

How Much of It is Real? Analysis of Paid Placement in Web Search Engine Results

Scott Nicholson (scott@scottnicholson.com), Assistant Professor
Syracuse University School of Information Studies
4-206 Center for Science and Technology
Syracuse, NY 13244

Tito Sierra, Digital Technologies Development Librarian
North Carolina State University Libraries
Raleigh, NC

U. Yeliz Eseryel, Ji-Hong Park, Philip Barkow, Erika J. Pozo, Jane Ward
Syracuse University School of Information Studies
Syracuse, NY

Abstract

Most Web search tools integrate sponsored results with results from their internal editorial database in providing results to users. The goal of this research is to get a better idea of how much of the screen real estate displays “real” editorial results as compared to sponsored results. The overall average results are that 40% of all results presented on the first screen are “real” results, and when the entire first Web page is considered, 67% of the results are non-sponsored results. For general search tools like Google, 56% of the first screen and 82% of the first Web page contain non-sponsored results. Other results include that query structure makes a significant difference in the percentage of non-sponsored results returned by a search. Similarly, the topic of the query can also have a significant effect on the percentage of sponsored results displayed by most Web search tools.

Introduction

Since 1996, there has been a growing trend toward reliance on the Internet for information needs. In order to find this information, many users start their exploration with a Web search tool. Most Web search tools are for-profit enterprises and use advertising to bring in enough income to allow them to continue offering their services for free to the public. Most search tools offer some type of targeted marketing where advertisements are connected to certain search terms. Some ads are placed on the top, side, and bottom of search results, while other ads are placed within the context of the “real” results. The result is that a searcher may be faced with more ads than actual editorial results. The goal of this research work is to gain a better understanding of the amount and type of advertising presented to search tool users.

This type of advertising is not necessarily bad or inappropriate. A user looking for products or services may have their information need met through this type of contextual advertising. This model of large eye-catching advertisements to support an information source has been employed for decades through our commercial telephone directories. The food and beverage industry sells

logo space on tabletop displays and menus, which can draw the eye from the text listing of the products available. Product placement outside of commercials has become more rampant in television shows and radio personalities plug products during their normal monologues. This model of contextual advertising has therefore been used and accepted in other domains.

There are some concerns about this type of contextual advertising. In some cases, the advertisements look similar to or are indistinguishable from non-paid listings relevant to the search terms (Wouters, 2004). Users expect these tools to provide information that is the best match to their query; instead they receive sponsored results from the highest bidder. As search tools become publicly-owned companies, they have to answer to a profit-seeking ownership. Thus, the temptation to provide more contextual advertising and fewer non-paid results grows.

In order to understand to what level search tools are providing advertisements over non-paid results, the research team submitted a series of controlled queries to eight different search tools. We then determine the amount of paid placement in the first screen and entire first page of results and identify patterns based on topic, query formulation, and type of search tool.

Research Questions

In order to explore paid placement in Web search tools, the following research questions are used:

1. Over all of the queries and tools, what mixture of content is displayed on the screen? Do different search tools present different amounts of advertising for the same searches? How does examining only the first screen of results differ from viewing the entire first page?
2. How much advertising is displayed for different searching topics? Do academic search topics produce less advertising than commercial search topics? Are there differences in the common topical searching areas?
3. How much advertising is displayed for different query structures? Do longer queries produce less advertising?
4. Are there patterns of interest in the data across facets of searches and tools that lead to future research questions?

Definitions and Assumptions

The primary definition to be discussed is that of an “editorial listing,” or the “real” results from a search tool. A search tool sends out spidering programs to Web pages that return with words and other information located on that Web page. These terms are integrated into an inverted index and an editorial listing is created to represent that Web page. This index is used in conjunction with matching and weighting algorithms to determine the pages returned for a query. In some cases, the pages in this primary index are there through paid inclusion, where a Web site can pay for rapid consideration for inclusion into a Web search tool index. There is no way to discern between pages that are in the primary index because a Web exploration program (a.k.a. spider) discovered them or because of paid inclusion. Thus, all of these results are grouped into the

category of “editorial listings” that are presented to a user purely because they are the most relevant items in the primary database.

The term “editorial listing,” used by others in the literature (Phillips, 2001, for example), does not imply that there were any decisions made by a human about the quality of the site. However, each search tool does have its own policies as to what pages to accept into the database. Most search tools will not index every page on the Internet; there are some basic quality standards that guide their selection. Many times, search tools will not index pages that are using tricks to raise their ranking, such as very small text or text that is the same color as the background. In addition, some tools do not index every page on a site, and index only those pages in the top levels of the site. Therefore, while the criteria a tool uses for page inclusion may not be as strict as a human’s selection, there is usually some type of editorial policy guiding what pages are indexed in the tool. The term “editorial listing” (or “real” result) reflects that these pages are part of the Web page database because they fit the editorial policies of the search tool, as automatically detected by a Web spider.

The other type of result provided to users is that of a sponsored advertisement. These listings are not visited by the traditional search tool spider, and the sponsor determines the content of the listing. Many times, sponsors pay to be listed when a user searches on specified terms, and the amount they are willing to pay determines their prominence. These listings may be graphical advertisements on the top or side of the screen or may be textual results placed within the context of the editorial results. Both of these are considered sponsored results, as they present choices to the user that represent other Web pages. The contrast in the current study is between the surrogates for pages chosen for relevancy, i.e. editorial listings, or surrogates chosen because someone paid the search tool to have them included, i.e. sponsored listings.

As with any search tool research, one assumption is that patterns found in Web search tool results are stable over time. This assumption is short-lived - as the search tools change, the results of the study will also change; this is the nature of Web research. The goal of this study is to provide patterns found from a snapshot of time, with hopes that these patterns will be useful in analyzing future situations.

Another assumption is derived from the measurement tool used for the study, which is based on one used by Nielsen and Tahir (2002). They used “screen real estate” as a measure in examining the content components of Web pages. The use of this measure is based on the assumption that the amount of screen real estate used to display results in a Web search tool is correlated to the amount of information provided by those results. This is not necessarily true in all cases. However, in this study this assumption allows us to discuss the concepts using a unit of measure that is comparable between tools – the percentage of a screen dedicated to a type of information. Throughout this paper, terms such as “amount of information” do not refer to relevance, the user’s perception of Web links, the number of links, or the content of those links. All of these possible measures represent areas for other research projects. In this paper, however, the measure used is the percentage of screen real estate.

Literature Review

There has been considerable discussion of paid placement in Web search engines in popular literature, but very little in academic literature. Trade articles and opinion pieces citing the

dangers of paid placement for users are plentiful, though the extent to which paid placement influences results of Web search engines has not been comprehensively measured. This literature review focuses on the user issues related to paid placement in Web search tools such as bias, consumer awareness, and disclosure. Additionally, it highlights prior research supporting the methodology choices made in the present study.

User Issues

In their paper on bias and search engines, Mowshowitz and Kawaguchi (2002a) argue that bias is present when some items occur more frequently or prominently with respect to the norm, while others occur less frequently or prominently with respect to the norm. In their definition, bias may be introduced into a retrieval system at any stage of its operation. Sources of bias include the rules used for document inclusion, the way items are indexed, the manipulation of text in Web pages to increase retrieval, query formulation constraints or predispositions, and the search algorithms used to retrieve and rank items. Although Mowshowitz and Kawaguchi do not mention the role of paid placement explicitly, it is clear that sponsored advertising introduces bias in Web search engine results.

