

Quality Uncertainty and the Performance of Online Sponsored Search Markets: An Empirical Investigation*

Animesh Animesh

Vandana Ramachandran

Siva Viswanathan

Department of Decision and Information Technologies

Robert H. Smith School of Business

University of Maryland, College Park, MD 20742

{ animesh@umd.edu, vramacha@rhsmith.umd.edu, sviswana@rhsmith.umd.edu }

ABSTRACT

Sponsored search mechanisms, where advertisers bid for better placement in the listing of search results on search services such as Yahoo! and Google integrate the benefits of online search with advertising, and have emerged as the dominant revenue model for online search engines. Interestingly, Yahoo! and Google employ different mechanisms to determine the placement of bidders' advertisements. This provides an unprecedented opportunity to not only test some of the predictions of earlier research relating advertiser's quality and their advertising intensity, in the online setting, but to also examine whether intervention by the search intermediary impacts the performance of these markets. Using data from online sponsored search auctions this study examines whether the relationship between advertisers' quality and their advertising intensity varies across product categories as well as across the different search mechanisms. Our results highlight significant and interesting differences in the quality-advertising relationships across the two market mechanisms as well as across products characterized by differing degrees of quality uncertainty.

*The authors would like to thank the NET institute for generous financial support. This paper has greatly benefited from comments from anonymous reviewers at ICIS 2006, CIST 2006, and participants at the Workshop on Sponsored Search Auctions, Vancouver, CA. We are also grateful to Carolyn Gideon, Sanjay Gosain, Nathan Larson, Ravi Mantena, and the participants at the 2006 NET Institute Conference for their insights and feedback. Danny Jamal provided valuable research assistance.

Quality Uncertainty and the Performance of Online Sponsored Search Markets: An Empirical Investigation*

1. Introduction

Sponsored search mechanisms, where advertisers bid for better placement in the listing of search results on search services such as Yahoo! and Google, have emerged as the dominant revenue model for online search engines. The emergence of *sponsored search* as a viable alternative not only to *organic* (algorithm-based) search but also to traditional advertising raises several issues of interest to academicians as well as practitioners. The presence of paid results in these online search engines presents a new kind of informational problem in the digital realm. An inherent conflict of interest arises in the sponsored search (also known as paid-placement or keyword advertising) context, where information intermediaries deliver information about sellers and their offerings, but are paid by those same sellers they “certify” (Gaudeul 2004). On the one hand “sponsored search” can potentially introduce a bias in the search results, thereby reducing the value of online search to consumers. On the other hand, the validity of the sponsored search model is evinced by the increased revenues it has bestowed upon hosting search engines since its inception, and the new complementary markets it has spawned in the wake of its success, such as search engine marketing and optimization. Given the controversies surrounding sponsored search mechanisms, and in light of their expanding importance for the emerging economic and competitive landscape online, it is important to gain a better understanding of the implications of sponsored search mechanisms for their participants as well as for policy makers.

Sponsored search advertising accounts for more than 40% of the total online advertising dollars spent by companies in the United States (Markoff and Ives 2005). This is projected to more than double over the next

The authors would like to thank the NET institute for generous financial support. This paper has greatly benefited from comments from anonymous reviewers at ICIS 2006, CIST 2006, and participants at the Workshop on Sponsored Search Auctions, Vancouver, CA. We are also grateful to Carolyn Gideon, Sanjay Gosain, Nathan Larson, Ravi Mantena, and the participants at the 2006 NET Institute Conference for their insights and feedback. Danny Jamal provided valuable research assistance.

few years. Despite the rapid growth of sponsored search, a vast majority of online consumers (62%) are unaware of the distinction between *organic search* results and *sponsored search* results (Fallows 2005). Even among consumers who are aware of sponsored search results, the majority believe that an advertiser ranked higher in the sponsored search results is of higher quality than those appearing lower in the search listings (iProspect 2006). These beliefs are also directly reflected in the substantially higher number of clickthroughs that firms listed on the top receive, compared to those appearing lower in the listings. According to an independent study by search marketing firms Enquiro and Did-it (using eye tracking tools), while nearly 50% of the users viewed the top ranked advertisement, this exposure was reduced to only 10% at rank 5 and below (Sherman 2005). In addition, as indicated by a recent study by DoubleClick (2006), more than 30% of total purchases are made from sellers listed at the top of the search listings. Given the proclivity of consumers to visit (and buy from) the advertisers appearing on top of the search listings, the highest bidders stand to gain disproportionately more than their counterparts placed lower down the search listings. Clearly being on top of the sponsored listings is beneficial to *all* firms. However, consumers as well as search intermediaries would benefit only if the firms listed on top are also of higher quality. Given the current state of consumer awareness and beliefs, sponsored search mechanisms characterized by adverse selection - where low quality bidders are placed at the top of the search listings – can lower consumer welfare and reduce the utility of such mechanisms for consumers. On the other hand, the widespread popularity and burgeoning revenues associated with sponsored search markets suggest that the market mechanism may be self-correcting. Consequently, understanding the quality-rank correlation becomes imperative as it has significant implications for consumer welfare as well as the future of sponsored search advertising. The relationship between an advertiser's quality³ and its ranking in the sponsored search results can differ across product categories and also across the

³ Quality in our study refers to the quality of the seller/advertiser, which is typically a multi-dimensional construct. We use multiple measures of advertiser-quality as described in section 2.

different online sponsored search markets. The *relative performance* of sponsored search mechanisms therefore remains an empirical question that needs to be validated, and is one that we examine in this study.

There exist a few different online sponsored search mechanisms, the two most popular ones being Yahoo! and Google. While the former determines the advertiser's rank/position in the sponsored listings based only on the bid amount (a pure market mechanism), the latter provides a more "regulated" environment where the ranking of an advertiser is based on a complex score that takes into account the "effectiveness of the advertisement" (its ability to attract clicks), in addition to the bid amount. The existence of these different sponsored search mechanisms, one with, and the other without intervention by the search intermediary, provides a unique opportunity to compare their relative performance and welfare implications. Specifically, we seek to address the following questions: (1) Are firms that are ranked higher in the sponsored search listings of higher quality than firms appearing lower down the listings? (2) More importantly, does this relationship between an advertiser's rank and its quality *differ* across products characterized by differing degrees of quality uncertainty? (3) Finally, are these *differences* if any, consistent across the two primary sponsored search markets?

2. Research Context

It is well known that electronic marketplaces have dramatically lowered the cost of consumers obtaining information about product offerings and prices as well as the cost to sellers communicating such information – with search engines such as Google and Yahoo! playing a pivotal role in this process. Traditionally these search services have used "crawler-based" search mechanisms to create their results lists (also called *organic* listings), based on algorithms that attempt to maximize the relevance of their search results to a user's query. The perceived unbiasedness and relevance of organic search results have contributed to the popularity of search engines and helped establish them as gateways to the Internet. Search engines have historically served

organic listings for free, while being supported by revenues from advertisers and from selling consumer search-related information.

Sponsored search

The sponsored search advertising model employs a continuous open bid auction mechanism that allows advertisers to bid for placement (to be as close to the top) in the listing of the paid search results which are displayed at the top and to the side of organic listings in response to keyword searches conducted by search engine users. Each advertiser pays the bid amount for each click received, but is not charged for the exposures in this pay-per-click (PPC) model. The higher the bid, the higher the seller's advertisement (usually a short text message and a link to the advertiser's URL) appears in the results. The paid listings (particularly, the highest bidders appearing at the top of the search results) on Yahoo! and Google are also distributed to a network of search websites, such as AltaVista, Ask, Lycos, and Netscape Search, among others. For example, until MSN developed its home-grown sponsored search listings, both MSN Search and MSN websites featured the top ranked sponsored listings generated by Yahoo! for user queried keywords.

Sponsored Search Mechanisms on Google and Yahoo!

Google, with a U.S. market share of nearly 37%, followed by Yahoo! at 30%, are the two premier search engines online. Both search engines routinely display sponsored listings alongside their organic search results. However interestingly, the two search intermediaries adopt different mechanisms in ranking their sponsored search results. Advertisers that wish to be listed on Yahoo!'s or Google's sponsored search listings (in response to a keyword search initiated by a consumer) can bid on the keywords related to their offerings. On Yahoo!, the higher the advertiser's bid per-click in the auction, the higher the placement the advertisement receives in the listing of sponsored search results triggered by a keyword query. However in the case of Google, the position of a firm's/advertisement's listing is a function of the advertiser's bid per-click as well as its clickthrough rate (CTR), i.e. the number of clicks the advertisement gets when displayed. If a seller

(advertisement) fails to generate sufficient clicks from users, it is penalized and is moved lower down the list. When the CTR consistently falls below a minimum threshold, then Google considers it as a signal indicating that the advertisement is irrelevant to the query, and drops the advertisement. The differences between Yahoo! and Google's sponsored search mechanisms therefore brings to forefront the question related to the effectiveness of the search engine's intervention (or the lack thereof) in determining the placement of paid advertisements. Our study investigates this issue by comparing the performance outcomes of the two mechanisms.

