/exec/obidos/subst/misc/music-launch-press-release.html](http://www.amazon.com/exec/obidos/subst/misc/music-launch-press-release.html).
- (1999), "Amazon.com Sues BarnesandNoble.com for Patent Infringement," Press Release, Seattle, WA, October 21, 1999; see [www.amazon.com/exec/obidos/subst/misc/press-releases/1-click-suit.html](http://www.amazon.com/exec/obidos/subst/misc/press-releases/1-click-suit.html).
- 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. (1903), "The Practice Curve," *Psychological Monographs*, 5, 5-70.
- Bakos, J. Yannis (1997), "Reducing Buyer Search Costs: Implications for Electronic Marketplaces," *Management Science*, 43(December), 1676-1692.
- Brynjolfsson, Erik, and Michael D. Smith (1999), "Frictionless Commerce? A Comparison of Internet and Conventional Retailers," Working Paper, Sloan School, Massachusetts Institute of Technology; see [ecommerce.mit.edu/papers/friction](http://ecommerce.mit.edu/papers/friction).
- , and ————— (2000), "The Great Equalizer? The Role of Shopbots in Electronic Markets," Working paper, Sloan School, Massachusetts Institute of Technology.
- 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), 601-613.
- , 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 (1999), "The Nature of Competition in Electronic Markets: An Empirical Investigation of Online Travel Agencies Offerings," Working Paper, Department of Operations and Information Management, The Wharton School, University of Pennsylvania.
- Crossman, E. R. F. W. (1959), "A Theory Of The Acquisition Of Speed-Skill," *Ergonomics*, 2, 153-166.
- 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.
- Diamond, P. A. (1985), "Search Theory," Working Paper 389, Department of Economics, Massachusetts Institute of Technology.
- 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, W. Z. (1952), "Manufacturing Progress Functions," *Review of Economics and Statistics*, 34, 143-155.
- Hunter, John E., and Frank L. Schmidt (1990), *Methods of Meta-Analysis: Correcting Error and Bias in Research Findings*, Thousand Oaks, CA: Sage.
- Jaccard, James, Turrisi, R., and Choi K. Wan (1990), *Interaction Effects in Multiple Regression*, Newbury Park, CA: Sage.

- Johnson, Eric J., Wendy Moe, Pete Fader, Steve Bellman, and Jerry Lohse (2001), "Modeling the Depth and Dynamics of Online Search Behavior," Working Paper, Columbia Business School, New York.
- 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), 515-539.
- Kolers, P. A. (1975), "Memorial Consequences Of Automatized Encoding," *Journal of Experimental Psychology: Human Learning and Memory*, 1(6), 689-701.
- 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(1), 149-157.
- Lal, R. and M. Sarvary (1999), "When and how is the Internet likely to decrease price competition?," *Marketing Science*, 18 (4), 485-503.
- Miller, George A. (1956), "The Magical Number Seven, Plus Or Minus Two: Some Limits On Our Capacity For Processing Information," *Psychological Review*, 63, 81-97
- (1958), "Free Recall Of Redundant Strings Of Letters," *Journal of Experimental Psychology*, 56, 485-491
- Morrisette, Shelley, James L. McQuivey, Nicki Maraganore, Gordon Lanpher (1999), "Are Net shoppers Loyal?" Forrester Report, March 1999; see [www.forrester.com/ER/Research/0,1338,7158,FF.html](http://www.forrester.com/ER/Research/0,1338,7158,FF.html)
- Neilsen, Jakob (1999), "Ten Good Deeds in Web Design," Alertbox, October 3; see <http://www.useit.com/alertbox/991003.html>
- (2000), "Why Doc Searls Doesn't Sell Any Books," Alertbox, August 6; see <http://www.useit.com/alertbox/20000806.html>
- Neisser, U., R. Novick, and R. Lazar (1963), "Searching For Ten Targets Simultaneously," *Perceptual and Motor Skills*, 17, 955-961.
- Newell, Allen, and Paul S. Rosenbloom (1981), "Mechanisms Of Skill Acquisition And The Law Of Practice," in *Cognitive Skills and Their Acquisition*, ed. John R. Anderson, Hillsdale, NJ: Lawrence Erlbaum Associates, 1-55.
- Neves, D. M., and John R. Anderson (1981), "Knowledge Compilation: Mechanisms For The Automatization Of Cognitive Skills," in *Cognitive Skills and Their Acquisition*, ed. John R. Anderson, Hillsdale, NJ: Lawrence Erlbaum Associates.
- Niebel, B. W. (1962), *Motion and time study* (3rd ed.), Homewood, IL: Richard D. Irwin.
- 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), 213-246.
- , ———, and Yiu-Fai Yung (2000), "Measuring the Customer Experience in Online Environments: A Structural Modeling Approach," *Marketing Science*, 19(Winter), 22-42.
- Payne, John W. (1976), "Task complexity and contingent processing in decision-making: An information search and protocol analysis," *Organizational Behavior and Human Performance*, 16, 366-387.
- , James R. Bettman, and Eric J. Johnson (1993), *The Adaptive Decision Maker*, New York, NY: 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*, ed. David Klahr, Pat Langley, and Robert Neches, Cambridge, MA: MIT Press, 221-286.
- Salop, S. (1979), "Monopolistic Competition with Outside Goods," *Bell Journal of Economics*, 10 (Spring), 141-156.
- Seibel, R. (1963), "Discrimination Reaction Time for a 1,023 Alternative Task," *Journal of Experimental Psychology*, 66(3), 215-226.
- 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.
- Snoddy, G. S. (1926), "Learning and Stability," *Journal of Applied Psychology*, 10, 1-36.
- Stiglitz, J. E. (1989), "Imperfect Information in the Product Market," in *Handbook of Industrial Organization*, ed. R. Schmalensee and R. D. Willig, 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, 228-232.
- 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.
- Zeff, Robbin, and Brad Aronson (1999), *Advertising on the Internet*, 2nd ed., New York, NY: John Wiley & Sons, Inc.