In a subsequent study, Mowshowitz and Kawaguchi (2002b) measure the bias present in fifteen commercial search engines using statistical analyses of variance in results returned. Bias values are computed by analyzing the top thirty URLs returned by each search engine for search terms across eight subject areas. The results suggest that their measure of bias can discriminate between search engines, but for most search engines bias does not depend on the subject domain searched, or the search terms used to represent that subject domain. The current research explores this finding to see if bias from sponsored placement depends upon the search topic.

Bhargava and Feng (2002) analyze the tension between benefits (to providers) and disutility (to users) inherent in search engine paid placement. An assumption in their analysis is that the perceived disutility to users, caused by search engine provider bias, can negatively impact search engine market share and user-based revenues. Their analysis leads them to develop a mathematical model for optimal design of a paid placement strategy.

However, not everyone agrees that paid placement inherently creates a disutility to users of search tools. Bill Gross, whose company Idealab founded Overture, recalls telling his Idealab colleagues in 1997 that “the best way to clean up search results was to use money as a filter” (Hansell, 2001). Gross finds that, at the time, search engines return a disproportionate number of results from porn Web sites; therefore, paid placement could operate as a filtering mechanism. Gross’ idea eventually led to the creation of the GoTo.com Web search engine where advertisers would bid for top placement in search results. The bid for placement model that GoTo.com helped popularize is akin to the Yellow Pages model, where the advertisers willing to pay the most get the most prominent placement. Yellow Pages combine textual listings of telephone numbers for companies with larger graphical advertisements for some of those companies. The Yellow Pages model suggests that some types of commercial searches (such as searches for products or services) can benefit from paid placement.

A study by Graham and Metaxas (2003) suggests that college students rely heavily on the Internet for their research needs. Participants in the study completed a survey of six questions

conducted over e-mail during the 2000-2001 school year. The results of the study reveal students' extreme confidence in search engines as a research tool, despite their having little awareness of how search engines select results. Additionally, the study suggests that students remain faithful to one search engine, even if it does not immediately provide the answers sought. Although not mentioned in the study, these findings also suggest a need to understand what impact paid placement has on the use of search engines in academia. The present study of paid placement considers academic and non-academic queries separately.

An ethnographic study commissioned by Consumer WebWatch in 2003 shows that participants had little understanding of how search engines ranked pages, and many were surprised to learn many search engines employed paid placement strategies (Marable, 2003). The study asked seventeen participants from varying demographic backgrounds to perform two online searches on five pre-assigned search sites. A total of fifteen major search and navigation Web sites were tested in the study. This study suggests a consumer awareness problem with regards to paid placement in search engines, and recommends that search engines take steps to fully disclose which results are ranked higher because of paid placement.

This same organization, Consumer Web Watch, subsequently released a report looking at the policies of the search tools regarding sponsored listings (Wouters, 2004). Their report shows that search tools do not do an adequate job of informing customers about sponsorship in their results in their help files or policy documents. Moreover, many tools do not strive to provide a significant separation between sponsored results and editorial results. Meta search tools proved to be a consistent offender by taking sponsored results from different tools and mixing them in with editorial results.

Two screen captures from 2003 demonstrate this concern. Figure 1 is from Overture.com, which allows Web site marketers to bid for placement. The highest bid receives first ranking and after all paid rankings have been displayed, Overture displays editorial results. The search topic was "basketball," and all of these listings are properly identified as sponsored listings. This is appropriate behavior for a search tool, as their advertising is labeled.

1. **New Balance Web Express**
Shop online for New Balance basketball shoes available in large sizes and widths (AA-6E). Discount closeouts, with up to 70% off. Official shoe source for the GusMacker.
www.newbalancewebexpress.com (*sponsored listing*)
2. **SportsPilot ScorePad for Basketball**
ScorePad is a Pocket PC-based scorekeeping, stats and shot chart program supporting basketball, baseball, softball, soccer and hockey. Includes integrated Web site.
www.sportspilot.com (*sponsored listing*)
3. **Hypnosis, Subliminal and Sleep Learning**
Master basketball. Beginner to advanced. Easy, fast, effective. Money back guarantee. Refined as a full time hypnotherapist since 1981. This is my calling, I will make sure you are happy.
www.hypnotictapes.com (*sponsored listing*)
4. **College Basketball on Your Video Screen**
A selection of College Basketball games played from the 1940s till now. Available on VHS or DVD.
www.footballvideos.com (*sponsored listing*)
5. **Licensed Broker Selling NBA Tickets**
Average ticket price is about \$100, but this site offers tremendous availability. TicketsNow.com is a collection of the nations top ticket brokers, all posting their tickets for sale here.
www.TicketsNow.com (*sponsored listing*)

Figure 1. Screen shot from Overture (2003)

Figure 2 shows a screen shot taken the same day from Excite, which is a meta-search tool. Users are led to believe by the top line that Excite has combined results from all of the search tools into one listing.

Excite searches these sites

[About Results](#)

Google • Looksmart • Inktomi • Ask Jeeves • About • Overture • FindWhat • FAST • Teoma NEW!

[Send
these
results
to a
friend](#)

1. [New Balance Web Express](#)

Shop online for New Balance basketball shoes available in large sizes and widths (AA-6E).

Discount closeouts, with up to 70% off. Official shoe source for the GusMacker.

<http://www.newbalancewebexpress.com>

2. [SportsPilot ScorePad for Basketball](#)

ScorePad is a Pocket PC-based scorekeeping, stats and shot chart program supporting basketball, baseball, softball, soccer and hockey. Includes integrated Web site.

<http://www.sportspilot.com>

3. [Hypnosis, Subliminal and Sleep Learning](#)

Master basketball. Beginner to advanced. Easy, fast, effective. Money back guarantee. Refined as a full time hypnotherapist since 1981. This is my calling, I will make sure you are happy.

<http://www.hypnotictapes.com>

4. [College Basketball on Your Video Screen](#)

A selection of College Basketball games played from the 1940s till now. Available on VHS or DVD.

<http://www.footballvideos.com>

5. [Licensed Broker Selling NBA Tickets](#)

Average ticket price is about \$100, but this site offers tremendous availability. TicketsNow.com is a collection of the nations top ticket brokers, all posting their tickets for sale here.

<http://www.TicketsNow.com>

6. [The Official Web Site of the Naismith Memorial Basketball Hall of](#)

Figure 2. Screen shot from Excite (2003)

The same paid results in Overture are listed in Excite as the top-ranked results. The problem is that these results are not labeled as sponsored. In fact, Overture owned Excite and was using it as another way to distribute their results. Users have no way of knowing that these top results are there simply because they paid top dollar for those spots. This problem was the impetus behind the study. As an aside, many search tools have improved their labeling of advertisements since these screen captures were taken.

It is possible that a search tool has both an editorial listing and a sponsored listing for the same commercial Web site. For this study, these would be counted as different types of listings, as the methods by which they were ranked and listed were different. We used contextual clues and information from the pay-for-placement tools, as in Figure 1 and Figure 2, to identify sponsored links that were not labeled. We were cautious in labeling the results this way, so the results related to the amount of unlabeled sponsorship should be taken as the minimum for that measure.