Performance and Quality of Sponsored Search Listings

In order to maintain credibility amongst their users, search engines employ a number of threshold criteria to determine which advertisements are placed in higher ranked positions in the sponsored search listings. Of primary importance is the *relevance* of the advertisement. Most of existing research in the domain of online search examines the *relevance* of the "advertisement" (search listing) to the user's keyword/query. Feng et al. (forthcoming) develop an analytical model to examine the performance of alternate sponsored search models in existence, and find that they perform nearly as well, provided that advertisers' willingness to pay for paid-placement is positively correlated with the true relevance of the advertisement to the searched keyword. This implies that the success of paid search relies heavily on the ability of the ads to generate clicks from users, and this in turn, depends on how closely the ad matches the searcher's needs (as depicted by the keyword(s) searched for). Not surprisingly then, search engines strive to attain appropriate matches between keyword queries and search results displayed.

While relevance ensures that non-matching results are not displayed in top-ranked positions on listings, it does not guarantee the quality of the seller submitting the advertisement. For instance, both "low-quality" and "high-quality" sellers may submit relevant advertisements for keywords of their choice. Our focus therefore is different — we are interested in understanding the relationship between the quality of the advertisers and their

rankings in the search listings. While relevance focuses on the *quality of the ad*, our focus is on the *quality of the bidders/advertisers* in the sponsored listings. To the best of our knowledge ours is the first empirical study to compare the relationship between advertisers' quality and advertising intensity as indicated by their rankings *across* product categories as well as *across* different sponsored search markets.

3. Theoretical Background

Fundamental to sponsored search advertising is the “keyword” or “query” that advertisers bid for. In the context of our study, these keywords or queries represent the underlying product, service, or offerings from sellers that consumers search for. This integration and co-evolution of *search* and *advertising* is a salient feature of sponsored search business models. As noted earlier, the nascence of online markets combined with the lack of salient cues exacerbates the problem of information asymmetries between buyers and sellers. Of particular importance for online purchases is the amount of information available to consumers regarding sellers and their offerings and the uncertainty resulting from any information asymmetries. Given our focus on quality uncertainty, we seek to examine if sellers of products characterized by varying degrees of quality uncertainty exhibit *differences* in their bidding behavior that manifests in the outcomes of sponsored search auctions.

A Framework for Product Categories

Nelson's (1970) classification of products into *Search*, *Experience* and *Credence* (SEC) goods is particularly relevant in this context, as it captures the underlying uncertainty consumers face in purchasing these products. The SEC goods framework has been widely used in the economics and marketing literature to examine consumer search behavior as well as firms' advertising strategies (Ekelund et al. 1995). According to the SEC framework, attributes of goods can be analyzed in terms of three properties – *search*, *experience*, and *credence* (Darby and Karni 1973; Nelson 1970). “These properties are used to categorize the point in the purchase process when, if ever, consumers can accurately assess whether a good possesses the level of an attribute claimed in advertising” (Ford et al. 1990, p. 433). *Search* goods have characteristics that are

identifiable through inspection and prior to purchase. *Experience* goods, on the other hand, have features that are revealed only through consumption. What differentiates *Credence* goods from *Experience* and *Search* goods is the fact that consumers can never be certain of the (long-term) quality and value of *Credence* goods purchased even from ex post observations and use. Typical examples of each include books, music CDs, and television for *Search*; brokerages, cruises, and healthcare for *Experience*; and psychics, tax services and counseling for *Credence* goods. It is pertinent to note here that the boundaries between these categories can be fuzzy, and the categories are best considered to represent regions in a continuum. What is important though is that the defining characteristic underlying this segmentation — pre-purchase quality uncertainty — increases from *Search* to *Experience* to *Credence* goods, as quality becomes more expensive to judge (Darby and Karni 1973).

Quality Uncertainty and Online Markets

There is a significant body of research examining issues related to adverse selection resulting from informational asymmetries between buyers and sellers regarding the reliability of the seller and the quality of goods sold (see for instance, Ausubel 1999 [credit cards], Genesove 1993 [used cars]). These informational asymmetries typically lead to a loss of efficiency stemming from the unmatched allocation of goods whereby consumers that value high quality products receive lower quality ones, and from exchanges that don't happen because buyers can't distinguish good products from bad and decide to stay out of the market altogether (Akerlof 1970). The uncertainties faced by consumers in their purchase decisions are only further exacerbated in online channels due to the lack of sensory cues.

Recognizing the potential of adverse selection in online markets, researchers have begun to examine various online markets for the possibility of adverse selection. In a study of financial brokerages, Bakos et al. (2005) find that investors using online brokerages can be burdened with hidden costs stemming from lower execution quality, compared to traditional brokerages. In a study comparing stamp auctions in online and

offline settings, Dewan and Hsu (2004) find that the lack of information regarding quality in online auctions leads to adverse selection and lower prices compared to traditional auctions. Jin and Kato (2004) analyze the extent of consumer frauds across graded and ungraded online auction markets. They find that uninformed consumers are trusting of, and willing to pay a premium to sellers who make vacuous overstatements of quality. More disturbingly, they find that it is the less reputable sellers that are likely to make such outright claims of quality in the ungraded market to overcome the otherwise lower probability of sale. These results highlight the risks faced by naïve online buyers.

To overcome these problems of adverse selection, online markets have resorted to technology-aided mechanisms such as online reputation systems that collect user feedback and product-merchant reviews. However, findings regarding the effectiveness of such mechanisms are inconclusive. Resnick et al. (forthcoming) find that reputation mechanisms are valuable to consumers as revealed by the price premiums commanded by sellers with better reputation scores. Lucking-Reiley et al. (2000) and Melnik and Alm (2002) find that the deleterious effects of negative reviews and reputations may unfavorably affect purchase intentions of potential buyers and deter cheating in a repeated game setting. Similarly, Garicano and Kaplan (2001) study used car auction sites, and find that third-party certification services help to reduce information asymmetry and adverse selection. However, other studies find that such mechanisms fail to live up to their promise as fraudulent sellers devise new loopholes to circumvent them (Brown and Morgan 2006; Resnick et al. 2000).

Our study contributes to this growing body of work that examines the effectiveness of various online markets by investigating a relatively new phenomenon on the Internet – sponsored search. The future of sponsored search advertising rests on how beneficial this mechanism is to consumers, and the promotion of low quality sellers in paid listings may adversely affect the value of such services to consumers. In particular, the comparison of the performance of different sponsored search mechanisms across the three different product categories promises to shed light on their relative effectiveness. Therefore, we seek to understand how advertiser (bidding) behavior differs across the three product categories characterized by different levels of

quality uncertainty, and the need and effectiveness of interventions in sponsored search mechanisms in alleviating any potential problems of adverse selection.

Advertising and Seller Quality

As mentioned earlier, sponsored search mechanisms are one of the fastest growing online advertising models. Advertising is clearly one important mechanism which can serve to reduce information asymmetries and help improve the efficiency of the market (Ekelund et al. 1995). Most of the existing work in advertising has focused on traditional media such as televisions, radio, magazines and newspapers. Results from analytical models suggest that advertising expenditures should be positively related to quality (Kihlstrom and Riordan 1984; Linnemer 2002; Nelson 1974). However, Schmalensee (1978) and Comanor and Wilson (1979) show that lower quality firms, under certain conditions will advertise more as compared to high quality sellers. Empirical research (for instance see, Moorthy and Zhao 2000) examining advertising in traditional media is also inconclusive about the relationship between seller quality and advertising intensity. Our research extends this stream of research to online markets. Since online markets are characterized by higher uncertainty and thus higher risk of adverse selection, we seek to examine the relationship between the quality of online advertisers and their level of advertising as indicated by their bidding intensity for search keywords.

Although online advertising has seen impressive growth and continues to grow rapidly, research in this field is lacking. Extant research on online advertising can be broadly classified into two categories: the first, examining the impact on online advertising on consumer behavior and attitudes (Schlosser et al. 1999), and the second, identifying factors that impact the effectiveness of online advertising (Chatterjee, Hoffman, and Novak 2003; Shamdasani et al. 2001). Our research complements these papers by comparing two competing markets for online advertising.

