RICE UNIVERSITY

Searching and Browsing on E-commerce Sites: Frequency, Efficiency and Rationale

by

Michael A. Katz

A THESIS SUBMITTED IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE DEGREE

Doctor of Philosophy

HOUSTON, TEXAS

MAY, 2001

ABSTRACT

Searching and Browsing on E-commerce Sites:

Frequency, Efficiency and Rationale

by

Michael A. Katz

Two experiments were conducted to investigate product finding behavior on e-commerce Web sites, focusing on the use of product menus or site search functions to locate products. Experiment 1 sought to characterize typical product finding behavior on e-commerce sites, and to shed light on the user's decision to browse a product menu or use a site search function to locate products. The results yielded a wide distribution of searching and browsing behavior and found that use of the site search functions did not yield faster or more accurate performance in locating products. Questionnaire data suggested that perceived effectiveness of the menu structure, menu and search function prominence, and the user's disposition toward using search functions all influenced the decision of whether to browse or search a site for a product, with the principle known as "information scent" playing a particularly important role.

Experiment 2 used experimentally controlled novel e-commerce sites to investigate the factors suggested to be important to product finding behavior by Experiment 1. The goal was to determine whether product finding behavior could be influenced via the manipulation of site design factors. The results provided evidence that such behavior could be manipulated through site design and supported the role of Experiment 1 factors in the decision process of the user. As a whole, the results suggest that product finding behavior is dependent upon both the characteristics of the site as well as the user. They also suggest that the efficiency of the menu structure of a site appears to be as critical to its usability as the fidelity of its search function. Further implications are also discussed.

ACKNOWLEDGMENTS

I wish to thank Mike Byrne for expressing interest in my research ideas and for his willingness to be my advisor. This dissertation would not have been possible without his support and guidance. I am also grateful to my committee for their helpful comments. Last but definitely not least, I wish to thank my friends and family for their constant support throughout my time at Rice.

TABLE OF CONTENTS

1. Introduc	ction1
1.1.	Factors that Affect Site Search / Menu Traversal Behavior9
	1.1.1. The Cost-Benefit Perspective
	1.1.2. The Attentional Capture Perspective24
1.2.	The Experiments
2. Experim	ent 1
2.1.	Method42
	2.1.1. Participants
	2.1.2. Materials
	2.1.3. Design
	2.1.4. Procedure
2.2.	Results51
	2.2.1. Comparison of Behavior Between Sets
	2.2.2. Comparison of General Versus Specific Items63
	2.2.3. Questionnaire Analyses
	2.2.4. Probability Distributions

	2.2.5. Frequency and Time Analyses	74
	2.2.6. Search Rates Across Participants	87
2.3.	Discussion	89
3. Experim	nent 2	102
3.1.	Method	105
	3.1.1. Participants	105
	3.1.2. Materials	105
	3.1.3. Design	108
	3.1.4. Procedure	115
3.2.	Results	117
	3.2.1. Overall Analysis Excluding Item Specificity	121
	3.2.2. Overall Analysis Including Item Specificity and	
	Excluding Search Prominence	125
	3.2.3. Analysis of General Search Predisposition of	
	Participants	131
3.3.	Discussion	138
4. General	Discussion	144

4.1.	Search Behavior and Prevalence147
4.2.	Efficiency of Searching Versus Browsing148
4.3.	The Cost-Benefit Perspective149
	4.3.1. Menu Breadth151
	4.3.2. Information Scent155
4.4.	Attentional Capture Perspective158
4.5.	Item Specificity164
4.6.	The Interplay Between Cost-Benefit and Attentional
	Capture Factors
4.7.	General Search Predisposition of Users168
4.8.	Design Recommendations171
4.9.	Conclusions174
References.	
Appendix A	: Selected Sites Used in Experiment 1
Appendix B	: Decision and User Profile Questionnaires
Appendix C	: Predictors Used in Logistic Regression and
	Corresponding Questionnaire Items

LIST OF TABLES

Table 1: Search frequency and accuracy by set in Experiment 1
Table 2: Logistic regression output from Experiment 1
Table 3: Analysis of qualitative responses from Experiment 170
Table 4: Frequency and accuracy of all browsing and searching
attempts in Experiment 179
Table 5: Switching behavior in Experiment 1
Table 6: Sites used in Experiment 2106

LIST OF FIGURES

Figure 1a: Distribution of searching behavior by site in Experiment 1
(Histogram)54
Figure 1b: Distribution of searching behavior by site in Experiment 1
(Box plot)55
Figure 2: Search rate by serial position of sites in Experiment 160
Figure 3a: Distribution of searching behavior by site and item
specificity in Experiment 1 (Histograms)64
Figure 3b: Distribution of searching behavior by site and item
specificity in Experiment 1 (Box plots)65
Figure 4: Probability distributions of time estimates for
Bluelight.com in Experiment 173
Figure 5: Frequency and accuracy of (all) browsing and searching
attempts in Experiment 177
Figure 6: Mean attempt duration as a function of location method and
accuracy in Experiment 178

Figure 7a: Distribution of searching behavior by participant in
Experiment 1 (Histogram)88
Figure 7b: Distribution of searching behavior by participant in
Experiment 1 (Box plot)
Figure 8a: Distribution of searching behavior by site in Experiment 2
(Histogram)118
Figure 8b: Distribution of searching behavior by site in Experiment 2
(Box plot)119
Figure 9a: Distribution of searching behavior by site and item
specificity in Experiment 2 (Histograms)120
Figure 9b: Distribution of searching behavior by site and item
specificity in Experiment 2 (Box plots)120
Figure 10: Main effects from Experiment 2123
Figure 11: Search prominence x menu breadth interaction in
Experiment 2124
Figure 12: Menu breadth x information scent interaction in
Experiment 2125

Figure 13: Information scent x item specificity interaction in
Experiment 2127
Figure 14: Menu prominence x information scent interaction for
specific items only (in Experiment 2)128
Figure 15: Menu prominence x information scent interaction for
general items only (in Experiment 2)129
Figure 16: Menu breadth x information scent interaction for
specific items only (in Experiment 2)130
Figure 17: Menu breadth x information scent interaction for
general items only (in Experiment 2)131
Figure 18: Search predisposition index as a function of participant in
Experiment 2133
Figure 19: Information scent x item specificity interaction for
browsers only (in Experiment 2)136
Figure 20: Information scent x item specificity interaction for
searchers only (in Experiment 2)137

1. INTRODUCTION

Five years is generally an insufficient incubation period for any new service or application to mature, but in such time the Internet has grown from a relatively little used service to a staple of popular culture. Although companies exist that track the number of online users and Web sites (such as MediaMetrix), there is relatively little data available regarding the behavior of users who interact with these sites. With the Web becoming an important part of the world economy and the lives of individuals, it is imperative to understand how Web site users view and exploit Internet sites. An example in the world of e-commerce will illustrate this point.

Users who choose to visit an Internet retailer are often bombarded with literally dozens of products on the site home page. The assumption held by many e-commerce site developers is that users will painstakingly scan the product lists until they find a product that interests them. However, an Ernst & Young survey revealed that a major flaw of e-commerce sites during the 1999 holiday shopping season was the inability of users to find desired products (Ernst & Young, Press Release, January 3, 2000; see also Lohse &

Spiller, 1998a), an ability highly valued by online users (The Conference Board, Inc., 1996). Unfortunately for those sites, it is becoming increasingly clear that users value time savings and convenience very highly while shopping (Tilson et al., 1998; The Conference Board, Inc., 1996; Donthu & Garcia, 1999), and they are much more likely to shop elsewhere than wade through dense product lists (Bellman, Lohse, & Johnson, 1999; Lohse & Spiller, 1998a). Despite these suggestive findings, however, the lack of any conclusive research on the best way to maximize user exposure to products has left the e-commerce community with the preferred strategy of maximizing the number of products presented to the user on a site's home page.

Factors that determine exposure to products online are plentiful, such as weekly time spent shopping online or weekly time spent surfing the Web (see Donthu & Garcia, 1999). However, an arguably key set of factors influencing exposure to products online deals with the *way* products are presented on e-commerce sites. Specifically, the way products are presented via the Web site design can potentially play a major role in the interaction between the site and the user. For instance, the use of interface controls (such as dropdown menus) in an inappropriate way has been shown to result in unnecessary extraneous steps for a user to complete a task (Tilson et al., 1998).

As a (perhaps unintentional) solution to the problem of users being unable to find sought after products online, many commercial sites now feature a search function so that users can find the products they seek without wasting valuable time traversing the menus for a site (potentially replete with products). ("Menu Traversal" (MT) is hereby defined as product finding behavior of a user *independent* of available search functions that enable location of a product through the specification of user-defined criteria. "Site Search" (SS) is hereby defined as product finding behavior through use of a search function that enables location of a product through the specification of user-defined criteria.) In fact, installation (or improvement) of a search function is a common recommendation for increasing the likelihood that a shopper will find what she seeks (Lohse &

Spiller, 1998a), and it has even been argued that search functions should be mandatory for all large Web sites (Lohse & Spiller, 1998b).

One researcher has suggested that sites with regularly formatted data, complex-to-digest data, or sites that contain more than 100 pages should include a local search engine (Powell, 2000). Such a recommendation is sensible in that sites with hundreds if not thousands of available documents would be placing a tremendous burden on users if a search function were not available (Sano, 1996). Arguably, the main advantage of a search function is that it delivers specific information *to the user* rather than leaving the user to browse coarse-grained hierarchies to locate it (Sano, 1996).

Furthermore, the types of visitors and tasks performed at the site influence the need for a search engine (Powell, 2000). Search is commonly used by power users who wish to bypass hierarchical navigation schemes, and frequent visitors to a site may wish to quickly locate items that they already know to exist (Powell, 2000).

In fact, searching occurs for a variety of reasons. Users may be looking for something known to exist, checking to see whether it exists, or simply exploring to determine the extent of something (Powell, 2000). Furthermore, users may consider a search function to be a shortcut, or alternatively, search may give users the illusion of control as they can specify what they want (Scanlon, 2000). Finally, users may consider searching to be a last resort after link-following has failed (Scanlon, 2000). Although a search function is usually meant to serve as a supplemental aid to MT for a site, for many sites it has become nothing less than a crutch for the exasperated online shopper. Why should a user bother to browse a dense site to find a product if she can simply type the product name into a search field on the company home page?

Although such a reliance on the search function may actually increase the likelihood that a user will locate the product originally sought, use of SS (in lieu of MT) may lead to decreased exposure in the long run for products available on the site. Employing SS does nothing less than prevent the user from being exposed to a multitude of products available throughout the Web site. Therefore, it is important to determine factors that influence the occurrence of SS and MT behavior online.

SS and MT behavior is not only relevant to e-commerce sites, but to academic and personal arenas as well. The decision to employ SS or MT arises in the environments of technical support, product information, university Web sites, and educational sites (i.e., courseware). Understanding the factors that determine how a user behaves on a particular site can have far-reaching implications. If such behavior can be modeled and predicted (based on user-specific, task-specific, and design-specific factors), site designers will have a substantial influence over how their site is used. A major problem in the usability of Web sites, as any usability engineer will attest, is the frequent discord between the way a site is intended to be used and the way it actually is used by the population with which it interacts.

Imagine a new Internet company named TopShopEquipment.com, as an example, that sells plumbing supplies with an e-commerce storefront. Being a startup company, TopShopEquipment has little financial resources available to design and develop its Web site. Nonetheless, company management would like potential customers to explore the many different types of products they have to offer, from toilet snakes to pipe wrenches, so that they can maximize their sales. Furthermore, they would like to cater to amateur plumbers in addition to professionals, so they would like to make it easy for a customer to find a product that may be somewhat unfamiliar to him or her.

How should TopShopEquipment.com design their Web site? Should they present a multitude of categories on their home page to encourage users to explore the many different products they have to offer? Alternatively, should they just present as many products as possible on the home page itself with the hope that the user will notice what he is looking for? Perhaps a search function should also be implemented for the site so that the customer can simply enter the name of the product that he seeks into a text field in case he happens not to know what category the item would be listed under in the product menu? But what if the customer knows what a product is used for but does not know its name?

Clearly, the design process of a Web site, particularly an e-commerce Web site, is a complicated matter. As discussed above, previous research has suggested that users will not be willing to spend a great deal of time attempting to find a product on an Internet storefront, as they value convenience very highly. Therefore, if a customer knows what type of product he needs but is unfamiliar with its name, then an online storefront that encourages users to use SS to find items will not be helpful for this customer. Similarly, if an office clerk unfamiliar with plumbing supplies is required to order an item with a specific name, a site that downplays its search function in favor of MT will be less than optimal.

Understanding and predicting how an Internet user will interact with an e-commerce site (and Web sites in general) is thus a critical issue if an online company manager with products to sell or an individual with information to disseminate wants to reach the largest possible audience. If the management of TopShopEquipment.com, for example, knows that their primary customer base is amateur plumbers who may be unfamiliar with the names of many plumbing supplies, then it is in their best interest to design a site that will encourage a user to use the product menus (i.e., use MT) rather than a search field (i.e., SS). It is the aim of this set of studies to determine if such behavior can be predicted and manipulated through Web site design factors, and two theoretical perspectives will be examined that can potentially account for the decision process of the user who must choose between SS and MT.

1.1. Factors that Affect Site Search / Menu Traversal Behavior

1.1.1. The Cost-Benefit Perspective

Decision-making behavior has been investigated extensively for many years (e.g., Payne et al., 1993; Hastie, 1991; Kleinmuntz, 1991). One approach to explaining human decision-making is through a cognitive costbenefit analysis, involving the consideration of various positive and negative dimensions of alternative strategies for a task (Kleinmuntz & Schkade, 1993; Payne et al., 1993). This approach focuses primarily on two cost-benefit dimensions: the cognitive effort required to use a strategy and the ability of a strategy to produce an accurate response. Individuals select the particular strategy that represents the best accuracy-effort tradeoff for the task at hand, but as has been pointed out, all else being equal, people are motivated to use as little effort as necessary to solve a problem (Payne et al., 1993). In fact, people often behave in accordance with the principle of least effort (Zipf,

1949), in which a strategy is selected that ensures that the minimum effort will be involved in reaching a specific desired result. In the context of how individuals gather information, Information Foraging Theory (Pirolli & Card, 1999; Pirolli & Card, 1995) argues that people modify their strategies to maximize their rate of gaining valuable information while minimizing cost. In this context, "cost" refers to the costs of accessing, rendering, and interpreting information-bearing items (Pirolli & Card, 1999).

Information displays define a cognitive incentive system for decision makers, whereby displays influence the effort and accuracy of each available strategy, and therefore, induce decision makers to use different strategies. Unfortunately, information about effort and accuracy is often limited or unavailable. Therefore, the more appropriate construct is *anticipated* effort and accuracy (Kleinmuntz & Schkade, 1993). As a first-time user on an ecommerce Web site is unable to foresee all the necessary steps required to locate (or purchase) a product, the initial home page design of the site will play a disproportionately large role in the decision process of the user.

Anticipated effort, for a number of reasons such as immediacy and accuracy of outcome feedback over multiple experiences, will likely play a more prominent role in the decision process than anticipated accuracy (Kleinmuntz & Schkade, 1993), and myriad factors are present on a Web site which can influence the anticipated effort to the user of a particular strategy. One such factor, referred to as "information scent" (Pirolli, 1997), describes the amount of remote information a user can derive regarding the location of information based on the design or labeling of the information structure. Formally, it is defined as "the (imperfect) perception of the value, cost, or access path of information sources obtained from proximal cues" (Pirolli & Card, 1999, p. 10). Despite the obvious argument that retrieval tasks characterized by high information scent should lead to better performance than those with low information scent (Pirolli et al., 2000), the importance of information scent in information retrieval paradigms and the design of information browsers is only recently becoming fully appreciated. For instance, in a comparison of the Hyperbolic Tree browser with a more conventional browser, it was found that strong information scent made

hyperbolic search faster than a conventional browser while weak information scent had the reverse effect (Pirolli et al., 2000). Further, computational modeling of human information foraging has suggested that users' browsing choices are based on the evaluation of information scent (Pirolli & Card, 1999).

In the context of shopping online, information scent could potentially refer to the amount of information a user could attain regarding the location of a product in a site based solely on the design of the site's home page. For instance, category labels (serving as product headings) on a site can be more or less distinctive resulting in differing degrees of information scent (Larson & Czerwinski, 1998; Tilson et al., 1998), with distinctiveness referring to the semantic aspects of an alternative that enhances its difference from other alternatives in the set (Norman, 1991). As has been noted, selection from a menu of options involves many aspects of human information processing including decision-making (Norman, 1991), and information scent can have a major impact on user decisions. Presumably, when a user must select one of many options to locate specific information, she evaluates each option for

its subjective likelihood of supplying the answer (Norman, 1991). If the likelihood that an option will provide the correct answer is perceived to be low, then other options will be preferred.

Semantic similarity has been offered as a driving principle in prior attempts to explain and model novice exploratory behavior of software applications, specifically the decision process for selecting from a list of available options (see Soto, 1999). According to many such models, users attempt to find the best semantic match between the task description (or goal) and the available labels on display objects (such as menus and tool bars). For instance, it has been argued that a novel user to an application will tend to explore those menu labels that share one or more words with the experimenter-supplied description of the tasks (or with the user's goal) (Engelbeck, 1986, cited in Soto, 1999), a heuristic known as the "labelfollowing strategy" (Polson & Lewis, 1990). Unfortunately for interface designers, the necessity of using hierarchical menus (with broad top-level categories) often makes such close semantic matching difficult (Soto, 1999). However, as has been argued, a good interface should nonetheless guarantee

that the correct label is always the one with the highest semantic similarity among the available labels (Soto, 1999). It should be noted, though, that the assessment of the value of a cue in revealing characteristics of an information source is not independent of the value of other cues. Value assessment of cues proceeds in an interactive manner, with the value of one cue affected by the value of the others (Pirolli & Card, 1999).

Latent semantic analysis, a tool for measuring semantic similarity, has been used to provide an objective and reliable method of providing semantic similarity estimates (Deerwester et al., 1990; Landauer et al., 1998; Soto, 1999). Its utility was convincingly shown as it was used in one study to demonstrate that semantic similarity between the labels in the menus of a graphing application and the task descriptions provided by the experimenter (for that application) predicted the ease of discovering the solutions of tasks (Soto, 1999). It therefore appears that the issues of label distinctiveness and semantic similarity must both be considered in the context of implementing effective category labels on a Web site. (For the sake of convenience, the term "information scent" will be used to encompass both principles.)

To illustrate the problem of labeling, one pair of researchers noted that "cute or humorous department names do not provide cues as to where products are located" (Lohse & Spiller, 1998b, p. 85). An information structure with poor information scent due to overly general top-level categories has also been suggested to cause forgetting of the correct traversal paths to a target (Snowberry et al., 1983). Most importantly, however, if after assessing all alternatives of a menu and no option proves satisfactory, the user may feel limited and frustrated. At this point, the user may attempt to find the necessary information in a totally different manner (such as through SS) or she may choose to abandon the system altogether (Norman, 1991).

Presumably, information scent can play a large role in a user's decision to traverse the menus or use SS. A site with poor information scent (due to uninformative product categories) will result in the user anticipating MT to be a more costly strategy than SS. As a result, the user will choose to use SS even though her cost-benefit analysis may be inaccurate. For example, although intuitively less costly than MT, employing SS for a site is often a complex undertaking (e.g., requiring strict search parameters or syntax). Further, as has been pointed out, search engines have no inherent scent. Rather, they only give users the ability to create their own scent by typing in words they think they might find on the site (Scanlon, 2000).

Equally compelling as information scent to the user's cost-benefit analysis of a site is the breadth and depth of the product headings (i.e., product categories). ("Breadth" refers to the number of options at a level in a hierarchy, while "depth" refers to the number of levels in a hierarchy.) In fact, the two issues are related in that hierarchies with general top-level categories tend to require a greater number of lower levels. A review of breadth versus depth research suggested that breadth was favored over depth in most studies of the organization of menu contents (Larson & Czerwinski, 1998). With increased breadth, there are increased demands on visual search processes along with corollary increases in time required to make a decision. Conversely, with increased depth, there are more decisions required by the user, and perhaps more importantly, there is greater uncertainty as to the location of target items. Although a deep menu may require less processing

time than a broad menu for each level of the hierarchy, this advantage is often outweighed by the greater number of hierarchy levels in a deep menu (Norman, 1991). Empirical studies of information retrieval using various menu structures have found evidence that high breadth menus lead to faster and more accurate performance than high depth menus (Snowberry et al., 1983; Kiger, 1984).

