

Googling Bing: Reassessing the impact of brand on the perceived quality of two contemporary search engines

Ahmed A. Atallah and Edward Lank

David R. Cheriton School of Computer Science
University of Waterloo, Ontario, Canada
{aataulla, lank}@cs.uwaterloo.ca

ABSTRACT

Given the high value of the online search market, whether brand perception or quality of search results matters more for users is a highly salient question. This paper presents findings of the largest controlled, systematic preference elicitation study of search quality versus brand perception. We examine a total of 548 instances of sponsored and organic search results from the Google and Bing search engines as rated by 25 participants. We find that, if users are not aware of the source of a set of search results, they will consistently rate Google results as better. However, the presence of the Google brand strongly influences perceived quality, essentially over-riding differences in search result quality. Together, these results demonstrate that, while Google may outperform Bing in blind searching, trust in the Google brand is a much more significant factor in users' search preferences.

Author Keywords

Search, Branding, Google, Microsoft, Bing

ACM Classification Keywords

H5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous.

INTRODUCTION

Despite a short history, internet search has a legacy of repeatedly producing well known brands. AOL, Yahoo and AltaVista each dominated the online search market at their respective peaks. Today, the online search is dominated by Google: By some estimates (Statscounter: gs.statscounter.com), Google controls in excess of 90% of the worldwide online search market. As a result of this dominance, the Google brand name has also become synonymous with online search. Similar to products such as Kleenex or Xerox, the word "google-ing" is part of several languages and is colloquially inseparable from the notion of online search.

While Google (the company) currently dominates the online search market, there has been very little open (publicly available) and conclusive research conducted into how the trust embedded in Google – and the recognition of its brand name – impact the perceived quality of online search results provided by Google. Are Google search results objectively better than those of its competitors? Does the Google brand influence the perception of quality of search results? Which is presently a stronger determinant of user preference, Google's quality or its brand? Given the recent integration of Microsoft Bing and Yahoo search engines [1] there are few major competitors left in the online search market. However the question of whether the competition between Google and alternative providers is going to be one of quality of search results (and the underlying technology) or simply a competition of branding remains largely unanswered.

Researchers have argued that brand perception can significantly impact design and evaluation of technology [10, 11]. However, Marcus [10] claims that branding, through the lens of human factors research, has typically not been investigated. He notes that branding causes prospective users to form clear expectations of the user experience even before interacting with a product or device. Other researchers have noted that expectations of technology are shaped by passive or active encounters with related systems, by messages from other people and from companies, and by other symbolic embodiments connected to any strongly branded system [11].

While researchers have argued that examining brand effects is important, answering questions about the role of brand versus quality is difficult in an area like online search. From a user perspective, the notion of one search engine being "better" than another is largely abstract. Terms like *precision* and *recall*, from information retrieval, are based on the notion of a ground truth for search, some absolute authoritative ranked list of results that should be returned by a perfect search engine. In online search, absolute ground-truth does not exist. This lack of authoritative ground-truth confounds questions of quality within online search, making it difficult to disentangle quality from brand.

In this paper, we disentangle the role of branding and quality by examining preferences for search results expressed by users over a large and diverse set of search

queries. Our work aims to address the question of whether the results from the market dominant engine of today (Google) are better in quality than those from a current competitor (Microsoft Bing), or whether it is the brand image and pre-established preferences of online consumers that influences their choice of search engine/result preference. Through a set of experiments we show that users still ascribe a higher quality to Google search results when presented with a blind (unbranded) set of competing results from the Google and Bing search engines. However, when we introduce branding, we show that brand preference overrides perceived quality differences within search.

There are two implications of the results described in this paper. More narrowly, they demonstrate that Google's brand recognition provides it with a strong competitive advantage within online search. This provides some credence to current concerns about antitrust regulation¹. More broadly, this study is one of only a very small set of academic studies that explores the role of brand in user perception and behaviour within computing environments [10].

The remainder of this paper is organized as follows. First, we summarize past work in analyzing quality and brand effects in search results. Next, we describe our experimental design to disentangle search result quality (or perceived search result quality) from brand perception. We then summarize the results of our experimental studies, and conclude by examining more deeply the implications of this work.