4. Data Description and Analysis

Description of Data

We collect data from two different sponsored search mechanisms – AdWords and Overture – used respectively by Google and Yahoo!. Following the Search, Experience and Credence framework commonly used in marketing literature, we selected a total of 36 keywords representing products, twelve each in the three categories. The classification of keywords into *Search*, *Experience*, and *Credence* categories is adapted from prior research (for instance, see Ekelund et al. 1995 and Nelson, 1970)⁴. For each of these keywords, we collected data on advertisers' positions or ranks achieved on listings (for both Yahoo! and Google) from the sponsored search results, once every day for a period of 60 consecutive days in late 2004. Of these 36 keywords, not all of them received sufficient bids from keyword advertisers across both these search mechanisms. We therefore restricted our focus to keywords that had a sufficient number of advertisers bidding for keywords representing the specific products, and also discarded any bidders that bid less than 30 days for each of the keywords. This helps to ensure that our data is devoid of noisy or sporadic behavior patterns exhibited by sellers. After maintaining the same number of keywords across good types, our final dataset consists of 9 keywords in each category, as listed in Table 1.

<<Insert TABLE 1 about here>>

The data on advertiser/seller quality was gathered from Alexa.com, which collects detailed site usage data from the millions of users that participate and contribute this information by using the online Alexa toolbar. Alexa then maintains and makes publicly available data on aggregate website statistics that we tap into for the purposes of our study. Prior research (see, Brynjolfsson and Smith (2000) has employed Alexa data as a proxy of website quality (Palmer 2002) as well as a proxy for firm's brand equity or social capital (Palmer et al. 2000).

Measures

⁴ The labels — *Search*, *Experience*, and *Credence* — are suggestive at best. The three categories are best thought of as representing products with increasing pre-purchase quality uncertainty for consumers.

The description of the measures is presented in Table 2. The measures are classified into three groups: a) the outcome of sponsored search auctions or the POSITION obtained by firms, b) the QUALITY of the firms, and c) the PRODUCT_TYPE.

<<Insert TABLE 2 about here>>

The *dependent variable* of interest is the POSITION (rank) advertisers receive on paid search lists on Yahoo!'s and Google's sponsored search results. In concert with industry research studies that find that consumers typically do not search beyond the first page of search results (BBC News 2006), we restrict our focus to the top fifteen search listings for each keyword. Within each keyword category, advertising firms are first ordered by their average rank/position in the sponsored search listings over the period of our data collection (not including the days that they did not bid) and the top fifteen ranked firms are then selected to form a smaller subset. We use two measures to depict the POSITION achieved by advertisers. The continuous variable *AvgRank* is a measure of the average POSITION obtained by the advertising firm in the paid search listing over all the days it bid. *OrdRank* represents the POSITION measure in discrete form, i.e. firms are first ordered based on the rank averaged by them over the 60-day observation period, and then ranked from 1(top) to 15 (bottom).

The main *independent variable* of interest in the study is the QUALITY of the advertiser. As noted earlier, advertiser-quality is a multi-dimensional construct. To account for this, we use three different measures/proxies of advertiser QUALITY⁵, which are particularly relevant in online settings. The greater the number of page views a seller's website attracts, and higher the proportion of all web users that visit the website, the higher we may regard the seller's quality to be. Alexa refers to these two numbers as page view (fraction of all the page views by toolbar users that go to a particular site, per million) and page reach (percentage of all Internet users who visit a given website), respectively. A new variable, *TrafficRank*, is then computed by Alexa that collectively

⁵ We also use factor-scores derived from these, as alternative measures of advertiser-quality, the details of which are provided in Section 6.

represents both page view and reach. For *TrafficRank*, a lower value indicates higher quality and vice versa. It is important to note that search engines like Google and Yahoo! are just one of the several sources that drive traffic to a website, the others being links from other websites, word-of-mouth, referrals by friends, and random surfing, among others. Consequently, a seller's ranking (POSITION) on Google or Yahoo! sponsored listings is neither synonymous with, nor the primary determinant of its *TrafficRank* on Alexa. This is further confirmed by our tests for endogeneity and the robustness checks (for more details see Section 6). A second measure is provided by the number of incoming links to a website, *Inlinks*. Originally popularized by search engines such as Google and Yahoo!, links pointing to a website are now commonly used as a measure of quality. An incoming link is considered as a positive recommendation by the originator of that link, so the more the *Inlinks* for an advertiser the more important it is considered to be. The third variable, *Ratings*, is calculated by averaging over the scores provided by customers who visit sellers' websites and rate them on their purchase and shopping experiences. *Ratings* are measured on a scale from 1 to 5. The use of three distinct measures reinforces the robustness of our findings. It is pertinent to note here that a positive (negative) relationship between QUALITY and POSITION depicts the presence of higher quality firms in higher (lower) positions.

Search, *Experience*, and *Credence* are binary (dummy) variables that represent PRODUCT_TYPE increasing in pre-purchase quality uncertainty. Finally, we collect information on the *age of the firm* from Alexa. AGE is measured as the number of days the firm has existed online and serves as a control.

Methodology and Analysis

Our primary goal in this study is to examine how the relationship between seller QUALITY and POSITION varies across the different product categories and across the two sponsored search markets. We conduct a series of increasingly sophisticated analyses. We use an OLS model as a benchmark in our analysis.

Interaction terms are created using centered main effects variables, QUALITY and PRODUCT_TYPE to minimize multi-collinearity. The OLS model uses a continuous dependent variable, *AvgRank*. Our dependent

variable is however naturally ordered as it measures the outcome or position in sponsored listings, and we therefore also repeat the analysis for an ordered dependent variable, *OrdRank* using ordered probit regressions. Further, we control for the age of the firm in order to account for the possibility that newly established, not necessarily low quality, firms may have lower *TrafficRank*. After accounting for missing values from the top fifteen ranked firms for the 27 product keywords, our total sample for the analysis of Yahoo! data is 353 (for OLS) and 350 (for ordered probit) observations. The corresponding numbers for Google data are 274 and 272. We specify the following equations, 1a and 1b, with age normalized quality measures, *N_QUALITY*, for both Yahoo! and Google.

$$(1a) \text{AVGRANK} = \gamma_1 + \gamma_2 \text{N_QUALITY}_i + \gamma_3 \text{PRODUCT_TYPE}_i + \gamma_4 \text{N_QUALITY}_i * \text{PRODUCT_TYPE}_i + \varepsilon_i$$

$$(1b) \text{ORDRANK} = \delta_1 + \delta_2 \text{N_QUALITY}_i + \delta_3 \text{PRODUCT_TYPE}_i + \delta_4 \text{N_QUALITY}_i * \text{PRODUCT_TYPE}_i + \varepsilon_i$$

It is further possible that unobserved variables relating to each keyword affect the outcomes observed in the above analyses. While the above analyses assume that the observations on the independent variables are not systematically correlated with the error terms, the observations, and subsequently residuals, within each keyword may not be independent. We examine three additional models to deal with this structural complexity. In model 2, we use cluster robust standard errors in a standard analysis to account for the fact that the observations are clustered into keywords and that they may be correlated within, but would be independent across keywords (White 1980; Wooldridge 2002).

$$(2) \text{AVGRANK} = \gamma_1 + \gamma_2 \text{N_QUALITY}_i + \gamma_3 \text{PRODUCT_TYPE}_i + \gamma_4 \text{N_QUALITY}_i * \text{PRODUCT_TYPE}_i + \varepsilon_i$$

In model 3, we conduct a Least Squares Dummy variable (LSDV) regression by including dummies for keywords in the baseline model to capture and control for keyword-level effects.

$$(3) \text{AVGRANK} = \delta_1 + \delta_2 \text{N_QUALITY}_i + \delta_3 \text{PRODUCT_TYPE}_i + \delta_4 \text{N_QUALITY}_i * \text{PRODUCT_TYPE}_i + \delta_5 \text{KEYWORD1} + \dots + \delta_{31} \text{KEYWORD26} + \varepsilon_i$$

Finally, in model 4 we also estimate a fixed-effects (FE) regression model treating keyword as the grouping variable over all sellers. We conduct split sample analyses for each of the three product types, while controlling for the fixed effects of keywords.

$$(4) \text{AVGRANK} = \zeta_1 + \zeta_2 \text{N_QUALITY}_i + \zeta_3 \text{KEYWORD1} + \dots + \zeta_{10} \text{KEYWORD8} + \varepsilon_i$$

5. Results

The correlation results are reported in Table 3. Significant correlations exist among the three measures of QUALITY for Yahoo and Google respectively: *TrafficRank* and *Inlinks* (0.84 and 0.85), *TrafficRank* and *Ratings* (0.62 and 0.62), and finally, *Inlinks* and *Ratings* (0.60 and 0.57), all of which are significant at the $p < 0.001$ level. Among the three however, the most direct and comprehensive quality measure available is *TrafficRank*, and therefore we present regression results using it as our primary measure of seller quality. The results remain consistent across *Inlinks* and *Ratings*, as discussed in Section 6.

<<Insert TABLE 3 about here>>

The results from the various regression analyses are presented in Table 4.