The number of (top-level) product categories and their level of specificity can therefore significantly impact the user's anticipated cost of traversing the product hierarchy. Previous research has suggested that the depth of a hierarchical menu can significantly influence the perceived complexity of a menu retrieval task (Jacko & Salvendy, 1994), and users have been shown to prefer menu structures with limited depth (Larson & Czerwinski, 1998). Further, it has been argued that increased depth leads to participants forgetting the target or the correct traversal path to the target (Snowberry et al., 1983). As an example relevant to online retailers, a site that presents only two product categories on its home page, namely "Men's Clothes" and "Women's Clothes" will imply to the user that many more

levels of the product hierarchy will necessarily be covered before a sought after product is located. Conversely, a site with many top level product categories such as "Scarves," "Hats", "Blouses," and "Pants" implies to the user that relatively few levels of the hierarchy remain to be explored (see also Lohse & Spiller, 1998a). Further, it is not unreasonable to argue that users are cognizant of the additional decision-making burden carried by additional levels of a hierarchy. In fact, the user may be aware that each level of a hierarchy represents a unique bundle of information that needs to be recognized and scanned. Presumably, the user is also aware that with a large number of levels, it would be time-consuming to learn and become familiar with all of them (Norman, 1991).

Coherent categorization of items is also an important factor as it has been shown to increase the speed and accuracy of item location (Snowberry et al., 1983). Product categories that contain unusual or unexpected items may lead a user to doubt her cognitive model (or cognitive map, Chase, 1986) of the product categorization structure of the site, increasing her anticipated cost of navigating through the product hierarchy to find a product. As many have argued, well-designed systems promote an effective mental model for the user on how the system operates, and it has even been argued that the issue of depth versus breadth of menus is transcended by the importance of revealing menu *organization* to users (Norman, 1991). For example, if a site lists "Socks" and "Ties" as a sampling of products available in the category of "Intimate Apparel," a user may be more inclined to use SS than to traverse the menus.

Limiting the number of available options on a Web site to facilitate quick and accurate decisions has also received theoretical and empirical support, as exceeding the short-term memory limitations of users can have detrimental effects on user performance (Miller, 1956; Larson & Czerwinski, 1998). Regarding breadth and depth of menu structures, it has been argued that depth should be limited by increasing menu breadth up to the capacity of short-term memory set by the seven plus or minus two guideline put forth nearly a half-century ago (Miller, 1956; Kiger, 1984).

Aside from the organization of menu structure, the cost-benefit analysis of users is likely to be influenced by the sheer number of options

available to them on a site (see Larson & Czerwinski, 1998). Whether the options are in the form of product lists, links to other sites, available audio and video, or company information, overloading the capacity of human short-term memory will likely lead to an increase in the anticipated cost of choosing any one option. In fact, there is no shortage of sites on the Internet that have enormous information breadth, and it has been demonstrated that users tend to view information structures with enormous breadth as confusing (Larson & Czerwinski, 1998). Increases in the amount of available information, it has been argued, can lead to the problem of maximizing the allocation of attention to information that will be useful (Pirolli & Card, 1999).

If a user is faced with a plethora of options on a Web site, the confidence of the user in a chosen option being the correct one (to satisfy set goals) may decline given time constraints. It is important to remember that although less time-intensive a cognitive process than decision-making (Sternberg, 1966), visual search nonetheless has a time and effort cost associated with it. In a product finding scenario, a user may view SS as a much more cost-effective (and confident) choice than an attempt to use MT for a site with an overwhelming selection of options.

As support for such an argument, it has been found that decision makers respond to increases in the number of alternatives for a task by switching to simpler strategies (Payne et al., 1993; Payne, 1982). Specifically, when faced with many alternatives, people consider the available information more selectively. For example, one factor in the decision (such as time) may take precedence over other factors with the result that any alternative that does not meet the required criterion is excluded, regardless of how that alternative fares on other dimensions (see Tversky, 1972). In support of this notion, it has been argued that choice tasks tend to elicit qualitative types of reasoning strategies that focus on a single attribute (Tversky et al., 1988), and that such attribute-based processing is cognitively easier (Russo & Dosher, 1983).

Arguably, the main lure of online shopping is the reduction in costs of finding products and product-related information (Alba et al., 1997). However, selling products over the Internet does not guarantee such a reduction. Site design factors such as information scent will be the ultimate determinants of (anticipated) search costs and possibly online sales as well. As has been known for many years, the way information is presented on a visual interface has a tremendous impact on the user's ability to process the information, the time required to interpret the display (Tullis, 1981, 1988; Norman, 1991), and the amount of cognitive effort required to implement various decision strategies (Payne et al., 1993). In fact, one researcher found that graphical format differences accounted for a large proportion of the variance in information acquisition and evaluation (Jarvenpaa, 1989, 1990).

Furthermore, it has been argued that one cost of a poorly designed system is that it will not be used (Ledgard et al., 1981), suggesting that a poorly designed menu structure on an e-commerce site will be avoided in favor of SS. Supporting this argument was a study that found that the use of unit price information by shoppers (in a supermarket) increased when the information was brought together for shoppers in the form of organized lists (Russo, 1977). Simply put, information must not only be available but easy to process if it is to be used (Payne et al., 1993).

As an example of a critical interface design principle, many have advocated the need to establish or reinforce logical relationships among components by establishing spatial relationships within the design (Mullet & Sano, 1995; Norman, 1991). One means of creating such relationships is through alignment of visual elements, and the virtues of alignment have been touted by many designers and researchers (e.g., Mullet & Sano, 1995). Failure to align visual elements in the interface design is likely to increase the anticipated costs of using MT for a site, as meaningful relationships among the elements will be more difficult to determine. Furthermore, it has been argued that the physical format of a menu of options should highlight the options and organize them in a meaningful way to help in visual search (Norman, 1991). Generally speaking, it is reasonable to argue that an ecommerce site that is difficult to process will lead to high anticipated costs of using MT to find a product; the user will quickly find SS to be a more beneficial strategy.

1.1.2. The Attentional Capture Perspective

Another general approach to determining factors that affect searching behavior can be termed the "attentional capture" perspective. Properties of the visual environment have been argued to be major determinants of whether an object in the environment draws (or "captures") attention. For instance, unique objects in terms of color or brightness have been found to capture attention in a visual display, as have objects that lie on or near visual boundaries (Todd & Kramer, 1993).

This attentional capture perspective is directly related to the notion of "visual momentum" (Hochberg & Brooks, 1978). Visual momentum describes an observer's desire to gather information from her visual environment in an attempt to form a stable understanding of her surroundings. It has been argued that visual momentum contains both an early, rapid component during which attention is drawn to visual landmarks followed by a more cognitively-driven visual analysis or "inquiry" (Hochberg & Brooks, 1978; Hochberg & Gellman, 1977).
Using this theoretical framework, it is possible to argue that attentional capture (encompassing visual momentum) is a major determinant of user behavior on e-commerce sites. As has been argued, increases in the amount of available information lead to the problem of maximizing the allocation of human attention to information that will be useful (Pirolli & Card, 1999). Generally speaking, the importance of the prominence of display elements on choice behavior has been well-known in the decisionmaking literature. It has been argued that the assessment of task factors may be biased by a variety of information display variables that cause individuals to pay more attention to less important factors simply because they are more salient in the display (Payne et al., 1993). Regarding the use (or disuse) of SS on an e-commerce site, if the design of a site is such that attention is immediately drawn to the search function (e.g., through brightness contrast of the text field with the page background), then users may be more likely to use SS even though no information regarding the effectiveness of doing so is available. Conversely, if attention is immediately drawn to product headings (of a menu), then an MT strategy may be adopted. In terms of the visual

momentum approach, the search function (or product headings) may serve as a landmark which influences the user's subsequent visual inquiry and information acquisition.

Upon reaching the phase of cognitively-driven visual inquiry, a user will be likely to choose a strategy to the extent that she is aware of its existence. For instance, if an option (such as SS) is available only after scrolling of the page, it is imperative that the user be signaled to this requirement (through a well-organized and consistent design). As an example, one study reported that a user was unaware that she needed to scroll down a page on an e-commerce site to find the product order links (Tilson et al., 1998).

Similarly, the sequencing of interface elements such as the search function and menus will affect the likelihood that an option is selected. As has been argued, if a decision-maker evaluates alternatives in the order they are presented on a display, then the first acceptable alternative is more likely to be chosen over others (Kleinmuntz & Schkade, 1993). Under cases of time pressure, in particular, users will curtail their processing and select the first alternative that exceeds a preset criterion value (Beach & Mitchell, 1987). As an example, the American Airlines reservation system (known as "Sabre") listed American Airlines flights first, resulting in more bookings for American Airlines (in comparison to other airlines; Phillips & Thomas, 1988). Therefore, if an online retailer presents several options on the home page of a site (such as product category menus) prior to the search function, then the user will be less likely to consider using SS as she will likely have yet to learn of its presence (see Lohse, 1997).

Distractions present on an e-commerce site can also detrimentally affect the attentional capture of important elements such as menus or a search function (Lohse & Spiller, 1998b). Although it has been argued that information irrelevant to the user's needs clutters up the screen and makes it more difficult and time-consuming for the user to locate the option that she desires (Norman, 1991), current commercial Web sites often present advertisements with flashing graphics or scrolling text. It has been argued that such special effects can have an overpowering effect on human peripheral vision affecting the processing of information elsewhere on the page (Nielsen, 1996). Further, as has been pointed out, "When no perceptual cues are available to support...information acquisition, visual momentum is absent" (Woods, 1984, p. 233). Therefore, if the design of a site is such that the attention of the user is not drawn to any particular product-oriented location or feature, the user will be more likely to select the most obvious (and simplest) strategy to locate a product, namely SS.

Initial orientation to a new environment, as many have argued, is supported by landmarks (Chase, 1986; Vinson, 1999; Evans, 1980; Golledge et al., 1985). Whether the landmarks are direct representations of a natural scene (such as a street corner) or a more abstract element existing as part of a Web site interface (e.g., a menu of options), they can exert a non-trivial influence on participant behavior. Relevant to the notion of attentional capture, the behavior of an e-commerce site user will arguably be heavily influenced by the existence of elements that she perceives as landmarks. If a search function is perceived as a landmark by the user due to brightness contrast with the background of the page, for example, then the user may perceive alternative strategies for finding a product to be subordinate (such

as traversing the product hierarchy). Conversely, if a menu of product categories captures attention and serves as the landmark, then use of SS may be subordinated. Where the user's attention is drawn may play a critical role in the decision process of the user. In fact, distractions on Web sites (discussed above) such as elaborate graphics with realistic detail can serve as spurious landmarks leaving the user wondering what action to take (Lohse & Spiller, 1998b).

It has been argued that users begin to look for information on a Web page by scanning the page trying to find the words they're looking for, referred to as "trigger words" (Scanlon, 2000). If users are unsuccessful at locating trigger words or synonyms, it is argued that many will use SS. While such an argument is not inconsistent with the attentional capture viewpoint, it is viewed as incomplete. While users may indeed be scanning a Web page for trigger words, attentional capture dictates that the scanning of the page is to a large extent pre-determined by the design of the site. As a result, the design of a Web page has a major impact on whether and how quickly a user will find the desired information. Furthermore, the attentional capture perspective predicts that users will often terminate their scanning for trigger words once their attention is drawn to an option on the page that is viewed to be highly cost-effective, even if more cost-effective trigger words are present.

1.2. The Experiments

Research with hypertext systems that predate the (present day) World Wide Web has attempted to address the search versus browse relationship albeit for particular information retrieval systems (e.g., Rada & Murphy, 1992; Campagnoni & Ehrlich, 1989; Marchionini, 1989). One such study (Rada & Murphy, 1992) tested the accuracy and speed of novice and expert users on various hypertext book systems (such as "Guide," "HyperTies," and "MaxiBook") when asked to perform various "search" or "browse" tasks. In their study, a task was defined as a search task if the answer were contained in one paragraph of the book while a browse task required looking at several parts of the book to locate the answer. The results of the study suggested that different hypertext systems were better suited to either search or browse tasks. For instance, the MaxiBook hypertext system was found to be superior to the Guide system for novices performing browse tasks due to the advantage of being able to enter a search command for any word in the book. Upon doing so, occurrences of that word were displayed in the book outline encouraging the user to visit more than one location in the book, a requirement for locating the appropriate answer.

Much research on decision making has examined how decision makers choose among available well-known strategies in a choice task, such as the lexicographic, elimination by aspects, or satisficing heuristics; see Payne et al. (1993) for discussions of these strategies. Research along these lines has even been conducted with e-commerce sites investigating user decision making processes for a Web-based product selection task (Fasolo & McClelland, 1999). As few studies have investigated how users locate information on e-commerce sites, this study will be discussed in detail.

In this study, users were presented with products and corresponding features (e.g., computers with different processors, varying amounts of memory, etc.) in a matrix form with products along the columns and features along the rows. Participants were asked to select a product under different circumstances in which number of products and corresponding features were manipulated. The goal of the study was to determine how users locate information in this environment. For example, do they seek information more by product or by feature? To investigate such behavior, the information in a cell of the matrix of products and features was not visible unless the mouse pointer was placed on that cell.

They concluded that the number of products available (in the matrix) had a substantial impact on the decision process. With more products, information was more difficult to explore and the information location process was less consistent leading to reduced accuracy. A greater number of available products also led to a smaller proportion of information (i.e., features) that was sought per product. Also, the number of features affected the decision process as well. With more features, people made more accurate choices, regardless of the number of products. Further, with more time available, decision makers paid more attention to features. Based on their findings, they argued that matrices of products and features should be designed with more features and fewer products, and that users should be

given the option of determining the most important feature and ordering the products along that feature. In their view, Payne et al.'s (1993) "adaptive decision maker hypothesis" was found to be valid on the Web, namely that decision makers appear to be very efficient "jugglers" of effort and accuracy.

While informative to how decision makers process comparison tables on the Web, this research provides little insight into users' searching and browsing behavior outside of matrices considering the e-commerce sites in their entirety. Furthermore, although Fasolo and McClelland concluded that users probably utilize an elimination-by-feature or lexicographic strategy in the matrix task, it is highly unlikely that such a simple and straightforward strategy would be suitable to characterize the user decision making process in deciding to search or browse to find a product on a site. Therefore, the goal of the current experiments was to determine the factors that influence product finding behavior based on original product finding data. Whereas a great deal of prior research has focused on determining the strategy used by individuals to accomplish certain tasks, the current research project aimed to take more of a bottom-up approach to determining the factors involved in the decision to search or browse a site.

Experiment 1 investigated the relative occurrence rates of SS and MT behavior on a sample of actual e-commerce sites in an attempt to characterize typical product finding behavior for online retailers. In addition, an attempt was made to understand the factors taken into consideration by users in deciding to use SS versus MT. Experiment 2 manipulated key factors in a controlled experimental setting to complement the results of Experiment 1. Specifically, Experiment 2 determined whether SS and MT rates could be manipulated as a function of changes in interface design.

As stressed earlier, the anticipated cost and benefit of using a particular strategy to find a product is at the heart of user product finding behavior. Such assessments can be further understood in terms of a series of questions (Payne et al., 1993): What cues are used to make the assessments of likely accuracy and likely effort? How are those cues combined into perceptions of anticipated accuracy and effort? What biases exist in such judgments? The experiments sought to find answers for these questions.

Aside from interface design (or more generally, properties of the decision task (Payne et al., 1993)), many different factors influence SS and MT behavior. Task-specific factors such as goal-directed versus exploratory behavior (e.g., looking for a specific product versus any product that meets certain requirements (Steiger et al., 1998), similar to the distinction of simple versus complex retrieval tasks (Pirolli et al., 2000)) are accompanied by a multitude of user-specific traits that likely play an enormous role on product finding behavior and online shopping behavior in general. Regarding task-specific factors, the intuitive prediction that specific items (or other unit of information) are more likely to elicit searching behavior and general items are more likely to elicit browsing behavior is not new. Such an hypothesis was put forth in earlier hypertext research (e.g., Campagnoni & Ehrlich, 1989).

Generally speaking, users vary in their repertoire of generating solutions, their ability and willingness to plan ahead, and the degree with which they will pursue a particular course of action before they quit (Norman, 1991). Furthermore, prior task knowledge and expertise in a problem domain can influence how information is processed (Alba & Hutchinson, 1987; Chi et al., 1988). For instance, experience in a problem domain may impact the frequency and recency with which available strategies have been used, thus affecting the accessibility of various strategies (Payne et al., 1993). As an example, users who are adept at using SS may be more inclined to use it in product finding scenarios. Similarly, users already familiar with a particular site may be more likely to use MT to find a product as they will be knowledgeable of the site's organization. However, the reverse may also be found if the user has decided based on previous exposure that using MT is not productive for that particular site.

Similarly, people with prior experience in dealing with a problem may simply draw a solution from memory that is based on prior evaluations of the alternatives (Payne et al., 1993); no processing of the characteristics of the current problem occurs, a strategy known as "affect referral" (Wright, 1975). Finally, an even simpler strategy known as the "habitual heuristic" may be used that encourages individuals to choose what he or she chose last time (Payne et al., 1993). Specific to shopping, numerous user-specific classifications of buying behavior have been outlined by previous research including rational, habitual, emotional, and social buyers (Steiger et al., 1998). Further, user perception of e-commerce site search functions themselves has also been reported, often with less than favorable results (Jarvenpaa & Todd, 1997). The focus of the current study was to take all of the potential factors into account in an attempt to understand product finding behavior.

Switching behavior in which a user begins to look for the product using SS but then decides to use MT or vice-versa was also of interest to the study. This issue is similar to between-patch versus within-patch foraging issues discussed in the Information Foraging literature (Pirolli & Card, 1999). Within-patch foraging refers to the gathering of information from a single source whereas between-patch foraging refers to the decision to use other sources. Relevant models assume that information foragers allocate their time to between-patch versus within-patch foraging activities with the goal of optimizing information gain while minimizing cost. Such models suggest that users will continue to look for a product using the method with which they started as long as they consider the continued use of the method to be more cost-effective than the act of switching to the alternative method. For example, if a user begins with SS but finds multiple attempts to find the product fruitless, she may decide that it is more cost-effective to explore the menu structure than attempt another revision of the search parameters.

2. EXPERIMENT 1:

The Relative Occurrence Rates of Site Search / Menu Traversal Behavior

The first step in understanding online product finding behavior was to investigate SS and MT behavior with existing corporate Web sites. Although it has been reported that MT behavior is slightly more common than SS behavior on some systems (e.g., an Internet phone number directory; Neal, 1995), it was important to take a fresh look at this issue as it pertained to ecommerce sites. Despite limited (and often non-existent) empirical data on the topic, many researchers are willing to make claims regarding search behavior and prevalence. For example, Spool et al. (1997) claimed that approximately one-third of users searched as their initial strategy at finding information on sites. More recently, Scanlon (2000) claimed that users rely on the search engine approximately half of the time, while Nielsen (2000) provided a more specific breakdown of user types.

According to Nielsen, roughly 50% of users are search-dominant users, about 20% are link-dominant, and the remainder exhibit mixed behavior. Search-dominant users, Nielsen claims, usually go straight for the search function when they enter a site as they are interested in finding specific information as fast as possible. Link-dominant users, on the other hand, will attempt to find information by following links on the homepage of a site. Finally, mixed-behavior users switch between searching and browsing depending on what seems to be the most efficient information-finding strategy at the time.

Despite the many claims made regarding the motivation for a search function, the reasons behind search function use, and the existence of patterns of searching behavior, there is surprisingly little data used to back these claims. Many researchers make statements such as, "My usability studies show" (Nielsen, 2000, p. 224) yet empirical data is rarely (if ever) presented. For example, one researcher supports her claim on search engine usage by stating, "Users seem to go to a search engine about half the time, at least in our tests" (Scanlon, 2000, p. 2). It is unfortunate that claims of this kind have garnered such credibility without firm empirical backing. In fact, many (if not most) of the claims made regarding searching and browsing are often taken as common sense or are general impressions formed over years of watching users on the Web. The lack of publicly-available data on this issue makes it impossible to generate justifiable guidelines for design; empirical studies are clearly warranted here.

For instance, it is widely accepted that searching is a faster means of finding information than browsing (e.g., Nielsen, 2000; Powell, 2000). Many of us can easily identify with entering a product name into a search field on an e-commerce site and getting a fast and accurate listing of products. However, is this indeed a rule of the Web? Are there cases where such a pattern does not hold? There is some evidence that such cases exist. During informal usability testing, Scanlon (2000) found that users who were successful using local search functions took more clicks to find their answer than those who followed links. Clearly, the superiority of search should not simply be taken for granted.

Furthermore, can users truly be characterized as search-dominant or link-dominant? Can situations (or sites) be found where the search-dominant user will consistently browse and vice-versa? If so, then what is it about such situations that leads to the change in behavior? It has been claimed that link-dominant users will use a search function "only when they get hopelessly lost" (Nielsen, 2000, p. 224) and that mixed-behavior users switch strategies depending on "what seems most promising to them at any given time" (Nielsen, 2000, p. 224). What is meant by "hopelessly lost" and what determines whether a strategy is "promising"? These are hardly trivial questions (see Sawyer, 2000 for one approach to this question).

Experiment 1 was intended to provide empirical data on many of these issues using existing e-commerce sites. Without such data, it is inappropriate to base design decisions and our understanding of Web usage in general on such weakly supported claims, regardless of the intuitive appeal of the claims. While documenting the relative occurrence rates of SS and MT behavior, questionnaires were used to get at the heart of the user decision process.

2.1. Method

2.1.1. Participants

Twenty Rice University undergraduate students participated to fulfill a course requirement. As part of the "decision" questionnaire administered for each site in set 2 (see Materials below), participants were asked whether they had used the experimental sites previously. They reported having previously used a mean of 1.5% of the sites.