RELATED WORK

The impact of brand on the perceived quality of on-line search results has been examined in the past. However the results from various marketing and academic studies have been inconclusive. Synopses of surveys by independent analysts and marketing firms such as Advertising-Age and iProspect are particularly difficult to assess because of the somewhat limited scope of the data and the lack of details on analysis techniques. For example, in a study of 16 participants, the Catalyst Group [2] concluded that, even though users did not see a difference in search result quality and preferred the Microsoft Bing search engine interface, they still would not switch to Bing for day to day use. Advertising-Age [3] made a somewhat orthogonal claim, suggesting that users preferred the actual quality of Bing and Yahoo search results as long as they were branded as Google. Unfortunately, the lack of transparency in marketing studies makes them easy to dismiss as unreliable.

There have been two recent controlled academic studies to identify the impact of branding on perceived search result

quality. Jansen et al. [4] claimed that branding plays a profound role in helping users assess the quality of search and sponsored search results. However the scope of their work was very limited as they only evaluated the impact of 4 pre-processed query results ("laser removal", "techno music camping", "Mexico" and "manufactured home") retrieved using Google but laid out in 4 different page interfaces representing 4 brands of search engines. We note that since users did not specify their own queries, they may have had less interest in search result quality and perhaps focused on interface elements of search engines. Thus we believe that results were unable to distinguish interface elements from branding.

Bailey et al. [5] improved upon the work of Jansen et al. by (i) conducting an interface neutral study [7], (ii) having a much larger sample size and (iii) allowing users to input their own queries. They presented search results, dynamically retrieved from Yahoo and Google, side by side in two columns (an approach that we too adopt in our experiments). They mislabeled brand of search results 50% of the time to determine if brand loyalty would convince users to prefer the results labeled with a better known brand or not. The methodology used in this study was very different approach from that in the work of Jansen et al. as it essentially stripped out all interface components in the web-forms presented to the participants, forcing them to analyze the actual results, instead of interface elements on the page. The conclusions of Bailey et al. contradicted the results of Jansen et al., revealing no statistical correlation between user preference and brand.

While Bailey et al. [5] showed no significant brand preference, there are aspects of their study design that make this conclusion suspect. Foremost, the design of the preference elicitation study conducted by Bailey et al. allowed users to pick a "No difference" option when brand names were displayed alongside search results. Presenting a "No difference" option allows users to default to a superficial response without looking deeply at results. It also limits the statistical power of the study. We expand on these issues in more depth here.

It is well accepted in the context of marketing and preference elicitation studies that the forcing users to choose between the presented option leads to finer discrimination among alternatives. The technique of forced choice was first discussed by Ferris [8] and Irwin [9]. A comprehensive examination of forced choice in preference elicitation is presented by Buchanan and Morrison [6]. Essentially, when a user is forced to make a choice between two results, s/he scrutinizes the options in greater detail. If there is no statistical difference in perceived quality, there should be a 50-50 split in user preference. A statistically significant deviation from this split implies a non-random skew (preferential treatment shown towards one option). Consequently if we want to see whether branding plays a role in consumer preference for search results, especially given that the quality of search engines for typical day-to-

¹ See, for example,

<http://www.techcrunch.com/2009/03/01/what-an-antitrust-case-against-google-might-look-like/>

day queries is often postulated to be roughly equivalent, we should force users to pick one search engine's results as better. Any difference (if any exhibited by either quality or branding) will automatically become apparent in the relevant conditions by the presence of a skew away from a pure 50-50 split.

The second significant shortcoming of prior work is that it fails to take into account pre-existing brand preferences for typical users [6]. Bailey et al. published their work in 2007, based on research done prior to publication. In 2005-2006, Google and Yahoo both controlled about 1/3 of the online search market. Bailey et al. did not either bias their brand study, or control for brand biases in analysis, so it is difficult to determine whether brand biases are absent, or whether brand biases are balancing each other.

Because of both contradictory results and methodological limitations, what effect, if any, brand has on search quality perceptions is unknown. We now describe a set of studies that aims to answer the question of brand effects on perceived search result quality.

EXPERIMENTAL DESIGN

Search engines present two types of results to users: *organic search results*, which represent content retrieved from the internet based on an indexing scheme used by the search engine; and *sponsored search results*, which are placed advertisements, usually from companies who purchase search placement from the search engine provider. In our study we tested the impact of brand on both these types of results.