In 2001, the non-profit group Commercial Alert filed a complaint to the FCC requesting that it investigate whether seven major search engine companies were violating existing advertising disclosure laws in their use of paid placement and paid inclusion. In response, the FCC stated that it would not take formal action against the seven search engines, but would send a letter to each of the seven search engine companies outlining the need for clear and conspicuous

disclosures of paid placement. The FCC letter to search engine companies outlined three specific recommendations: 1. paid ranking search results are clearly distinguished from non-paid results, 2. the use of paid inclusion is explained, and 3. no affirmative statement is made that might mislead consumers as to the basis on which a search result is generated (FTC, 2001).

Introna and Nissenbaum (2000) examined the political issues raised by search engines from a qualitative perspective, addressing the specific issue of whether the market mechanism can serve to correct search engine bias. They argued that since most users of the Web lack critical information about search engine alternatives and little understanding of how search engines work, users are not truly exercising free choice. Their recommendations include the need for full and truthful disclosure on the part of search engines regarding the rules governing indexing, searching, and prioritization, and consideration of public support for egalitarian and inclusive non-commercial search tools.

Methodological Review

Prior research about Web searching and use was examined to determine what topics were selected for searching, what types of queries were formed, the structure of those queries, which tools were selected, and how the results were judged. These studies served as inspiration to create a study that is representative of the most common searches of the Web.

Query and Tool Selection

A 2002 study by Whitmire using the Biglan model of disciplinary differences distinguish information-seeking behavior patterns among various academic disciplines. Data for their study were obtained from a 1996 College Student Experiences Questionnaire and grouped students according to the three Biglan dimensions (hard-soft, pure-applied, life-non-life). Although this study aims to understand how different disciplines compare with one another in their use of library resources, there may be value in applying the Biglan framework specifically to academic use of Web search tools. In the present study, these dimensions are used to formulate six different academic query topics, found in Table 1.

Spink, Wolfram, Jansen and Saracevic studied a log of over one million Web queries by users of the Excite search engine to understand the general public's practices and choices in Web searching (2001). A product of this study is a classification of 11 major search subject categories. In a subsequent analysis of the Excite data set, Ross and Wolfram (2000) categorized the 1,054 most frequently co-occurring term pairs into one or more of 30 subject areas. The five most popular subject categories in the Spink et al. study have been used as the basis for search categories representing non-academic users for the current study. These categories are summarized in "Non-academic" column heading in Table 1.

A query log analysis by Silverstein, Marais, Henzinger, and Moricz of approximately 1 billion queries submitted to the Altavista Web search engine found that the public used an average of 2.35 terms per search query (1999). 20.6% of AltaVista queries were submitted with no terms (an empty query), 25.8% had one term only, 26% had two terms, and 15% had three terms. The Spink et al. Excite study cited earlier found that the public used an average of 2.4 terms per distinct query. Excluding the 9.7% of queries submitted with zero terms, 26.6% of Excite queries

had one term only, 31.5% had two terms, and 18.2% had three terms. In the current study, we test the impact of paid placement on one, two, and three term queries separately. Additionally, the single-term query tested for each category is retested with an additional related term to test how the impact of paid placement changes with search terms that are not phrases. The four different categories for query term length are represented as column headings in Table 1.

We differentiate between academic and non-academic searches in the current study. For each class of search, we use a single query to represent the intersection of subject category (e.g. “engineering,” “people, places, things”) and query term lengths (e.g. “1 query term,” “2 query terms”). These queries appear in the boxes in Table 1 below. We used two different methods to select specific queries for academic and non-academic searches.

To derive queries for each of the six academic categories, we selected academic subject terms appearing in course titles in the 2003-2004 Syracuse University undergraduate course catalog. These queries are meant to be representative of possible subject-specific queries submitted by undergraduate students, but not necessarily representative of the most popular academic queries submitted by students.

For non-academic queries, known popular search queries were selected this study. The primary resource was a set of yearly statistical reports published by the Lycos search engine based on their ongoing “The Lycos Top 50” report (<http://50.lycos.com/>). As the name implies, The Lycos Top 50 is a list of the fifty most popular query terms submitted by Lycos users for a given time period. We used the Lycos top 50 early reports for 2000, 2001, and 2002 for our query selection process. To be included in our test list, the popular queries had to appear in the top 100 for at least 2 of the 3 years surveyed, to minimize the inclusion of short-lived topics of interest. Lycos excludes queries identified as “prurient” (e.g., 4-letter words, pornography terms) from its lists. To find specific queries for the “sex, pornography, preferences” category, we constructed queries based on the most popular query terms and query term pairs identified by the Spink et al. Excite study (2001).

Academic	1 term	w/ additional term	2 term phrase	3+ term phrase
Physical sciences	genetics	statistics genetics	organic chemistry	solar system astronomy
Engineering	propulsion	momentum propulsion	soil mechanics	computer aided design
Humanities	Writing	literature writing	architectural history	cultural theories of representation
Business	securities	securities privatization	consumer behavior	supply chain management
Social sciences	psychology	psychology marriage	behavior disorders	criminal justice system
Education	pedagogy	pedagogy religion	teacher development	adult literacy education
Non-academic				
Entertainment, recreation	baseball	baseball women	star wars	buffy the vampire slayer
Sex, pornography, preferences	sex	sex disease	xxx pics	free nude pictures
Commerce, travel, employment, economy	taxes	taxes bankruptcy	mortgage rates	world trade center
Computers, the Internet	napster	napster mp3	final fantasy	grand theft auto
Health, the sciences	diabetes	diabetes glaucoma	atkins diet	science fair projects
People, places, things	paris	paris london	britney spears	martin luther king

Table 1. Matrix of query terms used

These queries are used with three general Web search tools, three meta search tools, and two search tools primarily based on paid listings.

Sullivan summarizes search engine popularity in terms of “Audience Reach” and “Time Spent on Site” based on data from a January 2003 Nielsen NetRatings search-specific report (Sullivan, 2003b). Based on audience reach, the three most popular Web search tools are Google, Yahoo, and MSN, accounting for a combined total of 86% audience reach.

Meta search tools are those that collect results from several different search tools in response to a query. Wouters (2004) found that meta search tools are the worst offenders in displaying sponsored results that are not appropriately labeled or disclosed. Dogpile, Mamma, and ixquick are used for the meta search tests. Dogpile is identified as one of the “Big Four” meta search engines (Sherman, 2002); however, all of the “Big Four” metasearch engines are owned by the same parent company, InfoSpace, and return similar results for the same search. Therefore, two other meta search tools listed at Searchenginewatch.com were selected to supplement Dogpile.

We also tested our matrix of user queries against Overture and FindWhat.com, which provide paid listings for many other Web search tools.