**Table 1: Sites Analysed.**

---

**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)**

Acscs.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*
CDUniverse.com*	MusicBoulevard.com	Tunes.com*
CdUSA.com		

---

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

**Table 2. Estimated Power Law Functions.**

<i>Site</i>	<i>N</i>	<i>Individual-Level OLS</i>		<i>Empirical Bayes</i>		<i>Empirical Bayes</i>	
		<i>Power Law Estimates</i>		<i>Power Law Estimates</i>		<i>Linear Model Estimates</i>	
		$\beta$	$\alpha$	$\beta$	$\alpha$	$\beta$	$\alpha$
<b>Travel Sites</b>	<b>6146</b>						
Map Quest.com	1482	5.41	-.153***	5.37	-.041*	5.31	-.006
Travelocity.com	1394	5.44	-.211***	5.60	-.110***	5.47	-.028***
Expedia.com	1227	5.40	-.121***	5.41	-.040*	5.35	-.006
PreviewTravel.com	1167	5.18	-.165***	5.18	-.086***	5.07	-.019**
City.net	1005	4.88	-.221***	4.93	-.113***	4.80	-.031***
Southwest.com	620	5.36	-.303***	5.73	-.197***	5.50	-.050***
AA.com	595	5.36	-.173***	5.47	-.141***	5.29	-.030**
Delta-Air.com	425	5.02	.029	4.97	.048	5.00	.021*
NWA.com	402	5.35	-.260***	6.01	-.320***	5.64	-.082***
Continental.com	331	5.29	-.235***	5.37	-.156***	5.17	-.033**
UAL.com	326	5.17	-.205***	5.32	-.132**	5.17	-.036*
ITN.net	326	5.30	-.159*	5.24	-.234***	5.01	-.073***
Priceline.com	292	5.45	-.283***	5.59	-.241***	5.35	-.074***
USAirways.com	284	5.34	-.339***	5.36	-.353***	4.96	-.091***
TravelWeb.com	261	5.53	-.152*	5.37	-.274***	5.08	-.077***
TheTrip.com	213	5.21	-.257***	5.42	-.217***	5.14	-.047**
BestFares.com	203	5.55	-.455***	5.67	-.228***	5.36	-.045**
Amtrak.com	198	5.38	-.587***	5.77	-.462***	5.25	-.123***
MapBlast.com	181	5.41	-.086	5.31	-.002	5.26	.013
TWA.com	151	5.07	-.416***	5.44	-.14*	5.23	-.020
TravelZoo.com	150	5.46	-.131	5.32	-.223***	5.05	-.052*
AAA.com	104	5.13	-.147	5.54	-.317**	5.26	-.108***
LowestFare.com	99	5.51	-.190*	4.50	-.202	4.26	-.048
CheapTickets.com	95	5.04	-.296**	5.75	-.482***	5.32	-.162***
Avis.com	79	5.37	-.128	5.53	-.120	5.43	-.041
1096HOTEL.com	77	4.78	-.179	5.39	-.296**	5.09	-.089**
AlaskaAir.com	49	5.13	-.280*	5.28	-.186	5.09	-.055
ETN.nl	43	5.16	-.302*	5.41	-.428**	5.01	-.144*
LVRS.com	43	5.57	-.331**	5.53	-.366**	5.15	-.107*
HotelDiscount.com	39	5.11	-.215	4.57	.005	4.60	-.005
<i>BIC</i>					257,471		257,708
<b>Book Sites</b>	<b>1282</b>						
Amazon.com	1044	5.15	-.158***	5.24	-.106***	5.05	-.009**
BarnesandNoble	370	4.82	-.117*	4.74	.005	4.73	.005
<i>BIC</i>					30,796		30,816
<b>Music Sites</b>	<b>534</b>						
CD Now	256	5.24	-.164*	5.24	-.022	5.18	.004
Music Boulevard	206	5.09	-.230**	5.16	-.083*	5.01	-.006*
Best Buy	75	4.89	-.273*	5.15	-.231***	4.76	-.019**
CD Universe	42	5.09	.057	4.99	-.190*	4.70	-.027
<i>BIC</i>					11,706		11,730