<<Insert TABLE 4 about here>>

It should be noted here that since the left out category among product types is search, the coefficient of QUALITY in the regression equations represents the effect of one unit of change in QUALITY on the POSITION in listings for *Search* goods. The two interaction terms between QUALITY and *Experience* and QUALITY and *Credence* goods then are a measure of how much this association changes for *Experience* and *Credence* goods, *relative to* Search goods.

We first compare the corresponding baseline models across Yahoo! (Y1a, Y1b) and Google (G1a, G1b) depicted in Table 4. Across all the aforementioned models, we find that QUALITY is *positively correlated* with average POSITION obtained by the firm in the sponsored search listings for *Search* goods on both search mechanisms. More interestingly, we find that the coefficients of the interaction between PRODUCT_TYPE and

seller QUALITY are *negative* and significant in the case of *Experience* and *Credence* goods in all the four models for Yahoo!, but not for Google. These findings suggest that the relationship between seller QUALITY and POSITION achieved in listings are significantly different for *Experience* and *Credence* goods, as compared to *Search* goods for Yahoo!. In the case of Google, however, there appears to be no significant differences across the three product categories.

Results of our analyses of models 2 and 3 indicate that the results obtained above are robust. As shown in Table 4, the coefficient of QUALITY (for *Search* goods) is positive and significant across all models, while the coefficients of the interaction terms for *Credence* goods are negative and significant for Yahoo! in Y2 and Y3. The interaction coefficients for *Experience* goods are negative; but only significant in Y2 and not in Y3. In contrast, the corresponding interaction coefficients are not significant for Google in G2 and G3, reinforcing our earlier findings that there are no significant differences among *Search*, *Experience* and *Credence* goods in the QUALITY-POSITION relationship on Google's sponsored search outcomes.

While the above results focus on the *relative* differences in the relationship between QUALITY and POSITION in listings across the three product categories, it is also important to examine the *absolute* relationship between QUALITY and POSITION for each of the three product categories. We test for such a relationship across all three PRODUCT_TYPES using tests of linear combinations (as depicted in Table 5) for Yahoo! and Google, corresponding to the models in Table 4. These tests assess whether the absolute coefficient of quality on position in listings is different from zero across *Search*, *Experience*, and *Credence* categories.

<<Insert TABLE 5 about here>>

These results for Yahoo! in Table 5, which are mostly consistent across models 1a and 1b, suggest that the coefficient for QUALITY is positive and strongly significant for *Search* goods. On the hand, the coefficient for *Credence* goods is negative and significant in the linear regressions with *AvgRank*. The corresponding coefficients for *Experience* goods lie in between those of *Search* and *Credence*, and are not significantly

different from zero. These results are further reinforced in our analyses of Y2 and Y3. Table 5 suggests that different outcomes are evident on Google's sponsored listings. The coefficients for all three product types are positive and significant across all models G1-G3. The fixed effects models corresponding to equation 4 with split samples also provide consistent results. However, since the Hausman (1978) null is not rejected, the random effects model is a better choice for our setting⁶. These results are also shown graphically in Figure 1.

<<Insert FIGURE 1 about here>>

6. Robustness Checks

Tests for Endogeneity

The issue of endogeneity arises as an artifact of our particular measure of seller quality- *TrafficRank*. It is possible that *AvgRank*, the measure of the advertiser's POSITION in sponsored search listings may affect its *TrafficRank*, the measure of seller's QUALITY in our study. In other words, better positions in the search listings could increase traffic to the advertiser's site. However, our key findings of a negative correlation between QUALITY and POSITION for *Experience* and *Credence* goods in the case of Yahoo! imply that this relationship is in fact, reversed – i.e. *better* positions in the search listings are actually correlated with *lower* traffic. Thus, these results are likely to be strengthened in the absence of any potential endogeneity.

If our model suffers from recursive endogeneity (i.e. *AvgRank* affects *TrafficRank* and *TrafficRank* affects *AvgRank*), OLS would be insufficient. We therefore test for the endogeneity of *TrafficRank* using the Wu-Hausman F test (Wu 1973) and the Durbin-Wu-Hausman chi-sq test (Hausman 1978). These tests examine the null hypothesis that *TrafficRank* is exogenous by checking for a statistically significant difference between the OLS and 2SLS estimates of its beta coefficient when regressed on *AvgRank*. The OLS model is as specified earlier. In the 2SLS model, we use the variable AGE as the instrument for *TrafficRank*. Theoretically, the AGE

⁶ The Hausman test examines the hypothesis H0: the difference in coefficient across FE and RE is not systematic. In our data, it is not rejected, suggesting that there are no significant group or keyword effects in the model, and the criteria of efficiency suggests that we estimate a random effects model.

of a website would be correlated with its traffic (QUALITY), and therefore can be used to predict the latter. It is unlikely though, that AGE has a direct impact on the POSITION of the seller on sponsored listings.

From the 2SLS analyses, we find that the first-stage F is highly significant for both Google ($F(5,266)=51.34$, $p<0.01$) and Yahoo ($F(5,344)=39.62$, $p<0.01$); the F-statistic is much higher than the minimum value of 10 (Staiger and Stock 1997). The corresponding coefficients of AGE are also significant ($p<0.01$) indicating that AGE is both a valid and relevant instrument. However, since neither the Wu-Hausman F test (Wu 1973) or the Durbin-Wu-Hausman χ^2 test (Hausman 1978) is rejected, we fail to reject the null that *TrafficRank* is exogenous. Based on these analyses, we find that OLS is unbiased, consistent and the more efficient estimator for our model. Therefore we focus on the OLS estimates in our discussions. Further, the point estimates are qualitatively unaffected if we use 2SLS.

Additional Tests for Robustness

In this section, we address several sensitivity concerns that might arise from the measures and models used in our analyses, by conducting appropriate checks of robustness⁷. First, we examine the robustness of our findings by using several different measures of *TrafficRank* computed over different time periods. We then re-assess our models using several combinations of the three quality measures described earlier. We also control for the presence of different types of sellers in the sponsored search listings, and finally, we analyze the sensitivity of our results to subsets of keywords in each product category.

Alexa provides for each seller *TrafficRank* measures computed over 3-months, 1-month, 1-week and 1-day, along with an instantaneous measure, which is the one we used in the analyses presented above. We repeat our analyses using these measures of *TrafficRank* collected over different time periods, and find that our results are robust. The coefficient for the QUALITY for *Search* goods on POSITION is directionally consistent in all cases. *Experience* interaction coefficients for Yahoo across the 3-months, 1-month, 1-week and 1-day models are all negative and significant ($p<0.10$ or better); *Credence* interaction coefficients for Yahoo! are also negative

⁷ The results from these additional analyses are available from the authors upon request.

and significant ($p < 0.01$ or better) respectively. The corresponding interaction coefficients for Google are all insignificant. This suggests that our findings of a negative correlation in the case of *Experience* and *Credence* goods for Yahoo! are robust to changes in the *TrafficRank* measure. We also obtain consistent results when we test for clustering and keyword effects using the 3-month *TrafficRank* measure.

Next, we consider the two other available measures of seller quality - *Inlinks*, and *Ratings*. We first use *Inlinks* and *Ratings* as separate measures of QUALITY. We then create three different factor scores- a) consisting of *TrafficRank*, *Inlinks* and *Ratings*, b) consisting of *TrafficRank* and *Inlinks*, and c) consisting of *TrafficRank* and *Ratings*⁸. Across all the five models, the coefficient of QUALITY on POSITION is significant for *Search* goods for Yahoo ($p < 0.10$ or better) and for Google ($p < 0.05$ or better). Moreover, the QUALITY coefficients for *Credence* goods is significant and directionally opposite to *Search* ($p < 0.01$ or better); while the corresponding QUALITY coefficients for *Experience* goods are directionally consistent, but mostly insignificant in the case of Yahoo!. Both interaction coefficients are not significantly different from that for *Search* goods in the case of Google. These results reinforce the validity of our original findings.

Another check we conduct ensures that our results are not affected by the presence of different types of sellers such as retailers/manufacturers and infomediaries, in the sponsored search listings for a given keyword. Despite the presence of a mix of seller types, it is important to note that our results are based on an ordinal ranking of the sellers in the sponsored search listings; this ordinal ranking of sellers is still preserved with each type, and our results would therefore still hold. However, to test this, we include dummies to control for the presence of a mix of seller types. We find that the dummy for infomediaries is insignificant for Yahoo, but significant for Google ($p < 0.05$). After controlling for their effects, we observe that the main QUALITY coefficient for *Search* goods continues to be significant for Google ($p < 0.01$) but is interestingly insignificant for Yahoo. Further, while the product type interactions are insignificant for Google, both the *Experience* and *Credence*

⁸ The items load very well on the factor, with reliabilities well over 0.70, suggesting that measures reflect an underlying common quality factor

interaction coefficients are still directionally consistent for Yahoo, but only the latter coefficient is significant ($p < 0.01$). This suggests that in the worst case, the QUALITY-POSITION relationship may be reversed even for Search goods in Yahoo!'s listings.

