Information Systems Research
Vol. 21, No. 4, December 2010, pp. 760–772
ISSN 1047-7047 | EISSN 1526-5536 | 10 | 2104 | 0760

DOI 10.1287/isre.1100.0315

Research Commentary

Sponsored Search and Market Efficiency

Vasant Dhar, Anindya Ghose
Stern School of Business, New York University, New York, New York 10012
{vdhar@stern.nyu.edu, aghose@stern.nyu.edu}

Sponsored search is the mechanism whereby advertisers pay a fee to Internet search engines to be displayed alongside organic (nonsponsored) web search results. Based on prior literature, we draw an analogy between these markets and financial markets. We use the analogy as well as the key differences to present a theoretical framework consisting of a set of research questions about the pricing of keywords and design choices available to firms in sponsored search markets. These questions define an agenda for future research in sponsored search markets. They also have practical implications for advertisers and online marketplaces such as search engines and social media sites that support advertising.

Key words: social media; social commerce; sponsored search; financial markets; user-generated content; market efficiency
History: Vallabh Sambamurthy, Senior Editor. This paper was received on August 9, 2009, and was with the authors $7\frac{1}{2}$ months for 2 revisions. Published online in *Articles in Advance* November 18, 2010.

1. Introduction

Our central thesis in this article is that the increasing availability of information about products, brands, or services through IT-enabled social media platforms on the Internet will make sponsored search markets more efficient in the same way that greater access to information through IT made financial markets more efficient over time.

The concept of "information" or "price" efficiency says that any new information is quickly and accurately impounded in prices that provide an unbiased estimate of true asset value (Damodaran 2003). Over the last few decades, financial markets became increasingly efficient as reaction to new information quickened, driven by larger numbers of participants with greater access to information (Butman 1998). We are witnessing a similar trend in the online world where it is becoming easier for people to find relevant content and participate more easily in its generation (Dhar and Sundararajan 2007). We argue that the increasing observability of keyword pricing information in search markets and volume of buzz generated about brands or products related to those keywords in social media and social commerce forums present a rich set of opportunities to study the future evolution of sponsored search advertising markets. A practical implication of this research is that lessons and concepts from finance can be applied by firms in online search advertising to maximize their returns on investments by exploiting inefficiencies in the current market.

Sponsored search is the mechanism whereby advertisers pay a fee to Internet search engines to be displayed alongside organic (nonsponsored) Web search results. The search engine serves as a marketplace. Multiple sellers compete to be displayed to potential buyers, and hence bid for a display. Activity gravitates towards this marketplace if users find *relevant* answers to their queries. Considering the success of search engines in terms of attracting traffic, their relevance is clearly "good enough" and can only get better, which should in turn lead to greater traffic to them or social media sites that use them in the future.

Financial markets and sponsored search markets share several common characteristics, as well as differences. Financial exchanges and sponsored search markets are similar in that both are auctions with similar business models. Participation in financial exchanges requires membership, through which an exchange controls risk, such as the ability of members to make good on their obligations in a transaction. A central objective of financial exchanges is to provide deep liquidity, typically measured in terms of volumes transacted. Indeed, their profitability is driven by volume since they earn commission based on the number and size of transactions. Providing a liquidity pool creates two-sided network effects, attracting buyers and sellers in a virtuous cycle. Like financial exchanges, sponsored search markets also attempt to create network effects by attracting search query volume and thereby sponsored search advertisements, which form the basis for their business model. They try to accomplish this goal by making searches more relevant to their visitors, who are the potential "buyers." The "sellers" are advertisers, whose goals are to reach the right segment of buyers while minimizing their bidding costs and maximizing their revenues.[1] Price discovery in both markets occurs because of the exogenous information generated about the item being traded, which is reflected in its aggregate demand. Finally, financial portfolio managers typically attempt to optimize portfolio performance through a careful selection of assets, given a budget. This is similar to what online search advertisers do, namely, optimize marketing spend through the selection of keywords subject to a budget constraint.

The differences between the two types of markets are equally interesting. In financial exchanges, the entity being bid on is a stock as opposed to a keyword. In sponsored search markets, advertisers bid for rank and not just on "winning" the auction by being the highest bidder (although being towards the top of the auction is considered desirable). Viewed in this way, there are only bidders in sponsored search markets with a subset of the top bidders being displayed if their bids exceed the minimum reservation price established by the marketplace.

Another key difference between the two markets is their contrasting transparency. In financial markets the algorithm for winning is known to everyone and is based on price and time—the highest bidder wins the auction and ties are resolved by who gets there first. The high level of transparency is achieved via a considerable amount of regulation aimed at providing a level playing field for buyers and sellers. In contrast, sponsored search markets are relatively opaque. Price is only one factor in ranking. The ranking also takes into account the past "performance" of the keyword being auctioned such as previous click-through rates and other factors related to landing page quality that are not completely transparent to bidders. Because of this lack of transparency of the ranking algorithm, sponsored search markets are correspondingly far less transparent than financial exchanges. This creates an incentive for sellers to try to estimate the ranking criteria. In fact, the emerging industry of search engine marketing/optimization is largely dedicated to inferring the ranking criteria in sponsored search advertising even though the ability of search engines to change the criteria at little or no cost is a significant disincentive for advertisers to game them.

Financial markets and sponsored search markets have emerged during very different stages of technological maturity. Historically, the high costs of maintaining and running a continuous auction in financial exchanges necessitated that they limit themselves to listing stocks that had a significant trading interest or liquidity. Being listed on an exchange required extensive scrutiny. In contrast, there are relatively minor risks involved in what keywords can be bid on in sponsored search, so the universe of symbols in sponsored search markets is potentially much larger, with no regard to the inherent interest in them. While sponsored search auctions could in principle be continuous similar to financial markets, they are mostly periodic. While rankings for a given keyword fluctuate a bit, they do not change significantly in real time. In this respect, sponsored search is in its early days, not unlike financial markets in their early days. Just as IT facilitated information access in financial markets, we expect the same to occur with sponsored search keywords, especially considering the ease with which information can be created and accessed in Web 2.0-based tools such as social media platforms and sites that harness the wisdom of the crowds.