Category	Name	URL
Paid Listings Provider	1. Overture	http://www.overture.com
	2. FindWhat.com	http://www.findwhat.com
General Web search	3. Google	http://www.google.com
	4. Yahoo! Search	http://www.yahoo.com
	5. MSN Search	http://www.msn.com
Meta search	6. Dogpile	http://www.dogpile.com
	7. Mamma	http://www.mamma.com
	8. ixquick	http://www.ixquick.com

Table 2. Web search tools tested

Methods for Analysis of Results from Previous Literature

In a limited study of paid placement in meta-search engines, Danny Sullivan of SearchEngineWatch.com (2003) quantified paid placement in terms of the percentage of links on the first page of results that were “paid links.” Sullivan tested for paid links by examining the results for the single query term “canada” across eight well-known meta search engines. The range of paid links ranged from 0% to 86% depending on the search engine. Sullivan admits that the results he observed could vary, since only a single query term and time of day were tested. The problem with using a number of links is that a link can take many forms – a banner advertisement, a small listing, a colored and prominent listing with the words “sponsored link” at the bottom of the page. For this study, the goal was to measure what the user sees on the screen, which is more than just the link.

In their study of newspaper Web sites, Dewan, Freimer and Zhang quantified the amount of advertising in terms of the percentage of screen real estate devoted to advertising (2002). The percentage of advertising space was calculated as a percentage of an 800x600 pixel screen, as that was the most common resolution when this study was designed. A similar screen real estate approach is used by Nielsen and Tahir, in their study of large-scale Web site homepages, to calculate the percentage of homepage real estate devoted to navigation, advertising, site identity, etc. (2002). In our study we have selected the “screen real estate” approach to calculate the impact of advertising over the percentage of links approach, as it is a more accurate measure of the visual impact of advertising on users.

There are two general types of advertising commonly used in Web search tools. The first type, Non-integrated Advertisements, comprises paid advertisements that are presented differently from editorial listings. These are listings for which no payment for preferential placement has occurred and are clearly separated from the editorial listings on the search results page. Non-integrated Advertisement can be either text-based or graphical in nature. The second, Integrated Paid Listings, comprises search result listings for which payment for preferential placement has occurred.

Integrated Paid Listings can be conspicuous or inconspicuous, depending on whether or not they resemble editorial listings, or depending on the extent to which they clearly labeled as paid listings. In this study, we used the pay-for-placement search tools such as Overture and

Find what to identify inconspicuous paid listings. As in the example in Figures 1 and 2, it is clear when a tool has taken listings from one of these sponsored tools and listed them as editorial listing. It is probable that there are listings we were not able to identify as unlabeled paid listings when more devious methods of accepting payment for listing were employed. Therefore, this measure should be taken as a minimum number of listings, which may be actually be higher.

The FCC has recognized the use of ambiguous terminology in the labeling of paid listings (2002). In the present study we distinguish between ambiguous and unambiguous labeling of paid listings based on guidance provided by the FCC. Labels using the terms “paid” or “sponsored” to describe their paid listings are considered unambiguous. Labels using alternative terms such as “featured listings” or “recommended sites” are considered ambiguous.

Table 3 lists the different categories of information available on a search tool result page.

Category	Component	Definition
1. Search Interface	Interface	Search engine interface elements such as search boxes, submit buttons, and related categories; includes logos and help links
2. Editorial Listings	Editorial listings	Search result listings for which no payment for preferential placement has occurred
3. Non-integrated Advertisements	Banner ads	Graphical ads that are clearly separated from editorial listings
	Text ads	Text-based ads that are clearly separated from editorial listings
	Internal ads	Ads provided by the search tool to promote other parts of its site; may be graphical or text-based
4. Integrated Paid Listings	Conspicuous paid listings	Paid listings that are integrated with editorial listings, but appear unlike editorial listings because of a clear difference in presentation style
	Unambiguously labeled paid listings	Paid listings that are integrated with editorial listings, that appear similar to editorial listings, but are clearly labeled as paid listings
	Ambiguously labeled paid listings	Paid listings that are integrated with editorial listings, that appear similar to editorial listings, but are labeled using ambiguous language
	Inconspicuous paid listings	Paid listings that are integrated with editorial listings that appear similar to editorial listings and are unlabeled
5. Whitespace	Whitespace	Any unoccupied (blank) space on the screen

Table 3. Components of Search Results Pages Measured

The Silverstein et al. AltaVista study cited earlier reported an average of 1.39 screens were accessed per user query, with 85.2% of users accessing only one screen of results. The Spink et al. Excite study cited earlier found that 28.6% of users examined only one page of search results, and 19% looked at two pages only. The Fidel et al. observational study cited earlier found that “students often did not read beyond the first screen [of results].” The Consumer WebWatch ethnographic study cited earlier found that most people did not go beyond the first page of search results (2003). Based on the weight of evidence suggesting that most users only look at the first page of search results, only the first Web page of search results for each query is analyzed. To understand the impact of examining only the first screen as compared to the entire first page, both the amount of advertising both in the first 800X600 pixel screen of results and in the first Web page of results (full-page results) are analyzed.

Study Methodology

Based upon the literature review, each of the 48 queries listed in Table 1 was entered in the eight search tools listed in Table 2, leading to 384 searches during the fall of 2003. To calculate the amount of screen real estate devoted to advertising the first page of search results for each query were printed on 8.5 x 11 inch graph paper. The full page of screen results usually required several printed pages.

The goal of the study was to mirror the most common user behavior and determine how search tools results appeared to users. The decisions made in the study reflect this fact. Most users skim and browse, and therefore we chose the measure of screen real estate instead of just the number of hits. At the time, most users used an 800X600 size browser window with Internet Explorer 6.0. Other decisions were also based upon the typical user experience. This is important for anyone wishing to replicate this method; to be comparable, the study parameters must be adjusted to match the typical user behavior.

Queries were tested using all default settings of the search tool with the default Internet Explorer 6.0 browser sized to 800x600 pixels. To calculate results for the first screen, the point where the 600 pixel screen height ends was marked on the printed page. The last horizontal grid line occupied by page content is marked to indicate the end of the first Web page. The printed page margins were reduced to 1/4 inch to reduce additional whitespace introduced by printing Web pages. Each square on the graph paper was then coded to indicate the majority of content of the search page components discussed in Table 3.

To normalize the data, the number of 1/4 inch squares for each category was divided by the total number of squares that made up the first 800X600 screen of results and the first entire page of results. These percentages then became comparable measures across different interfaces and tools, and screen real-estate is a more representative measure of how typical users quickly browse the Web than the number of results. While the entire sites were categorized, the important types of information for this study are the amount of editorial listings and the amount of advertising and paid placement.

In order to combine the measurement into a single value that is easier to compare and conceptualize, the traditional concept of precision from information retrieval evaluation is used as a basis for a new measure, called Editorial Precision. *Editorial Precision (EPrec)* is the

percentage of screen real estate dedicated to editorial results (ED) divided by the sum of the percentage of real estate dedicated to editorial results and the percentage of the screen dedicated to all sponsored results(ED +AD):

$$EPrec = ED / (ED + AD)$$

The resulting value, which will be from 0 to 1, indicates what portion of the results provided by a search tool are “real” results as compared to results displayed to the user because the author of the page paid to have them displayed. This single measure focuses on the choices presented to the user as possible resolutions for their information need, regardless of the form of those choices. In the case that there are no editorial listings and no advertisements, such as a null searching result, then Editorial Precision is set at 0. This value factors out whitespace and interface, and narrows the measures down to a single value that represents the portion of search results that are “real,” i.e., non-paid placement, and therefore is the primary measure discussed in this research.

