



# **Advertising Disclosures: Measuring Labeling Alternatives in Internet Search Engines**

**Benjamin Edelman  
Duncan S. Gilchrist**

**Working Paper**

**11-048**

**January 13, 2012**

Copyright © 2010, 2011, 2012 by Benjamin Edelman and Duncan S. Gilchrist

Working papers are in draft form. This working paper is distributed for purposes of comment and discussion only. It may not be reproduced without permission of the copyright holder. Copies of working papers are available from the author.

# **Advertising Disclosures: Measuring Labeling Alternatives in Internet Search Engines**

Benjamin Edelman<sup>a</sup> and Duncan S. Gilchrist<sup>b</sup>

<sup>a</sup> -- corresponding author - Assistant Professor, Harvard Business School - Soldiers Field Rd, Boston MA 02163 - [bedelman@hbs.edu](mailto:bedelman@hbs.edu)

<sup>b</sup> -- Harvard Business School and Department of Economics, Harvard University - Soldiers Field Rd, Boston MA 02163 - [dgilchrist@hbs.edu](mailto:dgilchrist@hbs.edu)

# Advertising Disclosures: Measuring Labeling Alternatives in Internet Search Engines

Benjamin Edelman and Duncan S. Gilchrist  
January 13, 2012

In an online experiment, we measure users' interactions with search engines, both in standard configurations and in modified versions with clearer labels identifying search engine advertisements. In particular, for a random subset of users, we change "Sponsored links" or "Ads" labels to instead read "Paid Advertisements." Relative to users receiving the "Sponsored link" or "Ad" labels, users receiving the "Paid Advertisement" label click 25% and 27% fewer advertisements, respectively. Users seeing "Paid Advertisement" labels also correctly report that they click fewer advertisements, controlling for the number of advertisements they actually click. Results are most pronounced for commercial searches, and for vulnerable users with low education and little online experience.

Keywords: Internet advertising, search engines, disclosure, regulation, consumer protection, survey methods

JEL codes: L51, L86, M37, M38, C83

## 1 Introduction

Search engines combine two kinds of links: So-called "algorithmic" links (also known as "organic" links) present the results a search engine deems most relevant for a given search phrase, selected based on page contents, keywords, links, and other factors. So-called "sponsored" links give the results a search engine is paid to show, selected based on an advertiser's payment along with an assessment of the match between the advertisement and the user's search. (Edelman, 2007)

Search advertisements have become a huge business---the mainstay of Google's \$25+ billion of annual revenue. Many search advertisements are helpful---promoting products directly responsive to users' requests. Other search advertisements are less helpful; some promote counterfeits (Thompson, 2010), others trick users into installing spyware/adware (Edelman, 2006a), and still others resort to alternative methods of deception (Edelman, 2006b).

Pursuant to longstanding legal principles and recent FTC instructions, all leading search engines now include a label near advertisements. However, no search engine uses the "Paid Advertisement" label that legal precedent has required in other media. Seeing a divergence between industry practice and applicable legal precedent, we seek to measure the effects of labeling advertisements in the way that prior cases specify.

Our motivations are several. For one, with more than two billion people now using the Internet, search engine advertisements are a strikingly important and sizeable market. In this context, even a small change in user understanding could result in large changes in aggregate browsing behavior. That said, other compulsory disclosures provide cause for concern; consider the ever-present “objects in mirror are closer than they appear”---often a subject of mockery. Would improved labels on search engine advertisements make any real difference? Or are users already well-informed when it comes to search engine advertisements? Our online experiment lets us address these questions, while simultaneously providing evidence for which users respond most to changes in advertisement labels.

We proceed in six parts. In section 2, we review related research on search engine, advertisements, and user understanding. In section 3, we present applicable regulation. In section 4, we examine current labels at leading search engines. In section 5, we explain our methodology, including our online experiment. In section 6, we present results, and in section 7 we consider policy implications.

## **2 Prior Research**

An early literature suggests that users struggle to distinguish search engines’ algorithmic results from advertisements. A 2003 Consumer Web Watch study interviewed users about their searching patterns, finding that users tended to overlook disclosures associated with advertisements. (Marable, 2003) After learning that “sponsored” links are actually advertisements, one participant reported that “it wasn’t clear to me what Sponsored links were or what [the term meant]” (p.24), and the majority of participants reported that “sponsored” labels were too vague, easily misinterpreted, and confusing (pp.26-27). When CWW asked users to suggest alternatives, users suggested “advertisement” and “paid advertisement” (p.26), the same terms we test in the experiment in this paper.

A 2005 Pew Internet & American Life report found ongoing confusion among search users: Just 38% of searchers were aware of a distinction between paid and unpaid results. (Fallows, 2005) Even among users who are aware of the distinction between paid and unpaid results, only 47% said they could always tell which results are paid (p.ii). Of various options to improve advertisement disclosures, the label “paid” was most popular among all participating users (p.19).

Our research question is closest to Jansen et al. (2007). Randomly swapping algorithmic and advertisement links, Jansen finds that searchers prefer to click links that searchers believe are algorithmic results rather than advertisements. However, our implementation employs live page rewriting: where Jansen presented participants with static screenshots depicting modified results pages, we present working interactive search results---an important advance, we argue in section 5. Our page rewriting technique is most similar to similar Pan et al. (2007), though Pan addresses a different question (the relationship between result ordering and user attention).

Several articles consider the possibility that users dislike clicking on advertisements and will seek to avoid advertisements if they can. For example, Rayo and Segal (2010) posit an advertising platform which is able to bait-and-switch---fooling users into clicking advertisements in which they had no genuine interest. Of course such a strategy would be difficult in equilibrium; users might realize that they are being tricked. But our context presents a mechanism by which search engines could extract more clicks than users intended to perform---using labels that are less than fully informative in order to increase clicks from whatever subset of users do not understand the labels.

A parallel line of research posits that advertisements are less effective in circumstances that prompt consumers' suspicion. Kirmani and Zhu (2007) suggest that advertisers avoid such placements when deciding where to advertise, while Friestad and Wright (1994) suggest that advertisers refine their tactics in light of consumer skepticism. Putting aside the obvious legal and ethical concerns, advertisers would sidestep the concerns flagged in these articles if they could present their offers to consumers without consumers recognizing the offers as advertisements.

Separately, a sizable literature evaluates interactions between advertisements and certain kinds of users. For example, Biener and Alberts (2004) evaluate various demographic groups' relative receptiveness to tobacco marketing, finding patterns by age and race. Similarly, Jacoby et al. (1982) measure user understanding of pharmaceutical disclosures, finding differences by age, education, and race. Goldfarb and Tucker (2011) consider patterns in users' response to certain online advertisements, finding relationships between users' propensity to click on advertisements and users' stated privacy concerns.

Our findings also extend the literature on shrouded attributes. For example, Brown et al. (2010) compare the influence of item price and shipping cost in users' purchases at eBay, and Gabaix and Laibson (2006) consider users' choice of printers when ink costs vary. On one view, it is no surprise that consumers fail to consider costs of ancillary services in those contexts: in the period Brown et al. consider, users needed to visit an extra page to see shipping costs at eBay; users evaluating long-term printing costs would need to estimate cartridge lifespan, cartridge price, and their future printing needs. In contrast, search engines' disclosures appear adjacent to advertisements and require considering no additional information---a context where users might be expected to process the information particularly easily. In finding evidence to the contrary, we extend the literature on shrouded attributes, showing how information can be difficult for consumers to process even when visible and in plain view.