Why do we expect Web 2.0 technologies such as social media and social commerce to be significant

[1] Note that there is also another demand-supply curve in these markets wherein the buyers are advertisers who wish to purchase a slot on the search engine results page while the seller is the search engine who auctions these slots. However, this distinction is less relevant to our article and hence, we refrain from discussing this in more detail.

drivers of sponsored search market efficiency in the future? One basis for this expectation is that much of the current research in this area shows that for certain types of content, users value opinions of other users or the "wisdom of the crowds" more than those of "experts" or authoritative sources (Dhar and Chang 2009). This phenomenon threatens to break the hegemony of traditional content creators since it motivates large numbers of people to become high-quality publishers of information. User-generated content in social media spaces is an unprecedented source of rich information about a very wide variety and continuously evolving set of subjects, some mainstream and others very specialized. Users are creating and contributing useful content, covering both popular concepts as well as niche, obscure subjects that many firms and marketers do not anticipate. Indeed, users are increasingly rating, recommending, reviewing, tagging, and commenting on textual content such as in blogs, product reviews, social networks, and reputation systems. By connecting customers to customers, online social commerce is leveraging those connections for commercial purpose by moving the word-of-mouth needle from "buzz" to "buy" (Decker 2010). For example, online retailers have found that customer reviews are the most effective user-generated social tactic for driving sales because there is high demand for this content by shoppers. In addition, there is evidence that social media content (UGC) can drive increases in search engine traffic itself (McCarthy 2010). In May 2010, the website that got the highest amount of traffic from social media sites like Facebook was Google (Heseltine 2010).

It could legitimately be hypothesized that there is a significant amount of "noise" in social media content compared to financial markets, which would dampen their legitimacy. Unlike financial markets where information came largely from known agencies with established degrees of credibility, the Web provides significantly more diversity in terms of information sources and their credibility. For example, there could be thousands of reviews for products, all of which constitute relatively raw data that need to be researched and analyzed, making data interpretation costly. This is true. However, we anticipate these costs going down over time because of better systems for information filtering and interpretation. Several systems for aggregating and summarizing content from social media and social commerce spaces are being made available to users. For example, Microsoft has incorporated this kind of system in its search engine Bing, wherein it summarizes reviews for computers and electronics products. We would expect the content to be summarized, interpreted, and digested faster in the future, and hence reflected in the bidding behavior of advertisers in sponsored search. By this logic, we would expect the time lag between a change in the popularity of conversations involving certain keywords in social media sites and a corresponding change in the demand for those keywords in search engines to go down over time, all other factors being constant. In other words, we expect the efficiency of sponsored search markets to increase over time.

2. Prior Research in Financial Markets and Sponsored Search Markets

There is a large body of literature that has studied the impacts of information technologies on financial markets in terms of the impacts of lower costs, higher speed, more transparency, and faster reaction times of market participants (see Butman 1998). The most direct evidence regarding increasing efficiency of financial markets is from Butman (1998), who compared market reactions of thousands of stocks to negative and positive earnings surprises between 1995 and 1998 to those between 1983 and 1989. The price reactions that took three to four weeks in the 1980s took only two days in the 1990s (Butman 1998). For a detailed review of the literature on financial and search markets that is relevant to this paper, please refer to the online appendix.[2]

The explanation for these findings was that the use of advances in IT such as faxes and the Internet had contributed to faster dissemination of information and resulted in stocks adjusting much faster to news. Since then, price adjustments in response to new information have continued to become faster. Reaction times in the current world of rapid and

[2] Additional information is contained in an online appendix to this paper that is available on the *Information System Research* website (http://isr.pubs.informs.org/ecompanion.html).

ubiquitous information access and algorithmic trading are often measurable in seconds or less.

The focus in much of the literature on empirical finance has been on asset pricing. The capital asset pricing model (CAPM) has been the simplest and most enduring baseline model for asset pricing (Sharpe 1964). The general form of this model is the following:

$$E(R_i) = R_f + \beta_i(E(R_m) - R_f),$$

where $E(R)$ is the expected return of the asset, R is the risk free rate of interest, and β_i is the sensitivity of the expected asset returns to market returns. In this model, the market is the major "risk factor" that explains returns of all stocks to degrees signified by their respective betas. Much of the subsequent research in empirical finance focused on extensions to the CAPM to account for the large number of additional (to the market in the above equation) risk factors that impact security returns. These types of models are known as "factor models" because they attempt to account for an asset's returns using factors that are relevant to that asset. For example, "size" is a typical risk factor because large capitalization stocks have been observed to behave differently from small capitalization stocks. Another commonly used risk factor is "industry" because stocks in the same industry, say automobiles, would be subject to similar risks that are different from those factors impacting, say, oil stocks. The well-known BARRA and other factor models that grew out of the work of Rosenberg (1974) use these types of risk factors in estimating returns of stocks. In these models, each asset has a beta (its sensitivity) with respect to each factor that is used to estimate its expected returns.