Apparatus and Interface

The experiments were web-based and conducted on participants' personal systems and preferred browsers. The experimental interface we used is shown in Figure 1. Figure 1a shows a snapshot of results displayed in the *blind condition*, where all branding has been removed and the results have been reformatted to be visually indistinguishable from each other. Figure 1b shows results displayed in the *branded condition* where search results are labeled with a search engine brand using corporate logos. In the brand condition, results were incorrectly branded 50% of the time to measure presence or absence of a brand bias. Participants in our study were unaware that brands were being mislabeled and were specifically told that the intent of the study was to compare only the quality of retrieved results. We also counterbalanced placement to avoid a left-right column bias in result presentation. Participants selected their preferred search results by clicking the appropriate "(these) results are better" button.

Participants

25 participants (recruited by word-of-mouth) participated in this study. Participants filled out a brief survey on search engine preference and number of web searches performed during a one week period. All participants reported that Google was their preferred search engine, and they all performed at least 100 search queries in a one week period.

For a variety of reasons, the sample consisting of only Google users is not a drawback in our study; instead, it is a critical requirement to achieve meaningful results. First, as we note in the introduction, the goal of this study is to measure the effect of the Google brand on search result quality. Therefore we need to test Google users' perception of result quality with regard to their result preferences and branding. The effect of any brand will vary; for example, if we were to conduct a similar study on operating systems using the Microsoft brand, we would expect a different branding effect, one specific to the Microsoft brand on operating systems. Second, when designing brand effect studies, it is necessary to use a sample that prefers the stronger brand being tested. Incorporating a weaker brand in testing alters the magnitude of brand effect (impact of the branding, its induced loyalty and its perception), and is considered poor experimental practice [6]. Choosing a weaker brand also tells us little about the overall strength of the most market dominant brand. Finally, we are interested in the effect of the Google brand because Google controls the vast majority of online search. Biasing toward Google is an accurate portrayal of the real world, and creates a more ecologically valid study of the status quo in online search.

Methodology

Two identical experiments were conducted in this study. In our initial experiment we focus on organic search results, i.e. web content retrieved from the internet with both blind (Figure 1a) and brand (Figure 1b) conditions. We then contrast this with the results of sponsored search results in blind and brand conditions.

Experiment 1, our organic search experiment, was a two-block experiment. Our first block tested the blind condition, and was designed as a 2X2 counterbalanced experiment with source (Google or Bing) and location (left or right) as factors. Our second block was a 2X2X2 counterbalanced design with source (Google or Bing), location (left or right), and brand label (Google or Bing) as factors. Experiment 2 was identical, but used sponsored search results instead of organic search results. Participants completed the individual blocks and experiments in no particular order.

Tasks

In our organic search experiment, users input a series of self-specified queries based on information that we asked them to retrieve. We asked participants free form question that represent typical scenarios where users would use a search engine, and asked them to find answers using the blind/branded search tool we provided. These questions were divided into three different categories (a) informational question where a direct answer is required, for example "What is the capital of New Zealand?" (b) Subjective questions, "Try to find a good carrot cake recipe for me" and (c) Informal Queries, where users were asked to retrieve general information for a broad variety of topics, such as "Michael Jackson's fan site" or "Transformers 2 DVD". Users were free to construct any type of query using search operators such as +, - and quotations.

[Human Capital Institute | The Global Association for Strategic ...](#)

Day Three of **HCI's** Summit began with Bill Taylor explaining why the only way to stand out from the crowd is to stand for something special. ...

<http://www.hci.org/>

[HCI Bibliography : Human-Computer Interaction Resources](#)

The **HCI** Bibliography (**HCIBIB**) is a free-access bibliography on Human-Computer Interaction, with over 50,000 records in a searchable database. **HCIBIB.ORG** is also one of the premier ...

<http://www.hcibib.org/>

[HCI - Wikipedia, the free encyclopedia](#)

HCI (with the last character being uppercase " i " not lowercase " L ") may refer to:

<http://en.wikipedia.org/wiki/HCI>

[Human-computer interaction - Wikipedia, the free ...](#)

Human-computer interaction (**HCI**) is the study of interaction between people (users) and computers. It is often regarded as the intersection of computer science, behavioral ...

http://en.wikipedia.org/wiki/Human-Computer_Interaction

[Human-computer interaction - Wikipedia, the free encyclopedia](#)

Human-computer interaction (**HCI**) is the study of interaction between people (users) and computers. It is often regarded as the intersection of computer ...

http://en.wikipedia.org/wiki/Human%25E2%2580%2593computer_interaction

[Human Concern International](#)

HCI is in a great position to serve humanity for years to come. For the 2009 Annual Dinner, close to 300 guests joined **HCI's** staff, representatives and volunteers to raise ...