Finally, we examine the sensitiveness of our models to our choice of keywords. We reran our analyses presented in Table 4 (with 9 keywords each) using several combinations of 7 and 5 keywords each across Search, Experience and Credence goods. We obtain consistent results that reinforce the differences in the relationships for Experience and Credence goods across Yahoo! and Google. Models YR1& YR2, and GR1 & GR2, in Tables 4 and 5 provide the results for one such analysis for Yahoo! and Google respectively.

7. Post-Hoc Analysis

Quantile Regression Analysis

We are further interested in examining if and how the relationship between seller QUALITY and POSITION in the search listings differs *across* POSITIONS, or levels of the dependent variable. We do so by employing *quantile regression* analysis (Koenker and Bassett 1978). While OLS regression estimates the regression coefficient at the conditional mean of the regressor's distribution, quantile regression can provide parameter estimates at *different quantiles of the dependent variable*. This enables us to examine the variation in the effect of independent variables on the dependent variable at different quantiles (Buchinsky 1998; Koenker and Hallock 2001). Thus, quantile regression allows for possible heterogeneity that may be present in the QUALITY-POSITION relationship across different POSITIONS in the sponsored search listings.

Our findings show that the results from the linear regression hold at almost all the position ranges (i.e. from position 1 to position 15). We display the quantile graphs in Figure 2 for the two extreme cases for Yahoo! – *Search* (Fig. 2a) and *Credence* goods (Fig. 2b), and for *Credence* goods for Google (Fig. 2c), since the other two are very similar. The coefficients of Quality-Position are positive (the graph in 2a lies above 0) across all quantiles for search goods for Yahoo!, and the interaction coefficient (depicting the difference between *Search*

and *Credence* goods) is consistently negative (the graph in 2b lies below 0) across all quantiles. The advantage of quantile regression here lies in indicating how the strength of the negative correlation between QUALITY and POSITION for *Credence* goods changes across quantiles of sponsored search listings. As shown in Figure 2, the negative correlation appears to be the strongest in the top positions of sponsored listings, reinforcing our earlier findings. Figure 3 illustrates the inverse cumulative density function (CDF) for representative cases. These graphs depict the relationship between QUALITY and POSITION across different quality and quantiles of the position variable. As is evident from these figures, higher quality firms appear in top positions for *Search* goods (Fig. 3a), but not for *Credence* goods in Yahoo! (Fig. 3b); whereas the relationship holds for *Credence* goods in Google (Fig. 3c).

<<Insert FIGURE 2 about here>>

<<Insert FIGURE 3 about here>>

Non-linearities in the Quality-Position Relationship

Our analysis thus far assumed a linear relationship between advertiser's QUALITY and POSITION in the search results. In this section we test for the presence of non-linear relationship between the seller QUALITY and POSITION in the search listings. We conduct separate regressions for each product category (i.e. *Search*, *Experience*, and *Credence*) for both Yahoo! and Google, as specified below in 6a and 6b.

$$(5a) \text{ AVGRANK} = \alpha_1 + \alpha_2 \text{N_QUALITY}_i + \alpha_3 (\text{N_QUALITY}_i)^2 + \varepsilon_i$$

$$(5b) \text{ ORDRANK} = \beta_1 + \beta_2 \text{N_QUALITY}_i + \beta_3 (\text{N_QUALITY}_i)^2 + \varepsilon_i$$

The results of this analysis are depicted in Figure 4a for Yahoo! and Figure 4b for Google. As indicated by the graphs in Figures 4a and 4b, the quadratic terms are not significant for the two extreme product categories - *Search* and *Credence* goods, but the relationship is markedly non-linear for *Experience* goods. The coefficients for the squared quality variable in the model are not significant for search and credence goods, in both Google and Yahoo!. However, squared term for *Experience* goods is significant in both Yahoo! and Google, suggesting

a non-linear relationship. The coefficients of the squared QUALITY term in the equation 5a are 0.69 ($p < 0.10$) and 0.47 ($p < 0.05$) for Yahoo! and Google respectively. In equation 5b, the coefficients of squared QUALITY term are 0.21 ($p < 0.05$) for both Yahoo! and Google.

<<Insert FIGURE 4 about here>>

In the case of *Experience* goods we find that neither the low quality firms, nor the high quality firms bid the highest or appear on the top of the search listings. It is the medium quality firms that acquire the top rankings in the search listings of Yahoo! as well as Google. To explain this interesting difference in bidding patterns across the three categories it is useful to consider the economics of bidding for customer clickthroughs given the differences in cost structures of high, medium, and low quality firms.

The expected revenue for a firm bidding for keywords in a sponsored search context is a function of two factors – the probability of attracting clickthroughs and the conversion ratio. While the probability of attracting clickthroughs is a function of the advertiser's bid (the higher the bid, the higher the probability of attracting clickthroughs, irrespective of quality), the conversion ratio is dependent on the type of good (i.e. consumer's ability to distinguish between high, medium, and low quality providers). In the case of *Search goods*, since consumers have lower uncertainty about quality, low quality sellers will find it difficult to improve their conversion ratio through their bidding strategies and consequently will have less of an incentive to bid higher than higher quality firms. In the case of *Credence* goods, the conversion ratio is less related to quality of the sellers as consumers find it difficult to distinguish between low, medium, and high quality firms. Given that low quality firms have a lower cost of production compared to higher quality firms, lower quality firms can outbid higher quality firms and improve customer acquisition. Sellers of low-quality *Credence* goods thus, not only have the incentive but also the ability to bid higher and gain sufficient clickthroughs by capitalizing on consumers' inability to easily distinguish between high quality and lower quality sellers.

While conversion ratio is closely related to quality in the case of *Search* goods, and unrelated to quality in the case of *Credence* goods, the case for *Experience* goods lies in between these two extremes. While there

still exist information-asymmetries between buyers and sellers of *Experience* goods, the resulting quality uncertainty is not as high as in the case of *Credence* goods. Low-quality sellers, despite their lower costs would therefore not have a sufficiently high incentive to bid high in these markets. Medium quality firms though, stand to benefit more, given the significant (although lower than in the case of *Credence* goods) uncertainties surround the quality of their offerings in comparison to higher quality firms. Interestingly we find that this effect is consistent across Yahoo! and Google's results.

8. Discussion and Implications

Our results indicate that the sponsored search markets suffer from adverse selection; however the intensity of adverse selection differs across markets as well as product characteristics. While adverse selection was almost non-existent in the market for *Search* goods, the unregulated sponsored search mechanism used by Yahoo! suffered from problems of adverse selection for *Experience* and *Credence* goods. However, Google's intervention mechanisms of ranking bidder advertisements (by moderating the advertiser's willingness to pay with its performance measured by clickthrough rates) seem to be capable of circumventing the problem of adverse selection for *Search* as well as *Credence* goods. While adverse selection issues in the case of *Experience* goods are not as severe as in the case of *Credence* goods, Google's intervention mechanism does not seem to alleviate this issue.

This could adversely impact consumer welfare particularly for uninformed consumers and consumers who trust the search results provided by these search engines. This risk has been identified by consumer advocacy groups such as Consumer Reports WebWatch, as is evident from the following excerpt from their report (Marable 2003):

"... trust in search engines may make them (online consumers) vulnerable while online, as they are largely unaware such navigation sites often accept fees in exchange for giving advertiser Web pages prominent placement on their search results pages."

While online markets can improve consumer welfare by lowering search costs, in the presence of uncertainty regarding unobservable quality characteristics, sellers can distort or hide information, leading to adverse selection. Thus, it is possible that the higher costs of adverse selection counteract the benefits gained from lowered search costs for consumers (Fabel and Lehmann 2002). The sponsored search market provides an excellent test bed to examine issues of adverse selection. As advertising channels, online search mechanisms such as Yahoo! and Google lack traditional differentiators of firm quality, forcing consumers to seek out alternate sources of information, such as ranking of advertisements on paid listings. Our study is among the first to examine the performance outcomes of these important and powerful online advertising mechanisms.

Our findings also add to existing work that examines the efficacy of online markets for different product categories. In particular, the comparison of the two most popular sponsored search mechanisms allows us to illustrate their differential effectiveness in abating adverse selection. Our study also contributes to the literature on advertising by testing traditional theories in emerging channels. Just as eBay resorts to user-feedbacks and Amazon to reviews and ratings to alleviate adverse selection problems, our findings suggest that online sponsored search mechanisms may be able to decrease the negative impacts of adverse selection in markets with high pre-purchase uncertainty by providing alternate signals of quality about advertisers.