### **3 Applicable Regulation**

For more than two decades, the FTC has regulated advertisements that consumers might mistake for editorial or other non-advertising content. In literally dozens of cases, the FTC has pursued deceptive infomercials that purported to be independent programming rather than paid advertisements. (Ruskin, 2001, note 7) The FTC has specifically enjoined any "advertisement that misrepresents, directly or by implication, that it is not a paid advertisement." (FTC v. Levey et al.)

The FTC's call for appropriately-labeled advertisements applies equally online. In July 2001, Commercial Alert sent a letter to search engines complaining about the lack of "clear and conspicuous disclosure that the advertisements are, in fact, advertisements." (Ruskin, 2001) The FTC's June 2002 response noted that "many search engine companies do attempt some disclosure of paid placement" but concluded that "their current disclosures may not be sufficiently clear." (Hippesley, 2002)

The FTC's 2002 response called for search engines to present "clear and conspicuous disclosures" labeling advertisements. To assess whether a disclosure is clear and conspicuous, FTC guidelines call for considering whether the language of the disclosure is understandable to the intended audience, the placement of the disclosure, the disclosure's proximity to the information it qualifies, the prominence of the disclosure, and whether other portions of the advertisement distract attention from the disclosure. (FTC, 2000)

To date, in the context of search engine advertising, the FTC has not offered a view, one way or the other, as to whether any particular label satisfies the clear and conspicuous standard.

## 4 Practice at Leading Search Engines

Through October 2010, leading search engines presented most of their advertisements with labels reading “Sponsored links” (Google), “Sponsored Results” (Yahoo), and “Sponsored sites” (Bing).<sup>1</sup> At first glance, these labels might seem to be straightforward. That said, there is reason to doubt whether these labels are effective at conveying key information to consumers.

For one, when we talk to users about search engine advertising, we are struck by their misconceptions. We have interviewed various users, all with significant online experience and most with advanced degrees, who mistakenly believe all advertisements are on the right side of the page; it seems that they are unfamiliar with the top-of-page advertisements shown at all search engines. Indeed, Google’s “Search Basics” help page makes the same error: A diagram purports to show “*all* the elements ... on a Google search results page” (emphasis added), but the diagram incorrectly depicts advertisements appearing only at right, and never above search results.

In light of the FTC’s requirement that advertisement labels be “clear and conspicuous,” the next two sections assess search engines’ labels under those standards.

### 4.1 *Substance and Clarity of the Disclosure*

All three leading search engines have used the word “sponsored” to describe the paid advertisements they display. The word “sponsored” creates considerable ambiguity. For one, “sponsored” uses the passive voice---leaving unstated the subject of the sentence, i.e. *who* exactly did the sponsoring. If a user searches for “Priceline” at Google and sees a link to Expedia, who is the “sponsor” of that link? Sophisticated users are likely to understand that “sponsored” means “advertisement.” But the disclosure does not specifically inform anyone who does not already know. Indeed, when CWW’s professional ethnographers interviewed ordinary Internet users, many reported exactly these ambiguities in interpreting the word “sponsored.” (Marable at p.27.) One participant stated that “sponsored sounds like charity” and therefore did not seem to indicate that a result was a commercial advertisement.

Meanwhile, the English language offers a shorter label that avoids these ambiguities: The word “advertisement” refers to a commercial offer a media property is paid to present. That single word requires 13 characters, compared to 15-17 for the “Sponsored links” (and similar) labels search engines historically used, so “advertisement” cannot be rejected as overly verbose. At least in interviews, users seem to consider “advertisement” appropriate: When CWW asked its subjects to suggest alternative labels, “advertisement” and “paid advertisement” were their most frequent suggestions. (Marable at p.26.)

Our experiments evaluate user response to, and perceptions of, advertisements labeled as “Paid Advertisement.” In section 6.1 we find that the “Paid Advertisement” label yields a 25% to 27% reduction in clicks on advertisements compared to “Sponsored link” and “Ad,” and in section 6.3 we find that the “Paid Advertisement” label causes users to more accurately report how many advertisements they clicked.

---

<sup>1</sup> We quote these advertisement labels with the capitalization in which they appeared on the corresponding sites.

#### 4.2 *Format and Conspicuousness of the Disclosure*

A careful review of search engines' ad labels suggests that their format may limit their effectiveness. Figures 1-3 present cropped search engine labels as presented to users during October 2010. At each search engine, the "Sponsored" label is the smallest type on the page, and the capital "S" in "Sponsored" is typically similar in size to a lower-case letter in an ad title. Yahoo's site uses the smallest label, with the "S" just seven pixels tall. In fact, at Yahoo, the ® symbol (denoting a registered trademark) measures 8 pixels in height---taller than the capital "S" in "Sponsored." The following table compares the size of ad labels vis-à-vis other text on the page, at each of the three search engines.

	Google	Yahoo	Bing
1: Height of "S" in "Sponsored"	8 pixels	7 pixels	9 pixels
2: Height of a capital letter in advertisement title	12 pixels	12 pixels	12 pixels
3: Height of a capital letter in search engine's logo <sup>2</sup>	45 pixels	31 pixels	28 pixels
Ratio of 1 to 2	67%	58%	75%
Ratio of 1 to 3	17%	23%	32%

In addition, search engines present top-of-page ad disclosures against pastel background colors (pink, yellow, and blue at Google, Yahoo, and Bing, respectively). On one hand, the background color serves to distinguish top-of-page advertisements from the algorithmic results below. However, the colored background also reduces contrast with text---making disclosures harder to read and less likely to be noticed. As detailed below, our experiments were limited to alternative versions of disclosure text, but a similar methodology could test alternative background colors and other format adjustments.

#### 4.3 *Google's November 2010 Labeling Change*

As we finished an initial version of this article, Google made a notable change in its advertisement labels: On November 4, 2010, Google began to use the label "Ads" in lieu of "Sponsored links." To our knowledge, save for this recent change, Google was otherwise consistent in using the "Sponsored links" label since 2002. (4crests, 2002) In mid-2011, Bing also switched to the "Ads" label, though Yahoo continues to use "Sponsored Results" as of January 2012.

On one hand, "Ads" seems to embody a step towards greater clarity---a notch closer to the standard English word "advertisement," eliminating the passive voice ambiguity of "sponsored." A more formal label such as "advertisement" might better credit the importance of a legally-mandated disclosure, but "Ads" at least has immediate unambiguous meaning to a native speaker of American English. Thus, under the FTC's call for "clear" disclosures ("understandable to the intended audience" (FTC, 2000)), "Ads" seems to be a step forward.

While the new "Ads" label offers greater clarity, it seems to entail a reduction in conspicuousness. For one, the "Ads" label is strikingly small, measuring just 20x8=160 pixels, compared to 92x7=644 pixels for "Sponsored Results" at Yahoo. (See additional comparisons in the table below.) Indeed, the "Ads" label substantially fits inside an "o" in Google's logo, and "Ads" is both the shortest and the narrowest of all text on Google's search result page. (Edelman, 2010) In contrast, in identifying characteristics of "conspicuous" disclosure, the FTC calls for disclosure that is "prominen[t]" (FTC, 2000) ---a standard which small type is less likely to satisfy. Meanwhile, the placement of Google's "Ads" label---in a top corner, removed from the advertisements themselves or from anything else users have cause to look at---

<sup>2</sup> Bing presents its logo in lower-case: "bing". We therefore report the height of the full-height letter "b".

increases the likelihood that users will overlook the label. This corner placement is in tension with the FTC’s requirement that a disclosure be “proxim[ate] to the claim it is qualifying.” (FTC, 2000)

	“Sponsored links” (or similar)			“Ads”
	Google	Yahoo	Bing	Google
Height of capital letter in advertisement label	8 pixels	7 pixels	9 pixels	8 pixels
Width of advertisement label	77 pixels	92 pixels	90 pixels	20 pixels
Advertisement label area	616 pixels	644 pixels	810 pixels	160 pixels

Reviewing prior caselaw and research, we found no authority suggesting use of the label “Ads.” FTC litigation has consistently sought the label “Paid Advertisement,” never “ads.” (See e.g. *FTC v. Michael Levey*.) When CWW evaluated possible labels, they recommended the exact label “Paid Advertisement.” (See Marable at p.43.)