Measuring the screen real estate dedicated to real results as the percentage of the screen display allocated to all results eliminates some bias related to differences in presentation styles of the search engines. Calculating EPrec this way also eliminates the effect of the varieties of the search engine’s interface. Determining EPrec as a percentage of the screen dedicated to results measures what percentage of the listings are “real” content and creates a more realistic measure representing user behavior for comparing very different tools than counting the number of results and advertisements.

Editorial Precision (EPrec): The percentage of search results displayed due to a relevancy-based match to the query that are not sponsored, e.g., the percentage of all search results that are “real” listings.

Research Question 1: Overall Analysis and Differences Between Search Tools

The purpose of the first research question is twofold: what are the overall patterns of advertising and what differences can be found between search tools. The tools are compared both on an individual basis and as members of a class of search tool (as detailed in Table 2).

Overall Analysis

The first analysis is an average of all queries and tools in order to provide a general understanding of what appears to a user and to provide a baseline for comparison.

Type of Information	First Screen	Entire First Page
Interface	28.1%	12.8%
Whitespace	29.0%	39.8%
Total ads	6.0%	2.4%
<i>Banner ads</i>	1.3%	0.3%
<i>Text ads</i>	3.7%	1.5%
<i>Internal ads</i>	0.9%	0.7%
Total paid listings	18.2%	11.3%
<i>Conspicuous paid listings</i>	0.6%	0.2%
<i>Unambiguously labeled paid listings</i>	15.7%	10.1%
<i>Ambiguously labeled paid listings</i>	1.9%	1.0%
<i>Inconspicuous paid listings</i>	0.1%	0.03%
Total ads + paid listings	24.2%	13.8%
Editorial listings	18.7%	33.7%
Editorial Precision (editorial results / all results)	40.27%	67.39%

Table 4. Percent of Information Displayed across All Queries and All Tools

On average, 19% of the user’s first screen display contains results from the search tool’s primary index, and 24% of the screen contains advertisements or paid listings, so the user is being shown more sponsored material than non-sponsored content. However, over the entire first page of results, the amount of editorial material goes up to 33% while the amount of sponsored information drops to 14% of the entire first Web page. The search tools therefore front-load their advertising and sponsorship in order to encourage more people to see the listings; as the model of payment for these listings is per-click, this result is not surprising.

The Editorial Precision is 40% for the first screen, which means out of the entire screen of search results presented, 40% of it consists of “real” listings; in comparison, 67% of the results provided on the entire first page are “real” results. This provides a baseline upon which to compare other tools and searches.

Comparing Tool Categories

Table 5 compares the Editorial Precision (EPrec) scores for the three major categories of tools (from Table 2) for both the first screen of results and the first Web page of results over all queries.

Search Tool Category	Number of Searches	First Screen Editorial Precision	First Page Editorial Precision	Difference between First Page and First Screen
General	144	56%	82%	26%
Meta	144	48%	77%	29%
Paid Listing	96	5%	31%	26%

Table 5: Percentage of non-sponsored information in types of search tools

The general search tools provide the highest percentage of editorial listings, and the meta-search tools provide slightly fewer listings. The paid listing tools, as would be expected by the category name, provide much less non-sponsored information. An interesting finding is that all three

categories provide similar increases in editorial listings when moving from the first screen displayed to the entire first page of listings.

As with many aggregate measures, these aggregates over categories hide interesting differences between the individual search tools.

Comparing Individual Tools

Comparing individual tools reveals that the overarching statements about tool categories do not always hold true. Table 6 contains the average Editorial Precision scores for each of the tools, which indicates what percentage of the returned listings are editorial listings (as compared to paid placement and other sponsored pointers).

Search Tool	Search Tool Category	First Screen Editorial Precision	First Page Editorial Precision	Difference between First Page and First Screen
Mamma	Meta	83%	93%	10%
Google	General	82%	88%	6%
MSN	General	47%	73%	26%
Yahoo	General	40%	86%	46%
Dogpile	Meta	35%	54%	19%
ixquick	Meta	25%	84%	59%
Overture	Paid Listing	9%	61%	52%
Findwhat	Paid Listing	0%	0%	0%

Table 6: Percentage of non-sponsored information in different search tools

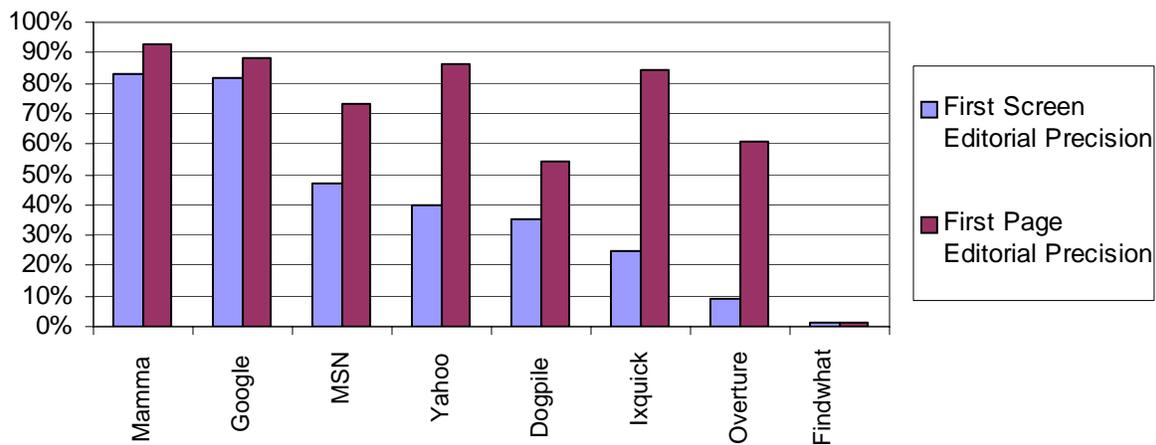


Figure 3: Percent of editorial information in different search tools

Several trends can be spotted looking at Table 6 and Figure 3. All of the search tools except Mamma and Google presented more sponsored content than editorial content on the first screen displayed. This is a clear indication of the financial goals of the search tools. When the entire first Web page of results is examined, however, the relative amount of editorial information displayed usually rises, sometimes dramatically.

Two of the search tools are designed to sell listings to the highest bidder. Overture and Findwhat both take bids for their primary listings. This chart shows a significant difference between them: Overture supplements paid listings with editorial listings when there aren't enough bidders for a search term, while Findwhat supplies only sponsored links as search results. These differences are explored in further detail under Research Question 4.

One overarching lesson to be taken from this is that users need to review the entire first Web page of results returned and not just the first screen displayed. The first screen may show users screens consisting mainly of sponsored results, but by scrolling down, users can access many more non-sponsored links.

Research Question 2: Differences in Topics

The next area of exploration is to understand differences between search topics. These differences are first compared at a broad level and then explored between different topics in those broad categories.

Comparing Academic and Non Academic Topics

Table 7 shows the differences in academic and non-academic searches. As expected, when a search is made in an academic field, significantly more editorial listings are displayed on the first page and screen than the sponsored information.

Academic / Non-Academic	Number of Searches	First Screen Editorial Precision	First Page Editorial Precision	Difference between First Page and First Screen
Academic	192	47%	71%	24%
Non Academic	192	33%	64%	30%

Table 7: Percentage of non-sponsored information in Academic and Non-Academic Topics

Just as observed before, average Editorial Precision significantly increases when one views the full page instead of just the first screen. The average first screen Editorial Precision is 33% for non-academic topics and 47% for academic topics. The average precision increases to 64% for non-academic and 71% for academic when the full screen is considered.