<http://www.humanconcern.org/>

[HCI Books](#)

Online Bookstore Self-Help,Chicken Soup for the Soul,Recovery,Addiction, Parenting,Love,Relationships,Spirituality, & Pets.

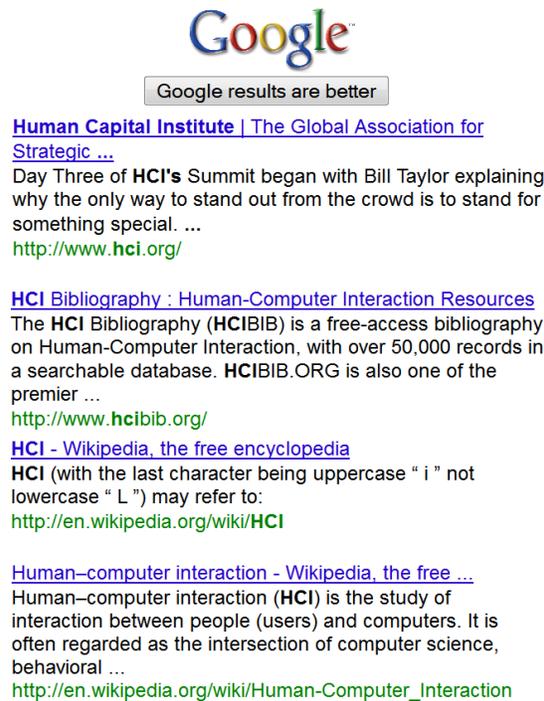
<http://www.hcibooks.com/>

[HCI Bibliography : Human-Computer Interaction Resources](#)

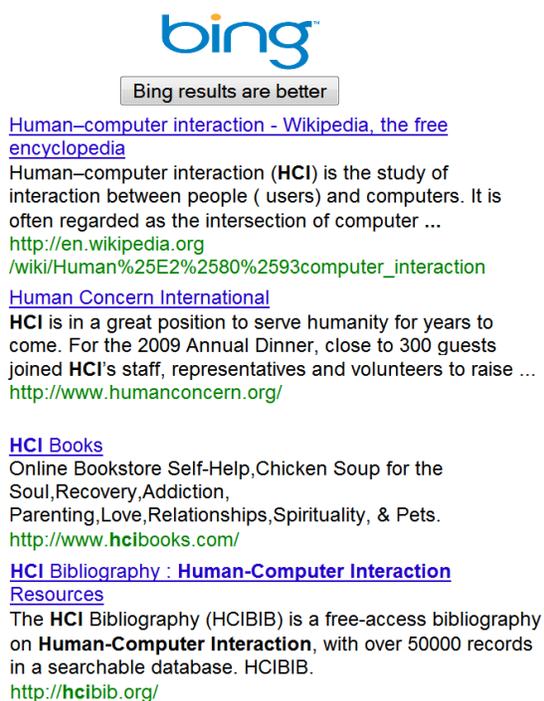
The **HCI** Bibliography (**HCIBIB**) is a free-access bibliography on **Human-Computer Interaction**, with over 50000 records in a searchable database. **HCIBIB**.

<http://hcibib.org/>

Figure 1A: Organic search results of a sample query ("HCI") as they would appear in our experimental interface without any branding.



The image shows a search engine interface for Google. At the top is the Google logo. Below it is a search bar with the text "Google results are better". The search results for the query "HCI" are displayed in two columns. The first result is from the Human Capital Institute, the second is from the HCI Bibliography, and the third is from Wikipedia. Each result includes a title, a short description, and a URL.



The image shows a search engine interface for Bing. At the top is the Bing logo. Below it is a search bar with the text "Bing results are better". The search results for the query "HCI" are displayed in two columns. The first result is from Wikipedia, the second is from Human Concern International, the third is from HCI Books, and the fourth is from the HCI Bibliography. Each result includes a title, a short description, and a URL.

Figure 1B: Our branding test interface, results were mislabeled 50% of the time. In both branding and non-branding conditions placement of logos and results (left or right column) was random.

Users were also free to click on results and browse through them before making a decision about which set of results were better.

For, experiment 2, the sponsored search experiment, queries were informational. Users were asked about recent purchases that they had made (online or offline) and were

asked to search for items relevant to that purchase. For example, two participants had rented cars recently and were asked to type queries that would get them the best car rental deal had they shopped online.