Our experiments evaluate Google’s prior and longstanding “Sponsored links” label as well as the new “Ads” label. As discussed in section 6.4, we find that “Ads” is at most a small step forward, but the effect, if any, is not statistically significant from zero.

## 5 Methodology

We seek to measure the effects of “Sponsored link,” “Ad,” and similar advertisement labels at leading search engines. Search engines could readily test such changes by presenting modified labels to a random subset of users and measuring results. But if search engines have run these tests, they have not chosen to publish the results. To prepare our own estimates of the effects of these labels, we built a software system to modify advertisement labels and measure subject response. In particular, in an online experiment, we invited participants to answer questions we posed, and we asked them to do their online research through a web browser window we provided. Unbeknownst to participants, the window was configured to give some subjects a modified display wherein each reference to “Sponsored links” (or similar) was replaced by “Paid Advertisement” or “Ads.” Thus, even though participants were browsing from their own computers and seeing results that seemed to come from ordinary search engines, their screens presented advertisement labels adjusted to suit the requirements of our experiment.

We chose to test the label “Paid Advertisement” because that is the exact label the FTC sought and obtained in prior litigation (e.g. *FTC v. Michael Levey et al.*) and because that is the specific label CWW recommended (Marable at p.43). The “Paid Advertisement” label fits easily within the screen space search engines allocate to “Sponsored link” (and similar).

We asked participants to answer two questions calling for research using search engines. Our first question called for commercial research, asking subjects to find websites that sold mattresses; our second question inquired about health information, calling for research on cancer treatments. Figure 4 presents the exact questions we asked. When participants ran searches to collect the required information, we monitored their searches and clicks, including counting how many advertisements each participant saw, and how many algorithmic and advertisement results each participant clicked.

We obtained participants via a commercial online survey service which handled all details of recruitment and compensation. Users in most demographic groups were compensated with entries into a lottery, and



the service offered increased compensation (certain payment) for users in underrepresented demographic groups.

The commercial online survey service directed users to our server where we provided instructions, a browser window, and follow-up questions. We presented search results to participants via standard HTML and JavaScript (without reconfiguring users' browsers or installing special software on users' computers). We ran a server-side proxy to receive search results, make modifications as required, and collect information about users' clicks. Our experiment uses "between-subject" variation: Advertisement labels were randomly assigned to each participant, and each subject's labels remained consistent for the entirety of the study.

Figure 5 shows search results a participant could have seen if searching for mattresses through our site, in a treatment condition wherein advertisement labels were rewritten to read "Paid Advertisements."

After each participant finished conducting the research our task requested, we asked follow-up questions regarding perception of the process. In particular, we asked how many "advertisements, sponsored links, or other listings the search engine was paid to show" the subject saw and how many such listings the subject believed he clicked.

We never explicitly asked participants their understanding of advertisement labels: If we asked "What does 'Sponsored links' mean?" or even "Do you see any advertisements on this screen?", the act of asking such questions might prompt subjects to consider questions they had not previously thought about. Instead, we present realistic, interactive search results to measure user behavior in a more natural context. Similarly, our follow-up questions obtain information about participants' recollection of their online activities without revealing why we are asking--preventing our questions from biasing user responses.

We believe our general methodology could prove useful to others seeking to study alternative web site layouts, formats, or other functionality. Historically, site operators have been able to test alternatives through randomized experiments they present to some of their users. As Varian (2010) points out, these experiments reveal all manner of useful site adjustments. But the resulting data ordinarily remains internal to the companies that run the experiments; interested researchers have no easy way to analyze results, short of requesting data and cooperation from site operators. In contrast, our use of a rewriting proxy allows us to examine possible alternative site designs without needing cooperation from the sites (in our case, search engines) we seek to study. Pan et al. (2007) ran a similar proxy system in a lab setting, but our implementation shows that this approach can be used without individualized attention from an experimental administrator and on a substantially larger scale.

## **6 Results**

723 participants completed our experiment and follow-up questions: 257 in the "Sponsored links" control group, 228 in the "Ads" control, and 238 in the "Paid Advertisements" treatment. We had little control over the environment from which respondents completed their assigned tasks; respondents completed their assigned tasks from their homes or offices, outside of our direct supervision. We must therefore exclude responses that are facially incomplete or unreliable. For example, if, contrary to our instructions, a respondent used a second browser window to find answers to our questions, then we could neither

modify the user's advertisement labels nor track the user's clicks. The following four paragraphs present the responses we exclude from our analysis and the reasons for these exclusions.

We first exclude responses from 122 users (58 from the "Sponsored links" group, 28 from the "Ads" group, and 36 from the "Paid Advertisements" group) who ran no searches or clicked on no links because we have no results for such users. We have no reason to believe the exclusion of these users biases our results because it seems that these users simply used a different browser window. It is possible that some users in the "Paid Advertisements" group sensed something was different with our window after seeing the unexpected "Paid Advertisement" label. However, we believe this is unlikely to be the case because more users in the "Sponsored links" control condition ran no searches or clicked no links.

We also exclude two users whose browsers submitted dozens of requests within seconds of each search. We believe these users were running "browser accelerator" software that automatically requests and caches every link on a page. We cannot use responses from these users because do not know which links the users actually clicked.

Next, we exclude 122 users (30 from the "Sponsored links" group, 52 from the "Ads" group, and 40 from the "Paid Advertisements" group) whose searches did not return a single advertisement from a search engine. Even if such users embody some distinct group that runs searches where ads are rare, we see no reason why their exclusion would differentially affect the composition of the treatment groups. Furthermore, these users can contribute nothing to our examination of how advertisement labels affect click behavior since they did not see any advertisement labels. We similarly exclude one user whose response indicated that he used ad-blocking software and therefore saw no advertisements.

Finally, we exclude users who were above our sample's 95<sup>th</sup> percentile in the number of links clicked, which is 16. Given the exclusions detailed above, this means excluding an additional 3 users in the "Sponsored links" group, 11 from the "Ads" group, and 13 from the "Paid Advertisements". Heavy clickers are troublesome for our analysis because the distribution of clicks is positively-skewed and it is well known that OLS analyses are sensitive to outliers, especially in small samples. (See Wooldridge (2009) at p. 351.) Our results should be interpreted as excluding the heaviest of clickers.

Net of these exclusions, we retain a total of 449 responses---166 in the "Sponsored links" group, 137 in the "Ads" group, and 146 in the "Paid Advertisements" group. The excluded users are not statistically significantly different from those that remain on average age, gender, or income (table omitted).

Throughout, we distinguish between "actual" user data (as recorded by our server) and "reported" data drawn from users' responses to our questions.

All our regressions include constant terms. However, we omit constant coefficients from our reported tables because the constants do not aid interpretation.

Table 1 presents summary statistics of the number of searches each user ran and the number of links clicked while users were completing the online research portion of the survey, grouping users based on their treated/untreated status. The total number of links clicked is further divided into advertisement clicks and organic link clicks. Recall that the treatment condition entails replacing "Sponsored link" (or similar) or "Ad" with "Paid Advertisement."

Table 1 suggests that behavior of the treatment and control groups was similar, although the treatment group on average ran somewhat fewer searches and clicked on fewer advertisements and organic links. However, the decrease in organic clicks is small and statistically insignificant. Most striking is a difference between the commercial and health searches: users clicked approximately four times as many links during the commercial research question as during the cancer research question, and advertising clicks made up a larger portion of total clicks during commercial research.