Some of the more recent work in empirical finance has added "liquidity risk" as a risk factor to factor models, and has examined the impact of liquidity on valuation and on market behavior. This work is summarized in Acharya and Pedersen (2005). The genesis of research in liquidity goes back to a fundamental result by Amihud and Mendelson (1987) on NYSE stocks between 1960 and 1980, which showed that illiquid stocks are associated with higher returns, implying that investors pay a "liquidity premium" for trading in the more liquid stocks by accepting lower returns. Conversely, investors are "rewarded" with higher returns for holding illiquid stocks. This observation has had numerous explanations for its existence (see Bodie et al. 2005). More recent work by Amihud (2002) showed that illiquidity is associated with higher ex ante returns and that lower liquidity is associated with a serial correlation in returns.

In contrast to financial markets, sponsored search markets are relatively new, but online advertising has come a long way in the last decade. The ascent of search engines such as Google, Yahoo, and MSN is driven by their unique position to sell new forms of advertisements. The increasing ubiquity and popularity of this mechanism has motivated several streams of work from multiple disciplines to examine a wide set of issues. Despite the vast body of theory work (for a summary of the theoretical literature, see Ghose and Yang 2009), very few empirical results exist in online sponsored search advertising. One stream of work has so far focused on examining issues related to adverse selection (Animesh et al. 2009), and keyword pricing differences based on ad context (Goldfarb and Tucker 2009). Another stream of work has begun to use firm-level data from search engine advertisers and marketing firms to examine the drivers of click-through rates and conversion rates along with the profitability of different positions on the search engine (Rutz and Bucklin 2007, Agarwal et al. 2008, Ghose and Yang 2009, Yang and Ghose 2010). Notably, Ghose and Yang (2009) show that keywords that have more prominent positions on the search engine results page, and thus experience higher click-through or conversion rates, are not necessarily the most profitable ones—profits are often higher at the middle positions than at the top or the bottom ones. An emerging stream of work has begun to use structural modeling approaches to examine issues such as externalities, consumer surplus, and social welfare (Song and Mela 2009, Jeziorski and Segal 2009, Athey and Nekipelov 2009). The final stream of work related to this article considers the issue of identifying relevant keywords (Joshi and Motwani 2006, Abhishek and Hosanagar 2007) and producing algorithms on how much to bid given budget constraints (Rusmevichientong and Williamson 2006).

Despite the copious literature on sponsored search, our understanding of selecting the most profitable set of keywords and deciding on how much to bid on them in sponsored search is still very nascent. This is primarily because the drivers of keyword pricing have not been studied in depth. Indeed, keyword pricing in sponsored search can be a complex problem. For a given keyword, the advertiser faces a different search volume across different search engines and this is reflected in the difference in bid prices across the search engines for a given keyword-advertiser pair. Even within a given search engine, different advertisers bid differentially for the same keyword. In addition, the same advertiser may bid differentially for a given keyword at different times on the same search engine. These data remain to be analyzed in detail, and the empirical issues we outline in this article are intended to provide future researchers with a broad framework that can be examined using scientific methods on the vast amounts of real-world data that are being generated through electronic commerce and social media.

3. Framing a Future Research Agenda

Before presenting our research questions, we need to define the variables to which they refer. These are presented in Table 1.

Table 1 **Description of Terms Used in the Paper**

Market efficiency	Time lag between information dissemination through the Web and its impact on bid prices
Keyword volume	The frequency with which a keyword occurs in social media over a given time period
Keyword liquidity	The number of unique searches for a given keyword over a given time period
Keyword volatility	Standard deviation of bid prices for a given keyword over a given time period
Advertiser demand	The number of unique advertisers who compete in an auction by placing bids for a given keyword
Long tail keyword	A "niche" or relatively obscure keyword whose search volume is very low compared to the more popular or "short head" keywords
Generic keyword	Keyword that does not identify either an advertiser or a brand but simply refers to a generic product
Branded keyword	Keyword that does identifies the brand (manufacturer, retailer, or product brand)
Broad match	If any of the keywords in the ad match the words contained in the search query.
Phrase match	If the keywords in the ad are in the same order as they are in the search query
Exact match	If the keywords in the ad match the exact order as they are in the search query.

Our theory of sponsored search markets builds on findings from prior work that show how economic value is embedded in online social media content (Ghose et al. 2006, Forman et al. 2008, Archak et al. 2008). If there were economic value, we would expect firms to engage more actively in monetizing this content. Our objective is to state the research questions that will frame inquiry in this area. Answers to these questions will help us understand how we should expect advertisers and search engines to behave in the future as sponsored search markets mature. Our thinking is driven by how financial markets matured over time, becoming progressively more efficient. During this maturation, players who were able to understand the inefficiencies benefited enormously. There are similar opportunities for advertisers in sponsored search markets. We classify the research questions into two sets as follows.

(a) How should keywords be priced in sponsored search advertising? Should we expect mispricing to occur and be exploitable my advertisers?

(b) What kind of design choices exist for advertisers and search engines? Specifically what payment mechanisms are optimal for advertisers under different conditions? How should advertisers design their advertising strategy, choose the type of match to use, content and style of ad creative, and landing page content quality, and on which search engine platform to advertise?

3.1. Keyword Pricing

There is accumulating evidence in the literature that the level of UGC in different social media forums such as blogs, product review sites, and other kinds of sites promoting "wisdom of the crowds" and "social commerce" play a key role in influencing the level of buzz about firms, brands, and products in the online world (see, for example, Archak et al. 2008, Forman et al. 2008, Dhar and Chang 2009, Ghose and Ipeirotis 2010). A good surrogate for this level of buzz is consumer search volume using specific keywords. This in turn is associated with advertiser demand for those keywords on search engines. The impact of the differences in advertiser demand on search engines should be reflected in the varying bid prices for those keywords. Our central thesis is that the volume of content about subjects that is being created by users in various social media and social commerce forums is an indication of "demand" for those subjects and should translate into a direct economic value for them.[3] The analogy here with the capital asset pricing model is that the overall level of buzz across subjects in the aggregate is similar to the "market factor"—the higher the buzz, the higher the overall keyword prices, but different keywords would have sensitivities to "the market" and other factors relevant to them. Given that queries on search engines take the form of keywords, we would expect advertiser demand to be proportional to the level of volume (or popularity) of such user-generated content on social media and social commerce sites. The level of advertiser demand should be reflected in the bid prices. Accordingly, increased search volume should have an impact on advertiser-level competition for those keywords and a consequent impact on keyword prices.