In both experiments, the queries were input into a generic text field with no branding present. The results were

[Budget.com: Official Site](#)
Rent Online at Budget!
Check Our Latest Deals Now.
[www.Budget.com](#)

[Avis.com: Official Site](#)
Save 10% on your US **Car Rental**
with Avis - Book Online Today
[www.Avis.com](#)

[Thrifty Car Rental®](#)
Choose Thrifty for your next low
rate **rental car**. Thrifty saves!™
[www.Thrifty.com](#)

[National® Holiday Rentals](#)
Book now and bypass the
crowds this holiday season! Go
Like A Pro.
[www.NationalCar.ca](#)

[Car Rentals from \\$10/day](#)
Car Rentals from Multiple
Suppliers Bookingwiz
EasySearch®. Save Big!
[Bookingwiz.com/Car-
Rental-Deals](#)

[Alamo® Car Rental](#)
Find the perfect **car** for your
trip. Book online at affordable
rates!
[www.Alamo.ca](#)



Google results are better

[Budget.com: Official Site](#)
Rent Online at Budget!
Check Our Latest Deals Now.
[www.Budget.com](#)

[Avis.com: Official Site](#)
Save 10% on your US **Car Rental**
with Avis - Book Online Today
[www.Avis.com](#)

[Thrifty Car Rental®](#)
Choose Thrifty for your next low
rate **rental car**. Thrifty saves!™
[www.Thrifty.com](#)



Bing results are better

[National® Holiday Rentals](#)
Book now and bypass the
crowds this holiday season! Go
Like A Pro.
[www.NationalCar.ca](#)

[Car Rentals from \\$10/day](#)
Car Rentals from Multiple
Suppliers Bookingwiz
EasySearch®. Save Big!
[Bookingwiz.com/Car-
Rental-Deals](#)

[Alamo® Car Rental](#)
Find the perfect **car** for your
trip. Book online at affordable
rates!
[www.Alamo.ca](#)

Figure 2: Sponsored search results of a sample query (“car rentals”) shown in both unbranded (left) and branded (right) conditions

dynamically extracted from queries run in the background against Google and Bing, parsed, and displayed in a generic interface (with or without brand logos present depending on the condition). If a particular search engine produced less than 10 organic results or less than 3 sponsored results, the user was asked to try a different query without being shown any results. Only the first page of results (containing either 10 organic or 3 sponsored search results) was shown to users. No elaboration was provided on what "better" meant in the context of organic or sponsored search. Although users were not aware of results being misbranded during the study, after the experiments we discussed participants' impressions of Bing and Google, how they rated results, what made them pick one over the other and whether they believed that they were biased towards their preferred brand or not. Insights from these post-task interviews are presented in our discussion of the results.

We did not allow repeated querying for the same terms. We were also able to verify after the study that the results presented by Bing and Google, were, in fact, always distinct and never matched 100% in content and in ordering, implying that there was always some difference that the participant could observe between the two sets of search results. A total of 548 individual searches were collected and analyzed, 170 for blind organic search, 170 for brand organic search, 103 for blind sponsored search, and 105 for brand sponsored search.

Hypotheses

There are four hypotheses which we aim to test experimentally:

1. In organic search, there is no difference in quality between Google and competing (Bing) search results.
2. Brand has a significant effect on user’s perceived quality of organic search results.
3. In sponsored search, Google returns ads which are of higher perceived quality (due to its larger advertiser base).

4. Brand has a significant effect on user’s perceived quality of sponsored search results.

RESULTS

One potential confound in our experimental design was left-right column biases in users selecting results (see Figure 1 and 2). Fortunately, in our experimental results, no left-right column biases were revealed. Left/right values are ignored in the remainder of this paper due to their lack of effect on results.

Organic Search

Query Type	Bing is better	Google is better
Informational	19	33
Subjective	30	41
Informal	17	30
Total	66	104

Table 1: Organic Search with No Branding – Frequency of preferences by query type

In the blind condition (Google and Bing results indistinguishable, Figure 1a) users favored Google by a significant margin, as shown in Table 1. Of the 170 trials conducted, 104 preferences for Google and 66 for Bing results were recorded. These figures are statistically significant (two-tailed binomial sign test, $p < .01$), allowing us to reject hypothesis 1 that proposes that there is no difference ($\mu=.5$) between Google and Bing results.