### *6.1 Change in Advertisement Clicks*

Table 2 explores the effect of treatment on actual advertisement clicks compared to both control groups, offering least squares estimates of the effect of treatment in seven regression specifications. Panels A and B respectively present coefficients comparing the effect of the “Paid Advertisement” treatment with “Sponsored link” and “Ad” controls. The first specification simply regresses the number of times a user clicked on advertisements (in the specified portion of the study) on the treatment variable. In specifications (2) through (5) we incrementally add self-reported user age, and then user-level dummies for self-reported education, income, and weekly hours browsing the Internet. In specification (6) we include an online shopping dummy indexed as the sum of binary indicators for whether the user reported spending money online in the last year on books, clothing, electronics, travel, or groceries, and in specification (7) we control the user’s gender.

When users performed our commercial research task, the “Paid Advertisement” treatment yielded a reduction in the number of advertisements clicked. In particular, the effect of the “Paid Advertisement” label was statistically significantly different, at a level of 0.10 or less, from the “Sponsored links” control group. See Table 2 Panel A, top row of coefficients. Table 2 Panel B shows similar results for the effect of the “Paid Advertisement” label versus “Ads.” The negative effect of treatment on advertisement clicks in the purchasing research becomes somewhat stronger when controlling for additional factors. Note the increasing magnitude of coefficient estimates and steady standard errors in the progression from (1) through (7). Meanwhile, when users ran health searches, estimated coefficients were also negative, but the effect was never statistically significantly different from zero. We attribute the differing effect of the treatment at least in part to differing base rates: With many fewer advertisements clicked during health research, we are ill-equipped to precisely estimate the effect of advertisement labels in that context. That said, Rodgers and Thorson (2000) posit differing psychological processes for evaluating advertisements in commercial and noncommercial contexts, and our results are consistent with that model.

Overall, depending on control variables, we estimate that the “Paid Advertisement” label yielded a 25% and 27% reduction in clicks on advertisements respectively compared to “Sponsored link” and “Ad.”

Table 3 presents an analogous examination of the effect of treatment on organic clicks. As expected, the estimates are not statistically significantly different from zero. We conclude that while the “Paid Advertisement” label does reduce clicks on advertisements, it has a negligible effect on organic clicks.

### *6.2 Demographic Differences*

We follow Goldfarb and Tucker (2011) in considering the possibility that varying advertisement configurations may differentially affect certain kinds of users. Tables 4 and 5 consider patterns in advertisement clicks in light of available demographic information.

We begin by reviewing base rates of advertisement clicks, as shown in Table 4. On the whole, we would expect vulnerable users---those with lower income, less education, and less online experience---to be more likely to click on advertisements. The “vulnerable” hypothesis makes no clear prediction as to age; older people could be more vulnerable (suckers, bad eyesight, etc.), but young people could also be distinctively susceptible (overly trusting, unfamiliar with the reasons for separation between editorial content and advertising, etc.). In a simple specification that splits age at the sample median of 47, we find that older users click significantly more ads. This result is robust with p-values below 0.05 when we split age instead at 30, 40, 50, and 60.

While the “Paid Advertisement” label has a negative effect on clicks compared with both the “Sponsored links” and “Ads” labels, as shown in Table 2, the effect is more pronounced in certain demographic groups. See Table 5. Here too, vulnerable users show the largest differences: When shown the “Paid Advertisement” label, those without college degrees reduced their clicks, as did those with light Internet usage and those without online shopping experience (those who reported not buying books, clothing, electronics, travel, or groceries online in the last year). To ease interpretation of these findings, we present in Table 6 the average number of advertisements clicked by treated and untreated users in each of these demographic groups. For example, while users without college degrees tend to click more advertisements (Table 5, column (3)), the “Paid Advertisement” label causes less educated users to click advertisements as infrequently as those with more education (Table 6, section 3). Similarly, although users with no online shopping experience click somewhat more advertisements on the whole (Table 5, column (5)), the “Paid Advertisement” label causes such users to click even fewer advertisements than their more experienced counterparts (Table 6, section 5).

### *6.3 User Perceptions of Advertisements*

We now turn to follow-up questions that asked users about their perceptions of the advertisements they saw. Responses to these questions are limited for several reasons. Due to an error in our follow-up survey, we discarded the first 124 responses we received. (Note that users only received follow-up questions after completing the research portion of the project, so dropping these responses did not affect user behavior in the research task.) We also discarded 92 responses that were internally inconsistent, for example reporting “no” to seeing advertisements in a first question, but then answering a subsequent question as to the number of advertisements seen. These exclusions have the unfortunate effect of removing many respondents from our analysis of follow-up questions, but inconsistent responses simply add noise and would not advance understanding of user perceptions. We regard the exclusion of these respondents as part of the challenge of recruiting reliable respondents in an online project in which respondents are rewarded for completion rather than quality. We are left with 233 responses to follow-up questions: 87 in the “Sponsored links” control group, 65 in the “Ads” treatment, and 81 in the “Paid Advertisements” treatment.

Table 7 presents summary statistics of users’ reports as to the number of searches they performed, the numbers of advertisements they viewed, and the numbers of advertisements they clicked.

In Table 8, we estimate the effect of treatment on reported advertisement clicks. Using a differences-in-differences analysis, we condition on how many advertisements a treated user actually clicked on. We find that when a user was treated with the “Paid Advertisement” label, not only did the user click fewer

advertisements (per Table 2), but also the user *correctly reported* clicking on fewer advertisements. See Table 8, row (4) in both Panels A and B.

Section 6.1 indicates that the “Paid Advertisement” label causes users to click fewer ads, but this change could be either good (users recognizing listings as ads and declining to click) or bad (users scared away from ads they would have benefited from clicking). However, these two mechanisms offer opposite predictions for user recall of the ads they saw: The former distinctively implies an increase in recall. Our questions about user perceptions---finding an increase in recall---support the first theory and not the second. In particular, because the “Paid Advertisement” label causes users to more accurately report how many advertisements they clicked, our results support the hypothesis that the “Paid Advertisement” label provides a genuine increase in user understanding.

#### 6.4 Comparing Google’s Labels: “Sponsored links” versus “Ads”

Our experimental framework allows us to examine the effects of Google’s November 4, 2010 switch from “Sponsored links” to “Ads.” As a threshold matter, both labels are statistically significantly different from “Paid Advertisements”: as detailed in the preceding sections, users seeing “Paid Advertisements” click fewer advertisements than users seeing either “Sponsored links” or “Ads,” but are more accurate in reporting how many advertisements they clicked. Furthermore, for both “Sponsored links” and “Ads,” differences from “Paid Advertisement” are most pronounced for users who have less education and less online experience.

Our data indicate that the labels “Sponsored links” and “Ads” are similar to each other, but both quite distant from “Paid Advertisement.” Table 2 formalizes this assessment: Treating “Sponsored links” as the baseline (since all search engines used substantially this label through November 2010), the label “Paid Advertisement” offers an average reduction of 0.586 advertisement clicks per user (during our task sequence), whereas the gap from “Ads” to “Paid Advertisement” is 0.538 clicks. Thus, “Ads” progresses only 5% along the spectrum from “Sponsored links” to “Paid Advertisements.” That said, our experiments found no statistically significant difference between users’ advertisement click rates in the “Sponsored links” and “Ads” treatments; whatever difference the new label made, we cannot distinguish that difference from zero. (See Table 9, reporting no coefficients statistically significantly different from zero.) This finding is consistent with industry reports that the “Ads” label did not reduce clickthrough rates and may even have caused a small increase. (Schwarz, 2010b)

## 7 Policy Implications and Conclusion

Our data indicates that the label “Paid Advertisement” is more effective than “Sponsored link” at conveying to users that a given link is in fact an advertisement: When the former label appears in place of the latter, users click somewhat fewer advertisements, but users more accurately report how many advertisements they clicked, indicating that users better understand which results are advertisements.