The important finding is that there is a significant difference, using a T-test for percentages, between academic and non-academic search results at both the first screen ($p < .005$) and first page ($p < .05$) level. The difference is in the number of sponsored listings displayed; searches on non-academic topics retrieved more sponsored links. This is good news for those using the Web search tools for academic research, as the results are sensitive to the type of topic searched since there is less advertising for academic searching.

There is little difference in distribution of banner ads, text ads or internal ads across academic and non-academic topics, which implies that these ad types are static between search topics. The

difference between these search types is that searches on non-academic topics have more paid listings that look like editorial listings, which can be seen in the shaded rows in Table 8. The non-academic search results display about 3.4% more results that are labeled as sponsored and 3% more results that are sponsored but ambiguously labeled. Therefore, searchers looking for non-academic topics need to take extra care to examine results carefully for sponsorship.

Paid Placement in First Page of Results	Non- Academic	Academic
Editorial listings	14.22%	23.12%
Total ads	6.25%	5.68%
<i>Banner ads</i>	1.20%	1.48%
<i>Text ads</i>	4.10%	3.35%
<i>Internal ads</i>	0.95%	0.85%
Total paid listings	21.03%	15.48%
<i>Conspicuous paid listings</i>	0.70%	0.50%
<i>Unambiguously labeled paid listings</i>	17.38%	13.97%
<i>Ambiguously labeled paid listings</i>	2.95%	0.80%
<i>Inconspicuous paid listings</i>	0.00%	0.22%
Total ads + paid listings	27.28%	21.17%

Table 8: Percentage of sponsored information across Academic and Non-Academic Topics

Comparing Individual Topics

Figure 4 shows the Editorial Precision between specific topics within both academic and non-academic categories. Academic searches in Engineering, Education, and the Humanities produced the highest percentage of “real” results, while searches in the Sciences and Business produced a higher percentage of sponsored results in both the first screen and first full page of results. In fact, the portion of the screen dedicated to sponsored results in Business is very similar to the portion of the screen dedicated to sponsored results for non-academic searches on Sex and Pornography, both of which are lucrative areas of Internet advertising.

Surprisingly, however, searches on Sex and Pornography produced the least amount of sponsored listings out of all of the non-academic topics. This could be related to the policies against sexually explicit advertising that most search tools employ. More commercial topics, such as People, Places, and Things and Commerce, produced the highest amount of sponsored results.

Figure 4 also demonstrates that while there is a large difference between topics in the amount of “real” results presented on the first screen, the difference between topics when the entire first Web page is considered is more subtle. The overall conclusion, that academic searches produce fewer sponsored results than non-academic searches, is still true.

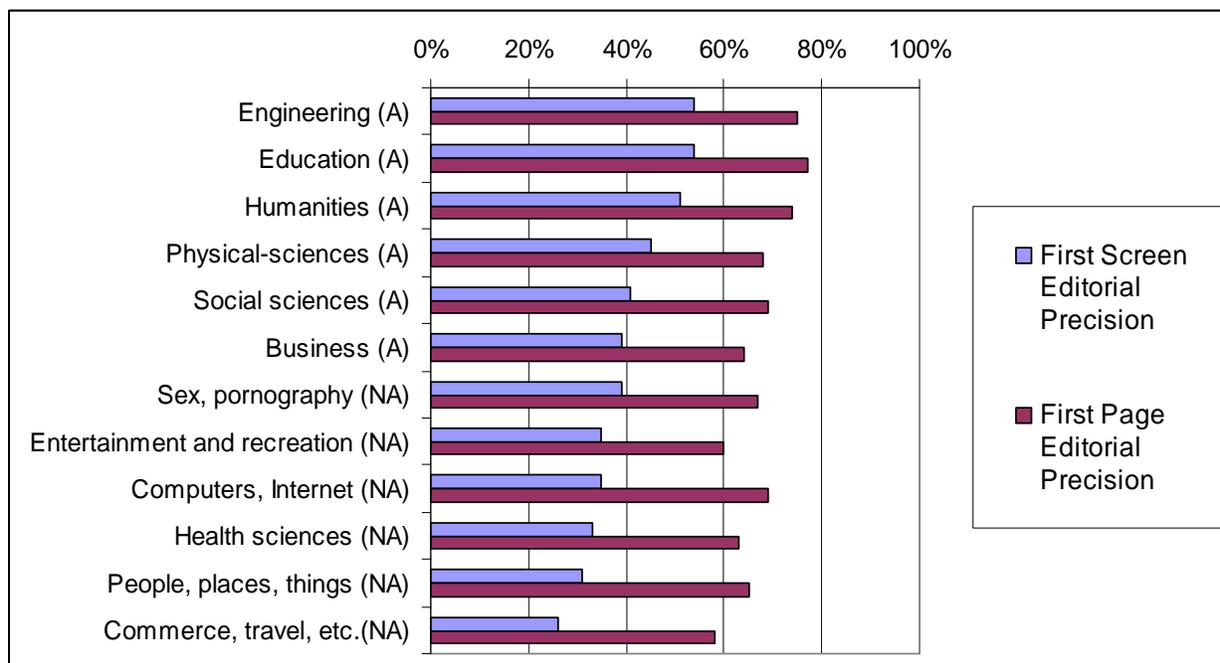


Figure 4: Percentage of non-sponsored information across Academic(A) and Non-Academic (NA) Topics

One overarching statement that can now be made in response to research question #2 is that there is a significant difference between academic and non-academic searches in the percentage of the screen containing sponsored results, and that that difference corresponds to the commercial nature of the search query. However, the amount of banner and other “static” advertising is similar across all search topics. Perhaps search tools would be willing to aid academic searchers and reduce the number of banner and other ads in correlation with the number of sponsored links to aid scholarly use of their tools.

Research Question 3: Differences in Query Structures

The third research question looks at the different query structures. Table 9 shows the difference in Editorial Precision for each of the four query structures used. The interesting finding from this examination of the data is that there is a significant difference (using T-test for percentages, $p < .005$) between searching on one word, a two-word phrase, or a three-word phrase and searching on two words that are not a phrase.

Query Structure	Number of Searches	First Screen Editorial Precision	First Page Editorial Precision	Difference between First Page and First Screen
1 word	96	32.1%	62.8%	30.7%
2 word phrase	96	32.6%	62.8%	30.2%
3 word phrase	96	38.8%	67.1%	28.3%
1 word + 1 word (non-phrase)	96	57.7%	76.8%	19.1%

Table 9: Percentage of sponsored information between different query structures

This finding is not surprising, as many of the sponsorship programs used by the search tools connect sponsored results and advertisements to specific terms or phrases. Sponsors select the words or phrases they want to be associated with and then pay based upon the popularity of those terms. The benefit to this plan is that searchers can take advantage of this fact by adding another synonym for their search terms that is not a phrase and greatly reduce the amount of sponsored information returned from their search. This search strategy is also sound for better retrieval, as adding more search terms give the tools a better chance of finding documents on the desired topic.

The pattern exhibited in Table 8 holds true for these data as well. The percentage of the screen dedicated to static advertisements remained relatively constant across different query structures, while the paid listings, since they are connected to the query, were responsible for the change in Editorial Precision.