The results of the branded condition (Figure 1b) are shown in Table 2. On the left side of the table, search result source is indicated. Of 170 search queries, we see an almost perfect split in terms of search preference (86 for Bing versus 84 for Google). Therefore, unlike in the blind condition, we do not see a significant perceived quality difference between Google and Bing. On the other hand a clear indication of bias (excessive preference) of Google

Result Source	Brand Shown		
	Bing	Google	Total
Bing	28	58	86
Google	39	45	84
Total	67	103	

Table 2: Organic Search with Branding – Frequency of preferences by brand and result source.

Ad Source	Brand Shown		
	Bing	Google	Total
Bing	14	21	35
Google	39	31	70
Total	53	52	

Table 3: Sponsored Search with Branding – Frequency of preferences by brand and ad source.

branded results was present (Table 2, bottom row). Users preferred Google branded results 103 times and Bing branded results only 67 times. These differences are statistically significant (two tailed binomial sign test for $\mu=.5$, $p < .005$). This confirms hypotheses 2.

Sponsored Search

During early pilot studies of our experiment, it seemed apparent that Google offered much more targeted advertisements compared to Bing. Google has a larger advertiser base and it therefore possess a larger number of ads in its repository than its competitors. Because of this dominance in the ad market, it is reasonable to assume that the quality (accuracy and relevance for the user) of the ads that Google will display will be higher. However, our goal in analyzing sponsored search results is not to prove that Google has more highly targeted ads. Instead, we wished to see whether brand preference persists despite Google’s significant targeted ads advantage.

To test this, we first validated Google’s quality advantage using our blind condition for sponsored search. Google ads were significantly better than Bing’s. For 103 blind trials (without branding) 75 resulted in Google ad preferences and 28 in Bing ad preferences (two tailed binomial sign test, $p < .0001$). This result validates hypothesis 3.

In contrast with the organic search results, when branding was introduced in the sponsored search case there was virtually no change in preferences. Users still preferred sponsored search results that originated from Google, even if they were branded as Bing (Table 3). Brand preferences, shown in the bottom row, were virtually identical (53 versus 52, not significant for $\mu=.5$). However, there was a statistically significant preference for ads sourced from Google regardless of the brand shown above them, 70 to 35 (two-tailed binomial sign test, $p < .0001$).

To test hypothesis 4, however, we need to determine whether the introduction of brand had an effect on the number of times a brand was chosen. In the unbranded condition, the likelihood of picking Google was 75 of 103 or 72.8%. The stated preference for the Google brand in the

branded condition (70 of 105 trials) was statistically unchanged (binomial test, $\mu=.728$, $p = 0.188$). This allows us to reject hypothesis 4, and to claim that brand does *not* impact user preference in sponsored search results.

DISCUSSION

Organic Search

The blind organic search results showed that Google results were, in general, rated significantly “better” than Bing results. In a forced decision study, if no quality difference exists, one would expect to see a lack of significance in the binomial sign test, i.e. that Google and Bing would be selected approximately an equal number of times. Instead, we see a significant preference for Google search results, allowing us to reject our first hypothesis. Our results show that, even in the lack of any branding and indication of the source of the results, users still preferred Google results over those produced by Bing. This trend was prevalent over all types of queries outlined in our Tasks section, i.e. informational, subjective and informal queries.

A much more interesting aspect of user behavior is seen when branding is introduced in the context of organic search. In our results, we also see a significant bias in user preference towards results labeled as “Google” regardless of their underlying data source. Of the four combinations of brand and source, users selected Bing search results that were branded as Google most often and Bing search results branded Bing least often.

The most interesting aspect of our organic search results is the strength of the Google brand. Consider the following. Search results are mislabeled 50% of the time in our second experiment, branded organic search. If *quality* dominates, then one would expect to see no brand preference, i.e. participants would always pick Google results, and 50% of the time these results would be branded Bing. If *brand* dominates, then one would, likewise, expect to see no quality difference, i.e. Bing and Google sourced queries would be selected an even number of times. If factors are confounded, one would expect to see effects in both brand and source.

In our results, for branded organic search, the Google brand preference was so strong amongst our participants that it eliminated any quality difference between search results. This is a highly significant result. First, it shows that when branding is introduced, the decision criteria used by people to judge which results are better significantly changes: they go with the brand they know and prefer. Furthermore this change in preference is solely triggered by the introduction of branding. Even though users could initially distinguish on the basis quality, they are unable or unmotivated to do so in the presence of branding. We are able to show that, even though a quality difference is present between Google and Bing search results, brand power is more pertinent than quality in organic search results. Google’s brand is sufficiently strong that it eliminates any perceived quality difference in organic search queries.