Our findings are particularly relevant for the providers of search services. Sponsored search listings that are biased can not only reduce consumer welfare, but also drive out higher quality firms, and eventually, reduce the profitability of the intermediary as well. Search intermediaries would do well to provide better information regarding their paid search mechanism and to incorporate reputation mechanisms to aid consumers in their decision-making for online purchases. Our findings indicate that additional signals of quality would serve to improve the efficiency and welfare properties of the sponsored search markets by reducing adverse selection. Provision of such additional quality information such as ratings and reviews from Bizrate.com and Epinions.com alongside the search listings can help reduce the risk faced by the consumer and improve consumer welfare.

9. Conclusions and Future Research

It is appropriate to discuss some of the limitations of our study. First, adverse selection in markets associated with unobservable seller characteristics such as quality has typically been known to be difficult to measure empirically. This is because reliable quality signals are hard to come by. Alexa is the only publicly available source for website *TrafficRank*. *TrafficRank* is calculated by aggregating the traffic generated by the site from among a subset of online users. Our measure of quality is therefore only accurate to the extent that this segment of consumers is representative of the broader online population. The use of two other measures of seller quality, namely the number of incoming links and ratings provided by a subset of online consumers provides us with a way to triangulate our findings. More importantly, our focus is on the *relative* performance of sponsored search and our key results highlight interesting differences across the three different product categories as well as across the two sponsored search mechanisms using the same measures of QUALITY and POSITION.

Second, traffic generated by the website is usually a function of how long the website has been in existence. The use of web-site traffic as a measure of quality would be problematic and confound the results of our study, if these new entrants and niche marketers were high quality sellers. We address this issue by normalizing quality by the age of seller, or the length of its existence online. Further, we find that all three measures of quality are highly and significantly correlated. While traffic rank and incoming links could be affected by the newness or niche focus of the website, the third measure – ratings, is less likely to be influenced by age. These measures therefore help reduce any potential confounding effects. Third, while the distinction among *Search*, *Experience*, and *Credence* goods is well established in theoretical literature, in reality all goods have search, experience and credence attributes, albeit to varying extents. While our classification largely conforms to prior literature, what is more important is that these products are characterized by increasing uncertainty for consumers purchasing them.

Our study opens doors for plenty of future research opportunities. Sponsored search auctions for keywords, though growing rapidly, are still in their infancy. Despite the nascence of sponsored search/advertising mechanisms, there exist significant differences between traditional advertising formats and sponsored search formats. Of particular interest is the fact that sponsored search advertising is a *performance-based* advertising model where firms incur an expense only when consumers click on the links to their websites. Thus firms' advertising expenses are closely linked to their revenues from potential sales to online consumers. In comparison, advertising in traditional print/broadcast media is characterized by fixed costs and further removed from any potential sales. Thus the two advertising formats (traditional vs. sponsored search) differ in their risk to advertisers. Future research should examine the implications of these differences on the incentives for (low vs. high quality) firms to advertise. The keywords in our study were chosen based on their popularity (as published on publicly available sources). It would be interesting to examine whether our results hold for "niche" or less popular keywords with fewer bidders. It would also be useful to examine the bidding patterns for keyword combinations as well as for keywords representing brands (such as "Sony Vaio" or "Dell Inspiron") rather than generic products.

Our study demonstrates the usefulness of SEC framework in studying the impact of different market mechanisms on the market effectiveness. Future research can employ the same SEC framework to study the dynamics of bidding strategies of the firms over a period of time. A more extensive analysis of such bidding dynamics promises to shed light on relative competition across different product categories as well as the existence of strategic groups within product categories.

Further, online seller quality is multidimensional, and possible extensions to our work may consider the impacts of several other alternate dimensions of quality such as website quality, seller trustworthiness, etc. Future studies should also examine consumer behavior in response to the sponsored search phenomena. Laboratory studies designed to analyze the differential search strategies adopted by consumers would help

understand how consumer search across different search formats. Studies of this nature are sparse, given the novelty of the phenomenon. Whether findings of studies relating to consumer behavior in traditional channels translate well to online settings is an empirical question yet to be answered.

Finally, as argued by Choi, Stahl, and Whinston (1997, p. 169), the future of the online markets “may depend on how non-technological but fundamentally economic issues as the lemons problem are solved”. As sponsored search mechanisms become more established and as the number of informed consumers using these markets increase, the markets with adverse selection will either breakdown or correct themselves due to the changes in the consumer, advertiser and/or search intermediaries behavior. It would be useful to examine how these sponsored search markets evolve over time.

TABLE 1
Classification of Goods and Services as per SEC Framework

Search	Experience	Credence
Apparel	Auto Insurance	Cosmetic Surgery
Books	Brokerages	Counseling
CD	Cruises	Home Security Systems
Cell Phones	Event Planning	Pest Control
Flight Tickets	Healthcare	Psychics
Laptops	Jewelers	Tax Services
Refrigerators	Martial Arts	Therapy
Television	Moving And Storage	Used Cars
Toys	Perfumes	Vacation

TABLE 2
Definition of Variables Used in Empirical Analysis

Type	Measure	Definition	Scale
POSITION	AVGRANK%	The sample average of the ranks obtained by the seller during the 60 day period	Continuous
	ORDRANK%	The position or ordered rank of the seller based on average rank	Discrete (1- 15)
PRODUCT_ TYPE	SEARCH EXPERIENCE CREDENCE	Dummies for search, experience and credence goods	Binary
QUALITY of advertiser or seller	TRAFFICRANK\$	A combination of historical page view rank (per million) and page reach (% of all Internet users).	Continuous Log Transformed
	INLINKS	Number of links pointing into seller's website.	Continuous Log Transformed
	RATING	A user generated score for the seller.	1-5 Log Transformed
Interactions	SEARCH* QUALITY	The product of SEARCH (mean-centered) and QUALITY (mean-centered)	
	EXPERIENCE* QUALITY	The product of EXPERIENCE (mean-centered) and QUALITY (mean-centered)	
	CREDENCE* QUALITY	The product of CREDENCE (mean-centered) and QUALITY (mean-centered)	
Control	AGE	The age of the Internet firm, as measured by the number of days since it was first established on the Internet	Continuous

% Lower Rank indicates top Position in the listings and vice versa

\$ Lower TrafficRank indicates higher Quality and vice versa

TABLE 3
Correlation Matrix for Yahoo! And Google

Yahoo! Correlations (pairwise N)							
	A	B	C	D	E	F	G
A. Rank	1						
	0.10						
C. INLINKS ^{\$}	(0.05)	1					
	0.09	0.60					
D. RATING [#]	(0.06)	(0.00)	1				
	0.10	0.84	0.62				
E. TRAFFICRANK ^{\$}	(0.07)	(0.00)	(0.00)	1			
	0.01	0.24	0.16	0.23			
F. Search	(0.89)	(0.00)	(0.00)	(0.00)	1		
	-0.02	-0.11	-0.06	-0.13	-0.50		
G. Experience	(0.64)	(0.04)	(0.20)	(0.02)	(0.00)	1	
	0.01	-0.14	-0.10	-0.11	-0.50	-0.50	
H. Credence	(0.75)	(0.01)	(0.05)	(0.04)	(0.00)	(0.00)	1
Google Sponsored Correlations (pairwise N)							
	A	B	C	D	E	F	G
A. Rank	1						
	0.10						
B. INLINKS ^{\$}	(0.10)	1					
	0.15	0.57					
C. RATING [#]	(0.01)	(0.00)	1				
	0.16	0.85	0.62				
D. TRAFFICRANK ^{\$}	(0.01)	(0.00)	(0.00)	1			
	0.05	0.43	0.27	0.52			
E. Search	(0.39)	(0.00)	(0.00)	(0.00)	1		
	-0.03	-0.25	-0.07	-0.28	-0.48		
F. Experience	(0.65)	(0.00)	(0.20)	(0.00)	(0.00)	1	
	-0.02	-0.18	-0.18	-0.23	-0.47	-0.55	
G. Credence	(0.72)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	1