Following the experiment, participants were asked to fill out the "user profile" questionnaire (see Materials below) to collect data regarding prior Web usage among other factors. Of the twenty participants, all but two reported more than five years of general computer usage, and all but one reported having used the Web at least two hours per week for more than one year. Regarding current Web usage, all but two participants reported at least five hours per week. Sixteen participants reported purchasing items online in the past, with eleven reporting having purchased at least five items online.

2.1.2. Materials

Twenty e-commerce sites were selected by the author using the criterion that they would be predicted to be highly effective at eliciting a range of SS and MT behavior. Such an assessment was made informally by the author based on the cost-benefit and attentional capture principles mentioned earlier. Any site that did not have both MT and SS capability was excluded from the pool of eligible sites. (See Appendix A for selected sites used in Experiment 1.)

For each site, the experimenter selected a broad sample of items available on the site for use in the study. Although participants were aware that items would always be present on the sites, there is no a priori reason to expect this to influence product location behavior. The number of items selected equaled the number of participants in the study (20) so that each participant was asked to find a different item. This requirement was meant to minimize the effect of the items chosen on user behavior; we were interested in user behavior with particular sites, not particular items. Furthermore, as the sites used covered a broad range of industries, it was not possible to ask the participants to locate the same items for each site. In fact, it is argued that such a strategy may have produced a confound as participants may have been influenced by previous success rates at locating a particular item using a particular method.

Besides being a broad sample of available products on each site, a concerted effort was made to select items that were not presented on the home page of the site at the time they were selected. This effort was made to attempt to minimize the likelihood that a participant would notice the item on the home page obviating the need to either traverse the menus or use SS.

To investigate the influence of task-specific factors such as goaldirected versus exploratory behavior (discussed above), 10 of the 20 items for each site met the requirement that only one item existed for that item specification (e.g., "Air Storm GPT Softball bat") while the other 10 item specifications were satisfied by a maximum of 20 items available on the site (e.g., "Softball bat"). The two sets of 10 were matched pairs as in the example so that the "specific" item would satisfy the "general" item requirement (i.e., "Air Storm GPT Softball bat" is one of the available "Softball bats").

Two questionnaires were also developed for use in the second set of stimuli given to each participant (see Procedure below). The first questionnaire, the "decision" questionnaire, consisted of 20 questions intended to reveal the identity and role of important factors involved in the decision process; the questions were based on the cost-benefit and attentional capture factors discussed earlier. Questions were developed with the aim of eliciting both interpretations of available information as well as predictions from participants as to what they expected to see once they proceeded. Furthermore, two versions of the questionnaire were created to accommodate participants who chose to use SS and those who chose to use MT. The second questionnaire, a "user profile" questionnaire, was used to obtain general user data regarding Web experience and typical online behavior patterns (see Appendix B for the "decision" and "user profile" questionnaires).

A Sony DCR-TRV10 Digital Camcorder was used to videotape the experimental sessions and audio was captured via a Sima Camcorder Lapel Microphone. Each participant was tested individually on an Apple iMac personal computer running Microsoft Internet Explorer 4.5 with an Ethernet / T1 Internet connection. The digital video (with audio) from each participant was captured onto an Apple G4 computer and analyzed with the task analysis software MacSHAPA (Sanderson et al., 1994).

2.1.3. Design

20 sites were used in Experiment 1. Each participant received two sets of sites, each set containing 10 unique sites, with set 1 always preceding set 2. For half of the participants, the assignment of sites into the two sets was constant. For the remaining half of participants, the assignment of sites to the two sets was reversed. In other words the sites used in set 1 for half of the participants were the sites used in set 2 for the other half. This counterbalancing scheme was necessary as a different procedure was used for set 1 and set 2 (see Procedure below). Each participant received 20 unique items, one unique item for each of the 20 sites, with the following requirements. Each participant received five general and five specific items per set, and for each half of participants (over all participants in that half), half of the items sought for each site were general and half specific.

2.1.4. Procedure

The sites were presented in a random order to each participant, and the item to be found on each site was randomly assigned to each participant (in accordance with the above requirements). The items to be found by the participants were presented one at a time on separate sheets of paper (one item on one sheet for each individual site), and participants were instructed that the purpose of the experiment was to investigate how people find products on Internet commerce sites. For each individual item, the experimenter stated, "You're looking for this item," as he placed the sheet of paper with the printed item in front of the participant. The remaining procedure was distinct for each stimulus set. For the first set of 10 sites, the participant located the item presented to her and wrote down the price of the product on the sheet of paper used to present the item. The experimenter recorded whether the user used MT or SS for the site as this was the dependent variable of primary interest. If both behaviors were exhibited, the experimenter noted the sequence of participant behaviors. Finally, if the participant wished to cease looking for a particular item, the experiment was continued with the next item.

For the second set of 10 sites, the participant was shown the item to be located on the sheet of paper as in the first set. The participant then wrote down on the sheet of paper used to present the item what she decided to do to locate the product. As an example, she might have written, "I would click on the 'Men's Clothes' menu option." Then the decision questionnaire was administered to the participant. During administration of the questionnaire, the site was left visible and the participant was permitted to scroll (but not click on) the site. Upon completion of the questionnaire, the participant was instructed to actually locate the item and provide the price information (as in the first set). Following this cycle of events, the next site was loaded by the experimenter and the same cycle of events was conducted. Following both sets of stimuli, the user profile questionnaire was administered to the participant and she was then debriefed and thanked. By administering the decision questionnaire in set 2 for each participant, the product finding behavior found in set 2 could be compared with that of set 1 to determine if there were any impact on natural SS and MT behavior due to the introduction of the questionnaire.

Participants were permitted to actually find the products not only so that product finding success rates could be determined, but also so that continuity could be maintained. As mentioned above, it also permitted the measurement of the impact on product finding behavior due to the questionnaire. Furthermore, allowing participants to find the products permitted the investigation of both information scent-finding and information scent-following behavior (Pirolli et al., 2000). Information scent-finding behavior refers to the decision made at the top-level menu, while information scent-following behavior refers to the following of a specific path in the menu structure. As mentioned above, switching behavior (from SS to MT or vice-versa) could also be investigated by permitting participants to find the products.

In the interest of gaining further insight into potential factors influencing SS and MT behavior, the participant was instructed to "think aloud" while looking for the items and the entire experimental session was videotaped (with audio). It was hoped that video and audio analysis would yield valuable data concerning the decision processes of the participants and the factors that influenced them. Think aloud instructions presented to participants were based on those relevant to concurrent verbalization (and not retrospective reports) from Ericsson and Simon (1993, p. 378). To familiarize participants with the process of thinking aloud, a verbal protocol practice task was administered consisting of the question, "How many windows are there in your parents' house?" (see Ericsson & Simon, 1993).

Furthermore, the experimenter verbally prompted participants to keep them on task when they became silent or if they began verbalizing irrelevant information. Each participant was reminded to think aloud if she remained silent for approximately 15 seconds (Soto, 1999). To keep participants focused on the task at hand, the question, "What are you thinking now?" was asked to them when necessary. As a procedural note, due to the potential influence of Web site loading behavior on participant impressions, each participant was asked to turn her head away from the screen until the site was fully loaded by the experimenter. For example, as a site loaded, the search function text field may have been one of the first graphical elements of the page to load, so the prolonged exposure to the search field absent other page elements may have biased participant behavior. (Although this may be an intentional aspect of page design instituted by developers, it was not a focus of the current research.)

As mentioned earlier, the main dependent variable of interest was whether the participants' initial attempt at locating each product was via MT or SS. Success rate data (for locating products) and time to complete the task were also recorded so that potential effects of SS and MT behavior on those dependent variables could be investigated. Participants were tested between March 23 and April 10, 2000.

2.2. Results

Participants were very accurate in locating items, with an overall accuracy rate of 94%. Due to such a high accuracy rate, such data will not be considered further. Five (1.3%) items were missing at the time of testing, five (1.3%) items were located on the homepage of the site, two (<1%) sites would not load, and one (<1%) site crashed in progress. For within-subjects statistical tests, missing data were replaced with the grand (search/browse) mean for that set (i.e., set 1 or set 2). Once the possibility of missing items was realized (due to site inventory changes), the presence of all future items was verified prior to further testing. As a result, 23 (5.8%) items were replaced with new items. For the two cases in which a site would not load, a replacement site was used (stressless.com) but the resulting data from that site was not included in any analyses due to the infrequent occurrence of such failures.

As can be seen in Figures 1a and 1b, a wide distribution of searching (i.e., SS) behavior was found for the 20 sites used (collapsed across set 1 and set 2). There was a highly significant main effect of site on search rates, F(19, 361) = 3.55, MSE = 0.185, p < 0.001, yet an attempt to cluster the sites based on prevalence of search usage yielded no clear-cut groupings. To locate items on Deerskin.com, participants utilized the search engine as the initial course of action only 5% of the time, while for Nike.com, they used the search engine 90% of the time. It is especially interesting to note that only four sites had a search rate greater than 50%, and the mean search rate across sites was 42%.



Figure 1a: Distribution of searching behavior by site in Experiment 1 (Histogram).



Figure 1b: Distribution of searching behavior by site in Experiment 1 (Box plot).

Deerskin.com was selected for use in the study for two reasons, namely it possessed a prominent search function in conjunction with very high menu breadth and specificity. It was predicted prior to conducting the experiment that the combination of these two factors would yield a high *search* rate as participants would be reluctant to read through the copious menu options in the presence of a highly available search function. Remarkably, however, participants were unfazed by the substantial menu breadth possibly due to the very high information scent associated with the options. Whether the participant was searching for a type of hat or pants, for instance, it was clear that the "Hat" or "Pants" menu option, respectively, would be an efficient strategy to locate the item. Supremevideo.com also showed a remarkable advantage of high information scent achieved through substantial menu breadth, albeit the search function was difficult to find on that site.

In contrast, Nike.com was selected for the simple reason that it possessed remarkably low self-evidence. Based on participants' comments and behavior, it is clear that they were generally unaware that the very prominent "STORE" heading just below the "Nike" heading was indeed a link to further menu options. The same held true for the product image links (e.g., an image of a sneaker) that would allow the participants to browse the Nike.com menus. Participants were thus left to use the search function ("product finder" on this site) as they felt they had no alternatives. Certain participants even remarked that there was really no product menu at all.

2.2.1. Comparison of Behavior Between Sets

Site	Set	Mean	Overall	Mean	Overall
		Search	Search	Correct	Correct
abccosmetics.com	1	0.70	0.63	0.90	0.95
	2	0.56		1.00	
bluefly.com	1	0.00	0.40	0.78	0.84
	2	0.80		0.90	
bluelight.com	1	0.10	0.40	1.00	0.95
	2	0.70		0.90	
computergear.com	1	0.13	0.41	0.90	0.95
	2	0.70		1.00	
crateandbarrel.com/store	1	0.20	0.40	1.00	0.95
	2	0.60		0.90	
deerskin.com	1	0.10	0.05	1.00	1.00
	2	0.00		1.00	
fogdog.com	1	0.20	0.20	0.90	0.90
	2	0.20		0.90	
furniture.com	1	0.60	0.50	0.70	0.75
	2	0.40		0.80	
intellesale com	1	0.30	0.30	1.00	0.95
	2	0.30	0.00	0.90	0.70

1 0 n anima and 1 . . 1 C

Items presented in set 2 yielded significantly more frequent searching

behavior than set 1 across participants, t(19) = 4.25, p < 0.001, see Table 1.

Site	Set	Mean	Overall	Mean	Overall
		Search	Search	Correct	Correct
jjill.com	1	0.50	0.40	1.00	1.00
(click on Online Store)					
	2	0.30		1.00	
nike.com	1	0.80	0.90	1.00	1.00
	2	1.00		1.00	
nordstrom.com	1	0.20	0.45	1.00	0.90
	2	0.70		0.80	
priorities.com	1	0.20	0.32	0.90	0.95
	2	0.44		1.00	
shopsports.com	1	0.56	0.48	0.78	0.84
	2	0.40		0.90	
smartbasics.com/products.htm	1	0.40	0.55	1.00	1.00
	2	0.70		1.00	
store.crabtree-evelyn.com	1	0.30	0.40	1.00	1.00
	2	0.50		1.00	
supremevideo.com	1	0.00	0.15	1.00	1.00
	2	0.30		1.00	
totalbodyandimage.com	1	0.30	0.30	0.80	0.90
(click on Go Shopping)					
	2	0.30		1.00	

Table 1 (continued): Search frequency and accuracy by set in Experiment 1.

Site	Set	Mean	Overall	Mean	Overall
		Search	Search	Correct	Correct
valueamerica.com	1	0.40	0.60	1.00	0.95
	2	0.80		0.90	
winona.com	1	0.33	0.47	1.00	0.95
(click on Shop Online)					
	2	0.60		0.90	

Table 1 (continued): Search frequency and accuracy by set in Experiment 1.

Although surprising, it suggests that the methodological differences between the two sets influenced searching behavior. However, the effect may be due to either the administration of a questionnaire for each item or fatigue as set 2 items always followed those of set 1. Support for a fatigue account was provided by a serial position analysis of search rates versus the order of presentation of a site to the participants (see Figure 2). As each participant received all twenty sites in random order, this analysis was performed to determine whether performance differed as a function of the order of presentation of the sites.



Figure 2: Search rate by serial position of sites in Experiment 1.

As shown in Figure 2, there was a significant linear relationship between search rate and serial position of the sites, F(1, 19) = 26.1, MSE = 0.206, p < 0.001, supporting the notion that fatigue led to increased searching behavior. Alternatively, some type of learning effect may be responsible for such a relationship, yet there is insufficient theoretical basis to make such a claim. Most importantly, the absence of a step function between the first ten and last ten sites strongly supports the notion that the change in methodology between set 1 and set 2, per se, was not responsible for the increase in
searching behavior between the sets. Finally, the aberrant serial position data for positions 6 and 20 are most likely due to a relatively high frequency of predominantly searched and browsed sites for each respective position (with browsing referring to MT behavior). In other words, it is possible that highly searched sites often appeared as the sixth site across participants by chance, and highly browsed sites often appeared as the twentieth site across participants by chance.

Based on analysis of the video (and audio) from the testing sessions, the difference between sets appears to be due to two main factors. First, participants were now cognizant (in set 2) of the fact that all sites possessed a search function. As a result, they would very often locate the search function on each site before making a decision on how to proceed. Compelling support for this latter hypothesis is the data for Computergear.com, a site with a prominent menu (but with very low information scent) and a very subtle search function. As can be seen in Table 1, participants used search only 13% of the time in set 1, whereas in set 2 search was used 70% of the time. It therefore seems reasonable to argue that participants were simply unaware of the search option in set 1 and were saddled with a very inefficient top-level menu. However, the same pattern holds for Bluelight.com, a site with better information scent and a very prominent search function.

Participants also became more analytical about the decision to search or browse in set 2. The presence of the questionnaires requiring detailed responses may have altered the participants' behavior by requiring greater thought on how to locate an item on each site. It is quite possible that participants acted according to their initial (immediate) evaluation in set 1 but became more analytical in set 2 out of the necessity to explain their actions in the questionnaires. Simply exposing participants to the concepts of prominence and information scent (indirectly) via the questionnaires may have led them to consider such factors in their decisions; factors that were not directly considered in set 1. One may thus draw the sensible conclusion that awareness of the search function and the more generally analytical approach of participants were both involved in the elevated set 2 search

rates, along with a fatigue or learning effect that influenced behavior as the experiment progressed.

Despite this significant difference between the sets, between-site differences remained fairly stable across sets, and the counterbalancing of sites across the two sets supports the use of collapsed (across sets) data for further analyses. Thus, the collapsed data will be used for further site analyses except where inappropriate or noted otherwise.

2.2.2. Comparison of General Versus Specific Items

Although more specific items were predicted to lead to a greater likelihood of searching behavior, such a pattern was not sufficiently consistent across sites to reach significance, t(19) = -1.07, p = 0.30, see Figures 3a and 3b.



Figure 3a: Distribution of searching behavior by site and item specificity in *Experiment 1 (Histograms)*.



Figure 3b: Distribution of searching behavior by site and item specificity in *Experiment 1 (Box plots).*

It appears that the varying levels of "generalness" for different items diminished the robustness of the effect. At the heart of the matter is *perceived* specificity rather than specificity as defined in the method. For example, the item, "Traditional Style Mahogany Bed" was classified as a "general" item as it could have been satisfied by more than one item (on Furniture.com), while "Women's Chanel Suit" was classified as a "specific" item as only one such item was available on the site (on Deerskin.com). The influence of specificity merits further investigation.

2.2.3. Questionnaire Analyses

Questionnaires were administered to participants for set 2 items following the initial decision to search or browse. These questionnaires (referred to earlier as "decision" questionnaires) sought to investigate an abundance of potential factors influencing the participants' selection of an initial product location strategy. Among the factors investigated were the predicted efficacy and prominence of the menu and search function. This data, along with response data from the user profile questionnaires were used as predictors in a logistic regression to determine which factors influenced the strategic decision on how best to locate an item.

The quantitative data collected from the questionnaire administered for each item (in set 2) was used to predict the likelihood of using search to find that item. To this end, a logistic regression was used to regress product location (i.e. SS / MT) behavior on the more than 50 variables collected from the questionnaires. Using stepwise predictor selection, a model was produced that yielded several significant predictors. Appendix C presents the predictors used in the regression and the questionnaire (or user profile) items used to create them. The resulting model's coefficients and significance levels are presented in Table 2.

Logistic Regression Predictor	Coefficient	p-Value
Number of menus expect to have to go through to find product	0.32	0.008
How long will it take to find products using menus	1.31	< 0.001
Considered time as a factor in searching vs. browsing	2.19	< 0.001
How clear as to what should type into search field	0.32	0.02
How noticeable is menu of options on page	- 0.48	< 0.001
How noticeable is search option on page	0.22	0.04
How many items would you expect search to return	- 0.04	0.02
How long will it take for search to find item	- 0.74	0.005
How often use search function to find products when shopping	4.99	< 0.001

Table 2: Logistic regression output from Experiment 1.

Based on the model, when a user draws any of the following

conclusions, she will be more likely to use the search function to locate an item: Using the menus will take a long period of time; I generally like to use search; time is important; menu options are subtle; it will take little time for the search function to produce the item; many menus will need to be traversed; I am clear as to what should be typed into the search field; search will return a manageable number of items; search function is prominent. Examples of factors that failed to reach significance were "whether there existed an appropriate number of menu options (on the home page)" and "the number of items expected to be found on the page containing the item."

Qualitative responses from the questionnaires were then sorted into meaningful categories and the frequencies of such responses were correlated with the search rates from set 2 by site. This way, data could be obtained regarding the factors that participants *believed* to be involved in their decision to search or browse. An example of a qualitative question from the decision questionnaire was "Did you consider time as a factor in choosing to use the Search function versus browse the menus? *If yes, in what way*? [italics added]" and a sample answer (to the follow-up question) was "Search engine is easier."

Appendix D presents the relevant questionnaire questions, the response categories developed, and a sample participant response from each category. Responses were placed into the most appropriate category (or categories) and the proportion of responses in each category for each site was tallied. So, for the above example, the response was sorted into a category labeled "Search is easier / faster / more direct." Then, the proportion of responses to the question "Did you consider time as a factor in choosing to use the Search function versus browse the menus? If yes, in what way?" that could be placed into the category "Search is easier / faster / more direct" was calculated for each site. For Valueamerica.com, the proportion of responses to that question that fit into that category was 30%. Finally, the response rates for each category (for each question) were correlated with search rates across sites.

As another example, for Crateandbarrel.com, the proportion of responses to the question, "Did you consider time as a factor in choosing to use the Search function versus browse the menus? If yes, in what way?" that were placed into the category "Unsure which menu to choose" was 40%. One such response (see Appendix D) was "I don't know exactly what this product is/does so it would be hard for me to know what menus to go to."

When the response rates were correlated with search rates (across sites), several meaningful correlations consistent with the quantitative

questionnaire findings were found (see Table 3): Comments classifying the search function as easier, faster or more direct than using the menus correlated positively with using search, r(19) = 0.55, p = 0.01; comments claiming a good match of item and menu label produced a significant negative correlation, r(19) = -0.57, p < 0.01; comments expressing doubt in the effectiveness or ease of use of the search function correlated negatively with frequency of search usage, r(19) = -0.51, p = 0.02; statements characterizing the menu options as vague or incomplete correlated positively with use of the search function, r(19) = 0.54, p = 0.01; finally, comments regarding the prominence of the menu, particularly due to high background contrast, were significantly negatively correlated with search usage, r(19) =-0.50, p = 0.02.

Factor Predicting Search Behavior p-Value r Search is easier / faster / more direct 0.55 0.01 Good match of item and menu - 0.57 0.009 Search may not work or will be confusing - 0.51 0.02 Menu options are vague / incomplete 0.54 0.01 Menu is highly noticeable due to good contrast with background - 0.50 0.02

Table 3: Analysis of qualitative responses from Experiment 1.