Furthermore, the “Paid Advertisement” label is especially beneficial in informing the users who otherwise least understand the meaning of “Sponsored links” (including older users, less educated users, and Internet novices). Indeed, as shown in Table 6, the insertion of “Paid Advertisement” labels makes non-college-graduates 31% less likely to click advertisements---transforming these users from frequent ad-clickers into users who equal educated users in their advertisement click rates.

The FTC has called for “clear” advertisement labels that are “easy for consumers to understand.” (Hipsley, 2002) Our results indicate that whatever the clarity and effectiveness of current advertisement labels, the “Paid Advertisement” alternative appears to be significantly more clear and significantly easier to understand. It seems search engines could change their labels easily; we see no technical costs in implementing such a change.

Search engines and advertisers might be concerned that improved advertisement labeling would reduce advertisement clicks. But FTC precedent does not indicate that company profits are a factor to consider when evaluating advertisement labels; the apparent purpose of regulation in this area is to protect users from advertisements they do not recognize as such, even if improved labels reduce company profits. That said, is it unclear whether a reduction in advertisement clicks would reduce search engine revenue in equilibrium. If advertisers seek to reach users who recognize their offers as advertisements---for example, users who want commercial messages because they are planning purchases---then improved advertisement labels could prompt higher bids and yield stable or increased revenue to a search engine.

We are struck by Google’s change from “Sponsored links” labels to “Ads.” We do not know what prompted Google to make this change; Google issued no public statement except to confirm that the change had occurred. (Schwartz, 2010a) Some news sources speculated that Google’s objective was to increase advertisement clicks (e.g. Catacchio, 2010) or to retain advertisement clicks while addressing possible regulatory concerns (Green, 2010). But if Google seeks to better comply with the requirements of FTC regulation and caselaw, our analysis suggests that the “Ads” label still falls short---in part for its small size and colloquial tone, as discussed in section 4.3, but also for apparent user response to the label, as measured by our experiments. Indeed, if “Ads” offers any improvement in user understanding compared to “Sponsored links,” the improvement is sufficiently small that our experiment could not distinguish it from zero. To clearly and conspicuously inform users, our experiments indicate that “Paid Advertisement” remains the better choice.

We view our contribution as threefold. First, we evaluate a widely-used disclosure and two plausible alternatives, assessing their relative merits and showing their distinctive effects on particular subsets of users. Second, we isolate the mechanism of labeling effects---showing that alternate disclosures reduce advertisement clicks not through user fear or confusion but through a genuine increase in understanding. Third, we offer a methodology of general applicability---showing how other researchers outside the lab setting can use online experiments to test changes to arbitrary web sites, without needing assistance from those sites.

## **Acknowledgements**

We thank Al Roth and two anonymous referees.

# Tables and Figures

Figure 1: "Sponsored links" at Google




Figure 2: "Sponsored Results" at Yahoo




Figure 3: "Sponsored sites" at Bing




Figure 4: Questions to Participants

### **Tell Us About Yourself**

1. What is your age? [5 choices]
2. What is your household's income? [7 choices]
3. Approximately how many hours do you spend browsing the Internet every week? [5 choices]
4. Do you use a social network (Facebook, LinkedIn, etc.)? [yes/no]
5. Have you bought the following online in the last year? [yes/no]
6. How would you rate your computer expertise? [4 choices]
7. What is the highest level of education you have completed? [7 choices]
8. What is your employment status? [7 choices]

### **Question 1: Online Mattress Purchasing**

In this study, we are trying to understand how people search for information on the Internet. To get started, we'd like you to answer the following question:

Find three websites that sell twin-size AeroBed® inflatable mattresses and a contact phone number for each website.

To answer this question, use our web-browser (accessible via the button below) to access websites you would ordinarily use for online research. You will notice that our web-browser has the functionality of any standard browser, including navigation buttons and an address bar. Please conduct all your research through our web-browser, without opening another tab or window.

Please jot down your answers on a piece of paper. When you're done with your research, press the Done With Research button and type your answers into the resulting screen. Not all websites selling twin-size AeroBed® inflatable mattresses will provide a contact phone number, but we ask that you find three that do.

### **Question 2: Online Medical Research**

As before, please answer the following question using our web-browser (accessible via the button below). Jot your answers down on a piece of paper before returning to the survey.

Imagine that a friend calls to tell you he or she has been diagnosed with leukemia (cancer of the blood or bone marrow). Your friend is panicking and knows little about the disease, so you decide to do some research. What methods are used to treat leukemia? (Name three methods.)

### **Follow-Up 1: Online Mattress Research Follow-Up**

The following questions refer to the Internet research you conducted regarding twin-size AeroBed® inflatable mattresses:

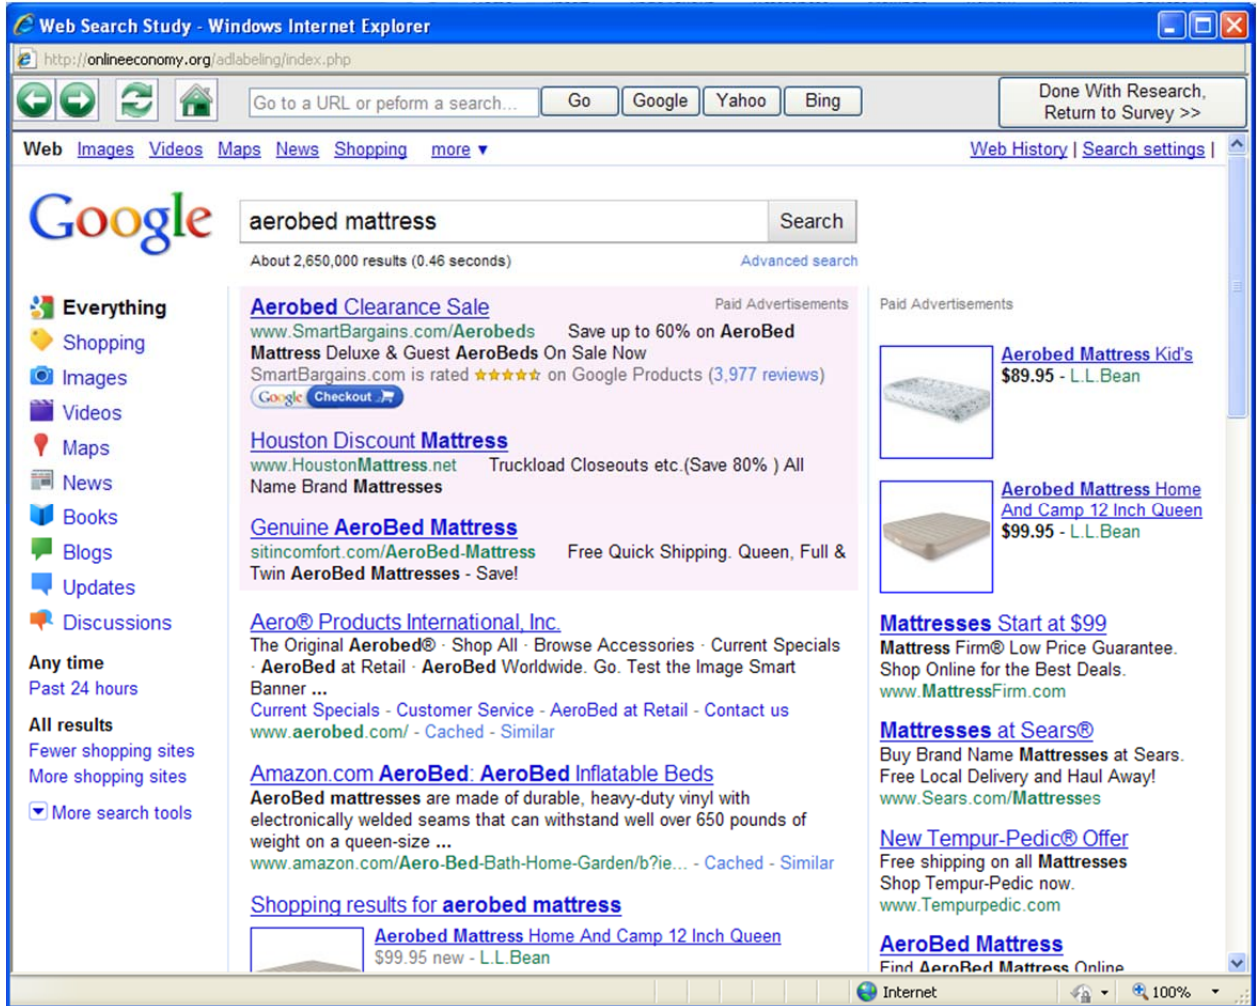
1. Did you use a search engine to do your research? [yes/no]
2. During your searches, you might have seen advertisements, Sponsored links, or other listings the search engine was paid to show you. Did you see any of these? [yes/no]
3. How many searches did you run? If you're not sure, enter your best guess. [numerical]
4. How many advertisements, Sponsored links, or other paid listings did you see? If you're not sure, enter your best guess. [numerical]
5. Did you click any advertisements, Sponsored links, or other paid listings? [yes/no]
6. How many advertisements, Sponsored links, or other paid listings did you click? If you're not sure, enter your best guess. [numerical]