Parentheses contain p values

[#] log transformed

^{\$} normalized by age of the firm, followed by a log transformation

TABLE 4
Regression Analyses

Model	Yahoo						Google					
	Y1a: Position (Cont)	Y1b: Position (Ord)	Y2: Position (Cont)	Y3: Position (Cont)	YR1: Position (Cont)	YR2: Position (Cont)	G1a: Position (Cont)	G1b: Position (Ord)	G2: Position (Cont)	G3: Position (Cont)	GR1: Position (Cont)	GR2: Position (Cont)
Constant	8.60 (0.23) ***		8.57 (0.20) ***	8.22 (0.71) ***	8.41 (0.26) ***	8.60 (0.32) ***	4.73 (0.15) ***		4.74 (0.13) ***	4.01 (0.55) ***	4.53 (0.18) ***	4.70 (0.20) ***
N_QUALITY	0.47 (0.24) *	0.14 (0.06) **	0.47 (0.21) **	0.62 (0.28) **	0.50 (0.26) *	0.62 (0.34) *	0.40 (0.15) ***	0.21 (0.07) ***	0.40 (0.15) **	0.48 (0.18) ***	0.37 (0.18) **	0.46 (0.21)**
EXPERIENCE	-0.25 (0.26)	-0.09 (0.06)	-0.25 (0.25)	-1.33 (0.79) *	-0.31 (0.29)	-0.41 (0.37)	-0.05 (0.17)	0.04 (0.08)	-0.05 (0.18)	0.74 (0.60)	-0.01 (0.19)	-0.06 (0.23)
CREDENCE	-0.16 (0.27)	-0.05 (0.06)	-0.16 (0.23)	-0.06 (0.81)	-0.30 (0.30)	-0.39 (0.36)	-0.18 (0.18)	-0.06 (0.09)	-0.18 (0.16)	0.15 (0.68)	-0.19 (0.23)	-0.24 (0.23)
QUALITY * EXPERIENCE	-0.50 (0.29) *	-0.13 (0.07) *	-0.51 (0.29) *	-0.48 (0.33)	-0.64 (0.30) **	-0.71 (0.41) *	0.12 (0.17)	0.06 (0.08)	0.12 (0.19)	0.19 (0.20)	0.19 (0.19)	0.29 (0.25)
QUALITY * CREDENCE	-1.30 (0.27) ***	-0.28 (0.07) ***	-1.29 (0.17) ***	-1.24 (0.29) ***	-1.29 (0.29) ***	-1.43 (0.36) ***	0.05 (0.17)	0.06 (0.08)	0.05 (0.18)	0.14 (0.20)	0.22 (0.22)	0.23 (0.23)
N	350	350	350	350	266	197	272	272	272	272	200	153
F or χ^2 (5)	5.53***	23.73 ***	13.89 ***	1.56**	4.81***	3.70***	1.89 *	13.18 **	2.74 **	0.75	1.12	1.53
\$R ²	0.075	0.013	0.075	0.124	0.085	0.088	0.034	0.009	0.034	0.082	0.028	0.050
Adj. R ²	0.061			0.045	0.067	0.064	0.016			0.028	0.003	0.017

* indicate significance at 0.1 level, ** at 0.05 level and *** at 0.01 level; parentheses contain standard errors

\$ Pseudo-R² and χ^2 values are reported for ordered probit analyses using ORDRANK

R1 and R2 are robustness models which include a set of 7 (excluding used cars, vacation, flight tickets, television, auto insurance, healthcare) and 5 (excluding cd, apparel, flight tickets, television, jewelers, event planning, perfumes, brokerages, counseling, pest control, psychics, used cars) keywords respectively

TABLE 5
Test of the Absolute Quality Coefficient Across Product Types

Type	Yahoo!						Google					
	Y1a: Position (Cont)	Y1b: Position (Ord)	Y2: Position (Cont)	Y3: Position (Cont)	YR1: Position (Cont)	YR2: Position (Cont)	G1a: Position (Cont)	G1b: Position (Ord)	G2: Position (Cont)	G3: Position (Cont)	GR1: Position (Cont)	GR2: Position (Cont)
Search	0.44 (0.24) *	0.14 (0.06) **	0.47 (0.21) **	0.62 (0.28) **	0.471 (0.245) *	0.616 (0.340) *	0.40 (0.15) ***	0.21 (0.07) ***	0.40 (0.15) **	0.48 (0.18) ***	0.37 (0.18) **	0.46 (0.21) **
Experience	-0.05 (0.40)	0.01 (0.10)	-0.03 (0.43)	0.14 (0.48)	-0.138 (0.418)	-0.090 (0.567)	0.53 (0.23) **	0.27 (0.11) **	0.53 (0.20) **	0.67 (0.28) **	0.56 (0.26) **	0.75 (0.34) **
Credence	-0.86 (0.37) **	-0.14 (0.09)	-0.82 (0.24) ***	-0.62 (0.41)	-0.793 (0.403) **	-0.810 (0.487) *	0.46 (0.22) **	0.27 (0.11) **	0.46 (0.19) **	0.62 (0.27) **	0.59 (0.32) *	0.68 (0.31) **
F-test for equality of interaction coefficients	F(3,344) = 9.21 ***	$\chi^2(3) =$ 23.63 ***	F(3,26)= 20.59 ***	F(3,320) =8.09 ***	F(3,260) = 7.73***	F(3,191) = 6.06***	F(3,266) = 2.66 **	$\chi^2(3) =$ 9.22 **	F(3,26)= 3.23 **	F(3,242) =2.60 *	F(3,194) =1.73	F(3,147) = 1.98

* indicate significance at 0.1 level, ** at 0.05 level and *** at 0.01 level; parentheses contain standard errors

\$ Pseudo-R² and χ^2 values are reported for ordered probit analyses using ORDRANK

FIGURE 1
Relationship Between Quality and Position Across SEC Goods
Figure 1a: Yahoo




Figure 1b: Google




FIGURE 2

Quantile Regression Coefficients
Figure 2a: Yahoo! Search Goods
Figure 2b: Yahoo! Credence Goods

Figure 2c: Google Credence Goods



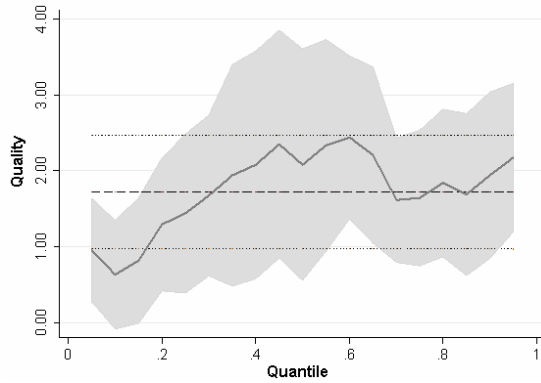



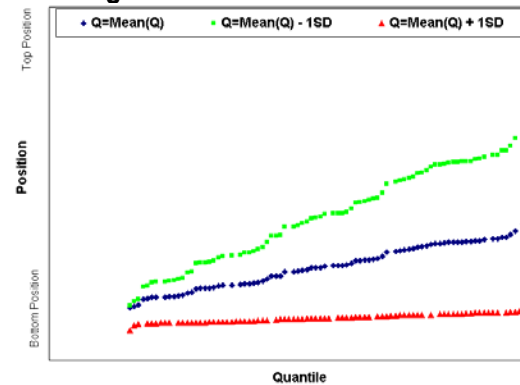
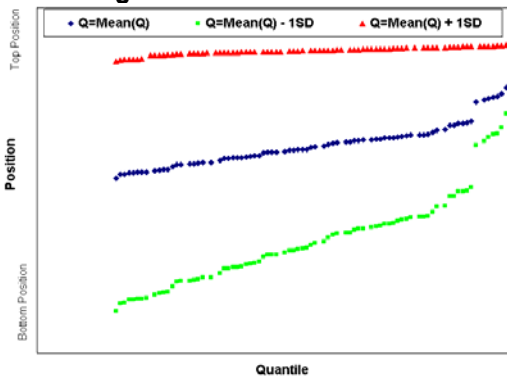
FIGURE 3

Quality Inverse CDF

Fig. 3a: Yahoo! Search Goods

Fig. 3b: Yahoo! Credence Goods

Fig. 3c: Google Credence Goods



Note: In Figures 3a-3c, we plot quality CDFs at 3 levels of seller quality– mean, mean +1 standard deviation (SD), and mean -1 standard deviation. When the CDF is shifted upwards for quality greater than mean, and subsequently shifted down for quality less than mean, it is indicative of greater probability of finding higher quality firms in top positions in sponsored search listings. This is the case for *Search* goods on Yahoo!, as seen in 3a. On the other hand in 3b, an opposite result is observed for *Credence* goods on Yahoo!, where the CDF for quality less than mean is shifted upwards, higher than the quality greater than mean. This is consistent with the linear regression results and suggests that the negative relationship between quality and position holds at every rank within our sample. Finally, the graph in 3c represents *Credence* goods on Google, where we obtain results similar to *Search* Goods on Yahoo! that adverse selection is averted.