Sponsored Search

While brand dominates in organic search, the sponsored search tests present a somewhat different story. First, as expected, the data verifies our hypothesis that Google ads have a significant quality advantage because of the large advertiser base that they possess. Our results show that, under blind conditions, where no indication of brand is present and the only discriminating factor is quality, participants consistently pick Google ads as better. Although this result is similar to the organic search result, when branding is introduced the results are dramatically different.

In contrast to the organic search results, the introduction of branding in sponsored search listings was unable to convince users to go with their preferred brand. This is an interesting observation because, even though our sample only consisted of Google users, they did not consistently rate their preferred brand for higher quality sponsored ad results. Instead their preference for ads actually taken from Google remained unchanged even though these ads were branded randomly. In other words, the participants did not hesitate to pick Bing branded ads as better. Thus we conclude that for the sponsored search results the actual quality is a much stronger factor at when Bing and Google are compared than is their brands.

Interpreting Branding Vs. Quality Effects

In our tests we attempted to first identify differences in quality between organic and sponsored search results taken from Google and Bing to establish a baseline criterion. We then measure if branding changes user preferences. One of the challenges that arise when interpreting results like these involves trying to rationalize observations. For example, are Google search results truly “better” than Bing search results (as shown in experiment 1) or are Google results just a better match to Google users’ expectations? Or in the context of ads, are user expectations of online ads different when they are examining ads from their preferred brand of search engine? These questions are difficult to address, and the answers depend on a definition of what constitutes “better” in web search. Given a set of search results that should be returned, i.e. given ground truth, “better” could be classified in terms of precision and recall. However, in web search, there is no ground truth. Instead, “better” as it pertains to search engine quality is a subjective response based on users’ perceptions of results that are returned.

To try to address these questions, we conducted a series of post-task interviews with participants to determine why they made the choices they did. However, the data from these interviews lead us to believe that there can be no universal definition of “better” when it comes to online search results. This is especially true in the case of subjective and informal queries where there is no right or wrong answer for the search results. Users in such complex situations pick one set of results that they find most appealing, but what they find appealing may change from one search query to another. For example a user looking for

a good cake recipe may have little knowledge of cakes and may choose to judge search engine quality based on the length of the first result page’s recipe (a short recipe is obviously better). Similarly some may like a particular type of cake so they might prefer the set of results that ranks those particular recipes higher. Participants in our study cited the reputation of the domain names from where the results originate, how “legitimate” results pages look, the accuracy of the search result snippet included with the page link, and the ranking of results as the most critical factors for their choice. Most could not express exactly their decision making criteria and noted that it varied depending on the query. Even when the answers required were direct facts, for example, trying to find out the location (city) of the world’s tallest building, participants often used varying criteria to rate one search result set as better. For example one participant mentioned that when s/he is looking for factual information, s/he has utmost faith in Wikipedia and would certainly prefer the search engine that gives preference (via ranking) to Wikipedia results.

One result of these complex, personal algorithms for judging search result quality is that it renders it difficult to determine whether participants chose Google because Google results were more accurate, or if they chose Google results because Google results matched their expectations in the blind organic search block. Because of this user-centric and fuzzy definition of “better”, it is also not clear how we can quantify “how much better” Google’s results are compared to those of Bing.

What is clear from our results is that, once branding is introduced, participants’ perceived overall quality of results is significantly different for organic search than it is for sponsored search. In the case of organic search, participants exhibit very strong brand loyalty, but in the case of ads, participants are not impacted by the presence of branding.

There are two possible factors at play here. The first and most intuitive is that sufficient quality difference can trump any brand loyalty. Google ads are sufficiently better to eliminate any brand bias that participants might have. This is analogous to how Google initially came to dominate online search, i.e. its results were sufficiently better than its predecessors (Yahoo, Altavista) that users switched because of higher quality. The other possible factor may be a difference in perception between organic and sponsored search. For organic search, participants were asking for tailored information from a query they specified, and a trusted brand may weigh more heavily in their preferences. For sponsored search, participants may care less about the results (viewing them as “just ads”) and may, therefore default to quality judgments.