The thesis above states a general relationship just as capital asset pricing models do, recognizing that there will be cases where search volume has no obvious commercial implications (i.e., zero betas). An increase in the use of the term "volcanic ash," for example, may be purely informational or fleeting, with no necessary economic consideration. Our focus, however, is on subjects such as products, brands, and firms, where there should be a relationship between buzz and demand. This leads to the first and primary research question, namely, what is the pricing model for keywords in sponsored search advertising?

To investigate that question, we propose a research framework that breaks down keywords into several attributes that we hypothesize are determinants of economic value. These are similar to factors in asset pricing models. As mentioned above for example, "the market" is a dominant factor in explaining prices of individual stocks, and we expect a similar market factor to impact keyword prices across the board. We are interested in highlighting both the speed and magnitude of changes in keyword prices as a result of changes in the volume of buzz around those keywords. Moreover, our interest is in specifying other first-order factors that we expect to impact keyword prices. Specifically, we consider the following: (i) branded versus generic keywords; (ii) more liquid versus less liquid keywords; (iii) short head versus long tail keywords; and (iv) futures and derivatives markets for keywords.

Social media content is rich and diverse. Much of this content often contains conversations regarding new products or brands, and can refer to the manufacturer, the retailer, or the actual product. Thus, buzz that is generated in social media and social commerce forums can be specific to a brand. For example, we might find the spike in buzz about the "Nikon D700" far greater than about "SLR cameras" when the camera is released or something noteworthy about that particular model is discovered by consumers, which can often result in conversations about that model among online customers.

Studies of consumer search patterns reveal that people use generic keywords (such as SLR cameras) when they are in the early part of their consideration life cycle and are seeking information, whereas branded keywords (such as the Nikon D700) are more likely to be used when consumers are closer to their purchase decision (Kaushik 2009). We might therefore expect the beta (the extent of price sensitivity) for branded keywords to be higher than those for generic keywords, with the expectation that an increase in buzz related to branded keywords will be reflected in purchase decisions. This motivates the following question that correlates price with keyword type:

***(a) Keyword specificity question (magnitude): Will bid prices for "branded" keywords within a search engine change more than bid prices for "generic" keywords for an equivalent change in buzz volume*?**

In addition to considering the relative magnitude of price changes, a related question is whether increases in volume of buzz related to branded keywords will be reflected *faster* in prices than for generic keywords. This leads to the following:

***(b) Keyword specificity question (speed): Will bid prices for "branded" keywords within a search engine change faster than bid prices for "generic" keywords for an equivalent change in buzz volume*?**

If the answer to the above question turns out to be positive, this could be an indication that the market for branded keywords is more efficient than the

[3] This association is well established in the literature on word-of-mouth where we have seen a positive relationship between product sales and the content of user-generated product reviews (see, for example, Archak et al. 2008, Ghose and Ipeirotis 2010).

market for generic keywords. There could be several reasons for this, but the message would be that social media content relating to brands has value and people pay more attention to it than generic content. A negative observation might suggest that the brand-related social media content is largely noise and thus ignored by advertisers.

The concept of "liquidity" has been studied extensively in the financial literature. It refers to the ease with which an asset can be converted to its theoretical economic value through a transaction. The more liquid an asset, the more easily it can be converted to, say, cash. Volume is a typical proxy for liquidity in that something that is transacted more is generally convertible more easily than if it is thinly traded. As mentioned earlier, an enduring observation in finance is that liquidity has economic value (Amihud and Mendelson 1987). People are willing to pay for more liquid assets relative to illiquid ones, all else being equal.

The concept of liquidity in sponsored search markets is similarly motivated. Higher volumes of keyword searches indicate higher levels of liquidity for those keywords in terms of number of impressions. The arrival of "new information" in the online world, however, can be quite different than in the financial arena where there is a lot of attention around anticipated events such as earnings announcements, and immediate reactions to unanticipated events such as investigations, resignations, dividend initiation, etc. In the online world, new information often permeates more slowly due to its diversity, inherent potential noise, and more amorphous nature. Our expectation is that the speed with which new information impacts pricing will be proportional to the liquidity of the keyword, the expectation being that a keyword that is already very prevalent in searches is likely to be impacted less in terms of pricing than an illiquid one. This motivates the following question, which correlates bid price with search volume:

(c) Liquidity question (magnitude): Are changes in bid prices of more liquid keywords within a search engine lower than the changes in the bid prices of the less liquid keywords for an equivalent change in buzz volume?

Thinking in terms of factor models, the above question asks whether the price changes of different types of words have different sensitivities to liquidity. This line of thinking is similar in spirit to that of liquidity as a risk factor impacting asset price changes. As with the keyword-specificity question, a natural question concerns the speed with which prices of liquid keywords will adjust relative to illiquid keywords for an equivalent change in search volume:

(d) Liquidity question (speed): Are changes in bid prices of more liquid keywords within a search engine faster than the changes in the bid prices of the less liquid keywords for an equivalent change in buzz volume?