Together, the quantitative and qualitative questionnaire analyses provided converging evidence regarding the rationale behind users' decisions to search or browse to find products. Similar results were found pertaining to the factors that govern the decision (quantitative analysis) and the factors that users' perceive to be important to the decision (qualitative analysis of user comments).

It therefore appears that many factors reliably predict whether a user will use SS or MT to find an item on an e-commerce site. Many such factors are directly related to the graphic design of the site, while others are related to the perceived or inferred underlying information structure of the site. Additional factors are pertinent to the perceived or inferred efficacy of the search function itself, while a final set of factors represent the user's disposition and preferences. Such factors will be elaborated upon in the discussion of Experiment 1.

A regression analysis was also performed to determine the proportion of variance in product location behavior explained by user-specific and sitespecific factors. While logistic regression would have been a more appropriate analysis due to the binary dependent variable (search / browse), a multiple regression would provide proportions of variance explained by each factor. When search behavior was regressed on participant and site, the model revealed that 17.6% of the total variance in product location behavior was accounted for by the participant. Similarly, the site used accounted for 13.3% of the total variance in behavior.

2.2.4. Probability Distributions

Among the questions asked of participants on the questionnaires was a request to estimate the time it would take to locate the item using menus versus using the search function. To achieve this end, participants were instructed to provide probabilities corresponding to different time periods. For example, a participant might indicate that there existed a 20% chance that the item would be found using the search function within one to three minutes. Of course, such estimates were made by the participants with only the benefit of having seen the home page of the site (i.e., prior to any interaction with the site). Based on this data, probability distributions were created plotting the probabilities on the ordinate and times on the abscissa for both browsing and searching for each site. The data for the probability distributions were normalized so that the mean probabilities summed to 1 (which was not necessarily the case for the original data).

The distributions suggested that participants were clearly sensitive to the differential designs of the various sites (see Figure 4 for the probability distributions for bluelight.com).



Figure 4: Probability distributions of time estimates for Bluelight.com in Experiment 1.

To explore this finding further, the expected search and browse times were computed for each participant for each site (using the normalized subjective probability data). Then the difference between expected search and browse time was computed and that measure was included in the logistic regression performed earlier (used to predict searching behavior). The model yielded a coefficient of 0.49 for the measure, p = 0.09, suggesting that participants considered subjective time estimates as part of their decision to search or browse a site. However, additional factors were clearly involved.

2.2.5. Frequency and Time Analyses

The frequency of searching and browsing behavior and corresponding time spent on each activity was recorded and analyzed, and success rates were also noted. Unlike all previous Experiment 1 analyses of searching and browsing behavior which involved only the initial product location decision, the current analyses included any subsequent browsing or searching attempts for a particular item. By doing so, a more complete picture of product finding behavior could be attained. Therefore, the overall accuracy rate reported earlier (94%) does not apply to this section as that success rate measured the accuracy of product locations after all attempts.

For this analysis, a browse attempt was defined as an attempt to locate a product without the use of the search function, and the attempt ended when participants either found the item, quit looking for the item, switched strategies (and began using the search function), backtracked to a higher menu level, or otherwise made a clear change in course of action (such as selecting an unrelated item on the site home page). A search attempt was defined as an attempt to locate the item using the search function and such an attempt ended when participants either found the item, quit looking for the item, switched strategies (and began browsing the menus), made a new search attempt, or otherwise made a clear change in course of action.

It appears that the effectiveness of the search function can have a substantial impact on whether a user will find the product he seeks. However, if one looks at the global frequency and time data (collapsing across all sites; Figures 5 and 6 respectively), it is clear that the use of the search function in an effort to locate a product does not lead to a greater rate of success. Although it failed to reach significance, participants showed a higher success rate when using the menus to find items as opposed to search, $C^2(1, N = 692) = 2.98, p = 0.08$ (see Figure 5). Such a finding is remarkable in light of the widely held belief that searching is more accurate than browsing (see Nielsen, 2000; Powell, 2000). Furthermore, a breakdown of the successful attempts (which were always final attempts) revealed a greater frequency of successful browsing attempts than searching attempts (207 vs. 160, respectively). Successful Unsuccessful



Figure 5: Frequency and accuracy of (all) browsing and searching attempts in Experiment 1.

Even more striking, the more intuitive notion that searching is

inherently faster than browsing the menus received no empirical support (see

Figure 6).



Figure 6: Mean attempt duration as a function of location method and accuracy in Experiment 1.

An ANOVA was conducted to investigate whether time to find the product varied as a function of success at finding the product and whether participants searched or browsed to find it. A 2 x 2 ANOVA with outcome (i.e., successful/unsuccessful) and product location method (i.e., browse/search) as between-subject factors and time as the dependent measure yielded no effect of product location method, F < 1. Further, no significant interaction was found, F < 1. In contrast, the outcome effect was highly significant, F(1, 73) = 11.10, MSE = 1134.13, p < 0.01, with successful product locations taking more time than unsuccessful attempts.

Furthermore, although participants presumably attempted to choose the most cost-effective product location strategy, their judgments were often flawed. As can be seen in Table 4, certain sites yielded discordant time and frequency data. For example, if one looks at the successful product locations for Fogdog.com, it appears that participants were much more likely to browse than search, yet browsing entailed a greater amount of time than searching (although the number of overall occurrences is small). The same pattern was true of Intellesale.com.

Site	Frequency			Time - Mean (SD)		
		(su	m)			
abccosmetics		Browse	Search		Browse	Search
	Successful	4	14	Successful	77 (48)	77 (38)
	Unsuccessful	6	2	Unsuccessful	104 (53)	85 (18)
		_	~ .		_	~ .
bluefly		Browse	Search		Browse	Search
	Successful	10	6	Successful	126 (66)	188 (80)
	Unsuccessful	25	20	Unsuccessful	70 (53)	65 (51)

Table 4: Frequency and accuracy of all browsing and searching attempts in Experiment 1.

Site		Frequ	ency		Time - M	lean (SD)
(sum)						
bluelight		Browse	Search		Browse	Search
	Successful	13	6	Successful	84 (32)	78 (24)
	Unsuccessful	1	4	Unsuccessful	43	119 (94)
computergeor		Browse	Search		Browse	Search
computergear	Successful	DIOWSC 0	10	Successful	14(26)	124(61)
	Junguagaaaful	9 12	10	Unguagastul	44(20)	124(01)
	Ulisuccessiui	15 December	U	Ulisuccessiui	43 (32) December 2	N/A Secret
cradiree-everyn	G (1	Browse	Search	G (1	Browse	Search
	Successful	11	8	Successful	// (42)	99 (39)
	Unsuccessful	6	1	Unsuccessful	58 (42)	64
crateandbarrel		Browse	Search		Browse	Search
crateandoarter	Successful	8	11	Successful	120(70)	04(58)
	Unguagaaful	0	2	Unguagastul	120(70)	74(30)
	Ulisuccessiui	11	5	Ulisuccessiui	01 (37)	11 (38)
deerskin		Browse	Search		Browse	Search
	Successful	19	1	Successful	74 (36)	30
	Unsuccessful	2	1	Unsuccessful	60 (17)	47
fogdog		Browse	Search		Browse	Search
	Successful	15	4	Successful	110 (41)	52 (41)
	Unsuccessful	2	3	Unsuccessful	170 (92)	57 (5)
framiteres		Drowse	Coorel-		Daorras	Soorah
Turniture	G 6.1	Browse	Search	a c 1	browse	Search
	Successful	8	1	Successful	182 (96)	113 (49)
	Unsuccessful	7	40	Unsuccessful	89 (64)	94 (131)

 Table 4 (continued): Frequency and accuracy of all browsing and searching

 attempts in Experiment 1.

Site	Frequency			Time - Mean (SD)		
		(su	m)			
intellesale		Browse	Search		Browse	Search
	Successful	11	8	Successful	99 (43)	70 (37)
	Unsuccessful	20	8	Unsuccessful	89 (84)	64 (39)
jjill		Browse	Search		Browse	Search
	Successful	8	9	Successful	138 (73)	102 (31)
	Unsuccessful	4	2	Unsuccessful	140 (93)	41 (4)
nike		Browse	Search		Browse	Search
	Successful	4	16	Successful	103 (49)	103 (71)
	Unsuccessful	3	11	Unsuccessful	41 (9)	118 (62)
nordstrom		Browse	Search		Browse	Search
	Successful	11	7	Successful	93 (48)	124 (47)
	Unsuccessful	11	7	Unsuccessful	82 (59)	102 (59)
priorities		Browse	Search		Browse	Search
	Successful	12	8	Successful	89 (50)	102 (67)
	Unsuccessful	4	0	Unsuccessful	47 (23)	N/A
shopsports		Browse	Search		Browse	Search
	Successful	8	8	Successful	188 (52)	79 (47)
	Unsuccessful	9	14	Unsuccessful	123 (105)	57 (40)
smartbasics		Browse	Search		Browse	Search
	Successful	8	11	Successful	38 (15)	126 (66)
	Unsuccessful	15	6	Unsuccessful	32 (29)	59 (36)

Table 4 (continued): Frequency and accuracy of all browsing and searching attempts in Experiment 1.

_

Site	Frequency			Time - Mean (SD)		
		(su	m)			
supremevideo		Browse	Search		Browse	Search
	Successful	18	2	Successful	102 (54)	106 (11)
	Unsuccessful	0	5	Unsuccessful	N/A	75 (61)
Totalbody and image		Browse	Search		Browse	Search
	Successful	14	3	Successful	83 (73)	89 (100)
	Unsuccessful	15	18	Unsuccessful	78 (51)	44 (35)
valueamerica		Browse	Search		Browse	Search
	Successful	6	13	Successful	163 (82)	79 (64)
	Unsuccessful	4	13	Unsuccessful	66 (60)	70 (45)
winona		Browse	Search		Browse	Search
	Successful	10	8	Successful	92 (50)	89 (40)
	Unsuccessful	4	5	Unsuccessful	72 (43)	55 (30)

Table 4 (continued): Frequency and accuracy of all browsing and searching attempts in Experiment 1.

Due to the low expected frequencies per cell, the potential dependence (for each site) of product location behavior (browse vs. search) with frequency of success was analyzed using Fisher's Exact Test (2-tailed). Family-wise alpha was set at 0.10 and error rate was controlled using a Bonferroni procedure. Significant relationships were found for three of the twenty sites, namely, Computergear.com, p < 0.01, Furniture.com, p < 0.01, and Supremevideo.com, p < 0.001. For these sites, success at finding the product was dependent upon the method used to locate it.

As can be seen in Table 4, for the most part, the significant effects found in the previous analyses were due to the extreme effectiveness (or ineffectiveness) of the search function. For example, Computergear.com had a very effective search function while Furniture.com fell prey to a highly ineffective search function. Supremevideo.com exhibited both an ineffective search function in tandem with a very effective menu structure.

Problems with search functions are common and well-known (e.g., Nielsen, 2000). User difficulty with query formation has been found in both the information retrieval literature (e.g., Borgman, 1986) and studies of the present-day Internet (e.g., Scanlon, 2000). One researcher has classified search function deficiency under four headings: Figuring out where to search (e.g., under which product section), entering the correct words, specifying the syntax, and interpreting the results (Scanlon, 2000). Of the four problems just mentioned, knowing what words to enter appeared to be the most common problem in our study. However, organization and presentation of search results caused much user difficulty (and frustration) as well.

It was also found that participants were willing to switch strategies (i.e., browsing to searching and vice-versa) when they felt the current strategy would not prove fruitful. Such behavior has been observed by others, with one researcher applying the label of "mixed-behavior users" (Nielsen, 2000). A tally was made of the frequency of such switches for each site and the results are presented in Table 5.

Site	Began by Browsing	Began by Searching	Total Number of Switches
abccosmetics	4	0	4
bluefly	15	12	27
bluelight	0	1	1
computergear	2	0	2
crabtree-evelyn	1	2	3
crateandbarrel	6	2	8
deerskin	1	1	2
fogdog	2	1	3
furniture	3	4	7
intellesale	6	4	10
jjill	7	1	8
nike	2	6	8
nordstrom	1	3	4
priorities	1	0	1
shopsports	4	1	5
smartbasics	4	2	6
supremevideo	0	1	1
totalbodyandimage	7	7	14
valueamerica	2	2	4
winona	5	1	6
TOTAL	73	51	124

Table 5: Switching behavior in Experiment 1.

Participants often made several attempts to find an item with the mean number of attempts per item being 1.75, but, as suggested by the data, if participants lost faith in the effectiveness of a particular strategy, they were willing to try a different method. Particularly prone to such switching behavior were Bluefly.com, Intellesale.com, and Totalbodyandimage.com, all with double digit switch rates. Also, as shown in Table 5, switching behavior was fairly consistent within sites regardless of whether participants began by using the menus or by using search. Noteworthy as well, it was discovered that participants would often switch strategies several times on a single site indicating that the decision on how to locate an item is a fluid and dynamic process. The mean number of switches while looking for each item was 0.31, but if one focuses only on those cases for which participants switched at least once, the mean number of switches per item rises to 1.58.

Bluefly.com's remarkably aberrant switch rate deserves some explanation. Traversing Bluefly.com's menu structure entails enduring a series of pull-down menus listing both brands and product categories, while the search function is without a text field for entering specific search parameters, instead relying on pull-down menus as well. As a result of this combination, participants were often confused as to which method would be most effective and were often quick to judge the initial method rather harshly only to find that the alternatives were no better. As an interesting aside, on first seeing Bluefly.com's search function page, participants often exclaimed, "What?!" Clearly not what they had in mind.

2.2.6. Search Rates Across Participants

Figures 7a and 7b represent the distribution of search rates across participants (looking at initial location behavior for each site). It is clear from the figures that search rate varied substantially across participants. It is also noteworthy that participant 3 rarely searched while participant 20 very often relied on the search function to locate items. This finding will be explored further in the Discussion.



Figure 7a: Distribution of searching behavior by participant in Experiment 1 (Histogram).



Figure 7b: Distribution of searching behavior by participant in Experiment 1 (Box plot).

2.3. Discussion

Essentially, Experiment 1 had two goals, namely to investigate the prevalence of searching versus browsing behavior on existing e-commerce sites, and to isolate factors that influence that strategic decision. Based on the results from Experiment 1, it is clear that search prevalence varies greatly among sites and that search rates are influenced by a multitude of factors, including site design factors. As presented above, nearly 20% of the variance in product location behavior was accounted for by the participants while nearly 15% was accounted for by the sites used. As argued throughout, numerous specific factors are implicated in a user's decision to use SS or MT to locate a product. Each will now be considered in turn.

Previous research is conflicting regarding the use of high menu breadth. On the one hand, high menu breadth suggests limited menu depth which has been shown to be preferred by users (e.g., Larson & Czerwinski, 1998). However, it has also been argued that users may not be willing to read through a large number of menu options which in turn could lead to the use of other alternatives (Norman, 1991). To this debate, Experiment 1 provided strong support for the virtues of high menu breadth. Deerskin.com and Supremevideo.com, the two most browsed sites, both possessed very broad and comprehensive menus. Apparently, participants were not bothered by the long list of menu options available to them as long as the correct choice became evident. On this point, Spool (reported in Koman, 1998) has

argued that "flattening" the information hierarchy by increasing menu breadth increases the probability that a user will find the content he seeks.

Information scent is most likely at the heart of the high menu breadth advantage. Although participants did consider the depth of the menus (and the associated time required to find the item) as an important factor when judging how to proceed, it is clear from their comments that information scent more strongly governed their behavior. The semantic match between the item name and menu label under which it resided on the site was a key determinant of the decision to use the menus. Deerskin.com and Supremevideo.com had a maximum amount of information scent on many occasions as products were organized into very discrete categories (e.g., Deerskin.com: "Hats," "Pants"; Supremevideo.com: "DVD Players," "TVs"), leading to high browse rates. Further, it is important to note that Deerskin possessed a very prominent search function, suggesting that very high information scent will win out in virtually every situation.

(Valueamerica.com also showed the same pattern albeit to a lesser degree.)

Relevant to the use of graphic menu options, Fogdog.com is an interesting case in point. It was predicted that the bold, colorful, and graphic menu options at the top of the page would lure users more than the highly prominent search function just above the menu. Generally, though, neither was as powerful at drawing users as expected. Akin to Supremevideo.com, Fogdog.com possessed a broad menu of options located lower on the page, far below the prominent main menu and search function, and it was these menu options that were often used by participants. While this supports the breadth arguments made earlier, it also casts doubt on the effectiveness of graphic menus. It is not unlikely that participants interpreted the graphical menu options as one large graphical heading or advertisement and paid it little mind, similar to the phenomenon of "banner blindness" (Benway & Lane, 1998), in which users often ignore large, colorful items that are clearly distinguished from other items on a Web page. However, one must not rule out the possibility that participants deliberately preferred the more numerous (and more specific) menu options to the prominent menu options that

possessed a good deal of scent themselves. If so, then it would serve as even more compelling evidence for the virtues of high menu breadth.

An excellent illustration of product menus with high prominence but low information scent can be found at Jjill.com, which has menu headings such as "Relaxed Wear" and "Weekend Wear." (Under which heading would a particular dress be located?) On this note, it is important to point out that the four main menu headings on the Jjill.com site including the two just mentioned were *never* used by the participants despite their prominent display on the site home page. (Jjill.com had a non-zero browse rate due to other available menus, most notably a pull-down menu organized by product category.) User comments indicated that for those users seeking a particular product on the Jjill.com site, attention was quickly drawn to a large and colorful area of the screen that was simply useless.

As argued throughout this paper, users value time very highly and they are reluctant to behave in a way that will make a less than efficient use of that time. Not surprisingly, therefore, it was found that users who considered time important were more likely to use the search function as search was widely believed to be faster than wading through menu options. However, as shown in Table 4 and discussed earlier, this popular belief was not supported by the data.

Prominence of the menu and search function were also deemed important by the participants in determining which option to choose. In accordance with predictions, the more prominent participants considered an option to be, the more likely it was to be used. However, some important caveats apply to this generalization and they concern information scent and whether one or both the menu and search are prominent. For instance, the prominence of the main menu for Fogdog.com rivaled (or perhaps even surpassed) that for its prominent search function and it is quite possible that this accounted for the surprisingly low search rate for that site. In addition, when information scent was very low, the prominence of a menu had little influence, with the Jjill.com main menus serving as the exemplar for this point.

An important qualification for this latter point, however, is that it applies only when a search function is available to the user and he is aware of that option. For instance, Computergear.com possessed a prominent menu with very low information scent along with a very subtle search function. In set 1, fully 87% of users used the menu despite its flaws, but for set 2, when they presumably became aware of the presence of a search function, usage of the menu fell to 30%.

As an interesting aside, being unaware of a subtle search function may lead to its disuse, but this is not the only possible interpretation of that correlation. Another interpretation, supported by occasional participant comments, is that participants interpreted a subtle search function as a warning from the site's designers. Some participants felt that a search function that was very difficult to locate would more likely be ineffective than one prominently displayed on a site, while others considered it irritating that they would be "coerced" into using menus so that they could be exposed to other products that they did not intend to buy. It appears that promoting menu usage by hiding the search function is not a wise idea. Instead, based on the results of Experiment 1, menu usage can be encouraged through the use of high information scent.

Furthermore, all designers should keep in mind that subtle search functions often go unnoticed and the result can be a great deal of frustration for the user if the menus prove unhelpful. Interestingly, the inverse of the prior dilemma can also be found. Nike.com served as a prime example of a site with a very subtle menu. As discussed earlier, participants most often used the search function simply because they felt they had no alternative.

Finally, and perhaps as critical to user behavior as information scent is the user's evaluation of and general predisposition toward the search function. As mentioned earlier, participants often considered the search function to be less time-consuming than browsing the menus, yet many participants had serious reservations with choosing search as their product location strategy. Concerns relevant to the search parameters that should be used, the number of items that the search function would be expected to return, and the time for the search function to locate the items all weighed heavily on the minds of users. More generally, users' general opinion toward using search functions and their experience with searching also played a substantial role in their strategic decisions. Some participants were content
to search on nearly every site while others were reluctant to do so unless absolutely necessary.

One researcher has suggested that roughly half of all users are searchdominant (i.e., they go straight for the search function when entering a site (Nielsen, 2000)). However, as shown in Figures 7a and 7b, far fewer than half of the users in our sample were search-dominant. In fact, if a searchdominant user is expected to use the search function in every situation, then none of the users in our study could be classified as search-dominant. Such a finding in this experiment can be interpreted as support for the notion that particular sites play a critical role in determining searching behavior.

"Loss aversion," the notion that losses loom larger than corresponding gains (Kahneman & Tversky, 1979) may have had an impact on the decision to use SS or MT. Specifically, users may have concluded (due to a priori biases or site-design factors) that SS was generally a more efficient way of locating products on Web sites. If that were the case, then a user may have surmised that the *probable* loss in time to find the product using the menus was a more important factor than the *possible* (but usually unlikely) gain in accuracy by using the menus (as opposed to using SS).

On a general note, the commonly held belief that SS is faster and more accurate than MT received no empirical support in this study. The frequency and time data, discussed above, clearly suggested that the participants in Experiment 1 were not irrational by often choosing to use the menus to locate items. It seems clear that they acquired no time or accuracy disadvantage by doing so.