The results when examined from this perspective paint a positive picture for the competing search engine. In our study, although Bing’s organic search results are shown to be of lesser “perceivable quality,” they are not “sufficiently worse” and can pass as Google results. If Google users are

fooled by rebranding of other search results as Google search results then we can claim that the brand overpowers minor differences in actual quality of results. Bing is a new brand, with limited following. If Bing can enhance its brand perception, then branding effects may begin to work in its favour.

LIMITATIONS

One potential limitation of our study is in search engine interface design. In our evaluation of search engine quality preferences, we present users with blind and branded search results in a generic interface that we design, one that ensures anonymity of the underlying source of data. While this allows us to make claims about the content returned by search engines, and the way that brand influences perception of that content, search engine interfaces are undergoing continual refinement. For example, Bing, alongside the presentation of search results, provides users with pop-out synopses of pages retrieved. Similarly both Microsoft and Google offer a wide number of other services such as free email, maps, and news that can add to the user experience and impact reputation of the brand. In our study, we focus solely on search results and perceived quality of results.

CONCLUSION

In this paper we presented the results of a quantitative study that captures the impact of brand on sponsored and organic search results. Organic search results show that branding in the online search market (specifically the Google brand) has a strong impact on user preferences in cases where the retrieved results are sufficiently close in quality. However, using sponsored search results, we show that brand preference is not sufficient to overcome the difference in quality for of advertisements returned.

There are real-world implications to the results presented in this academic study. More narrowly, these results have conflicting implications for how regulators view Google. For organic search, we see that the Google brand is sufficiently strong that it eliminates perceived differences in search engine quality. This creates the potential for abuse of a market leading position. As one example, Google could use its brand dominance to over-charge for advertising revenue in search, and then use this revenue to provide services such as word processors or email at a loss, forcing competitors out of the market. However, working against this concern is the lack of effect that brand has on sponsored search results in our study. This argues that, with sufficient technical innovation, Google's dominant search position may still be challenged.

More broadly, this paper represents a study of how branding can affect users' impressions of a computer

system or application. While it seems obvious that brand perception can impact design and evaluation of technology, relatively few studies have analyzed the effect of brand on perceptions of quality, and specifically how brand effects change the perception of quality. In this paper, we provide one such study.

While Google's search results are ranked higher by participants in unbranded conditions, when we introduce brand (and misbranding), we show that brand effects are so strong that they eliminate any perceived quality difference in organic search. As a concrete example of the potential impact brand effects can have, this work broadens our understanding of the potential factors that can influence users' subjective ratings of technology we design.

ACKNOWLEDGMENTS

We thank the participants in our study and our funding agencies.

REFERENCES

1. Microsoft Press Release: "Microsoft, Yahoo! Change Search Landscape", July 27, 2009. Available at <http://www.microsoft.com/Presspass/press/2009/jul09/07-29release.mspx>
2. The Catalyst Group: <http://www.catalystnyc.com>
3. Advertising Age: <http://www.adage.com>
4. B. Jansen, M. Zhang, C. Schltz, "Brand and Its Effect on User Perception of Search Engine Performance", In the Journal of the American Society for Information Science and Technology, Volume 60 : 8 (August 2009), pp. 1572-1595
5. P. Bailey, P. Thomas and D. Hawking, "Does brand name influence perceived search result quality?", Proceedings of the 12th Australasian Document Computing Symposium, online: <http://goanna.cs.rmit.edu.au/~aht/adcs2007/proceedings.html>
6. B. Buchanan and D. Morrison, "Measuring simple preferences: An approach to blind forced choice product testing", Marketing Science, Volume 4 : 2 (Spring 1985), pp. 93-109
7. P. Thomas, and D. Hawking, "Evaluation by comparing result sets in context", Proceedings of the 15th ACM international Conference on information and Knowledge Management (CIKM 2006)
8. G. E. Ferris, "The k-Visit Method of Consumer Testing", Biometrics, Vol 14, No.1, pp 39-49, Mar. 1958
9. F. W. Irwin, "An Analysis of the Concepts of Discrimination and Preference", The American Journal of Psychology, Vol. 71, No. 1, pp 152-163, Mar. 1958
10. Marcus, A. 2004. Branding 101. Interactions 11, 5 (Sep. 2004), 14-21.
11. Cockton, G. From quality in use to value in the world. In CHI '04 Extended Abstracts on Human Factors in Computing Systems (Vienna, Austria, April 24 - 29, 2004). CHI '04. ACM, New York, NY, 1287-1290.