The questions articulated up to this point raise another natural question, namely, when do bid prices behave symmetrically in response to increases and decreases in buzz volume. It is conceivable that for some keywords, after awareness is generated, prices remain high independent of subsequent buzz. That is, these keywords may display stickiness in prices such that even after a decline in the volume of buzz (and consequently in the volume of search), bid prices on search engines for those keywords remain at the same level as before. This calls attention to the need to examine in what situations or for what types of keywords we might expect prices to increase in response to volume but remain sticky, that is, not revert to previous levels when buzz decreases. The question can be framed as follows:

(e) Directional symmetry question: When is the relationship between keyword prices and buzz volume unidirectional versus bidirectional?

The above question seeks to identify situations in which once awareness is generated, bid prices remain unchanged with subsequent buzz. An equally interesting question is whether we would expect price stickiness to be inversely correlated with price sensitivity. Specifically, if prices are highly sensitive to buzz volume, is it likely to be indicative of a bidirectional relationship? Is this more likely to be true of branded versus generic keywords?

The next two research questions consider the behavior of "long-tail" versus "short-head" keywords. For some time now, we have seen that the Internet has been shifting consumer purchases from the more popular products to the relatively more niche and obscure products because the latter have become easier to find compared to before the Web.

Indeed, the term "long tail" was coined by Chris Anderson (Anderson 2006) to describe the emergence of niche markets that were not feasible in the physical world (such as physical stores selling "oversized shoes" for which there is a market that is geographically dispersed and therefore infeasible in the brick-and-mortar world but ideal for electronic commerce). The concept has been discussed subsequently in detail in Brynjolfsson et al. (2006). The emergence of such markets creates a demand for obscure keywords and strategies of targeting less competitive and niche markets in addition to the broad markets. Other things being equal, one would expect a greater level of buzz to be generated for the popular products relative to the niche ones.

If short-head keywords have greater demand and consequently higher liquidity on both sides of the market, one would expect the prices for short-head keywords to be less volatile than those of the long tail keywords. The higher the liquidity, the more trades and more frequent the changes in prices. However, in such cases (with higher liquidity) these price changes are likely to be smooth, as opposed to cases with lower liquidity where there might be fewer price changes (because of fewer trades), but the magnitude of these price changes are likely to be larger (higher volatility). Said simply, one would expect prices for the short-head keywords to be less volatile compared to the long-tail keywords where prices are likely to adjust more infrequently but the changes could be of higher magnitude (and therefore more volatile). This leads to the following research question that correlates bid price and volatility of different types of keywords, based on whether they are in the tail or in the head:

(f) Volatility question: Would the prices for short-head keywords within a search engine exhibit lower volatility than for long-tail keywords?

Long-tail keywords are also likely to be longer (for example, contain more than three words in length) and more specific in nature than frequently searched generic phrases.[4] These types of keywords are often used by people searching for a specific product or service and are more likely to convert into a sale or enquiry because user intentions are already quite specific (Kaushik 2009). Accordingly, the level of advertiser demand for these keywords could increase faster than that for the short-head keywords in response to increasing search volume for them.

Finally, in concluding this section, it is worth considering the possibility of futures markets emerging in sponsored search as they did in financial markets. Futures markets arise for commodity products due to the time difference between a contract and a delivery, which often requires buyers or sellers to hedge. At the current time, there are no futures markets in sponsored search. An advertiser cannot buy a ranking on a search engine for the keyword "Olympics 2012" for June 2012. We expect this to change, as some advertisers seek to lock in prices they are willing to pay in the future. The necessary conditions for these markets to emerge would be identical to those in financial markets: commoditization of products, a high level of interest or liquidity, and trust that market makers would indeed deliver on their obligations. If a highly trustworthy counterparty in such a transaction is the search engine or social media/commerce site itself, such a market could prove to be a big money maker for them by enabling them to "sell insurance" to advertisers, guaranteeing them future prices for keywords and allowing for the same type of hedging made possible by futures and options markets. This motivates the following research question.

(g) Futures/derivatives markets: What attributes of keywords would make them amenable to the emergence of futures and derivatives markets?

In financial markets, derivatives such as futures and options help "complete the market," which means that they enable the complete set of possible gambles on future states-of-the-world to be constructed. The most liquid futures markets are those that provide the most useful hedging potential or an expression of interest in the most covered parts of the market. For example, index futures tend to be heavily traded since they cover large parts of the market of interest to large numbers of investors. These include the S&P 500 Index futures, which represent a set of 500 liquid stocks, and other more specialized indices such as exchange-traded funds (ETFs) represent aggregate prices of certain sectors such as financials, pharmaceuticals, etc. Prices of these indices are based on

[4] See, for example, SEO terminology at www.seojunkies.com.

the prices of individual assets that comprise them. It is possible that in search markets, indices such as travel or automobiles could become the equivalent of ETFs, consisting of sets of keywords (equivalent to sets of assets) with some kind of weighted aggregate price representing the overall price of the set-of-words index. As in finance, these derivatives would enable sellers to express gambles on future states of the world that are not currently feasible.

To summarize, the above-mentioned questions present a basic understanding of keyword pricing in sponsored search markets in terms of well-studied concepts in the financial world such as efficiency, volatility, and elasticity. How have researchers of financial markets tested their hypotheses based on similar concepts? What is similar and different about the data emanating from sponsored search markets relative to financial markets?

In large part, researchers in financial markets have studied prices and bid-ask spreads of different assets across time and under a wide range of conditions, including the arrival of new information. Over the last decade, studies have occurred in increasingly granular time intervals, which have led to considerably more sophisticated theories than what existed during periods of low-frequency data. Despite the fact that data quality tends to be a pervasive problem in the world of finance, researchers have been fortunate enough to be able to obtain standard historical data from vendors for hypothesis testing and theory development.