It is also interesting to point out that many participants showed a general preference for browsing the menus. It was predicted (or perhaps feared) prior to conducting the experiment that participants would take the course of action that would most quickly release them from their experimental obligation. In other words, they would use the search function because they would consider it faster than using the menus and they would be permitted to leave the testing room sooner. As has been discussed throughout, however, this was not the case. In addition to the data indicating the contrary, participants often commented that they preferred to use the menus (even when a prominent search function was present) as it allowed them to see all of the items available on a site permitting them to select the best item for their money. It appears that the common expectation that users will use SS whenever possible is based more on myth than data.

We have presented evidence that suggests that users do not simply use the site search function whenever possible, and hypertext research conducted in the 1980s supported such a view. The hypertext-based help system study discussed above (Campagnoni & Ehrlich, 1989) found that despite the use of stimuli designed to elicit both browsing and searching behavior, most users preferred to browse. In that study, browsing was defined as scanning tables of contents and paging through relevant topics to find answers while searching was defined as using indexes to look up specific query terms and following the links to the appropriate topic and page (to find the answers). Users in their study strongly preferred browsing over searching to find the answers to the questions and the authors of the study offered two explanations for the results that are consistent with the results of the current study. To account for the browse over search advantage, they reported that

participants were often unable to formulate the query terms necessary to make effective use of the indexes (i.e., they were unclear as to the search parameters). However, for questions that resulted in uncommonly high searching behavior, they reported that the key words in the question did not appear in the tables of contents (i.e., there was a poor match of item and menu).

It is noteworthy that instances from prior hypertext research in which users prefer searching over browsing are sometimes tempered by caveats of search usage. For example, one study found that while "searching" was more common than "browsing" when using an electronic encyclopedia, participants failed to take full advantage of the search capabilities available within the system (Marchionini, 1989). Although users in the study preferred the more powerful full-text search option over the alternative, they commonly neglected to use Boolean connectives (AND, OR) or change scope or proximity settings.

Two factors which were not directly investigated in Experiment 1 deserve recognition. The *perceived* specificity of the item and the familiarity with the item are both likely to be important determinants of the decision to use SS or MT to locate the item. Although not directly measured in Experiment 1, participant comments may provide some insight into these factors. A more controlled experimental investigation of these issues is warranted.

Some readers may question the findings from Experiment 1 on the grounds that the experiment was an artificial situation. However, there is evidence to contradict this claim. While looking for items, participants frequently spent a considerable amount of time deciding between two items that equally satisfied the requirements of the product description, and several participants had to be reminded as they navigated to the product purchasing screens that they were not supposed to actually purchase the items. One participant even went so far as to investigate the shipping and handling costs associated with purchasing an item. Similar to this objection, readers may also claim that much of the time, e-commerce site users browse sites without a particular item in mind. While this may be true some of the time, it is argued that this is the minority of e-commerce site usage. In fact, one

researcher found that two-thirds of web users are looking for specific information (Koman, 1998). Granted, users may not be intending to buy an Air Storm GPT Men's Softball Bat, but it is not unlikely that they have some type of bat in mind.

3. EXPERIMENT 2:

Manipulating Site Search /

Menu Traversal Behavior in a Controlled Setting

Based on the results of Experiment 1, numerous factors presented themselves that may have affected the decision to use SS or MT. As pertains to the cost-benefit perspective, information scent appeared to be a critical factor in a user's decision to use the menus or search function on a site. In the discussion of cost-benefit issues in the Introduction, the focus was on the *anticipated* costs and benefits of using particular product finding strategies. Therefore, the actual categorization structure of products beneath the toplevel product headings was superfluous for the purposes of this study. Supporting the emphasis on the top level of the menu hierarchy, McEwen (1981) observed that most navigational errors are made at or near the top of the menu hierarchy.

Also important is the issue of menu breadth as Experiment 1 produced rather unintuitive findings. As discussed earlier, participants were not dissuaded from using menus with high breadth under the circumstance that the information scent of the menu was high. This surprising result merits further investigation in a controlled experiment.

As pertains to the attentional capture perspective, prominence of the search function and menu options both appeared to play important roles in the users' decisions to use SS and MT and both would be investigated in Experiment 2. As discussed earlier, if the design of a site is such that attention is immediately drawn to the search function (e.g., through brightness contrast of the text field with the page background), then users may be more likely to use SS. Conversely, if attention is immediately drawn to product headings, then an MT strategy may be adopted. In terms of the visual momentum approach, the search function (or product headings) may serve as a landmark which influences the user's subsequent visual inquiry

and information acquisition. The user may perceive alternative strategies for finding a product to be subordinate (such as traversing the product hierarchy). Conversely, if a menu of product categories captures attention and serves as the landmark, then the use of SS may be subordinated. As argued above, where the user's attention is drawn may play a critical role in her decision process.

Two of the factors not directly (or appropriately) investigated in Experiment 1 that were deemed important, namely familiarity and specificity of the item being sought, also deserved empirical evaluation. All items used in Experiment 2 were popular and familiar to all participants (unlike in Experiment 1). For instance, instead of having participants locate "jumpers," they were asked to locate "jeans." Item specificity, although examined in Experiment 1, was manipulated more effectively in Experiment 2 in an attempt to gauge the influence of the factor. Only clearly specific and clearly general items were used in Experiment 2.

Finally, the users' general predisposition toward using the search function was examined in Experiment 2. The results from Experiment 1 suggested strongly that users possess varied opinions of search functions and that such impressions affect behavior.

3.1. Method

3.1.1. Participants

Thirty-two Rice University undergraduate students participated to fulfill a course requirement. None participated in Experiment 1.

3.1.2. Materials

Sixteen Web sites were created representing every combination of four factors (menu breadth, information scent, prominence of menu, and prominence of search function; see Table 6). The created sites were portrayed as sixteen competing sites in the genre of "clothing retailer." Therefore, the content of the sites was equivalent (to the extent allowed by the varying experimental conditions). Primary variation among the sites was due to the variance in site home page design, the main focus of this Experiment. (See Appendix E for selected sites used in Experiment 2.)

Site	Site Name	Search	Menu	Breadth	Scent
1	Fashion Plus Clothing	Prominent	Prominent	High	High
2	J & M Apparel	Prominent	Prominent	High	Low
3	The Cave	Prominent	Prominent	Low	High
4	Thunder Alley Outfitters	Prominent	Prominent	Low	Low
5	Style Dynamics	Prominent	Subtle	High	High
6	Queenstown Clothing	Prominent	Subtle	High	Low
7	Blue Creek Clothier	Prominent	Subtle	Low	High
8	Radical Wear	Prominent	Subtle	Low	Low
9	Hip Trends, Inc.	Subtle	Prominent	High	High
10	Nite Owl Fashion	Subtle	Prominent	High	Low
11	TRI Apparel	Subtle	Prominent	Low	High
12	Fabulous Fabrics	Subtle	Prominent	Low	Low
13	New Wave Fashion	Subtle	Subtle	High	High
14	Trixie Wear	Subtle	Subtle	High	Low
15	Cougar Apparel	Subtle	Subtle	Low	High
16	Fireball Outfitters	Subtle	Subtle	Low	Low

Table 6: Sites used in Experiment 2.

All sites possessed a product menu and search capability. Further, each site contained eight unique JPEG images of models that were each 100 pixels x 100 pixels in area. The purpose of having images was to create natural variation among the sites. To further achieve this end, the fonts used were varied among sites (at least among company name headings) and alignment of company names at the top of the page was varied as well. In an effort to make the sites look more natural, each site contained five irrelevant (but live) links. These links were "Customer Service" or "Security Policy" links, for example, and were extraneous to the factors of interest. Such links were varied to the greatest extent possible to promote variation among the sites. Each site possessed a unique company name and an effort was made to have each site display a unique color scheme. Finally, to avoid confounding the factor of prominence, scrolling was not required for any of the sites.

Accompanying the sixteen Web sites were sixteen matched pairs of items to be located by participants, with each pair consisting of a general and specific version of an item. For example, the general version of an item might have been "sweater" while the specific version might have been "Kenneth Cole Diamond Pattern Sweater." As stated earlier, all items were very popular and therefore familiar to all participants (i.e., there was no doubt as to the comprehension of the item description). Finally, to investigate the influence of participants' general predisposition toward using search functions, eight additional (half general and half specific) items were created to gauge users' opinions of the search functions (see Procedure below).

Each participant was tested individually on an Apple iMac personal computer running Microsoft Internet Explorer 5 with the display set to a resolution of 1024 x 768 pixels. The created sites were located on the computer's hard drive so no Internet connection was necessary.

3.1.3. Design

Five factors were manipulated completely within subjects: prominence of search function, prominence of menu, menu breadth, item specificity, and information scent. Below is a depiction of the factors that were manipulated, the two levels for each factor, and the method of operationalizing each level.

Search Function Prominence:

PROMINENT: The "Search" text label had high brightness contrast with its background and there was a blank text field present, also with high brightness contrast with its background. High color contrast was achieved with a white text label and a white field on a colored background. The Search text label was presented in 18 pt. Arial font near the top of the page. (For example, see Site 1 in Appendix E.) SUBTLE: No text field was present (i.e., the Search text label served as a link). The Search label link was presented as the final option on the menu with the same font characteristics as other menu options. (For example, see Site 11 in Appendix E.)

Menu Prominence:

PROMINENT: The menu options consisted of relatively large text with high brightness (and/or color) contrast with the background. The menu items were presented in 14 pt. Times New Roman font on various colored backgrounds. Conditions that called for broad, high scent menus also possessed product category headings ("Men's Department" and "Women's Department.") For prominent menus, those headings were presented in an 18 pt. Arial font. The menus were presented vertically down the page occupying a large amount of screen real estate in the center of the page. (For example, see Site 11 in Appendix E.) SUBTLE: The menu options consisted of relatively small text. The menu options were presented in black 12 pt. Times New Roman font and the menus were presented horizontally (i.e., in a horizontal row or rows) at the bottom of the page. The menu text was always presented on a white background (the default background color of the sites). Sites with subtle menus also contained promotional text so as to occupy more screen real estate. (For example, see Site 6 in Appendix E.)

Menu Breadth:

HIGH: The menu contained 30 options (not including the search option for sites with subtle search functions). (For example, see Site 1 in Appendix E.)

LOW: The menu contained 9 options (not including the search option for sites with subtle search functions). (For example, see Site 11 in Appendix E.)

Item Specificity:

GENERAL: The item was very general (e.g., scarf).

SPECIFIC: The item was very specific (e.g., Elizabeth Gillett Silk Scarf).

Information Scent:

HIGH: An arbitrary criterion based on the Latent Semantic Analysis method described below. As there were both a broad and a narrow high scent menu, each menu was analyzed individually:

Broad-High Scent Menu:

Maximum Scent: 0.55 Minimum Scent: -0.05

Narrow-High Scent Menu:

Maximum Scent: 0.46 Minimum Scent: 0.02

(For an example of a site with a high scent menu, see Site 1 in

Appendix E.)

LOW: An arbitrary criterion based on the Latent Semantic Analysis method described below. As there were both a broad and a narrow low scent menu, each menu was analyzed individually:

Broad-Low Scent Menu:

Maximum Scent: 0.38 Minimum Scent: -0.07

Narrow-Low Scent Menu:

Maximum Scent: 0.36 Minimum Scent: -0.04 (For an example of a site with a low scent menu, see Site 16 in

Appendix E.)

While others have measured scent by asking users to rate how confident they are before they click on a link (Spool & Klee, 2000), the information scent measure used in Experiment 2 was based on a method of measuring semantic similarity known as Latent Semantic Analysis or LSA (Deerwester et al., 1990; Landauer et al., 1998). LSA is a mathematical technique for extracting and representing the similarity of meaning of words and passages by analysis of large bodies of text. It assumes that the information about all the word contexts in which a word does and does not appear provides a set of mutual constraints that largely determines the similarity of the meaning of words to a set of words (Landauer et al., 1998).

The specific method used is based on applications provided by the Science and Applications of Latent Semantic Analysis Group (SALSA) at the University of Colorado at Boulder (http://lsa.colorado.edu). The LSA application processes a large sample of language and represents the words (and sentences, paragraphs, or essays) used in it as points in a very high dimensional semantic space. A semantic space is a mathematical representation of a large body of text. Every term (i.e., word, sentence, paragraph, essay) and every novel combination of terms has a high dimensional vector representation. When two terms are compared, the cosine of the angle between the vectors representing the terms is compared within a given semantic space. (For further information on LSA, see Deerwester et al., 1990, Landauer et al., 1998, or the SALSA group Web site (http://lsa.colorado.edu).)

Using LSA, the following procedure was used to compute the scent measures (listed above) that were used to quantify the distinction between high scent and low scent menus. First, the scent of each item for each menu label (for all menus) was determined. Second, for each item, the scent of the "best" and "worst" menu labels for each of the four main menus was determined. Using the Broad-High Scent Menu as an example, the menu label "Sweaters" was the best label for the item "sweater" (i.e., "Sweaters" yielded the highest scent for "sweater") while "Scarves" was the worst menu label (i.e. yielded the lowest scent). Then, for each main menu, the mean scent of all the best matches was computed, and the same was done for the worst matches. The mean of the best matches corresponds to the "Maximum Scent" measure listed above while the mean of the worst matches corresponds to the "Minimum Scent" measure listed above. Looking at the values presented above for the high scent and low scent menus, it is clear that the high scent menus possessed a higher level of maximum scent than the low scent menus, while the low scent menus possessed a lower level of minimum scent than the high scent menus.

In addition to the semantic similarity of links with the sought information, other factors that determine scent have been put forth. For example, it has been argued that links with fewer, more general words possess less scent (Sawyer, 2000). While using more general words in links may be an acceptable means of reducing semantic similarity, link length per se was not manipulated as semantic similarity was considered far more important. Manipulation of link length may have also led to a more difficult evaluation of the information scent of the links.

The items were counterbalanced so that each item was sought on every site across participants and each site was viewed only once by each participant. Due to the fact that there were twice as many items as sites and participants were shown each site only once, four participant groups were required. The participants were assigned to one of two main groups (Group A and Group B) and one of two sub-groups (Sub-group I and Sub-group II). For Group A, sites 1-8 were assigned general items and sites 9-16 were assigned specific items. For Group B, sites 1-8 were assigned specific items and sites 9-16 were assigned general items. The two sub-groups for each main group were necessary to ensure that every item was sought for every site across all participants. For sub-group I, sites 1-8 were assigned items 1-8 while for sub-group II, sites 1-8 were assigned items 9-16.

3.1.4. Procedure

As in Experiment 1, participants were shown each site one at a time in random order and were presented with a single item provided on a sheet of paper for each site. However, unlike the previous experiment, the participants in this study were not actually asked to find the products (at any time). After a participant was shown the product to be found, she was instructed to decide what she would do first to locate the item and then perform that action. After performing that action, she was presented with a Javascript message box that simply stated what action she took (e.g., "You clicked on: Pants"). Then, the experimenter recorded the menu option clicked or the text typed into the search field and loaded the next site.

The reason behind this somewhat unnatural instruction was the lack of functionality behind the sixteen created sites. It would have been a massive undertaking to create sixteen sites with full product menu structures and, more importantly, fully operational search functions. If a participant were to learn that the search functions were non-functional, for example, it could have influenced her decision to use SS on future sites.

Following administration of all sixteen sites (and corresponding items), an additional eight (half general and half specific) items (new to the participants) were presented with eight of the original sixteen sites. The new

items were randomly assigned to the existing sites and each site-item pair was presented individually. For each site-item pair, participants were asked various questions pertaining to their opinion and evaluation of each search function. The questions were used to gather information including the clarity of the required search parameters, the number of expected search hits, and the predicted time required for the search function to return the item in question. The goal of administering such questions was to collect further data on the very important factor of the users' general search predisposition. Due to the fact that the additional items were to be used to compare participants to each other in terms of search predisposition, the stimuli and their order of presentation were held constant for all participants. As in Experiment 1, the dependent variable of primary interest in Experiment 2 was the initial course of action taken by participants to use SS or MT to locate the product.

3.2. Results

The overall search rate collapsing across all conditions was 27% and Figures 8a and 8b present the distribution of search rates across sites.



Figure 8a: Distribution of searching behavior by site in Experiment 2 (Histogram).



Figure 8b: Distribution of searching behavior by site in Experiment 2 (Box plot).

The distribution of searching behavior across sites taking item specificity into account (see Figures 9a and 9b) reveals that search behavior is dependent upon many factors including item specificity. The following analyses were conducted to investigate the influence of each manipulated factor on the likelihood that a user would search or browse to find a product.



Figure 9a: Distribution of searching behavior by site and item specificity in Experiment 2 (Histograms).



Figure 9b: Distribution of searching behavior by site and item specificity in *Experiment 2 (Box plots).*

Due to the fact that the specificity factor (i.e., general vs. specific items) was not completely crossed with the other four factors of interest, two separate repeated factor ANOVAs were conducted. The first analysis investigated the effects of the four completely crossed factors (search prominence, menu prominence, menu breadth, and information scent) without the item specificity factor. The first analysis revealed a nonsignificant effect of search prominence permitting us to replace the search prominence factor with the specificity factor in the second ANOVA. This way, specificity would be completely crossed with the three other remaining factors (menu prominence, menu breadth, and information scent) and its main effect and interactions with the other factors could be investigated.

3.2.1. Overall Analysis Excluding Item Specificity

A completely within-subjects ANOVA was conducted with search prominence, menu prominence, menu breadth, and information scent as factors. The analysis revealed a significant main effect of menu breadth, F(1, 31) = 27.59, MSE = 0.092, p < 0.001, with low breadth menus more likely to elicit searching behavior (see Figure 10 for all main effects). The main effect of information scent also reached significance, F(1, 31) = 12.31, MSE = 0.280, p < 0.01, with low information scent menus more likely to elicit searching behavior. Although subtle menus were more likely to elicit searching behavior than prominent menus, the effect failed to reach significance, F(1, 31) = 2.61, MSE = 0.108, p = 0.12. Finally, although a prominent search function was more likely to elicit searching behavior than a subtle one, the main effect of search prominence failed to reach significance, F < 1. The lack of a search prominence main effect prompted us to replace the factor with the item specificity factor in the follow-up ANOVA to be described later.



Figure 10: Main effects from Experiment 2.

The analysis also yielded two interesting interactions involving menu breadth that approached significance (see Figures 11 and 12). The search prominence x menu breadth interaction, F(1, 31) = 3.72, MSE = 0.076, p = 0.06, suggested that menus with higher breadth elicited greater browsing behavior, particularly for sites that possessed a subtle search function (see Figure 11).



Figure 11: Search prominence x menu breadth interaction in Experiment 2.

The information scent x menu breadth interaction, F(1, 31) = 2.64, MSE =

0.145, p = 0.11, suggested that menus with high information scent elicited greater browsing behavior, especially when the menu possessed high breadth (see Figure 12).



Figure 12: Menu breadth x information scent interaction in Experiment 2.

3.2.2. Overall Analysis Including Item Specificity and Excluding Search

Prominence

As described earlier, the first ANOVA revealed a non-significant effect of search prominence. Therefore, the factor was replaced with the item specificity factor so that the main effect of specificity and its interactions with the other remaining factors could be investigated. The resulting ANOVA analyzed the four factors (item specificity, menu prominence, menu breadth, and information scent) as completely-within factors. As this ANOVA merely replicated the main effects found in the first ANOVA, item specificity will be at the center of this discussion. A main effect of item specificity was found, F(1, 31) = 4.28, MSE = 0.308, p < 0.05, with specific items more likely to elicit searching behavior (see Figure 10).

Only one two-way interaction involving item specificity had an Fvalue greater than 1, namely specificity x information scent, F(1, 31) = 1.19, MSE = 0.105, p = 0.28 (see Figure 13). This interaction suggested that specific items elicited greater searching behavior, particularly for menus that possessed low information scent.



Figure 13: Information scent x item specificity interaction in Experiment 2.

Two three-way interactions involving item specificity were found that either reached or approached significance. The item specificity x menu prominence x information scent interaction (see Figures 14 and 15) reached significance, F(1, 31) = 4.16, MSE = 0.068, p < 0.05. Looking first at specific items (Figure 14), it appears that low scent menus yielded a high level of searching behavior regardless of menu prominence. However, the data for general items (Figure 15) reveals that prominent low-scent menus



were much more likely to be used than subtle low-scent menus.

Figure 14: Menu prominence x information scent interaction for SPECIFIC items only (in Experiment 2).



Figure 15: Menu prominence x information scent interaction for GENERAL items only (in Experiment 2).

The item specificity x menu breadth x information scent interaction (see Figures 16 and 17) approached significance, F(1, 31) = 3.07, MSE = 0.092, p = 0.09, and was even more intriguing than the previous three-way interaction. As can be seen in Figure 16, no interaction was present between menu breadth and information scent for specific items. Menus with high information scent elicited greater browsing behavior independent of menu breadth. However, as shown in Figure 17, a marked interaction was found between the two factors for general items. While menus with high information scent elicited greater browsing behavior, the effect was much more pronounced when the menu possessed high breadth. In fact, when general items were sought on sites containing a menu with low breadth, the browsing advantage usually found for high scent menus virtually disappeared.