### **Follow-Up 2: Online Medical Research Follow-Up**

Analogous to Follow-Up 1.



Figure 5: Our Online Experiment Interface



*Table 1: Summary Statistics for User Searching*

Condition:	Purchasing research			Health research		
	“Sponsored links”	“Ads”	“Paid Advertisements”	“Sponsored links”	“Ads”	“Paid Advertisements”
Actual searches	3.440 (3.125)	3.131 (2.499)	2.856 (2.178)	1.596 (0.997)	1.387 (1.009)	1.555 (1.842)
Actual link clicks	6.133 (3.056)	5.752 (2.879)	5.459 (2.891)	1.476 (0.939)	1.431 (1.056)	1.404 (0.965)
Actual ad clicks	2.042 (2.104)	1.781 (2.035)	1.637 (1.975)	0.271 (0.544)	0.204 (0.487)	0.260 (0.539)
Actual organic clicks	4.090 (2.778)	3.971 (2.898)	3.822 (2.575)	1.205 (0.957)	1.226 (1.078)	1.144 (1.037)
N	166	137	146	166	137	146

Standard deviations are reported in parentheses.

Table 2: Effect of Treatment On Actual Ad Clicks

Panel A: Effect of “Paid Advertisements” Treatment vs. “Sponsored links” Control								
Specification:	(1)	(2)	(3)	(4)	(5)	(6)	(7)	
Dependent variable: ad clicks during...	Purchasing research	-0.405* (0.232)	-0.434* (0.228)	-0.443* (0.230)	-0.467** (0.231)	-0.518** (0.234)	-0.565** (0.237)	-0.567** (0.237)
	Health research	-0.0108 (0.0614)	-0.0207 (0.0595)	-0.0191 (0.0603)	-0.0240 (0.0608)	-0.0182 (0.0614)	-0.0180 (0.0627)	-0.0186 (0.0627)
	All research	-0.416 (0.254)	-0.455* (0.247)	-0.462* (0.249)	-0.491* (0.251)	-0.536** (0.254)	-0.583** (0.257)	-0.586** (0.257)
Panel B: Effect of “Paid Advertisements” Treatment vs. “Ads” Control								
Specification:	(1)	(2)	(3)	(4)	(5)	(6)	(7)	
Dependent variable: ad clicks during...	Purchasing research	-0.144 (0.238)	-0.394 (0.244)	-0.409* (0.245)	-0.445* (0.247)	-0.477* (0.248)	-0.505** (0.250)	-0.500** (0.250)
	Health research	0.0559 (0.0612)	-0.0247 (0.0618)	-0.0291 (0.0624)	-0.0368 (0.0633)	-0.0373 (0.0634)	-0.0387 (0.0642)	-0.0378 (0.0642)
	All research	-0.0881 (0.257)	-0.419 (0.260)	-0.438* (0.261)	-0.482* (0.263)	-0.514* (0.265)	-0.544** (0.267)	-0.538** (0.267)
Covariates (for both panels)								
Age		X	X	X	X	X	X	
Dummies								
Education			X	X	X	X	X	
Income				X	X	X	X	
Weekly hours online					X	X	X	
Online shopping						X	X	
Gender							X	

\*Significant at the 10% level. \*\*Significant at the 5% level.

Number in “Sponsored links” condition = 166, “Ads” condition = 137, “Paid Advertisements” condition = 146. Standard errors are reported in parentheses. Each estimate/standard error pair results from a different regression.

Dummies split income into 7 categories, education into 7 categories, and weekly hours online into 6 categories. The online shopping dummy is split into 5 categories which are indexed as the sum of binary indicators for whether the user reported spending money online in the last year on books, clothing, electronics, travel, and groceries.

Table 3: Effect of Treatment on Organic Clicks

Panel A: Effect of “Paid Advertisements” Treatment vs. “Sponsored links” Control								
Specification:	(1)	(2)	(3)	(4)	(5)	(6)	(7)	
Dependent variable: organic clicks during...	Purchasing research	-0.268 (0.305)	-0.258 (0.305)	-0.254 (0.306)	-0.196 (0.302)	-0.176 (0.307)	-0.185 (0.311)	-0.181 (0.311)
	Health research	-0.0610 (0.113)	-0.0587 (0.113)	-0.0459 (0.114)	-0.0481 (0.115)	-0.0610 (0.117)	-0.0550 (0.118)	-0.0535 (0.118)
	All research	-0.267 (0.360)	-0.255 (0.361)	-0.256 (0.362)	-0.190 (0.357)	-0.181 (0.362)	-0.178 (0.368)	-0.176 (0.369)
Panel B: Effect of “Paid Advertisements” Treatment vs. “Ads” Control								
Specification:	(1)	(2)	(3)	(4)	(5)	(6)	(7)	
Dependent variable: organic clicks during...	Purchasing research	-0.149 (0.325)	-0.00556 (0.339)	0.0604 (0.341)	0.0603 (0.346)	0.120 (0.345)	0.143 (0.348)	0.132 (0.347)
	Health research	-0.0824 (0.126)	-0.0522 (0.131)	-0.0433 (0.132)	-0.0308 (0.135)	-0.00811 (0.136)	-0.00273 (0.136)	-0.00409 (0.136)
	All research	-0.0330 (0.378)	0.119 (0.395)	0.200 (0.397)	0.229 (0.403)	0.317 (0.402)	0.349 (0.406)	0.339 (0.405)
Covariates (for both panels)								
Age		X	X	X	X	X	X	
Dummies								
Education			X	X	X	X	X	
Income				X	X	X	X	
Weekly hours online					X	X	X	
Online shopping						X	X	
Gender							X	

\*Significant at the 10% level. \*\*Significant at the 5% level.

Number in “Sponsored links” condition = 166, “Ads” condition = 137, “Paid Advertisements” condition = 146. Standard errors are reported in parentheses. Each estimate/standard error pair results from a different regression. Dummies are as in Table 2.