In sponsored search markets, we have a deluge of time series data, which is growing exponentially over time. These data provide a rich set of observables for testing the types of propositions we have stated. At the same time, the vast swath of social media content in online forums is visible and available as well, often grouped by industry or subject. These two sets of data are required to map the interdependence between keyword prices and online buzz, making our questions testable. Fortunately, we are beginning to see the emergence of data sets available to researchers for theory development, which we mention in §4. Nevertheless, given the lack of standard data sources, researchers must be strategic about gathering some of these data themselves.

3.2. Design Choices in Keyword Advertising

There are a number of design choices available to both search engines and advertisers in sponsored search. We focus on two specific ones: choice of payment mechanism, which determines how the marketplace generates revenues and choice of advertising strategy by advertisers, which includes choosing the type of match to use, content, and style of ad creative, landing-page content quality, and which market to advertise on. We elaborate on these two questions below.

(a) Payment mechanism question: What payment mechanisms such as pay-per-impression, pay-per-click, or pay-per-conversion are optimal for advertisers?

A promising area for future research would be to examine the profit-maximizing payment mechanism for advertisers and search engines and the associated pricing policies. There is a considerable amount of debate in the trade press with regard to the issue of the optimal pricing mechanism in sponsored search and the use of hybrid auctions. A recent stream of work has begun to examine these issues using various game-theoretic models that incorporate different incentives for different entities involved (Agarwal et al. 2009). The problem is complex because of different payment options available to advertisers, such as pay-per-impression, pay-per-click, or pay-per-conversion (or action). The three payment models differ based on when the advertiser pays the search engine. In the pay-per-impression (PPI) model, the advertiser pays the search engine every time the ad is shown, irrespective of whether there are any clicks or conversions. In the pay-per-click (PPC) model, the advertiser pays when a user actually clicks on the advertisement. In the pay-per-action (PPA) model, the advertiser pays when an actual concrete action (such as a product sale, filling out a form, subscription to a service, etc.) takes place.

Depending on the product or service being sold, the returns on advertising investments from the three prevalent pricing models could be very different, and hence incentives for participation in these three different mechanisms vary for advertisers. These differences also necessitate different bid prices for each of them, which is a nontrivial decision for advertisers. Keeping in mind our first research question on keyword pricing, we believe that a comprehensive answer to the above question would necessitate a mix of analytical modeling followed by empirical analysis of experimental and secondary data. We would expect that the answer would be predicated on the nature of the product or the service that is being advertised. For example, for a product, which is not available for sale online, a click-through may be less meaningful than a simple exposure. Hence, it might make sense for advertisers to participate in a PPI mechanism only. On the other hand, for products whose entire transaction from search to purchase can be completed online, it would make relatively more sense for the advertiser to participate in PPC or PPA mechanisms. Since the volume of buzz that is generated will also be correlated with product characteristics, an answer to the above question necessitates that researchers take into account findings related to the first question.

(b) Keyword attribute question: How should advertisers choose various keyword-level attributes towards the objective of maximizing their ad budget portfolio?

What complicates the nature of the keyword pricing process is that keyword advertising often entails a multi-dimensional decision-making process from the advertiser's perspective with respect to choosing and bidding on keywords. This is because the advertiser has to decide which keywords to choose, and how much to bid on them. In addition, the advertiser needs to decide on the scope of the keyword (such as whether they should go for an exact match or a broad match), the ad creative (that is, the text appearing on the top of the ad on the search engine results page), and the design and content of the landing page (since it affects the quality score and thereby the final rank in the auction process). For a given keyword, the advertiser faces a different search volume across different search engines and this is reflected in the differences in bid prices across the search engines for a given keyword-advertiser pair. In addition, the same advertiser bids differentially for a given keyword at different times on the same search engine (Rutz and Bucklin 2007, Ghose and Yang 2009, Yang and Ghose 2010). These interdependencies across information elements also make the pricing of keywords in sponsored search a complex, nontrivial problem for researchers to model and analyze.

Advertisers typically have a daily budget for keyword ads and their main objective is to maximize return on their portfolio of keywords given budget constraints, much as we see in financial markets where investors aim to maximize risk-adjusted returns on their portfolio of stocks given a specific amount of capital. Indeed, the problem is akin to multi-period portfolio optimization where the problem is to decide how much to invest in each possible instrument at a point in time based on the covariance between the instruments and the expected return for each. In sponsored search, the decision is how much to invest in each keyword at a point in time given the expected return of each, the return correlations between each pair of keywords, and constraints on how much can be spent on each and as a whole. Despite the complexity involved in making these decisions, managers can put our questions into practice in the selection of data that they should be recording about keyword advertisements. This is as a first step towards the optimal design of advertising campaigns, akin to how financial portfolio managers design optimal portfolios.

4. Research Approach

There are several possible approaches for addressing the central research question related to the pricing of keywords. One simple measure of popularity of keywords, for example, is the frequency with which they appear on social media and Web 2.0 sites such as blogs, opinion forums, social tagging sites, social commerce sites, etc. The volume of social media content for any set of keywords on the Web can be tracked using publicly available social search tools. For a given sample of keywords, one could collect data on frequency of occurrence (volume) of those keywords from blog forums. Blog frequency data is now available through several social search tools such as the Blogscope project at the University of Toronto, which tracks over 30 million blogs across the world and provides the historical frequency of any keyword (the data set currently tracks data over the past 12 months). In a similar vein, for a given sample of keywords, one could collect data on frequency of occurrence (volume) of those keywords using social search tools on microblogging forums (for example, Twitter). The interactive social search tool tweetvolume.com enables us to determine how often any keyword or phrase has appeared on a Twitter post in a given day, week, or year from March 2001 to present. Tweetvolume generates the estimates by utilizing resources from Google's search engine and Twitter's open API to capture and compare data points. Another social search tool, Tweetscan enables us to determine the individual users who are talking about a keyword and what they are saying.