Figure 16: Menu breadth x information scent interaction for SPECIFIC items only (in Experiment 2.)



Figure 17: Menu breadth x information scent interaction for GENERAL items only (in Experiment 2.)

3.2.3. Analysis of General Search Predisposition of Participants

As mentioned earlier, participants were given an additional 8 site-item pairs following the main 16 trials to collect data regarding the general opinion and tendency participants possessed toward using the search function. These data were collected using questionnaires. For each site-item pair, participants were presented with questions aimed at determining their impression of the efficiency and cost-effectiveness of using search. For instance, they were asked to provide estimates of the number of expected search hits, the likelihood that the search function would return any hits, and the expected time to locate the product using the search function.

A "search predisposition index" score was then calculated for each participant using the questionnaire data (see Figure 18). To calculate the index score for each participant, the following procedure was used. The mean response (and standard deviation) for each question for each site was calculated (across all participants). Then z-scores were calculated for each participant's response to each question. The z-scores were then normalized so that positive scores represented a stronger preference (or more positive opinion) for using search functions. Finally, the mean of the z-scores was calculated for each participant producing the search predisposition index score for each participant as presented in Figure 18.


Figure 18: Search predisposition index as a function of participant in Experiment 2.

The index scores clearly suggest that users possess varied opinions of local search functions. Some view search functions as generally effective and valuable uses of time while others view them negatively. To explore this finding further, analyses were conducted to determine whether participants' general search predisposition influenced their behavior in Experiment 2. Although sufficient evidence exists to argue for variance among participants' views of search functions, the site-specific factors investigated in Experiment 2 may have been strong enough in their own right to influence behavior regardless of search predisposition.

Participants' search predisposition index scores were first correlated with the proportion of searching behavior in Experiment 2 for each participant yielding a coefficient of -.03. While not the most convincing evidence, the lack of a meaningful correlation suggested that site-specific factors were substantially more important determinants of behavior in Experiment 2.

To further examine the influence of the index scores, participants were separated into two groups (of equal size) using a median split of index scores with one group labeled "browsers" and the other group "searchers." Browsers were participants who held generally negative views of search functions while searchers held generally positive views of search. Then, the browser/searcher distinction was entered into the earlier ANOVAs for Experiment 2 data (presented above) as a between-subjects factor with the goal of exploring interactions of the new grouping factor with the other factors of Experiment 2 (such as scent, menu breadth, etc.). For these analyses, as in the previous ANOVAs, the site-specific (and item-specific) factors were manipulated within subjects, thereby controlling other betweensubjects differences. While the grouping factor failed to produce a main effect, F < 1, two higher-order interactions were found, both involving item specificity. The item specificity x information scent x group interaction reached significance, F(1, 30) = 5.45, MSE = 0.092, p = 0.03, as did the item specificity x menu prominence x menu breadth x group interaction, F(1, 30)= 4.20, MSE = 0.091, p < 0.05.

Figures 19 and 20 represent the item specificity x information scent interaction for browsers and searchers, respectively. The interaction for users who possessed an unfavorable opinion of the search function magnified the non-significant two-way interaction between specificity and scent (discussed above). Specifically, an increase in searching behavior was found for specific items particularly on sites with low scent menus (compare Figures 13 and 19). In contrast, participants who possessed a favorable opinion of search functions (in terms of efficiency and usefulness) were more likely to search for specific items regardless of the information scent of the menu (see Figure 20). Simply put, users with favorable opinions of searching were more likely to search for specific items even in the presence of a high scent (and presumably more efficient) menu.



Figure 19: Information scent x item specificity interaction for BROWSERS only (in Experiment 2).



Figure 20: Information scent x item specificity interaction for SEARCHERS only (in Experiment 2).

Finally, a regression analysis was performed (as in Experiment 1) to determine the proportion of variance explained by the participants and the sites used. Product location behavior (search / browse) was regressed on participant and site and yielded findings supporting the influence of user characteristics. The regression model revealed that 27.9% of the variance in product location behavior was accounted for by the participants while 7.2% was accounted for by the sites used.

3.3. Discussion

Many predictions based on data reported in Experiment 1 were supported in the more controlled environment of Experiment 2, particularly predictions relevant to the cost-benefit perspective. Sites that contained menus with high breadth or high information scent were significantly more likely to elicit browsing behavior than sites with low breadth or low scent, respectively. Furthermore, the breadth and scent factors interacted in that menus with high information scent were even more likely to elicit greater browsing behavior when the menu possessed high breadth. Although the information scent x menu breadth interaction failed to reach significance, the pattern of results is not surprising. When made up of high information scent items, broad menus allow an even greater level of comprehensiveness and semantic matching (of items with menu labels) than low breadth-high scent menus.

Item specificity played a very important role in determining product finding behavior on the sites. Specific items were found to be more likely to elicit use of the search function to find the item. Furthermore, two three-way interactions were found that further supported the influence of item specificity. While the three-way interactions are complex and difficult to interpret, they seem to suggest that cost-benefit and attentional capture factors do interact when item specificity is taken into account. While such an interplay was found directly via the marginal search prominence x menu breadth interaction, item specificity brought out the interaction of information scent with menu prominence (as well as menu breadth).

Variance among search predisposition index scores of participants (Figure 18) suggested that general opinions of the effectiveness and usefulness of local search functions varied among users. While the two higher-order interactions involving grouping of participants as "browsers" and "searchers" did not provide clear and convincing support for the influence of search predisposition, they suggested that such general opinions played a role under certain circumstances. In fact, a multiple regression model revealed that nearly 30% of the variance in product location behavior was explained by user characteristics. Together with the variance of the search predisposition index score among participants, the analyses provided support for its status as a meaningful factor that must be considered when investigating product finding behavior on e-commerce sites.

Unfortunately, the data were less supportive of the factors related to attentional capture, namely search function and menu prominence. Although prominent elements were more likely to be used than subtle ones, the effects were not strong enough to reach significance. Nevertheless, the attentional capture factors were not without influence as suggested by its interactions with other factors. For example, the browse advantage for menus with high breadth was further pronounced when the search function was subtle. Considered jointly, the main effects and interactions concerning prominence suggest that independent of other factors, search function and menu prominence are generally not critical to product finding behavior as long as the user is aware of the available options. The marginally significant search prominence x menu breadth interaction makes clear the possibility that subtle search functions may often go unnoticed when presented among broad menus.

One possible account of the three-way interaction found among item specificity, menu prominence, and information scent (Figures 14 and 15) centers on the relationship between item specificity and searching behavior. As discussed earlier, there was a substantially higher search rate for specific items than general ones and the three-way interaction suggests that this tendency affects menus of varying prominence differentially. When looking for general items, users apparently were much more willing to consider (and use) a low-scent menu if it were prominently displayed on the page. For specific items, however, the prominence of a low-scent menu did not lead to it being used more often relative to a subtle one. It appears that the tendency to search when asked to find a specific item is not mitigated by a prominent menu when that menu possesses low information scent.

As discussed earlier, the breadth x scent interaction that was so pronounced for general items was non-existent for specific items (see Figures 16 and 17). Higher-order interactions are often difficult to interpret and the (marginally significant) item specificity x menu breadth x information scent interaction is no exception. One possible explanation that can be put forth, however, is that users seeking a general item on a site with a low breadth menu fail to examine the menu options carefully relative to a site with a high breadth menu. One can only speculate why this would be the case. Perhaps if a user is unimpressed with the first few options in a low breadth menu, she may assume that the other options will be unhelpful as well. Conversely, the same user may be willing to examine more options in a high breadth menu even if the first few options are unhelpful in the hope that other options would prove fruitful.

While this seems counter to cost-benefit principles, an example may show how this might not be the case. If a site sells clothing items and represents such items with a top-level menu that contains thirty product categories, a user may then consider the ratio of number of perceived (or assumed) products available on the site to the number of menu labels present. If a user perceives a site to sell roughly twenty categories of items and the menu contains thirty categories, she may then assume that the appropriate category is present on the site *regardless of the information scent of the first few menu options*. Based on this assumption, she may be willing to explore the other menu options. Alternatively, if she perceives twenty categories of items to be available on the site but the menu contains only nine categories, one or two low information scent menu categories may be enough to dissuade the user from investigating the remaining options. The question still remains, however, as to why such an effect would be found for general but not specific items.

Overall, much supportive evidence was found in Experiment 2 for principles suggested in the exploratory first experiment. The results as a whole seem to suggest that the structure of information on e-commerce sites is more important than the way that information is presented. In fact, other studies have concluded that graphic design is completely unrelated to success at finding information on Web sites (Spool et al., 1997).

Specific to this experiment, the site design factors of information scent and menu breadth played key roles in determining whether a participant used the search function or browsed the menus to find a product. Furthermore, a regression model revealed that approximately 7% of variance in product location behavior was explained by the sites used in the experiment. Finally, it is noteworthy that the overall search rate of 27% echoed the pattern from Experiment 1 (42%), namely a much lower search rate than that predicted by many researchers.

4. GENERAL DISCUSSION

Collectively, the exploratory Experiment 1 and the experimentally controlled Experiment 2 provided support for the cost-benefit and attentional capture perspectives of how users decide to find products on e-commerce sites. Additionally, valuable data was collected that pertained to search prevalence and efficiency for users interacting with such sites. Each major topic will be discussed in turn, but it is important to make clear that data from this research unequivocally supports the implementation of search functions on e-commerce sites; the value of a search capability is wellknown. Rather, the main argument is that more attention be placed on site design and its interaction with user characteristics. A user's behavior on an e-commerce site is not solely determined by the characteristics of that user. Instead, the traits and goals of the user interact with the design of the site to

determine behavior, and the cost-benefit and attentional capture principles lay at the heart of the user's product finding strategy.

While the experiments provided empirical data on searching behavior and factors that influence the product location decision, conclusions must be tempered with the realization that college students may behave differently from the population at large. Rice University undergraduates are unique relative to the population at large in terms of age, intelligence, computer and Web experience, and socioeconomic level. Any one of these factors could have influenced behavior in these set of experiments potentially making generalizations to average Web usage invalid.

One could argue, for example, that younger users are more willing to browse sites for information as they may value time less highly than adults. Time is extremely valuable for adults who must juggle work, family, and financial responsibilities. While college students are often pressed for time as well, they arguably have much more leisure time than their older counterparts. Being a prestigious university, Rice attracts very intelligent students who are often very computer and Web savvy. Unlike many Web users outside of Rice, they may have a great deal of experience navigating Web sites and e-commerce sites and may have already even purchased items online. Such proficiency may have affected participant behavior in unforeseen ways. For example, more intelligent and experienced users may be more skilled at navigating menu structures or more adept at entering search queries than the average user. Future experiments must be conducted with a more heterogeneous sample of participants before any concrete conclusions can be drawn.

Furthermore, this experiment focused on a specific genre of Web sites, namely e-commerce sites. It is important to realize that the data provided by this set of experiments may not generalize to all (or even most) Web sites. Many different types of sites exist and they may use any number of design methods to disseminate information to users. How users interact with such sites to obtain the desired information may differ greatly depending on the genre of such sites.

4.1. Search Behavior and Prevalence

Despite the scarcity of empirical data on actual search behavior on Web sites, many researchers were willing to make claims regarding search prevalence. The proportion of users asserted to be "search-dominant," i.e., who head straight for the search function upon entering a site, ranged from 33% to 50%. Results from Experiment 1, however, demonstrated that searching behavior varied depending on the site being used. Rather than finding search rates for each site between 33 and 50 percent, a wide distribution of searching behavior was found for the 20 sites used (see Figures 1a and 1b). Furthermore, the distribution of searching behavior across participants (from Experiment 1; see Figures 7a and 7b) revealed that only 30% of users showed a search rate greater than 50%, with only one user reaching a search rate of 80%. (Overall search rates were 42% and 27% for Experiment 1 and Experiment 2, respectively.) It therefore appears that the claims made by other researchers regarding search behavior and prevalence neglect the influence of site design and place too much emphasis on the user himself.

It was also found (in Experiment 1) that participants switched strategies (i.e., browsing to searching and vice-versa) when they felt the current strategy would not be effective at locating the product. Further, they would often switch strategies multiple times on individual sites suggesting that the decision on how to locate an item is a dynamic process.

4.2. Efficiency of Searching Versus Browsing

It is a widely held belief among many researchers and users alike that searching is inherently faster and more accurate than browsing. However, data from Experiment 1 challenged many of these accepted truths of the Web. No empirical support was obtained for the notion that searching is faster than browsing the menus, and the global frequency and time data (Figures 5 and 6 respectively) provided striking evidence that the use of the search function in an effort to locate a product does not lead to a greater rate of success. Surprisingly, the trend was actually toward higher success rates when using the menus to find items as opposed to using search.

For three of the twenty sites, namely Computergear.com, Furniture.com, and Supremevideo.com, success at finding the product was found to be dependent upon the method used to locate it. At the heart of these effects was the extreme effectiveness (or ineffectiveness) of the search function (see Table 4). For example, Computergear.com had a very effective search function while Furniture.com fell prey to a highly ineffective search function. Overall, the frequency and time data clearly suggested that participants in Experiment 1 were not irrational by often choosing to use the menus to locate items. It seems clear that the efficiency of their performance was not diminished by doing so.

4.3. The Cost-Benefit Perspective

The cost-benefit perspective was introduced as a means of accounting for the decision process of users who must choose between traversing a product menu hierarchy or entering search parameters into a local search function. The approach focuses on two dimensions, namely the anticipated cognitive effort required to use a strategy and the ability of a strategy to produce an accurate response. Individuals select the particular strategy that represents the best accuracy-effort tradeoff for the task at hand. In terms of Information Foraging Theory, people modify their strategies to maximize their rate of gaining valuable information while minimizing the costs of accessing, rendering, and interpreting information-bearing items (Pirolli & Card, 1999).

While the data from the experiments provided direct support for the specific principles argued to be encompassed by the cost-benefit perspective, such as information scent or menu breadth (to be discussed below), support was found for the cost-benefit perspective as a general view as well. For instance, the logistic regression model from Experiment 1 that sought to find factors relevant to the search / browse decision found that users who felt that using the product menu on a site would take a long period of time were more likely to use the search function to locate the item.

The probability distributions calculated in Experiment 1 provided further support that participants considered subjective time estimates as part of their decision to search or browse a site (see Figure 4 for the probability distributions for bluelight.com). Among the questions asked of participants was a request to estimate the time it would take to locate the item using menus versus using the search function. The expected search and browse times were then computed for each participant for each site (using the normalized subjective probability data). Subsequently, the difference between expected search and browse time was computed and that measure was included in the logistic regression performed to predict searching behavior. The results suggested that subjective time estimates played a role in their product finding decision.

Furthermore, cost-benefit factors were found to interact in Experiment 2 suggesting that information scent and menu breadth do not act in isolation, but rather are pieces of a larger puzzle in which users weigh the costs and benefits of particular courses of action. Even more interestingly, as will be discussed below, cost-benefit and attentional capture factors were found to interact under certain circumstances suggesting that both perspectives are directly relevant to the search / browse decision on e-commerce sites.

4.3.1. Menu Breadth

The debate over the number of options at a single level of the menu (breadth) versus the number of levels in the hierarchy (depth) has been waged for several decades. Some designers believe that breadth should be maximized while others advocate equally strongly that breadth should be minimized, thus maximizing depth.

A review of breadth versus depth research suggested that breadth was favored over depth in most studies of the organization of menu contents (Larson & Czerwinski, 1998). The increased demands on visual and cognitive processes incurred for high breadth menus were viewed to be more acceptable than those accompanying high depth menus. To this debate, empirical studies of information retrieval using various menu structures have found evidence that high breadth menus lead to faster and more accurate performance than high depth menus (Snowberry et al., 1983; Kiger, 1984).

Support for the virtues of high menu breadth was found in the controlled setting of Experiment 2. A significant main effect of menu breadth was found that provided support for the argument that high breadth menus are more likely to be used than low breadth menus. Specifically, low breadth menus were more likely to elicit searching behavior than high breadth menus. It seems clear that when a user perceives that many menu levels will need to be traversed, she appears perfectly willing to avoid the menu entirely.

The logistic regression model used in Experiment 1 to isolate factors relevant to the search / browse decision also provided support for high menu breadth. It revealed that when users perceived a need to traverse many menus to find a product, they were more likely to use the search function to find it. As discussed in the Introduction, a site that has few top-level menu options implies to the user that many menus will need to be traversed. Based on the model, users apparently take this fact very seriously when deciding whether to use the menus as their product finding strategy, often preferring to bypass the menus entirely in favor of the search function.

Deerskin.com and Supremevideo.com were selected for use in Experiment 1 mainly due to their possession of menus with very high breadth. While participants were expected to avoid such menus, these sites yielded the lowest search rates of all the sites presumably due to the very high specificity and information scent (to be discussed below) afforded by their high breadth menus. Numerous sources of evidence therefore exist to advocate the virtues of high menu breadth. Notwithstanding the corollary increases in menu option specificity and information scent (to be elaborated below), high breadth menus permit awareness of the full scope of products (or services) available on a site (Nielsen, 2000). In fact, participants in Experiment 1 often commented that they preferred to use the menus (even when a prominent search function was present) as it allowed them to see all of the items available on a site permitting them to select the best item for their money.

It is especially interesting to note that many supporters of high menu breadth often insert a disclaimer that menu breadth should not be overly great as to overwhelm the user. However, the most browsed site, Deerskin.com, possessed over thirty menu options in its top-level menu. While modest upper limits on breadth may be appropriate for menus in other interfaces, such as traditional applications, the Web (particularly ecommerce sites) appears to be a special case in that more liberal limits are acceptable. (Experiment 2 also provided support for this argument in that the high breadth menu possessed thirty options.)

4.3.2. Information Scent

One factor argued to influence the anticipated effort to a user of a particular strategy is information scent (Pirolli, 1997). Scent describes the amount of remote information a user can derive regarding the location of information based on the design or labeling of the information structure. Formally, it is defined as "the (imperfect) perception of the value, cost, or access path of information sources obtained from proximal cues" (Pirolli & Card, 1999, p. 10).

On e-commerce sites, scent could refer to the amount of information a user would be able to attain regarding the location of a product in a site based solely on the design of the site's home page. For instance, product headings on a site can vary in distinctiveness resulting in differential degrees of information scent (Larson & Czerwinski, 1998; Tilson et al., 1998). Presumably, the key determinant of the scent of a menu option is the semantic similarity between the item sought and the menu option under which it resides (see Soto, 1999).

Experiment 2, in which factors such as information scent were stringently controlled and manipulated revealed a main effect of information scent with low scent menus more likely to elicit searching behavior than high scent menus. Furthermore, the breadth and scent factors interacted in that menus with high information scent were even more likely to elicit greater browsing behavior when the menu possessed high breadth. While not significant, the interaction is intuitive when one considers the improved semantic matching provided by a greater number of menu options. Therefore, the argument that high scent lies behind the advantage of high breadth menus received empirical support.

Further supporting the role of information scent as the foundation of the menu breadth advantage was the results of Experiment 1. Although participants considered the breadth of the menus as an important factor when judging how to proceed, it is clear from their comments that information scent more strongly governed their behavior. The semantic match between the item name and menu label under which it resided on the site served as a key determinant of the decision to use the menus. In Experiment 1, Deerskin.com and Supremevideo.com had a maximum amount of information scent on many occasions as products were organized into very discrete categories (e.g., Deerskin.com: "Hats," "Pants"; Supremevideo.com: "DVD Players," "TVs"), leading to high browse rates. Further, it is important to note that Deerskin possessed a very prominent search function, suggesting that very high information scent will dominate the product finding decision process even in the presence of a prominent search function.

Qualitative analysis of the comments made by participants in Experiment 1 (see Table 3) also revealed many meaningful correlations consistent with the postulated role of information scent. For instance, it was found that when participants noted a good (semantic) match of the item to be sought and the menu label they would choose, they were more likely to use that menu option (as opposed to the search function). Considering the fact that search engines have no inherent scent, relying instead on giving users the ability to create their own scent (Scanlon, 2000), such a finding seems intuitive. Also revealed in the qualitative analysis of user comments was the finding that users who characterized the menu options on a site as vague or incomplete were more likely to use the search function to find the item. Such a result provides additional support for semantic matching as the underlying construct behind information scent.

4.4. Attentional Capture Perspective

Another major perspective that was taken to account for the users' decision to browse or search was referred to as the attentional capture perspective. Properties of the visual environment have been argued to be major determinants of whether an object in the environment draws (or "captures") attention. For instance, unique objects in terms of color or brightness have been found to capture attention in a visual display, as have objects that lie on or near visual boundaries (Todd & Kramer, 1993).

This attentional capture perspective is directly related to the notion of "visual momentum" (Hochberg & Brooks, 1978). Visual momentum describes an observer's desire to gather information from her visual environment in an attempt to form a stable understanding of her surroundings. It has been argued that visual momentum contains both an early, rapid component during which attention is drawn to visual landmarks followed by a more cognitively-driven visual analysis or "inquiry" (Hochberg & Brooks, 1978; Hochberg & Gellman, 1977).