Table 4: Effect of Demographic Characteristics On Actual Ad Clicks

Specification:	Dependent variable: actual ad clicks during all research				
	(1)	(2)	(3)	(4)	(5)
“Paid Advertisements” vs. “Sponsored links”	-0.428*	-0.420*	-0.437*	-0.403	-0.393
	(0.246)	(0.252)	(0.252)	(0.252)	(0.250)
“Paid Advertisements” vs. “Ads”	-0.348	-0.069	-0.096	-0.100	-0.131
	(0.263)	(0.264)	(0.264)	(0.264)	(0.263)
Age ≥ 47	1.000***				
	(0.210)				
Annual household income <\$100k		0.284			
		(0.260)			
Not a college graduate			0.299		
			(0.211)		
Browses the Internet <4 hours/week				0.339	
				(0.297)	
No online shopping experience					0.696**
					(0.278)

\*Significant at the 10% level. \*\*Significant at the 5% level. \*\*\*Significant at the 1% level.  
N=449.

Table 5: Effect of Treatment On Actual Ad Clicks by Demographic Group

Dependent variable: actual ad clicks during all research					
Panel A: Effect of "Paid Advertisements" Treatment vs. "Sponsored links" Control					
Specification:	(1)	(2)	(3)	(4)	(5)
Treatment	-0.540 (0.394)	-0.173 (0.534)	0.050 (0.381)	-0.261 (0.277)	-0.118 (0.280)
Age ≥ 47	0.755** (0.348)				
Age ≥ 47 * Treatment	0.190 (0.510)				
Annual household income <\$100k		0.510 (0.410)			
Annual household income <\$100k * Treatment		-0.321 (0.607)			
Not a college graduate			0.701** (0.346)		
Not a college graduate * Treatment			-0.876* (0.510)		
Browses the Internet <4 hours/week				0.513 (0.451)	
Browses the Internet <4 hours/week * Treatment				-0.949 (0.694)	
No online shopping experience					1.124*** (0.422)
No online shopping experience * Treatment					-1.467** (0.639)
Panel B: Effect of "Paid Advertisements" Treatment vs. "Ads" Control					
Specification:	(1)	(2)	(3)	(4)	(5)
Treatment	-0.184 (0.352)	-0.202 (0.608)	0.183 (0.395)	0.091 (0.274)	0.140 (0.276)
Age ≥ 47	1.383*** (0.376)				
Age ≥ 47 * Treatment	-0.437 (0.516)				
Annual household income <\$100k		0.039 (0.513)			
Annual household income <\$100k * Treatment		0.150 (0.671)			
Not a college graduate			0.300 (0.372)		
Not a college graduate * Treatment			-0.475 (0.520)		
Browses the Internet <4 hours/week				1.064* (0.588)	
Browses the Internet <4 hours/week * Treatment				-1.501* (0.776)	
No online shopping experience					1.432** (0.569)
No online shopping experience * Treatment					-1.775** (0.733)

\*Significant at the 10% level. \*\*Significant at the 5% level. \*\*\*Significant at the 1% level.

Number in "Sponsored links" condition = 166, "Ads" condition = 137, "Paid Advertisements" condition = 146. Standard errors are reported in parentheses.

Table 6: Average Total Actual Ad Clicks by Treatment Condition and Demographics

	Condition:	“Sponsored links”	“Ads”	“Paid Advertisements”
Age $\geq$ 47		2.622	2.894	2.273
Age < 47		1.868	1.511	1.328
<hr/>				
	Condition:	“Sponsored links”	“Ads”	“Paid Advertisements”
Annual household income <\$100k		2.433	1.991	1.939
Annual household income $\geq$ \$100k		1.923	1.952	1.750
<hr/>				
	Condition:	“Sponsored links”	“Ads”	“Paid Advertisements”
No College Degree		2.651	2.117	1.826
College Graduate		1.950	1.817	2.000
<hr/>				
	Condition:	“Sponsored links”	“Ads”	“Paid Advertisements”
Browses the Internet <4 hours/week		2.733	2.933	1.524
Browses the Internet 4+ hours/week		2.221	1.869	1.960
<hr/>				
	Condition:	“Sponsored links”	“Ads”	“Paid Advertisements”
No online shopping experience		3.200	3.250	1.615
Online shopping experience		2.076	1.818	1.958

We classify a user as having no online shopping experience if the user reports not buying books, clothing, electronics, travel, or groceries online in the last year.

Table 7: Summary Statistics of User Perception of Searches and Advertisements

Condition:	Purchasing research			Health research		
	“Sponsored links”	“Ads”	“Paid Advertisements”	“Sponsored links”	“Ads”	“Paid Advertisements”
Reported searches	2.713 (1.982)	2.277 (1.654)	2.667 (2.012)	1.264 (0.769)	1.169 (0.627)	1.395 (1.057)
Reported ad views	4.701 (6.279)	3.400 (4.010)	4.200 (6.900)	2.130 (4.411)	1.215 (2.553)	2.160 (6.351)
Reported ad clicks	0.517 (1.190)	0.554 (1.000)	0.284 (0.794)	0.0690 (0.255)	0.0923 (0.458)	0.136 (0.494)
N	87	65	81	87	65	81

Standard deviations are reported in parentheses.

Table 8: Effect of Treatment on Reported Ad Clicks

Dependent variable:						
Ad clicks reported after	Purchasing research		Health research		All research	
Panel A: Effect of “Paid Advertisements” Treatment vs. “Sponsored links” Control						
Coefficient estimates						
Treatment	-0.308*	-0.350	0.0493	-0.0390	-0.259	-0.556**
	(0.161)	(0.215)	(0.0612)	(0.0609)	(0.197)	(0.261)
Actual ad clicks		0.0742		0.175**		0.0683
		(0.0510)		(0.0873)		(0.0570)
Treatment * Actual ad clicks		0.0510		0.245**		0.174**
		(0.0803)		(0.109)		(0.0863)
Diff-in-diff estimate of the effect of treatment <sup>c</sup>		-0.270*		0.0367		-0.229
		(0.161)		(0.0552)		(0.191)
Panel B: Effect of “Paid Advertisements” Treatment vs. “Ads” Control						
Coefficient estimates						
Treatment	-0.303*	-0.321	0.0290	-0.110	-0.274	-0.538*
	(0.158)	(0.200)	(0.0859)	(0.0898)	(0.217)	(0.275)
Actual ad clicks		0.115**		-0.0568		0.115*
		(0.0494)		(0.158)		(0.0656)
Treatment * Actual ad clicks		0.0338		0.491***		0.155*
		(0.0743)		(0.179)		(0.0921)
Diff-in-diff estimate of the effect of treatment <sup>c</sup>		-0.236*		0.0112		-0.236
		(0.153)		(0.0799)		(0.204)

\*Significant at the 10% level. \*\*Significant at the 5% level. Sample size is 168 in Panel A, 146 in Panel B. Regressions include an age covariate and dummies for education. Standard errors are reported in parentheses.

<sup>c</sup> Diff-in-diff estimate is calculated as follows:

[Treatment coefficient] + [Treatment \* Actual ad clicks coefficient] \* [Avg actual ad clicks | treated]



Table 9: Effect of “Ads” label vs. “Sponsored links” label

Specification:	(1)	(2)	(3)	(4)	(5)	(6)	(7)	
Dependent variable: ad clicks during...	Purchasing research	-0.261 (0.239)	-0.0266 (0.243)	-0.0273 (0.243)	-0.0551 (0.244)	-0.0470 (0.245)	-0.0174 (0.245)	-0.0195 (0.245)
	Health research	-0.0667 (0.0599)	0.00920 (0.0598)	0.00510 (0.0597)	0.00930 (0.0606)	0.0147 (0.0603)	0.0237 (0.0612)	0.0235 (0.0613)
	All research	-0.328 (0.260)	-0.0174 (0.261)	-0.0222 (0.261)	-0.0458 (0.262)	-0.0323 (0.263)	0.00635 (0.263)	0.00394 (0.264)
Covariates (for all panels)								
Age		X	X	X	X	X	X	
Dummies								
Education			X	X	X	X	X	
Income				X	X	X	X	
Weekly hours online					X	X	X	
Online shopping						X	X	
Gender							X	

\*Significant at the 10% level. \*\*Significant at the 5% level.