It is also possible to record data on various sponsored search advertising metrics for keywords such as the average cost per click, average number of clicks, average number of impressions, extent of advertiser competition, etc., on search engines such as Yahoo, MSN Live, and Google. This time-series data can be used to identify variation within keywords over time and test for the association between social media and sponsored search advertising. Aggregate-level sponsored search data of this kind are now available through the Microsoft Ad Center as well as through the keyword-pricing tool on Google. The Ad Center tool provides access to panel data that has both time-series and cross-sectional variation in the above-mentioned sponsored search metrics for a sample of keywords. Such data can be used to examine if the volume of social media for specific keywords shows an association with the prices advertisers are willing to pay for these keywords in search engines.

One could also examine if the volume of different kinds of UGC across different social media sites is correlated with the potential market size of these keywords, where market size is estimated on the basis of the number of clicks or the number of impressions for those keywords across advertisers. Based on our central thesis, we expect this correlation to increase on average as the market becomes more efficient, i.e., the time lag between the generation of new information and its impact on advertiser competition and bid prices goes down. A particularly useful website for publicly available data on the comparative volume of keywords appearing on different social media sites is Virtue.com. This company analyzes online conversations on a variety of social networks, blogs, status update sites (i.e., microblogs), and photo- and video-sharing sites on a daily basis. A series of algorithms are applied to reflect the frequency of usage, the size of the social media environment, and the magnitude of the conversation. The result is a numeric score for each brand known as the Vitrue Social Media Index (SMI). Future research can study the correlations between this score and various metrics (such as bid prices and click-through rates) in sponsored search.

It is important to note that in any empirical estimations that are conducted along the lines proposed above, it would be critical to control for past demand of that keyword on the search engine, and then draw conclusions. Said simply, given that past demand is likely to be an excellent predictor of current and future demand, a researcher would need to control for past popularity of that keyword on the search engine. One useful measure of popularity is the volume of search queries for a given keyword on a search engine. The idea is that a higher search volume for a given keyword reflects a higher level of consumer interest in that keyword, and this past demand is likely to be positively correlated with future demand for that keyword. If markets become more efficient, this would be reflected in a higher level of advertiser interest in sponsoring such keywords. Such keyword search volume data is now publicly available through the Google Trends tool in ready-to-use Excel spreadsheets. Thus, future academic research that aims to examine such time-series correlations about the impact of social media content on keyword pricing would need to directly control for past demand of that keyword on the search engine to make any meaningful inferences.

Are there specific industries where we would expect our propositions to hold particularly well? We expect these to be ones where the time lag between consumer conversations and their implications on firm sales and profitability is relevant. Industries such as media, healthcare, consumer packaged goods, and travel and hospitality industries are among the ones that have been strongly impacted by social media over the last several years and where users tend to influence each other. There have also been several examples from such industries where a small and adverse consumer reaction to a product or service had a snowballing effect on sales. In crisis control or disaster management situations with respect to public relations, tracking consumer conversations and using that data to shape an online advertising strategy for damage control would be very valuable. This is important given that spikes in user-generated content and consumer search queries in search engines can be for positive sentiments as well as for negative sentiments associated with brands or products. Hence, advertisers who use consumer (user-generated) language to shape their ad strategies on the Internet in lieu of or in addition to industry (professional) language will benefit significantly by extracting underpriced keywords.

5. Conclusion

Search has come a long way in the last decade. In the pre-Internet era, it was slow and laborious, with each search requiring time and resources. Good information was scarce and locating it was costly. In the Internet era and its associated explosion of information fueled by user-generated content in social media and social commerce platforms, information is plentiful, but the problem is one of filtering the torrent of results that comes at a user who merely type something into a search engine. Modern-day search engines are already quite good at determining user intent and are likely to get even better over time. Because of this, the Internet will continue to expand as a marketing channel where advertisers can be confident of displaying their offerings only when they are likely to match user intent and lead to desired outcomes. Because customer behavior can be tracked, it makes the impact of an advertising campaign or strategy measurable. This fact has transformed sponsored search into a continuous auction market, much like financial markets, where sellers are bidding for "space" on the channel.

A key practical implication of our thesis is that search advertising on the Internet is currently mispriced, because it takes time for the market to absorb new information about keywords—whether they have become more or less valuable. Because of this inefficiency, there are keywords that are cheap to buy relative to their level of interest (i.e., buzz volume) and others that may be overpriced for similar reasons. Due to the mispricing of keywords, there is potential for advertisers to improve their ROI once a good basis for mispricing is identified. In this context, it is worth noting that, similar to financial markets, where algorithmic and systematic trading now makes up a majority of trading activity, it is equally likely that sponsored search bidding will also become algorithmic as advertisers obtain reliable economic estimates associated with search terms, which feed into the bidding decision. In other words, we can expect trading activity in sponsored search markets to become increasingly sophisticated as advertisers tease out mispricing and value from the markets and are able to measure expected returns more accurately from past data.

Our article resonates well with the emerging phenomenon of social media keyword research (Nelson 2009) which is a new approach to choosing the right keywords and phrases for a firm's search engine optimization strategy. The main premise of this strategy is that a firm tracks the conversations that its customers are having about its brands, its products or its services using a variety of social search tools (for example those available on Facebook or Twitter) and mines them to extract keywords most suitable for their search engine marketing strategy. Said simply, firms track a number of websites where their online community would congregate, and look at the keyword phrases that appear most frequently in the websites' content on a search. Once patterns of usage for certain keywords are visible across websites, it can be informative about what language used in keywords would best resonate with their audience.

We have drawn a parallel between financial and sponsored search markets. The similarities are compelling, but there are also some important differences. Despite the differences, however, financial markets provide a stable and well-understood set of concepts that we have found useful in the analysis of sponsored search markets. The set of questions we have presented in this paper are testable with the data currently available. We expect future research will test these questions and refine them into a more complete theory of search markets.