Regarding usage of search functions on e-commerce sites, if the design of a site is such that attention is immediately drawn to a search function (text field) with high brightness contrast against the page background, then users may be more likely to search despite the unavailability of information regarding its effectiveness. Conversely, if attention is immediately drawn to menu headings, a browsing strategy may be adopted. In terms of the visual momentum approach, the search function or menu headings may serve as a landmark which influences the user's subsequent visual inquiry and information acquisition.

While the experiments provided unequivocal support for the costbenefit perspective, only qualified support was found for the attentional capture view. It seems clear that the role of graphic design factors such as the prominence of screen elements is less important than the influences of information scent or breadth. However, empirical data was collected that suggested a non-trivial role of attentional capture factors.

Initial support was obtained from the logistic regression model from Experiment 1. When participants perceived the search function as prominent, they were more likely to use it to find products. Likewise, when they perceived the menu options to be subtle, they were less likely to use them as their product finding strategy. Qualitative analysis of user comments also revealed that when users commented on the prominence of the menu on a site, referring particularly to the background contrast of the menu against the site, they were more likely to use the menu.

Although the quantitative and qualitative findings from Experiment 1 implicated attentional capture in the decision of the user to search or browse, Experiment 2 revealed that its role is more complex than originally thought. Despite the finding that subtle menus were more likely to elicit searching behavior than prominent menus, the effect failed to reach significance. Likewise, although a prominent search function was more likely to elicit searching behavior than a subtle one, the effect also failed to reach significance. Further support for attentional capture, albeit modest support, was found via interactions with other factors, as revealed by the results of Experiment 2. For example, the browse advantage for menus with high breadth (discussed above) was further pronounced when the search function was subtle (see Figure 11).

The results of Experiment 2 suggested that, independent of other factors, search function and menu prominence are generally not critical to product finding behavior *as long as the user is aware of the available options*. The marginally significant search prominence x menu breadth interaction made clear the possibility that subtle search functions may often go unnoticed when presented among broad menus. Despite the failure of search function and menu prominence main effects to reach significance in Experiment 2, however, Experiment 1 provided support for the importance of the prominence of interface elements. The sites used as stimuli in Experiment 2 were simplified relative to the vast majority of e-commerce sites on the Web and this may have diminished the potential influence of element prominence on product finding behavior.

A related finding that also supports the importance of user awareness of available options was discussed earlier and pertains to the difference found between set 1 and set 2 stimuli in Experiment 1. It was argued that set 2 stimuli yielded greater searching behavior than set 1 partially due to the fact that participants were made aware in set 2 that every site possessed a search function (due to the questionnaire). Supporting this line of reasoning, certain sites showed marked decreases in menu usage in set 2. For example, Computergear.com possessed a prominent menu with very low information scent along with a very subtle search function. In set 1, fully 87% of users used the menu despite its flaws, but for set 2, when they presumably became aware of the availability of a search function, usage of the menu fell to 30%.

Another less intuitive example of an interface element potentially going unnoticed by users is when a large graphical menu is used. A case in point is Fogdog.com from Experiment 1. It was predicted that the bold, colorful, and graphic menu options at the top of the page would lure users more than the highly prominent search function just above the menu. Generally, though, neither was as powerful at drawing users as expected. Fogdog.com also possessed a broad menu of options located lower on the page, far below the prominent main menu and search function, and it was these menu options that were often used by participants.

While this supports the breadth arguments made earlier, it also casts doubt on the effectiveness of graphic menus. It is likely that participants interpreted the graphical menu options as one large graphical heading or advertisement and paid it little mind, similar to the phenomenon of "banner blindness" (Benway & Lane, 1998), in which users often ignore large, colorful items that are clearly distinguished from other items on a Web page. However, one must not rule out the possibility that participants deliberately preferred the more numerous (and more specific) menu options to the prominent menu options that possessed a good deal of scent themselves. If so, then it would serve as even more compelling evidence for the virtues of high menu breadth.

The aforementioned interaction of cost-benefit and attentional capture factors found in Experiment 2 (menu breadth x search prominence) scratches the surface of the relationship between these two perspectives of user behavior. A later section serves to elaborate on the influence of search function and menu prominence by exploring their interactions with costbenefit factors more deeply. Before doing so, however, the role of item specificity must be discussed if one is to make sense of the interactions that were found.

4.5. Item Specificity

Task-specific factors such as goal-directed versus exploratory behavior (e.g., looking for a specific product versus any product that meets certain requirements (Steiger et al., 1998)) have been suggested to influence product finding behavior. To examine this claim, both Experiments 1 and 2 manipulated item specificity. Although the method used to operationalize item specificity in Experiment 1 was flawed, it suggested that *perceived* item specificity was more important than an arbitrary objective measure. Therefore, Experiment 2 used items that were clearly general (e.g., "sweater") or clearly specific (e.g., "Kenneth Cole Diamond Patterned Sweater") and revealed a main effect of item specificity with specific items more likely to elicit searching behavior than general items (see Figure 10).

Item specificity also played a role in two three-way interactions with one interaction involving only cost-benefit factors and the other involving an interplay between cost-benefit and attentional capture factors. The latter interaction will be discussed in a separate section below as it serves as an example of both cost-benefit and attentional capture factors interacting to explain user behavior. The remaining interaction involved item specificity and two cost-benefit factors (menu breadth and information scent) and produced intriguing results.

The item specificity x menu breadth x information scent interaction presented a formidable interpretation challenge. While completely unexpected, it suggested that item specificity may influence not only overt user behavior but also how users interpret available information presented on such sites. For some unclear reason, seeking a general item may have led users to compare the number of menu options with the number of perceived product categories available on a site. For cases in which the latter dominated the former, users may have possessed very low tolerance for menus lacking information scent.

One should keep in mind that the product finding task presented to participants is only one of many possible tasks in the realm of e-shopping. Claims made regarding product location behavior may not extend to situations in which users are asked to perform more ambiguous tasks such as "finding a gift for Bob." Such overly general item descriptions deserve special attention.

4.6. The Interplay Between Cost-Benefit and Attentional Capture Factors

The marginally significant search prominence x menu breadth interaction suggested that menus with higher breadth elicited greater browsing behavior, particularly for sites that possessed a subtle search function. It also served as evidence that cost-benefit and attentional capture factors can interact to influence product finding behavior on e-commerce sites.

Further evidence was provided by the significant three-way interaction between item specificity, menu prominence, and information scent (see Figures 14 and 15). One possible account of this three-way interaction centers on the relationship between item specificity and searching behavior. As discussed earlier, there was a substantially higher search rate for specific items than general ones and the three-way interaction suggests that this tendency affects menus of varying prominence differentially. When looking for general items, users apparently were much more willing to consider (and use) a low-scent menu if it were prominently displayed on the page. For specific items, however, the prominence of a low-scent menu did not lead to it being used more often relative to a subtle one. It appears that the strong tendency to search when asked to find a specific item is not mitigated by a prominent menu when that menu possesses low information scent.

The interactions found in Experiment 2 suggested that the cost-benefit and attentional capture perspectives are both relevant to product finding behavior and the strategic decisions users make. They also imply a complex relationship among the cost-benefit factors of menu breadth and information scent and the attentional capture factors of search function and menu prominence, a relationship made clearer by item specificity. While these experiments fail to provide a complete account of the relationships among these factors, they serve as an excellent first step toward that goal.

4.7. General Search Predisposition of Users

Besides site-specific and task-specific factors, characteristics of the user played an important role in determining product finding behavior. While there may also be effects of user gender, ethnic group, and socioeconomic status, this discussion will focus on the general predisposition that participants possessed toward the search function.

Participants in Experiment 1 often considered searching to be less time-consuming than browsing, yet many participants had qualms with choosing search as their product location strategy. Concerns relevant to the search parameters that should be used, the number of items that the search function would be expected to return, and the time for the search function to locate the items were all considered by users.
More broadly, users' general opinions and prior experiences with searching played a substantial role in their strategic decisions. Some participants were content to search on nearly every site (in Experiment 1) while others were reluctant to do so unless absolutely necessary.

The quantitative analysis from Experiment 1 revealed that users who generally liked using search functions, perceived search as an efficient method of producing a reasonable number of items (i.e., "hits"), or were confident as to the parameters required by the search function (i.e., the search term) were all more likely to use the search function to locate items. Furthermore, analysis of comments from participants in Experiment 1 yielded consistent results. Comments classifying the search function as easier, faster or more direct than using the menus was associated with more frequent search usage while comments expressing doubt in the effectiveness or ease of use of the search function were associated with less frequent search usage.

Experiment 2 sought to investigate more directly the influence of users' general opinions of using search functions. A search predisposition

index score was created for each participant and it yielded a wide distribution of scores across participants. Despite this finding, however, no clear link was found between the index scores and participant behavior in Experiment 2. While two higher-order interactions were found involving the grouping of participants as "browsers" or "searchers," the site-specific factors of information scent, menu breadth, and prominence of the menu and search function failed to interact with the grouping factor. Despite the failure to find a stronger connection between user opinions and site-specific factors, the data do not rule out such a connection.

Nevertheless, the results suggested that while users differ in their opinions of the effectiveness and value of local search functions, site design factors (particularly those relevant to information structure) exert a powerful influence over user product finding behavior that may be largely independent of individual user characteristics. In fact, regression models performed for each experiment revealed that the sites used accounted for 13% and 7% (Experiment 1 and 2, respectively) of the total variance in product location behavior. Such an argument adds further fuel to a central theme of this research, namely the overemphasis placed by many researchers on user characteristics as determinants of product finding behavior on e-commerce sites. While user-specific characteristics account for a large proportion of the variance in product location behavior (18% and 28% in Experiments 1 and 2, respectively), the data suggest that site design factors play a substantial role as well.

4.8. Design Recommendations

Our results suggest the following with respect to the design of sites with both menus and search functions. It is clear that menu usage can be encouraged by implementing a menu with options high in information scent. Incorporating menu options with close semantic matching to the items (or information) that users of a site will pursue will result in a higher proportion of users choosing the menu as their strategy.

To additionally foster menu usage, site designers can implement menus that contain many top-level options (and limited depth). Doing so can provide greater prominence to the menu and, more importantly, can substantially increase the information scent of the menu options. Users often used menus with high information scent even in the presence of a prominent search function.

The common belief that search functions are inherently faster and more accurate at locating items than product menus was not supported by our data. The message to site designers is that implementing a local search function on a site is no guarantee that users will successfully find the products (or information) they seek. Yet another reason exists for site designers and the companies that back them to get to know the target users of their sites and the tasks that they intend to use the site to accomplish.

However, search functions can be very useful aids to locating information and can serve as valuable alternatives to effective menus. Despite their limitations, (well designed) search functions should be implemented whenever possible (and appropriate). Most importantly, the decision of whether to implement a search function should be independent of the design of the site's information structure. Menus and search functions should not be considered substitutes for one another. As participants were willing to switch strategies while locating products, it would also be beneficial for designers to make such switching efficient and seamless. Tracking (and representing) past user behavior (in the current visit) may provide for a superior user experience as it will allow users to avoid redundant search queries or previously explored menu traversal paths.

Use of either menus or search functions can be encouraged by bestowing prominence to the element of interest. Large size relative to alternative options and high brightness (or color) contrast with the page background may lead to more frequent use of certain elements relative to others. However, it is strongly recommended that all major alternatives at the disposal of users be made clearly available to them. Making screen elements (such as search functions) overly subtle may result in users being unaware of their existence causing them much frustration. Furthermore, making both the search function and menus clearly available can preclude abandonment of a site if one product location method is not fruitful. Participants also occasionally commented that a subtle search function served as a warning to the user that it would be less effective than a more prominently displayed search function. Other participants went even further to assert that the site designers were attempting to coerce them into using menus so that they could be exposed to other products that they did not intend to purchase.

Perhaps the most important recommendation that can be made is one that has been advocated for many years, namely for site designers to become more knowledgeable of their users. Along these lines, it is important for designers to be aware of the general search predisposition of their user base. While users differ substantially in their views toward using search, there may be cases where users of a particular site will share certain characteristics. Some "quick-and-dirty" experiments may allow the site producers and designers to uncover such patterns.

4.9. Conclusions

As mentioned in the Introduction, searching and browsing behavior are not only relevant to e-commerce sites, but to academic and personal arenas as well. The decision to employ a search function or use a menu of options arises in many other areas such as technical support, product information, university Web sites, and educational sites (i.e., courseware). Understanding the factors that determine how a user interacts with a particular site will enable site designers to have a more substantial influence over how their site is used. This research is a first step in that direction.

Site-specific, task-specific, and user-specific factors all reliably predict whether a user will search or browse to find an item on an ecommerce site. Encompassed by these factors are graphic design elements of the site, the perceived underlying information structure of the site, the perceived efficacy of the search function itself, and more generally, the user's disposition. There are large individual differences in users' general attitudes toward using search functions that must be considered prior to implementing a search function on a site. Designers must therefore evaluate assumptions about the preferences of particular user populations before using them as the basis for design decisions. Together, Experiments 1 and 2 lead to the conclusion that the information structure of e-commerce sites is a far better predictor of user behavior than the manner in which that information is presented. Specifically, the graphic design of e-commerce sites appears to be given far greater emphasis than it deserves, consistent with other studies suggesting no link between graphic design and success at finding information (see Spool et al., 1997).

Based on the data, the following speculative view of "visual momentum" on e-commerce sites will be put forth. While the data do not provide direct support for the following view, it is intended to serve as a potential framework for further research. It is argued that a user who first arrives at an e-commerce site with a particular product in mind will first have her attention drawn to a particular element (or elements) of the page. This will occur based on size, color contrast, and brightness contrast of the element with the background of the page. Then, the user will evaluate that option in terms of its costs and benefits for finding the desired information, which is heavily determined by the information scent of the option. If the element meets some arbitrary cost-benefit cutoff, it will be pursued. If not, the next element to capture attention will be evaluated and so on until a satisfactory option presents itself. If the first element to capture attention fails to meet some minimum cost-benefit cutoff, the site may be abandoned altogether (as in cases when a large but worthless graphic covers the majority of screen real estate). Likewise, if each element meets the costbenefit minimum but all fail to meet the cutoff of further investigation, the site may be abandoned. According to this view, it is critical to draw attention quickly to a cost-effective option or the site may be at high risk of losing the user, perhaps permanently.

Of course, this view is undoubtedly an oversimplification of the interaction of users with e-commerce sites. Users of a site may be somewhat familiar with the site due to word of mouth, for example, and may have formed expectations of the information or products available on the site. Furthermore, there may be effects of user experience and expertise that affect product (or information) finding behavior on sites. Nevertheless, the version of visual momentum offered here is a reasonable framework for further investigation of user behavior on e-commerce sites.

Encouragement of menu usage through the implementation of menus with high scent and breadth may yield an additional unexpected benefit. Although lacking empirical support, there may exist many potential advantages for a company to consistently encourage menu traversal on their site. The increase in familiarity of the site organization and contents obtained by traversing the menus (as opposed to using a search function) may lead to time savings in using the menus in the future as well as increased user exposure to additional products. This may be an especially beneficial decision if the site designers have doubts about the effectiveness or ease-of-use of the local search function, or if they are unwilling (or unable) to implement one at all. Furthermore, repeated use of the menus may encourage the user to return to the site in the future (assuming the site organization is effective).

It is argued that it is in the best interest of Web site designers to assume that users would prefer to use product menus if it is reasonably effective for them to do so. As discussed throughout, a search function is often a complex method to locate a product. The user must enter the correct words, correctly specify the syntax, and properly interpret the results (Scanlon, 2000). It is possible that the a priori preference to use search functions possessed by many users may be due directly to negative experiences with using product menus in the past.

A more bold prediction on the advantages of encouraging menu usage on a site is that it may lead to increased impulse buying. Among the findings to emerge from previous modeling studies of online shopping behavior is that a major predictor of online buying is the user behavior of looking for product information online (Bellman, Lohse, & Johnson, 1999). The more likely users are to look for product information online, the more likely they are to buy products over the Internet. A sensible extrapolation of this finding is to suggest that increased exposure to products online may lead to a greater likelihood that a user will buy products over the Web. Such an hypothesis has been put forth by others (Lohse & Spiller, 1998a), but to date no significant relationship has been found. However, during a beta-test of an

electronic supermarket, in which the interface allowed browsing of menus but was without a search function, several users reported the desire to make impulse purchases (Henderson et al., 1998). Clearly, much more research is needed to determine the link, if any, between user product finding behavior and online sales. Hopefully, this set of studies will serve as an excellent springboard for such a research endeavor.

REFERENCES

- Alba, J., & Hutchinson, J. W. (1987). Dimensions of consumer expertise. Journal of Consumer Research, 13, 411-445.
- Alba, J., Lynch, J., Weitz, B., Janiszewski, C., Lutz, R., Sawyer, A., & Wood, S. (1997). Interactive home shopping: Incentives for consumers, retailers, and manufacturers to participate in electronic marketplaces. *Journal of Marketing*, 61(July), 38-53.
- Beach, L. R., & Mitchell, T. R. (1987). A contingency model for the selection of decision strategies. *Academy of Management Review*, 3, 439-449.
- Bellman, S., Lohse, G., & Johnson, E. (1999). Predictors of online buying behavior. *Communications of the ACM*, 42(12), 32-38.
- Benway, J. P., & Lane, D. M. (1998, December). Banner blindness: Web searchers often miss "obvious" links, *Internetworking* [On-line], **1.3**.
- Borgman, C. (1986). The user's mental model of an information retrieval system: An experiment on a prototype online catalog. *International Journal of Man-Machine Studies*, **24**, 47-64.
- Campagnoni, F. R., & Ehrlich, K. (1989). Information retrieval using a hypertext-based help system. *Proceedings of the Twelfth Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, 212-220.
- Chase, W. G. (1986). Visual information processing. In K. R. Boff, L. Kaufman, & J. P. Thomas (Eds.), *Handbook of perception and human performance, Vol. II: Cognitive processes and performance* (pp. 28-1 to 28-71). John Wiley and Sons.

- Chi, M. T. H., Glaser, R., & Farr, M. J. (Eds.). (1988). *The nature of expertise*. Hillsdale, NJ: Erlbaum.
- Deerwester, S., Dumais, S., Furnas, G., Landauer, T., & Harshman, R. (1990). Indexing by latent semantic analysis. *Journal of the American Society for Information Science*, **41**, 391-407.
- Donthu, N., & Garcia, A. (1999). The internet shopper. *Journal of Advertising Research*, **39**(3), 52-58.
- Ericsson, A. K., & Simon, H. A. (1993). *Protocol analysis*. Cambridge: MIT Press.
- Evans, G. W. (1980). Environmental cognition. *Psychological Bulletin*, 88, 259-287.
- Fasolo, B, & McClelland, G. (1999, November). *Tracing decision processes on the Web*. Poster presented at the annual meeting of the Society of Judgment and Decision Making.
- Golledge, R. G., Smith, T. R., Pellegrino, J. W., Doherty, S., & Marshall,
 S. P. (1985). A conceptual model and empirical analysis of children's acquisition of spatial knowledge. *Journal of Experimental Psychology*, 5, 125-152.
- Hastie, R. (1991). A review from a high place: The field of judgment and decision making as revealed in its current textbooks. *Psychological Science*, **2**, 135-138.
- Henderson, R. Rickwood, D., & Roberts, P. (1998). The beta test of an electronic supermarket. *Interacting with Computers: The Interdisciplinary Journal of Human-Computer Interaction*, **10**(4), 385-399.

- Hochberg, J., & Brooks, V. (1978). Film cutting and visual momentum. InJ. W. Senders, D. F. Fisher, & R. A. Monty (Eds.), *Eye movements* and the higher psychological functions. Hillsdale, NJ: Erlbaum.
- Hochberg, J., & Gellman, L. (1977). The effect of landmark features on mental rotation times. *Memory and Cognition*, 5, 23-26.
- Jacko, J. A., & Salvendy, G. (1994). Incorporating task complexity into hierarchical menu design. *Proceedings of the Human Factors and Ergonomics Society, 38th Annual Meeting*, 967.
- Jarvenpaa, S. L. (1989). The effect of task demands and graphical format on information processing strategies. *Management Science*, **35**, 285-303.
- Jarvenpaa, S. L. (1990). Graphic displays in decision making: The visual salience effect. *Journal of Behavioral Decision Making*, **3**, 247-262.
- Jarvenpaa, S. L., & Todd, P. A. (1997). Consumer reactions to electronic shopping on the World Wide Web. *International Journal of Electronic Commerce*, 1(2), 59-88.
- Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision making under risk. *Econometrica*, 47, 263-291.
- Kiger, J. I. (1984). The depth/breadth tradeoff in the design of menudriven interfaces. *International Journal of Man-Machine Studies*, 20, 201-213.
- Kleinmuntz, D. N. (1991). Decision making for professional decision makers. *Psychological Science*, 2, 138-141.