\*\*\*  $p < .001$  (one-tailed)    \*\*  $p < .01$  (one-tailed)    \*  $p < .05$  (one-tailed)    All  $\beta$  significantly  $> 0$ .,  $p < .001$

**Table 3: Binomial Test of Differences in Visit Duration**

<i>Site</i>	$N_1$	$\%(t_1-t_2)>(t_2-t_3)$	$N_2$	$\%(t_2-t_3)>(t_3-t_4)$
<b>Travel Sites</b>				
MapQuest.com	1482	55.7***	970	57.5***
Travelocity.com	1394	56.2***	967	53.1*
Expedia.com	1227	55.7***	860	56.7***
PreviewTravel.com	1167	55.5***	719	55.9***
City.net	1005	55.7***	603	58.4***
Southwest.com	620	53.9*	427	57.4***
AA.com	595	55.5**	379	57.8***
Delta-Air.com	425	60.7***	273	59.0***
NWA.com	402	56.0**	287	57.8**
Continental.com	331	58.9***	225	55.6*
ITN.net	326	58.0**	217	56.2*
UAL.com	326	62.3***	190	53.2
Priceline.com	292	57.2**	143	56.6*
USAirways.com	284	60.9***	175	59.4**
TravelWeb.com	261	57.1**	155	56.8*
TheTrip.com	213	53.5	157	61.2**
BestFares.com	203	61.6***	146	62.3***
Amtrak.com	198	57.6*	109	66.1***
MapBlast.com	181	58.0*	107	59.8*
TWA.com	151	64.2***	96	49.0
TravelZoo.com	150	67.3***	89	58.4*
AAA.com	104	58.7*	57	56.1
LowestFare.com	99	63.6**	56	44.6
CheapTickets.com	95	65.3***	61	55.7
Avis.com	79	59.5*	43	51.2
1096HOTEL.com	77	55.8	35	57.1
AlaskaAir.com	49	59.2	29	48.3
ETN.nl	43	58.1	24	70.8*
LVRS.com	43	60.5	25	36.0
HotelDiscount.com	39	69.2**	21	47.6
<b>Book Sites</b>				
Amazon.com	1044	60.3***	678	54.0*
BarnesandNoble.com	370	57.0**	209	50.2
<b>Music Sites</b>				
CD Now.com	256	58.6**	156	55.8
MusicBoulevard.com	206	53.4	128	53.9
BestBuy.com	75	56.0	50	54.0
CDUniverse.com	42	52.4	23	65.2*
Overall	13854	57.2***	8889	56.2***

**Table 4: Logistic Regression Predicting Buying on Visit  $N$  from  $\alpha$ ,  $\beta$ ,  $N$ , and the Interactions of  $\alpha$  and  $N$ , and  $\beta$  and  $N$ .**

	$n$	$\alpha$	$B$	$N$	$\alpha N$	$BN$	<i>Likelihood</i> <i>5 d.f.</i>
Books	2824	-5.54*	.035	-.004	.192	.005	14.23*
Music	1526	-5.52**	-.80***	.010***	.068*	.009***	13.02*
Travel	57639	-2.19***	-.45***	.001***	.000	-.002***	516.08***