FIGURE 4
Quadratic Relationship Between Quality and Position Across SEC Goods
Figure 4a: Yahoo


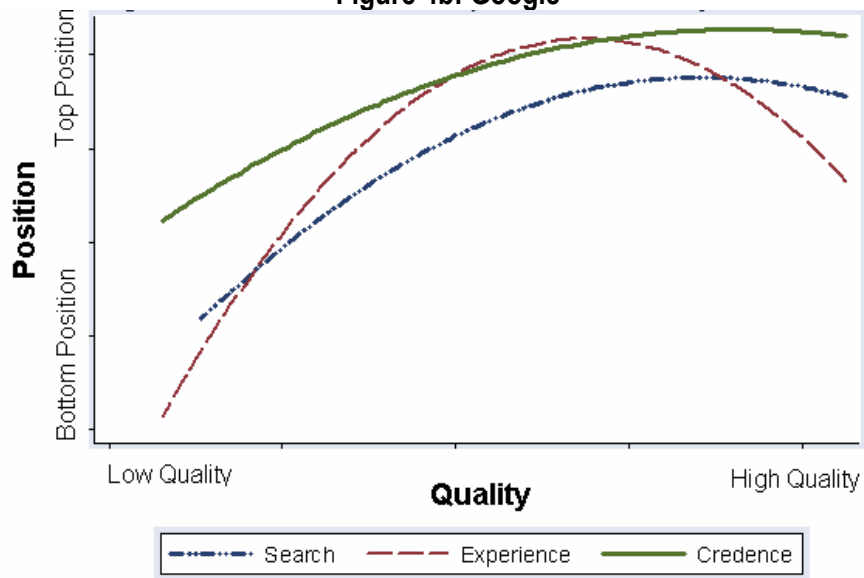


Figure 4b: Google



* Note: High Quality represent lower TrafficRank and Top Position represents lower AvgRank

REFERENCES

- Akerlof, George A. (1970), "The Market for Lemons: Quality Uncertainty and the Market Mechanism," *Quarterly Journal of Economics*, 84 (3), 488-500.
- Ausubel, Lawrence M. (1999), "Adverse Selection in the Credit Card Market," Working Paper, University of Maryland.
- Bakos, Yannis, Jr. Henry C. Lucas, Wonseok Oh, Gary Simon, Siva Viswanathan, and Bruce W. Weber (2005), "The Impact of E-Commerce on Competition in the Retail Brokerage Industry," *Information Systems Research*, 16 (4), 352-71.
- BBC News (2006), "Search Users 'Stop at Page Three'," April 12, 2006.
<http://www.webcitation.org/5FmwvAX3h>
- Brown, Jennifer and John Morgan (2006), "Reputation in Online Markets: Some Negative Feedback," Working Paper, University of California, Berkley.
- Brynjolfsson, Erik and Michael D. Smith (2000), "Frictionless Commerce? A Comparison of Internet and Conventional Retailers," *Management Science*, 46 (4), 563-85.
- Buchinsky, Moshe (1998), "Recent Advances in Quantile Regression Models: A Practical Guide for Empirical Research," *Journal of Human Resources*, 33, 88-126.
- Chatterjee, Patrali, Donna L Hoffman, and Thomas P Novak (2003), "Modeling the Clickstream: Implications for Web-Based Advertising Efforts," *Marketing Science*, 22 (4), 520-41.
- Choi, Soon-Yong, Dale O. Stahl, and Andrew B. Whinston (1997), *The Economics of Electronic Commerce*. Indianapolis: McMillan Technical Publishing.
- Comanor, William S. and Thomas A. Wilson (1979), "Advertising and Competition: A Survey," *Journal of Economic Literature*, 17 (2), 453-76.
- Darby, Michael R. and Edi Karni (1973), "Free Competition and the Optimal Amount of Fraud," *Journal of Law and Economics*, 16 (1), 67-88.
- Dewan, Sanjeev and Vernon Hsu (2004), "Adverse Selection in Electronic Markets: Evidence from Online Stamp Auctions," *Journal of Industrial Economics*, 52 (4), 497-516.
- DoubleClick (2006), "Search Trend Report Q4 2005," <http://www.webcitation.org/5Fmwx1gAZ>
- Ekelund, Robert B., Franklin G. Mixon, and Rand W. Ressler (1995), "Advertising and Information: An Empirical Study of Search, Experience and Credence Goods," *Journal of Economic Studies*, 22 (2), 33-43.

- Fabel, Oliver and Erik E. Lehmann (2002), "Adverse Selection and Market Substitution by Electronic Trade," *International Journal of the Economics of Business*, 9 (2), 175-93.
- Feng, Juan, Hemant K. Bhargava, and David M. Pennock (forthcoming), "Implementing Sponsored Search in Web Search Engines: Computational Evaluation of Alternative Mechanisms," *INFORMS Journal on Computing*.
- Ford, Gary T., Darlene B. Smith, and John L. Swasy (1990), "Consumer Skepticism of Advertising Claims: Testing Hypotheses from Economics of Information," *Journal of Consumer Research*, 16 (4), 433-41.
- Garicano, Luis and Steven N. Kaplan (2001), "The Effects of Business-to-Business E-Commerce on Transaction Costs," *Journal of Industrial Economics*, 49 (4), 463-85.
- Gaudeul, Alexandre (2004), "Internet Intermediaries' Editorial Content Quality," Working Paper, Industrial Organization, Economics Working Paper Archive at WUSTL.
- Genesove, David (1993), "Adverse Selection in the Wholesale Used Car Market," *Journal of Political Economy*, 101 (4), 644-65.
- Hausman, J. (1978), "Specification Tests in Econometrics," *Econometrica*, 46 (6), 1251-71.
- iProspect (2006), "Search Engine User Behavior," <http://www.webcitation.org/5FmxuJBfl>
- Jin, Ginger Z. and Andrew Kato (2004), "Consumer Frauds and the Uninformed: Evidence from an Online Field Experiment," Working Paper, University of Maryland.
- Kihlstrom, Richard E. and Michael H. Riordan (1984), "Advertising as a Signal," *Journal of Political Economy*, 92 (3), 427-50.
- Koenker, Roger and Gilbert Bassett (1978), "Regression Quantiles," *Econometrica*, 46 (1), 33-50.
- Koenker, Roger and Kevin F. Hallock (2001), "Quantile Regression," *Journal of Economic Perspectives*, 15 (4), 143-56.
- Linnemer, Laurent (2002), "Price and Advertising as Signals of Quality When Some Consumers Are Informed," *International Journal of Industrial Organization*, 20 (7), 931-47.
- Marable, Leslie (2003), "Consumer Reaction to Learning the Truth About How Search Engines Work," *Consumer Reports WebWatch*.
- Markoff, John and Nat Ives (2005), "Web Search Sites See Clicks Add up to Big Ad Dollars," *New York Times*, February 4, 2005. <http://www.webcitation.org/5FmwsRaq9>
- Moorthy, Sridhar and Hao Zhao (2000), "Advertising Spending and Perceived Quality," *Marketing Letters*, 11 (3), 221-33.
- Nelson, Phillip (1974), "Advertising as Information," *Journal of Political Economy*, 82 (4), 729-54.

- (1970), "Information and Consumer Behaviour," *Journal of Political Economy*, 78 (2), 311-29.
- Palmer, Jonathan W. (2002), "Web Site Usability, Design, and Performance Metrics," *Information Systems Research*, 13 (2), 150-67.
- Palmer, Jonathan W., Joseph P. Bailey, and Samer Faraj (2000), "The Role of Intermediaries in the Development of Trust on the Www: The Use and Prominence of Trusted Third Parties and Privacy Statements," *Journal of Computer-Mediated Communication*, 5 (3).
- Resnick, Paul, Richard Zeckhauser, Eric Friedman, and Ko Kuwabara (2000), "Reputation Systems," *Communications of the ACM*, 43 (12), 45-48.
- Resnick, Paul, Richard Zeckhauser, John Swanson, and Kate Lockwood (forthcoming), "The Value of Reputation on Ebay: A Controlled Experiment," *Experimental Economics*.
- Schlosser, Ann E., Sharon Shavitt, and Alaina Kanfer (1999), "Survey of Internet Users' Attitudes toward Internet Advertising," *Journal of Interactive Marketing*, 13 (3), 34-54.
- Schmalensee, Richard (1978), "A Model of Advertising and Product Quality," *Journal of Political Economy*, 86 (3), 485-503.
- Shamdasani, Prem N., Andrea J. S. Stanaland, and Juliana Tan (2001), "Location, Location, Location: Insights for Advertising Placement on the Web," *Journal of Advertising Research*, 41 (4), 7-22.
- Sherman, Chris (2005), "A New F-Word for Google Search Results," *Search Engine Watch*, March 8, 2005. <http://www.webcitation.org/5FmwyPgDv>
- Staiger, Douglas and James H. Stock (1997), "Instrumental Variables Regression with Weak Instruments," *Econometrica*, 65 (3), 557-86.
- White, Halbert (1980), "A Heteroskedasticity-Consistent Covariance Matrix Estimator and a Direct Test for Heteroskedasticity," *Econometrica*, 48, 817-38.
- Wooldridge, Jeffrey M. (2002), *Econometric Analysis of Cross Section and Panel Data*. Cambridge, MA: The MIT Press.
- Wu, De-Min (1973), "Alternative Tests of Independence between Stochastic Regressors and Disturbances," *Econometrica*, 41 (4), 733-50.