Research Question 4: Patterns in Facet Combinations

Finally, we look at combinations of these facets to better understand the differences between individual search tools. Looking at the combination of search tool type and general topic reveals nothing new of interest (Table 10); the trends established earlier hold true.

Search Tool Type	First Screen EPrec		First Page EPrec	
	Academic	Non-Academic	Academic	Non-Academic
General	67.23%	45.52%	86.83%	77.68%
Meta	54.05%	41.05%	80.49%	73.32%
Sponsored	7.34%	2.05%	34.04%	26.55%

Table 10. Average Editorial Precision: Search tool Type by Search Topic

Examining the search tool type by query structure reveals some interesting patterns. The overall structure of these results (from Table 9) is that one term, two-term phrase, and three-term phrase searches are all significantly lower in Editorial Precision than results from two terms not in a phrase. As can be seen in Table 11, this pattern is slightly different for each of the search tool types. The general search tools have similar results for one term and two-term phrases; then the Editorial Precision rises by about 10% for a 3-term phrase and rises much higher for the two-term non-phrase search. Meta search tools have a more even rise with a higher EPrec for each query type. Sponsored search tools have an erratic pattern, which is cause for more detailed exploration.

Search Tool Type	First Screen EPrec				First Page EPrec			
	1 term	2-term phrase	3-term phrase	2-term non-phrase	1 term	2-term phrase	3-term phrase	2-term non-phrase
General	45.67%	45.70%	54.92%	79.20%	78.87%	78.12%	81.97%	90.07%
Meta	36.46%	40.11%	45.69%	67.94%	72.38%	74.02%	74.97%	86.25%
Sponsored	3.16%	1.25%	4.38%	10.00%	22.27%	22.89%	33.20%	42.82%

Table 11: Average Editorial Precision: Search Tool Type by Query Structure

General Search Tools

There are some interesting differences between the General search tools, detailed in Table 12. Google is one of the more consistent tools, delivering a high level of Editorial Precision in both the first screen of results and over the entire first page. General search tools provided the non-academic searcher with slightly more advertisements. MSN provided better-than-average Editorial Precision for academic topics, but many more sponsored results for non-academic topics. It is clear from this that MSN works to customize the number of advertisements to the topic, plastering the first screen of non-academic results with almost 80% advertising. Yahoo provided more stable results than MSN across topics, so their advertising is not as heavily tied into the search term. The large difference between the first screen and first page of results indicates that users of Yahoo must scroll down after doing a search to reach the majority of the non-sponsored results.

Search tool	First Screen EPrec		First Page EPrec	
	Academic	Non-Academic	Academic	Non-Academic
Google	85.61%	78.85%	88.97%	87.59%
MSN	69.79%	23.25%	84.88%	60.78%
Yahoo	46.27%	34.45%	86.65%	84.66%

Table 12: Average Editorial Precision: General Search Tools by Search Topic

Looking at the search tools by query structure, no new information is unearthed; the same pattern of lower results for the one term, two-term phrase and three-term phrase as compared to the higher two-term non-phrase query holds true.

Search Tool	First Screen EPrec				First Page EPrec			
	1 term	2-term phrase	3-term phrase	2-term non-phrase	1 term	2-term phrase	3-term phrase	2-term non-phrase
Google	82.25%	74.80%	80.25%	91.62%	88.71%	85.28%	87.00%	92.15%
MSN	30.45%	35.92%	45.52%	74.20%	66.48%	65.52%	72.65%	86.66%
Yahoo	24.30%	26.38%	38.99%	71.77%	81.42%	83.55%	86.25%	91.39%

Table 13: Average Editorial Precision: General Search Tools by Query Structure

Meta Search Tools

Examining the meta search tools separately demonstrates the importance of not relying upon aggregates to understand a phenomenon (Table 14). Just as above, each of the search tools has a different pattern. Mamma is unusual for the meta search tools in that its Editorial Precision is often higher than that of Google. Dogpile is similar to MSN, in that the amount of sponsored information varies greatly by topic (although Dogpile's Editorial Precision was consistently lower than MSN's). ixquick, on the other hand, does not supply ads based on the topic of the query, as the results for ixquick are similar for academic and non-academic searches.

Search tool	First Screen EPrec		First Page EPrec	
	Academic	Non-Academic	Academic	Non-Academic
Dogpile	51.36%	18.95%	63.82%	44.06%
ixquick	26.74%	23.02%	85.03%	82.91%
Mamma	84.07%	81.18%	92.61%	92.99%

Table 14: Average Editorial Precision: Meta Search Tools by Topic

There are also interesting patterns in the query structures for the meta search tools (Table 15). When entering a single term, Dogpile piled on the advertisements, with only 6.5% of the first screen results and 40% of the first page results being non-sponsored. Entering a two- or three-word phrase helped the results somewhat, but a high Editorial Precision was reached only with a two-term non-phrase query. ixquick, as in the above example, had relatively stable results across query structures. While there was a difference in the Mamma searches, they were relatively small and Mamma continued to deliver the highest amount of non-sponsored results.

Search Tool	First Screen EPrec				First Page EPrec			
	1 term	2-term phrase	3-term phrase	2-term non-phrase	1 term	2-term phrase	3-term phrase	2-term non-phrase
Dogpile	6.55%	25.67%	28.82%	79.56%	39.59%	49.57%	49.79%	76.80%
ixquick	23.42%	15.80%	26.00%	34.30%	84.88%	80.66%	83.21%	87.14%
Mamma	79.41%	78.87%	82.24%	89.97%	92.65%	91.83%	91.91%	94.81%

Table 15: Average Editorial Precision: Meta Search Tools by Query Structure

Examining the Dogpile data more closely, some disturbing trends are revealed. For the one term, two-term phrase, and three-term phrase searching, Dogpile included listings to sponsored sources that were not explicitly labeled; in fact, between three and five percent of the entire first page was made up of these ambiguously labeled paid listings (on top of the sixteen to nineteen percent of the page dedicated to clearly labeled paid listings). Dogpile did not do this for the two-term non-phrase queries, which means the match of a keyword introduces some of these additional paid listings into the results. Users should be aware of search tools that are introducing paid listings that are not clearly labeled as such. There data support the same conclusion – that meta search tools are problematic in regard to sponsored links – as the report produced for Consumer Web Watch (Wouters, 2004).

Sponsored Search Tools

The final two tools examined are sponsored tools, which mean their raison d'être is to provide sponsored links in response to queries. Both of these tools will allow a person to search his or her search tool directly, although the typical way searchers get these results is by seeing them in conjunction with a search in another tool.

Tables 10 and 11 show that these tools have the lowest Editorial Precision of all tool types. Looking at the tools individually reveals more concerns. Tables 16 and 17 show that Findwhat provided no non-sponsored content. Twenty-five percent of the time, there were no results at all from the Findwhat tool, as the search did not match any of the purchased keywords.

Overture, on the other hand, supplements their sponsored links with non-sponsored links. They display as many advertisements as people have purchased for that query, and then provide results from Yahoo to fill the page. Therefore, they have a higher Editorial Precision than Findwhat in all search topics, although their scores are worse than other types of tools.