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

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

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

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

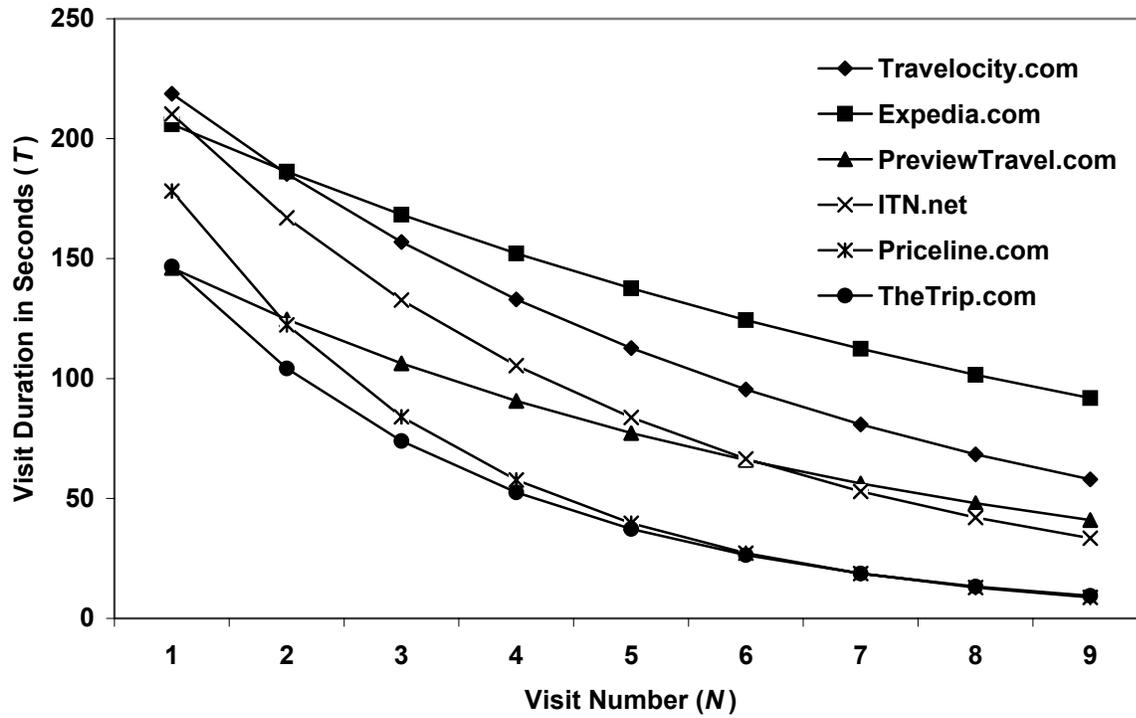


Figure 1. Power Law Learning Curves for Sites from the Travel, Music, and Books Categories.