Number in “Sponsored links” condition = 166, “Ads” condition = 137. Standard errors are reported in parentheses. Each estimate/standard error pair results from a different regression. Dummies are as in Table 2.

## References

- Biener, L., Albers, A., 2004. Young Adults: Vulnerable New Targets of Tobacco Marketing. *Amer. J. Pubic Health*, 94(2), Feb, pp.326-330.
- Brown, J., Hossain, T., Morgan, J., 2010. Shrouded Attributes and Information Suppression: Evidence from the Field. *Quart. J. Econ.*, 125(2), pp.859-876.
- Catachio, C., 2010. "Say Goodbye to 'Sponsored Links', Hello to 'Ads' on Google Results Pages." *The Next Web*. <http://thenextweb.com/google/2010/11/05/say-goodbye-to-sponsored-links-hello-to-ads-on-google-results-pages/> , accessed Jan 11, 2012.
- Edelman, B., 2006a. "Pushing Spyware through Search." January 26, 2006, <http://www.benedelman.org/news/012606-1.html> , accessed Jan 11, 2012.
- Edelman, B., 2006b. "False and Deceptive Pay-Per-Click Ads." <http://www.benedelman.org/ppc-scams/> , accessed Jan 11, 2012.
- Edelman, B. "A Closer Look at Google's Advertisement Labels." November 10, 2010. <http://www.benedelman.org/adlabeling/google-nov2010.html> , accessed Jan 11, 2012.
- Edelman, B., Ostrovsky, M., Schwarz, M., 2007. Internet Advertising and the Generalized Second Price Auction: Selling Billions of Dollars Worth of Keywords. *Amer. Econ. Rev.*, 97(1), Mar, pp.242-259.
- Fallows, D., 2005. "Search Engine Users." *Pew Internet & American Life Project*. <http://www.pewinternet.org/Reports/2005/Search-Engine-Users.aspx> , accessed Jan 11, 2012.
- Friestad, M., Wright, P., 1994. The Persuasion Knowledge Model: How People Cope with Persuasion Attempts. *J. Cons. Res.*, 21(1), pp.1-31.
- FTC, 2000. "Dot Com Disclosures." <http://business.ftc.gov/sites/default/files/pdf/bus41-dot-com-disclosures-information-about-online-advertising.pdf> , accessed Jan 11, 2012.
- FTC v. Michael S. Levey, et al., 1993. "Proposed Consent Agreement with Analysis to Aid Public Comment." 58 Fed. Reg. 38764.
- Gabaix, X., Laibson, D., 2006. Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets. *Quart. J. Econ*, 121(2), pp.505-540.
- Goldfarb, A., Tucker, C., 2011. Online Display Advertising: Targeting and Obtrusiveness. *Marketing Sci.*, 30(3), pp.389-404.
- "Google Search Basics: Search Results Page." *Google Help*. <http://www.google.com/support/websearch/bin/answer.py?hl=en&answer=35891> , accessed Jan 11, 2012.
- Green, A., 2010. "Click-Through Rate Jumps 11.4% After Google Changes 'Sponsored Links' to 'Ads'." *The Search Agents*. November 15, 2010. <http://www.thesearchagents.com/2010/11/click-through-rate-jumps-after-google-switches-from-sponsored-links-to-ads/> , accessed Jan 11, 2012.
- Hipsley, H., 2002. "RE: Commercial Alert Complaint Requesting Investigation of Various Internet Search Engine Companies for Paid Placement and Paid Inclusion Programs." <http://www.ftc.gov/os/closings/staff/commercialalertattach.shtm> , accessed Jan 11, 2012.
- Jacoby, J., Nelson, M., Hoyer, W., 1982. Corrective Advertising and Affirmative Disclosure Statements: Their Potential for Confusing and Misleading the Consumer. *J. Marketing*, 46(1), winter, pp.61-72.

- Jansen, B.J., Brown, A., Resnick, M., 2007. Factors Relating to the Decision to Click on a Sponsored Link. *Dec. Support Systems*, 44, pp.46-59.
- Kirmani, A., Zhu, R., 2007. Vigilant Against Manipulation: The Effect of Regulatory Focus on the Use of Persuasion Knowledge. *J. Marketing Res.*, 44(4), pp.688-701.
- Marable, L., 2003. "False Oracles: Consumer Reaction to Learning the Truth About How Search Engines Work. Results of an Ethnographic Study." *Consumer Web Watch*.  
<http://www.consumerwebwatch.org/news/searchengines/index.html> , accessed January 11, 2012.
- Pan, B., Hembrooke, H., Joachims, T., Lorigo, L., Gay, G., Granka, L., 2007. In Google We Trust; Users' Decisions on Rank, Position, and Relevance. *J. Computer-Mediated Communication*, 12, pp. 801-823.
- Rayo, L., Segal, I., 2010. Optimal Information Disclosure. *J. Pol. Econ.*, 118(5), pp.949-987.
- Rodgers, S., Thorson, E., 2000. The Interactive Advertising Model: How Users Perceive and Process Online Ads. *J. Interact. Advertising*, 1(1), pp. 42-61.
- Ruskin, G., 2001. "RE: Deceptive advertising complaint against AltaVista Co., AOL Time Warner Inc., Direct Hit Technologies, iWon Inc., LookSmart Ltd., Microsoft Corp. and Terra Lycos S.A." <http://www.commercialalert.org/PDFs/SearchEngines.pdf> , accessed January 11, 2012.
- Ruskin, G., 2003. "RE: Request for Investigation of Product Placement on Television and for Guidelines to Require Adequate Disclosure of TV Product Placement." <http://www.commercialalert.org/ftcall.pdf> , accessed January 11, 2012.
- Schwarz, B., 2010a. "Google Does Away With 'Sponsored Links' Label, Now Ads Are Labeled 'Ads'." *Search Engine Land*. November 5, 2010. <http://searchengineland.com/google-does-away-with-sponsored-links-label-now-ads-are-labeled-ads-54956> , accessed January 11, 2012.
- Schwartz, B., 2010b. "Google 'Ads' From 'Sponsored Links' Name Change Positive Impact on Click Through Rate." *Search Engine Roundtable*. November 17, 2010.  
<http://www.seroundtable.com/archives/023278.html> , accessed January 11, 2012.
- Thompson, L., 2010. "Google Making Millions from Advertising Counterfeit Goods." *The Times of London*. [http://technology.timesonline.co.uk/tol/news/tech\\_and\\_web/the\\_web/article6981227.ece](http://technology.timesonline.co.uk/tol/news/tech_and_web/the_web/article6981227.ece) , accessed January 11, 2012.
- Varian, H., 2010. Computer Mediated Transactions. *Amer. Econ. Rev.*, 100(2), pp.1-10.
- Wooldridge, J. M., 1994. *Introductory Economics*, fourth ed. South-Western Cengage Learning, Mason OH.
- 4crests, "Wow, look at the new 'Sponsored Link'," 2002. *Webmaster World*.  
<http://www.webmasterworld.com/forum81/85.htm> , accessed February 28, 2011.