Acknowledgments

The authors thank Yannis Bakos, Foster Provost, V. Sambamurthy, and Bruce Weber for detailed discussions and useful suggestions. Funding for this project was provided by an NYU-Poly Seed Fund grant. The second author also acknowledges the generous financial support from NSF CAREER Award IIS-0643847.

References

Abhishek, V., K. Hosanagar. 2007. Keyword generation in sponsored search. *Proc. Internat. Conf. Electronic Commerce, Minneapolis*.

Acharya, V., L. Pedersen. 2005. Asset pricing with liquidity risk. *J. Financial Econom.* **77**(2) 375–410.

Agarwal, N., S. Athey, D. Yang. 2009. Skewed bidding in pay-per-action auctions for online advertising. *Amer. Econom. Rev.* **99**(2) 441–447.

Agarwal, A., K. Hosanagar, M. D. Smith. 2008. Location, location and location: An analysis of profitability of position in online advertising markets. Working paper, SSRN.

Amihud, Y. 2002. Illiquidity and stock returns: Cross sections and time series effects. *J. Financial Markets* **5** 31–56.

Amihud, Y., H. Mendelson. 1987. Trading mechanisms and stock returns: An empirical investigation. *J. Finance* **42**(3) 533–553.

Anderson, C. 2006. *The Long Tail: Why the Future of Business Is Selling Less of More*. Hyperion, New York.

Animesh, A., V. Ramachandran, S. Viswanathan. 2009. Research Note—Quality uncertainty and the performance of online sponsored search markets: An empirical investigation. *Inform. Systems Res.* **21**(1) 190–201.

Archak, N., A. Ghose, P. Ipeirotis. 2008. Deriving the pricing power of product features by mining consumer reviews. Working paper, SSRN, and Working Paper NYU CeDER 07-05.

Athey, S., D. Nekipelov. 2009. Equilibrium and uncertainty in sponsored search auctions. Working paper, Harvard University, Cambridge, MA.

Bodie, Z., A. Kane, A. Marcus. 2005. *Investments*. Tata McGraw-Hill Inc., New York.

Brynjolfsson, E., Y. Hu, M. Smith. 2006. From niches to riches: Anatomy of the long tail. *SMR* **47**(4) 67–71.

Butman, R. 1998. Big news on your stock? Hold on to your hat. *Wall Street Journal* (April 27).

Damodaran, A. 2003. *Investment Philosophies: Successful Strategies and the Investors Who Made Them Work*. Wiley, New York.

Decker, S. 2010. Social commerce 101: Leverage word of mouth to boost sales. Click, New York. http://clickz.com.

Dhar, V., E. A. Chang. 2009. Does chatter matter? The impact of user-generated content on music sales. *J. Interactive Marketing* **23**(4) 300–307.

Dhar, V., A. Sundararajan. 2007. Issues and Opinions—Information technologies in business: A blueprint for education and research. *Inform. Systems Res.* **18**(2) 125–141.

Forman, C., A. Ghose, B. Wiesenfeld. 2008. Examining the relationship between reviews and sales: The role of reviewer identity information in electronic markets. *Inform. Systems Res.* **19**(3) 291–313.

Ghose, A., P. Ipeirotis. 2010. Estimating the helpfulness and economic impact of product reviews: Mining text and reviewer characteristics. *IEEE Trans. Knowledge Data Engrg.* Forthcoming.

Ghose, A., S. Yang. 2009. An empirical analysis of search engine advertising: Sponsored search in electronic markets. *Management Sci.* **55**(10) 1605–1622.

Ghose, A., P. Ipeirotis, A. Sundararajan. 2006. The dimensions of reputation in electronic markets. NYU CeDER working paper, New York.

Goldfarb, A., C. Tucker. 2009. Search engine advertising: Pricing ads to context. Working paper, SSRN.

Heseltine, S. 2010. Can Facebook drive great traffic to websites? Accessed May 19. http://searchenginewatch.com/3640379.

Jeziorski, P., I. Segal. 2009. What makes them click: Empirical analysis of consumer demand for search advertising. Working paper, Stanford University, Palo Alto, CA.

Joshi, A., R. Motwani. 2006. Keyword generation for search engine advertising. *Sixth IEEE Internat. Conf. on Data Mining Workshops*, IEEE Computer Society, Washington, DC, 490–496.

Kaushik, S. 2009. Paid search analytics: Measuring value of "upper funnel" keywords. Web Analytics Blog, Occam's Razor. (February).

McCarthy, D. 2010. Social media marketing drives search traffic 61% for one small business: A case study. (June 7), http://www.viralhousingflix.com.

Nelson, P. 2009. How to use social media keyword research to drive traffic to your website. (December 17), http://www.underdesign.co.uk.

Rosenberg, B. 1974. Extra-market components of covariance in security returns. *J. Financial Quant. Anal.* **9**(2) 263–274.

Rusmevichientong, P., D. Williamson. 2006. An adaptive algorithm for selecting profitable keywords for search-based advertising services. Working paper, Social Sciences Research Network (SSRN).

Rutz, O., R. Bucklin. 2007. A model of individual keyword performance in paid search advertising. Working paper, Social Sciences Research Network (SSRN).

Sharpe, W. 1964. Capital asset prices: A theory of market equilibrium under conditions of risk. *J. Finance* **19**(3) 425–442.

Song, Y., C. Mela. 2009. A dynamic model of sponsored search advertising. Working paper, Social Sciences Research Network (SSRN).

Yang, S., A. Ghose. 2010. Analyzing the relationship between organic and sponsored search advertising: Positive, negative, or zero interdependence? *Marketing Sci.* **29**(4) 602–623.