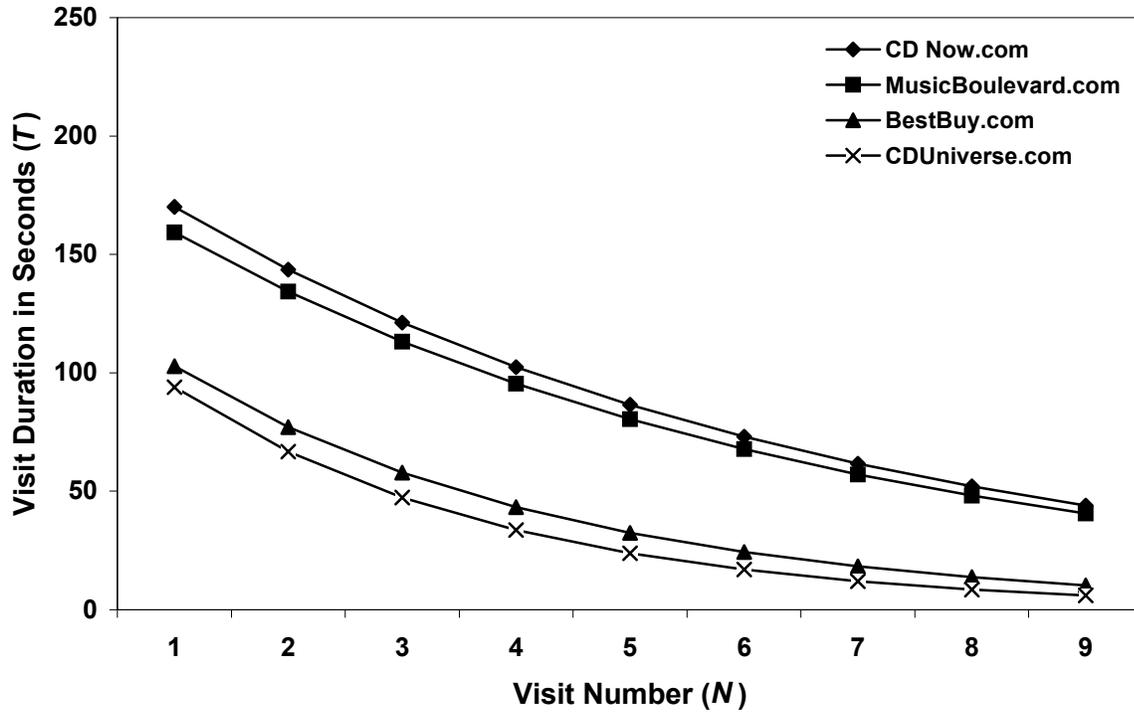
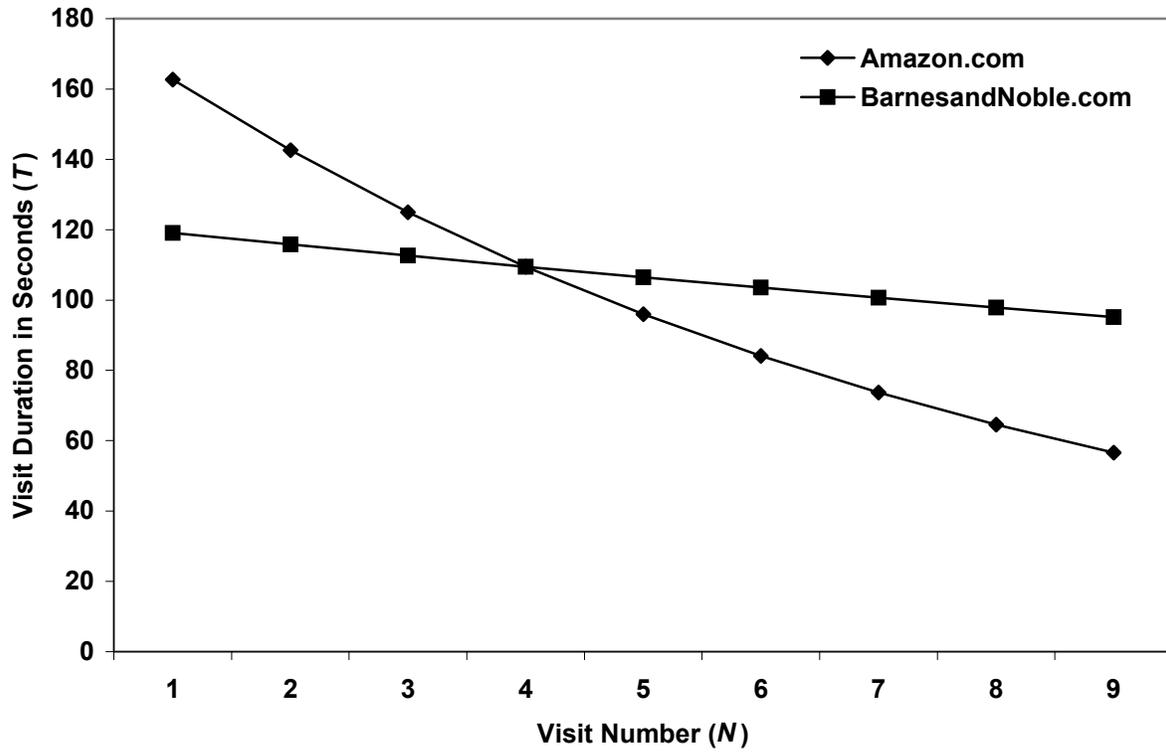
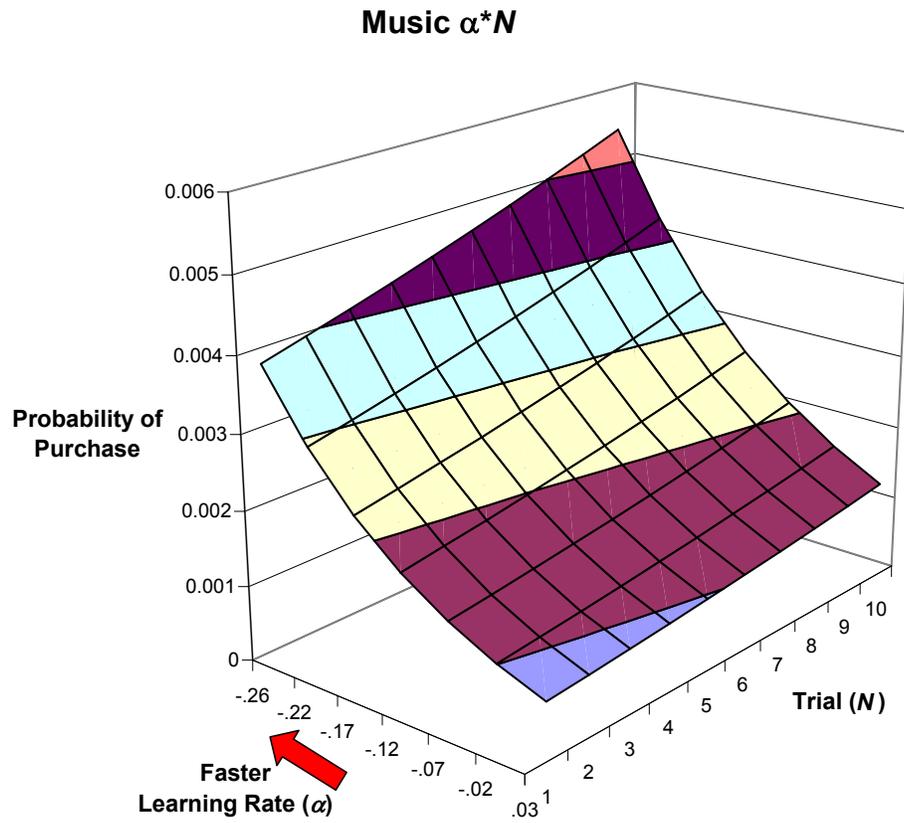