Search tool	First Screen EPrec		First Page EPrec	
	Academic	Non-Academic	Academic	Non-Academic
Findwhat	0.00%	0.00%	0.00%	0.00%
Overture	14.67%	4.10%	68.08%	53.09%

Table 16: Average Editorial Precision: Sponsored search tools by Topic

Search Tool	First Screen EPrec				First Page EPrec			
	1 term	2-term phrase	3-term phrase	2-term non-phrase	1 term	2-term phrase	3-term phrase	2-term non-phrase
Findwhat	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
Overture	6.32%	2.49%	8.75%	19.99%	44.54%	45.77%	66.39%	85.63%

Table 17: Average Editorial Precision: Sponsored search tools by Query structure

Discussion and Conclusions

This study has shown that, on average, 60% of the search results presented on the first screen are there because the author of the page paid for them to be in the listing. If the user scrolls to view the entire page, then 33% of the results presented on the page are sponsored. Other studies (Wouters, 2004) suggest that the search tools do an inadequate job of informing users about where these sponsored results come from. Most search tools have moved away from using non-labeled advertisements or at least are labeling those results coming from a pay-for-placement result.

The measure used in this result is sensitive to several facts. First, the 800X600 resolution selected at the time of the study is no longer the most common resolution used. It is expected that search tools are aware of this fact and have adjusted their results to match the higher resolutions commonly used. Therefore, others using this method should match the resolution used for data collection to the common resolution used at the time.

Another measure to explore for future exploration would be the number of editorial results and advertisements instead of the space taken up by the advertisements. We focused on screen real-estate, as that is more appropriate if one assumes that users typically skim screens and do not read them in detail.

This study also shows that generalizations about search tools mask important differences among the tools. Google and the meta search tool Mamma.com both provided a high level of “real” results, while sponsored tools such as Findwhat provided a consistent barrage of paid results. Most tools varied by topic, providing more “real” results for academic queries. In addition, most tools varied by query structure, providing significantly more “real” results for searches that were not standard phrases. Conversely, ixquick did not vary for either topic or query structure, providing the same mix of sponsored and real results on each search.

Sponsored content is not necessarily bad for information seekers. The most important future line of research to come from this study is to determine, using real users with real information needs, how relevant the sponsored and editorial links are to their information needs. While this study determines how much of the screen is taken up by sponsored results, it does not consider the relevance of those results. For certain types of non-academic searches, it is possible that sponsored links will be more relevant than editorial links. Even so, the model of “he who pays the most gets the top position” may still not be appropriate for information services to employ.

There are a few lessons that searchers can take away from this research. First, searchers should add an additional synonym to their queries. This will greatly reduce the number of sponsored links and serve to help the search tool provide better results. Second, searchers need to be aware that many of the results displayed on the first screen of results will be sponsored, and they need to scroll down and see the entire page (and even to continue on to other pages of results) before abandoning a search term. At this time, academic searches provide fewer sponsored links than non-academic searches due to the fact that fewer advertisers focus on keywords used in academic searching.

This study adds to a growing body of research that indicates that part of information literacy must be to train users to consider their search results, asking if results exist because of their relevance to the query or simply because the page author is the highest bidder. Teaching users to be more skeptical of search tools results will allow them to become savvy Web searchers.

References

- Bhargava, H.K. and Feng, J. (2002). Paid placement strategies for Internet search engines. *Proceedings of the Eleventh International Conference on World Wide Web*, 117-123.
- Dewan, R., Freimer, M., & Zhang, J. (2002, January). Managing Web sites for profitability: Balancing content and advertising. Paper presented at the Hawaii International Conference on Systems Sciences, Kona, HI. Retrieved September 28, 2003 from <http://rdewan.ssb.rochester.edu/papers/web2.pdf>
- Federal Trade Commission. (2002). Commercial Alert Letter. Retrieved November 23, 2004 from <http://www.ftc.gov/os/closings/staff/commercialalertletter.htm>
- Fisher, I.E., Smith, L.M. (2000). The role for web search engines. *The CPA Journal* 70(1), 42-48.
- Miyagishima, B. K., & Toney, C. D. (1999). A visit to the information mall: Web searching behavior of high school students. *Journal of the American Society for Information Science* 50(1), 24-37.
- Graham, L. & Metaxas, P. (2003). Of Course It's True; I Saw It on the Internet. *Communications of the ACM* 46(5), 70-75.
- Hansell, S. (2001, June 4). Paid Placement Is Catching On in Web Searches. *The New York Times*. p. C1.
- Introna, L. D., & Nissenbaum, H. (2000). Shaping the Web: why the politics of search engines matters. *The Information Society* 16(3), 169-185.
- Marable, L. (2003). False oracles: Consumer reaction to learning the truth about how search engines work. *Consumer WebWatch Research*. Retrieved January 30, 2003 from <http://www.consumerwebwatch.org/news/searchengines/ContextReport.pdf>
- Mowshowitz, A. and Kawaguchi, A. (2002a). Assessing bias in search engines. *Information Processing and Management* 38(1), 141-156.
- Mowshowitz, A. and Kawaguchi, A. (2002b). Bias on the web. *Communications of the ACM*. 45(9), 56-60.
- Nicholson, S. (2000). Raising reliability of Web search tool research through replication and chaos theory. *Journal of the American Society for Information Science*, 51(8), 724-729.
- Nielsen, J., & Tahir, M. (2002). *Homepage Usability: 50 Websites Deconstructed* Indianapolis, IN: New Riders Publishing.
- Phillips, T. (2001). Search and ye shall find. *Director* 55(2), 116-117.

- Ross, N. & Wolfram, D. (2000). End user searching on the Internet: An analysis of term pair topics submitted to the Excite search engine. *Journal of the American Society for Information Science* 51(10), 949-958.
- Sherman, C. (2002). The big four meta search engines. *SearchEngineWatch*. Retrieved September 17, 2002 from <http://searchenginewatch.com/searchday/article.php/2160781>
- Silverstein, C., Henzinger, M., Marais, H. & Moricz, M. (1999). Analysis of a very large web search engine query log. *ACM SIGIR Forum* 33(1), 6-12.
- Spink, A., Wolfram, D., Jansen, M.B.J., & Saracevic, T. (2001). Searching the Web: The public and their queries. *Journal of the American Society for Information Science and Technology* 52(3), 226–234.
- Sullivan, D. (2001). Meta search or meta ads? *SearchEngineWatch*. Retrieved November 22, 2004 from <http://www.searchenginewatch.com/sereport/article.php/2163821>
- Sullivan, D. (2003a). 2002 Search Engine Watch awards. *SearchEngineWatch*. Retrieved September 28, 2003 from <http://searchenginewatch.com/awards/article.php/2155921>
- Sullivan, D. (2003b). Nielsen Netratings search engine ratings. *SearchEngineWatch*. Retrieved February 25, 2003 from <http://www.searchenginewatch.com/reports/article.php/2156451>
- Whitmire, E. (2002). Disciplinary differences and undergraduates' information-seeking behavior. *Journal of the American Society for Information Science and Technology* 53(8), 631-638.
- Wouters, J. (2004, Nov. 8). Searching for Disclosure: How Search Engines Alert Consumers to the Presence of Advertising in Search Results. *Consumer WebWatch Research*. Available at <http://www.consumerwebwatch.org/news/paidsearch/finalreport.pdf> (retrieved Nov. 23, 2004).