- Kleinmuntz, D. N., & Schkade, D. A. (1993). Information displays and decision processes. *Psychological Science*, **4**(4), 221-227.
- Koman, R. (1998). The scent of information: Helping users find their way by making your site "smelly." [On–line]. Available: *Http://webreview.com/wr/pub/98/05/15/feature/index.html*.
- Landauer, T., Foltz, P., & Laham, D. (1998). Introduction to latent semantic analysis. *Discourse Processes*, **25**, 259-284.
- Larson, K., & Czerwinski, M. (1998). Web page design: Implications of memory, structure, and scent for information retrieval. *Proceedings of CHI* '98 Human Factors in Computing Systems, 25-32.
- Ledgard, H., Singer, A., & Whiteside, J. (1981). *Directions in human factors for interactive systems*. New York: Springer-Verlag.
- Lohse, G. L. (1997). Consumer eye movement patterns on Yellow Pages advertising. *Journal of Advertising*, **26**(1), 1-13.
- Lohse, G. L., & Spiller, P. (1998a). Quantifying the effect of user interface design features on cyberstore traffic and sales. *Proceedings of CHI '98 Human Factors in Computing Systems*, 211-218.
- Lohse, G., & Spiller, P. (1998b). Electronic shopping. *Communications of the ACM*, **41**(7), 81-87.
- Marchionini, G. (1989). Making the transition from print to electronic encyclopedias: Adaptation of mental models. *International Journal of Man-Machine Studies*, **30**, 591-618.

- McEwen, S. A. (1981). An investigation of user search performance on a Telidon information retrieval system. In *The design of videotex tree indexes* (pp. 35-64). Ottawa, Ontario: Behavioral Research and Evaluation, Department of Communications, Government of Canada.
- Miller, G. A. (1956). The magical number seven plus or minus two: Some limits on our capacity for processing information. *Psychological Review*, **63**, 81-97.
- Mullet, K., & Sano, D. (1995). *Designing visual interfaces: Communication oriented techniques*. Englewood Cliffs, NJ: Sunsoft Press.
- Neal, D. (1995). Electronic commerce: Design of the AT&T Internet tollfree 800 directory. *Proceedings of the Human Factors and Ergonomics Society, 39th Annual Meeting,* 215-216.
- Nielsen, J. (1996, May). Top ten mistakes in Web design. [On-line]. Available: www.useit.com/alertbox/9605.html.
- Nielsen, J. (2000). *Designing Web usability: The practice of simplicity*. Indianapolis, IN: New Riders Publishing.
- Norman, K. L. (1991). *The psychology of menu selection: Designing cognitive control at the human/computer interface*. Norwood, NJ: Ablex.
- Payne, J. W. (1982). Contingent decision behavior. *Psychological Bulletin*, 92, 382-402.
- Payne, J. W., Bettman, J. R., & Johnson, E. J. (1993). *The adaptive decision maker*. Cambridge: Cambridge University Press.

- Phillips, C., & Thomas, P. (1988, August 19). Texas Air asks transportation agency to examine rivals' reservation systems. *Wall Street Journal*, p. 13.
- Pirolli, P. (1997). Computational models of information scent-following in a very large browsable text collection. *Proceedings of CHI '97 Human Factors in Computing Systems*, 3-10.
- Pirolli, P., & Card, S. (1995). Information foraging in information access environments. *Proceedings of CHI '95 Human Factors in Computing Systems*, 51-58.
- Pirolli, P., & Card, S. (1999). Information foraging. *Psychological Review*, **106**(4), 643-675.
- Pirolli, P., Card, S., & Van Der Wege, M. (2000). The effect of information scent on searching information visualizations of large tree structures. In V. Gesu, S. Levialdi, & L. Tarantino (Eds.), Advanced Visual Interfaces 2000 (pp. 161-172). New York: ACM Press.
- Polson, P. G., & Lewis, C. H. (1990). Theory-based design for easily learned interfaces. *Human-Computer Interaction*, 5(2-3), 191-220.
- Powell, T. A. (2000). *The complete reference: Web design*. Berkeley, CA: Osborne/McGraw-Hill.
- Rada, R., & Murphy, C. (1992). Searching versus browsing in hypertext. *Hypermedia*, 4(1), 1-30.
- Russo, J. E. (1977). The value of unit price information. *Journal of Marketing Research*, **14**, 193-201.

- Russo, J. E., & Dosher, B. A. (1983). Strategies for multiattribute binary choice. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 9, 676-696.
- Sanderson, P. M., McNeese, M. D., & Zaff, B. S. (1994). Handling complex real-world data with two cognitive engineering tools:
 COGENT and MacSHAPA. *Behavior Research Methods, Instruments, & Computers,* 26(2), 117-124.
- Sano, D. (1996). *Designing large-scale Web sites: A visual design methodology*. New York: Wiley & Sons.
- Sawyer, P. (2000). *Links that give off scent*. North Andover, MA: User Interface Engineering, Inc. Report.
- Scanlon, T. (2000). *On-site searching and scent*. North Andover, MA: User Interface Engineering, Inc. Report.
- Science and Applications of Latent Semantic Analysis Group. (2001). [Online]. Available: *http://lsa.colorado.edu*.
- Snowberry, K., Parkinson, S. R., & Sisson, N. (1983). Computer display menus. *Ergonomics*, 26(7), 699-712.
- Soto, R. (1999). Learning and performing by exploration: Label quality measured by latent semantic analysis. *Proceedings of CHI '99 Human Factors in Computing Systems*, 418-425.
- Spool, J. M., & Klee, M. (2000). Designing for scent. North Andover, MA: User Interface Engineering, Inc. Report.

- Spool, J. M., Scanlon, T., Schroeder, W., Snyder, C., & DeAngelo, T. (1997). Web site usability: A designer's guide. North Andover, MA: User Interface Engineering.
- Steiger, P., Stolze, M., & Good, M. (1998). Beyond Internet business-asusual: A CHI 98 workshop. SIGCHI Bulletin, 30(4), 48-52.
- Sternberg, S. (1966). High-speed scanning in human memory. *Science*, **153**, 652-654.
- The Conference Board, Inc. (1996). *Commerce in Cyberspace: A Conference Report*. Edited by Joseph L. McCarthy. New York: The Conference Board, Inc.
- Tilson, R., Dong, J., Martin, S., & Kieke, E. (1998). A comparison of two current e-commerce sites. ACM 16th International Conference on Systems Documentation, 87-92.
- Todd, S., & Kramer, A. (1993). Attentional guidance in visual attention. Proceedings of the Human Factors and Ergonomics Society, 37th Annual Meeting, 1378-1382.
- Tullis, T. S. (1981). An evaluation of alphanumeric, graphic, and color information displays. *Human Factors*, 23, 541-550.
- Tullis, T. S. (1988). A system for evaluating screen formats: Research and application. In H. R. Hartson & D. Hix (Eds.), *Advances in Human-Computer Interaction, Volume 2* (pp. 214-286). Norwood, NJ: Ablex Publishing Corporation.
- Tversky, A. (1972). Elimination by aspects: A theory of choice. *Psychological Review*, **79**, 281-299.

- Tversky, A., Sattath, S., & Slovic, P. (1988). Contingent weighting in judgment and choice. *Psychological Review*, **95**, 371-384.
- Vinson, N. G. (1999). Design guidelines for landmarks to support navigation in virtual environments. *Proceedings of CHI '99 Human Factors in Computing Systems*, 278-285.
- Woods, D. D. (1984). Visual momentum: A concept to improve the cognitive coupling of person and computer. *International Journal of Man-Machine Studies*, **21**, 229-244.
- Wright, P. L. (1975). Consumer choice strategies: Simplifying vs. Optimizing. *Journal of Marketing Research*, **11**, 60-67.
- Zipf, G. K. (1949). *Human behavior and the principle of least effort*. Cambridge, MA: Addison-Wesley.

APPENDIX A: SELECTED SITES USED IN EXPERIMENT 1

Bluefly.com



Deerskin.com



Fogdog.com



Jjill.com



treatyourself



gossamer skirt An inspired creation of nearly weightless sik organza layered over a dyed-to-match lining. Seam detailing and no-waistband styling create a flattering line. Side-zip entry. 37" length, 33" petite, 40" tall. Dry clean. Imported. Lavender or black (not shown). Only at J. Jill.

\$98.00

<u>buy this item</u> see more must haves



Nike.com



APPENDIX B: DECISION AND USER PROFILE QUESTIONNAIRES

Decision Questionnaire

1) What is your estimate of the probability that the item you are looking for can be found under the menu selection you chose? _____

2) On a scale of 1 (not at all) to 10 (very), how related in meaning is the label of the menu option you chose to the item you were looking for?

3) When you reach a page with the item you're looking for, how many other items do you expect to see on the page? _____

4) After clicking on that menu option, how many other menus do you expect to have to go through to find the product? _____

5) How long do you expect it will take to find the product using the menus?

6) What amount of time would you consider to be too long to find a product?

7) Did you consider time as a factor in choosing to browse the menus versus using the Search function? _____

If yes, in what way?

8) Which of the following characterizations of the number of menu options is most appropriate (circle one): Not enough / About the right amount / Too Many

9) Do you think the designers of this page wanted you to use the menu or use the Search function to find this item? _____

Why?

10) On a scale of 1 (low) to 10 (high), rate how clear you would be as to what you should type in the Search field to find the item.

Did this affect your decision to not use the Search function?

What would you type into the Search field?

11) On a scale of 1 (not at all) to 10 (very much), how would you agree with the statement: I immediately knew what to do to find the product as soon as I saw the site.

12) On a scale of 1 (not at all) to 10 (very), how noticeable is the menu of options on the page? _____

Why?

Did that affect your decision to use the menu?

13) On a scale of 1 (not at all) to 10 (very), how noticeable is the Search option on the page?

Why?

14) Was either the menu or Search function the first thing you noticed on the page? _____

If not, what was the first thing you noticed?

Where in the order of things you noticed on the page was the menu and Search function (i.e., second, third, etc.)?

Why do you think that is?

15) How long do you think you will have to wait once you select a menu option for the next page to appear?

16) What is your estimate of the probability that the Search function would return the item you were looking for? _____

17) How many items would you expect the Search function to return?

18) How long do you think you would have to wait for the Search function to return the results of the search?

19) How long do you expect it would take to find the item using the Search function?

20) Have you seen and/or used this site before?

Have you been to other sites that sell the same kind of products?

If so, which ones?

User Profile Questionnaire

1) How much general computer experience do you have? (check one)
I have used computers for more than 10 years
I have used computers for between 5 and 10 years
I have used computers for between 1 and 5 years
I have used computers for less than 1 year
2) How much prior experience do you have using the Web?
More than 1 year of regular use (at least 2 hours per week)
Between 6 months and 1 year of regular use
Between 1 and 6 months of regular use
Less than 1 month of regular use
3) How often do you currently use the Web?
More than 10 hours per week
Between 5 and 10 hours per week
Between 1 and 5 hours per week
Less than 1 hour per week

4) Have you ever purchased anything online before?

____Yes

____ No

5) How many items have you ever purchased online?

_____ More than 20

_____ Between 10 and 20

_____ Between 5 and 10

_____ Less than 5

6) List every site that you have ever bought products from (that you can remember):

7) How often do you look online for items to buy?

_____ More than 10 times per week

_____ Between 5 and 10 times per week

_____ Between 1 and 5 times per week

_____ Less than 1 time per week

8) When you shop for an item online, how much time do you usually spend looking for the item?

_____ More than 1 hour

_____ Between 30 minutes and 1 hour

_____ Between 10 and 30 minutes

_____ Less than 10 minutes

9) How much Web programming experience do you have?

_____ I have created websites that include text and images, and I am familiar with the use of both tables and layers

_____ I have created websites that include text and images, but I am not familiar with the use of both tables and layers

_____ I have created websites that include only text

_____ I have never created a website

10) If you are shopping for a specific item online, how would you most likely find the site on which to buy it? (Please rank order the choices)

_____ Use a search engine (like Yahoo)

_____ Word of mouth

_____ Commercials or advertisements

_____ Other: _____

11) If you are shopping for a specific item online and you have found the site on which to buy the item, how do you usually find the item on the site? (Check all that apply)

_____ Browse the site using the product menus

_____ Use the Search function on the site

_____ Scan the homepage of the site to see if the item is featured there

_____ Other: _____

12) If you checked more than one choice on question #11, rank the order in which you would usually try each option. (First (1), Second (2), Third (3), etc.)

_____ Browse the site using the product menus

_____ Use the Search function on the site

_____ Scan the homepage of the site to see if the item is featured there

_____ Other: ______

13) Complete this sentence as it pertains to you: "When I am shopping online, I would say that I use the menus to find a product on a site _____% of the time and the Search function _____% of the time."

APPENDIX C: PREDICTORS USED IN LOGISTIC REGRESSION AND CORRESPONDING QUESTIONNAIRE ITEMS

Predictor used in Logistic Regression

Number of menus expect to have to go through to find product

How long it will take to find products using menus

Considered time as a factor in searching vs. browsing

How clear as to what should type into search field

How noticeable is menu of options on page

How noticeable is search option on the page

How many items would you expect search to return

Question on Questionnaire or User Profile

After clicking on the menu option you would choose, how many other menus would you expect to have to go through to find the product? How long would you expect it

will take to find the product using the menus?

Did you consider time as a factor in choosing to use the Search function versus browse the menus?

On a scale of 1 (low) to 10 (high), rate how clear you are as to what you should type in the Search field to find the item.

On a scale of 1 (not at all) to 10 (very), how noticeable is the menu of options on the page?

On a scale of 1 (not at all) to 10 (very), how noticeable is the Search option on the page?

How many items do you expect the Search function to return?

Predictor used in Logistic Regression

How long will it take for Search to find item

How often use search function to find products when shopping

Question on Questionnaire or User Profile

How long do you expect it will take to find the item using the Search function?

Complete this sentence as it pertains to you: "When I am shopping online, I would say that I use the search function to find a product on a site ____% of the time."

APPENDIX D: CATEGORIZATION OF QUALITATIVE QUESTIONNAIRE RESPONSES

Questionnaire Question	Response Categories	Sample Response
Did you consider time as a factor in choosing to use the Search function versus browse the menus? If yes, in what way?	Search is Easier / Faster / More Direct	Search engine is easier
	Search is Hidden / Obscure	Search is hidden; didn't consider
	Menus are Faster / Easier	I can find it faster using the menu
	Unfamiliar with product so browsing would take too long	I have no idea what I'm looking for so using the search function would probably be faster
	Want to see product, so browsed	Want to see the piece of furnituretime not important
Questionnaire Question

Response Categories

Sample Response

Did you consider time as a factor in choosing to use the Search function versus browse the menus? **If yes, in what way?** (*Continued*) Unsure which menu to choose

Good match of item and menu

Take to long to read all menu options

Search may not work or will be confusing

Search requires less wading through of options / screens I don't know exactly what this product is/does so it would be hard for me to know what menus to go to

I would choose the menu item on this one because it would likely take me to the product Well, I thought it would be faster not to have to read through all the detailed choices Don't want to type in, might not give what I want

Searching usually bypasses other menus

Specific product so faster to use Search

Questionnaire Question **Response Categories**

Did you consider time as a factor in choosing to use the Search function versus browse the menus? **If yes, in what way?** (*Continued*) Name of product so long, would rather browse than type in Search

Search will be unhelpful if enter whole product name

Promote Buying of /

Exposure to Other

Products

Do you think the designers of this page wanted you to use the Search function or the menu to find this item? **Why**?

Menus are comprehensive

I thought it would be more efficient to use search since I have a specific item

Sample Response

In this case, it seemed to be faster to click through a couple of menus than actually type in the name

The search may not give me what I want if I enter the whole name (or it might not give me anything) They want you to browse and see what they have so you'll buy more

There are a lot of menu items, and they are thorough Menu is more prominent

There is a great emphasis on menus

Questionnaire Question	Response Categories	Sample Response
Do you think the designers of this page wanted you to use the Search function or the menu to find this item? Why? (continued)	Search is prominent	Search function very prominentmenu headings smaller
(commutu)	Menu is organized well (clear and unambiguous)	The menu options were carefully divided and specific
	So many menu options	So many optionswho has patience to look at them all?
	Menu options are vague / incomplete	Menus don't seem to help identify the product
	Search is built into menu	[Search] built into the menu directly
On a scale of 1 (not at all) to 10 (very), how noticeable is the menu of options on the	Large (or relatively large)	They are bold and large

page? Why?

QuestionnaireResponse CategoQuestion		ries Sample Response	
On a scale of 1 (not at all) to 10 (very), how noticeable is the menu of options on the page? Why?	Good contrast with background	The menu is brighter than most of the page	
(Continued)	Different color	Standard position, different color	
	Good placement	On the left	
	Good Font Choice	Simple, relatively large font	
	Accompanied by Pics / Icons	The icons	
	Ads / images draw attention	The ads draw too much attention	

Questionnaire Question	Response Categories	Sample Response
On a scale of 1 (not at all) to 10 (very), how noticeable is the menu of options on the page? Why? (<i>Continued</i>)	Small	small
(Bad Font Choice or size	Pretty clear on the left frame, but small font
	Poor Contrast w/ Background	Small bad color and font, kind of blends with rest of page
	Unclear as to which elements were menu options	Menu options aren't very clear
	Bad placement	The options are in bold, but they are not at the top or in the center of the page
On a scale of 1 (not at all) to 10 (very), how noticeable is the search option on the page? Why?	Large	it is big and the word search is in clear, bright font

Questionnaire Question	Response Categories	Sample Response
On a scale of 1 (not at all) to 10 (very), how noticeable is the search option on the page? Why? (<i>Continued</i>)	Good Placement	It's at the top of the menu
	Good Contrast (color/font)	High contrast, large size compared to other text
	Icon present	Small, but top of page with own picture icon
	Text field present	Right under site icon - - actual text box
	No Text field / "icon"	It's not even a separate box like it usually is
	Bad Contrast (color/font)	The word "search" is small and the colors white and olive do not contrast enough

Questionnaire Question	Response Categories	Sample Response
On a scale of 1 (not at all) to 10 (very), how noticeable is the search option on the page? Why ?	Bad Placement	sits off to the side
(Continued)	(Relatively) Small	Too small, hidden at top of page
	Looks like an Ad or title	Though it is at the top of the pageit looks like an advertisement
	Not called "Search"	I don't see an obvious one, unless "product finder" is a search option
(What was the first thing you noticed?) Why do you think that is?	Animation / Scrolling	The pictures and flashing NIKE ID icon are distracting
	Designers want you to use / notice element	Because the designers wanted you to become distracted at first

Questionnaire Question	Response Categories	Sample Response
(What was the first thing you noticed?) Why do you think that is? (<i>Continued</i>)	Colors	colorsdark menu buttons
	Size	Menu is so enormously sized and centeredonly pics on the page are in the menu
	Placement	Order of appearance as you scroll down the site
	Contrast	The menu is very bright and obvious
	Icons / Graphics	My attention was drawn to all the graphics first. Lots of different menus
	Buy/see other products on page	They want me to see as many products as possible. Having pictures of products does that. The menu does that.

Questionnaire Question	Response Categories	Sample Response	
(What was the first	Was looking for the	Llooked for menu	
thing you noticed?)	element	T looked for menu	
Why do you think			
that is? (Continued)			
	Famous Face	Barkley is such a well-	
		known athlete	

APPENDIX E: SELECTED SITES USED IN EXPERIMENT 2

Site 1 - Fashion Plus Clothing

(Prominent search, prominent menu, high breadth, high scent)

Fashion Plus				
Catalog Request	Customer Service Returns & Excl	hanges Our Privacy Policy Our Gua	rantee	
	Search			
	Men's Department	Women's Department		
	Sweaters	Gloves		
1 60	Ties	Jeans	6	
0 13	Pants	Robes	Real and	
	Shorts	Shoes	12	
	Boxer Shorts	Jackets		
	v ests Suita	Scarves	14	
	Socks	Dresses		
	Belts	Bras		
	Blazers	Panty Hose	199 S	
	Watches	Handbags		
	Coats	Raincoats		
	Cuff Links	Cosmetics		
- 🧟 🗋		Earrings		
		Hair Accessories		

(Prominent search, subtle menu, high breadth, low scent)



Outerwear | Active Wear | Business Wear | Weekend Wear | Evening Wear | Relaxed Wear | Lounge Wear | Gifts for Him | Gifts for Her | Sale Items | What's Hot | What's New | Specialty Items | Cool Stuff | Best Sellers | Must Haves | Fashion by Season: Fall Styles | :Winter Styles | :Spring Styles | :Summer Styles | Basics | Classics | Our Top Picks | Rare Finds | Club Wear | Fashion by Region: L.A. Fashions | :New York Fashions | :Miami Fashions

Site 11 - TRI Apparel

(Subtle search, prominent menu, low breadth, high scent)



Site 16 - Fireball Outfitters

(Subtle search, subtle menu, low breadth, low scent)

		Fir	eball C	Dutfitters
Ordering Information	Shipping Options	Fireball Quarterly	Email Updates	Sign In
	Like a rocke Fireball Out E∨en befor are working	t straight from th fitters offers the re the models lea to bring the fres to	e runways of F styles of today ve the runway hest and hotte you.	Paris and Madrid, y and tomorrow. , our designers st styles straight
Just check out what we have offer. You'll be glad that you o Nothing is worse than being o of the loop, and fashion is n exception.	to lid. but o		- A	
			2	
Basics A Rela	.ctive Wear Business W scd Wear Classics Spe	ear Weekend Wear Ev cialty Items Must Haves	ening Wear Search	