Figure 1. Power Law Learning Curves for Sites from the Travel, Music, and Books Categories.



**Figure 2: Probability of Purchase: Significant Interaction Effects**



## Endnotes

<sup>1</sup> Data from Jupiter Media Metrix, June, 1998.

<sup>2</sup> Systematic 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, for example, a cigar-rolling machine (Crossman, 1959), and the observed data curve upward from a straight line as  $N$  increases. Secondly, when the baseline time is not observed for an individual, the empirically estimated power law curve is shifted horizontally and appears flatter than curves estimated from subjects for whom the first observed trial is in fact the baseline. Newell and Rosenbloom (1981) augmented 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 on which the individual learned to perform the task prior to observation.

<sup>3</sup> We 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 of them would have made enough trials to be hitting up against their personal asymptotic performance. It is unlikely, given the present state of the Internet, that a constant asymptote exists for physical performance of the site navigation task, given the typical variance in network delays across Web sessions experienced by most Web visitors. Since we have data from in-home Web surfing only, we may be missing many observations that occur when the panelists visit these sites from other locations. It is highly likely that many of our subjects have visited these sites many times before they joined the panel, so that the number of trials is underestimated. The number of prior trials,  $E$ , can be estimated using 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 don't have for most of our subjects.

<sup>4</sup> Many of these Web companies have a number of different Web sites, or a number of pseudonyms, which Jupiter Media Metrix identify with a single domain name. For example, Barnes and Noble have seven Web addresses for their site, six of which are hosted on AOL servers. Because it is important for our analysis that we identified all the related sites on which a visitor could learn a particular interface, we independently checked Jupiter Media Metrix's "roll-up" definitions of domain names for the sites we considered. We searched for all the sites that had similar words in their URLs in data from one month, June 1998, and checked whether these sites belonged to companies on our list and whether some were in fact 'pseudonyms' for identical storefronts. We verified the number of page views for our roll-up definitions with the Jupiter Media Metrix counts for the same domain names.

<sup>5</sup> We have also examined aggregate patterns for the power law, a method that is inferior because of heterogeneity across consumers. The power law results are qualitatively quite similar. For example, an analysis of Amazon.com shows an  $\alpha$  of  $-.31$ , 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.

<sup>6</sup> Similar analyses with an untransformed dependent measure showed a weaker pattern of results than the log-transformed visit times.

<sup>7</sup> We performed similar tests using the 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 six sites (16.7%) have more significant estimates of  $\alpha$  from the linear model than the power law model

<sup>8</sup> We thank an anonymous reviewer for this insightful suggestion..

<sup>9</sup> One explanation for the low level of comparison-shopping is that people are using one site to comparison shop, i.e., a pricebot. We found very little usage of pricebots in the Jupiter Media Metrix data.