

# New Twists on Auctions

**Abstract:** *With the advent of the Internet, traditional auction forms have evolved to fit into a plethora business niches, either integrating into traditional approaches or simply creating new opportunities. This chapter examines two novel uses for auctions in a business context, namely sponsored search auctions and prediction markets. Understanding the potential auction benefits and limitations can hopefully provide practitioners with a more informed and successful approach when employing these auction-based tools in their business.*

**Keywords:** Advertisements Pricing, Decision Support, E-Commerce Pricing, Electronic Bidding, Electronic Marketplaces, Electronic Markets, Group Decision Making, Internet Advertising.

## SPONSORED SEARCH AUCTIONS

### Introduction

Search engines are big business on the Internet. The two main players, Google and Yahoo!, made \$6 billion and \$5 billion in 2005 respectively. This bonanza mainly comes from *sponsored search*, where search engines charge advertisers to display their ads when users search for particular terms. Instead of large banner ads directed to everyone, sponsored search allows advertisers to target their audience. Real estate agents can display their ads only to users searching on the term “real estate”. In Figure 1 below, Google’s unpaid (or native or organic) search results are in the middle of the screen, while its advertising network’s paid results (or sponsored links) are on the right-hand side.

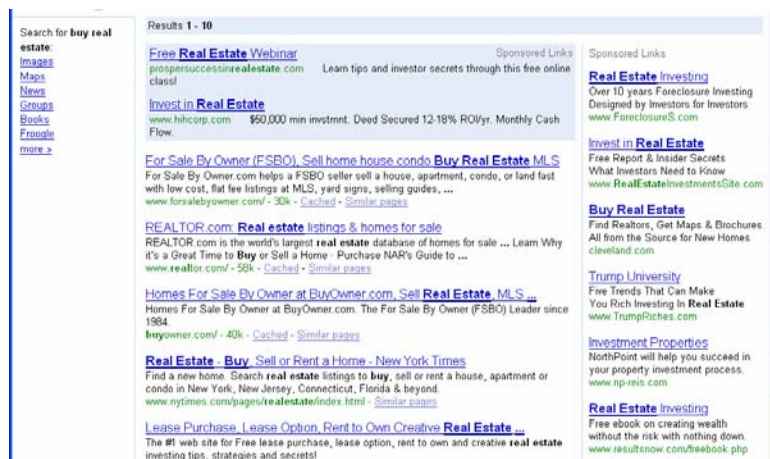


Figure 1: Search for “buy real estate” on Google

To advertise via sponsored search, advertisers first identify potential search terms of interest and then bid on the *cost per click* (CPC) that they are willing to pay to the search

engine if a user clicks on their advertisement. Each time a user searches, an auction occurs based what advertisement keywords match the search terms. The auction ranks advertisements based on bid amount, desired geographic location and other criteria, then displays the advertisements according to their ranking in the paid results slots (e.g., Google Adwords, Yahoo! Search Marketing). Alternatively, Google's AdSense serves ads when users visit websites enrolled in their program. Instead of matching using a user's search query, they match advertisement keywords according to the website's content.

Search engines make money from advertisers, who can choose to take their business to any one of several search engine companies. For a search engine to attract advertisers, they must attract or be able to target the users an advertiser desires, and to attract users, they must have good unpaid search results. So Yahoo! Search Marketing's main competitor in the sponsored search business is not Google Adwords, but Google's unpaid search results, and vice versa (Balcan, Chen, Devanur, & Kumar, 2006).

Behind the scenes in sponsored search, huge numbers of transient auctions are occurring, one for every time a user enters a query (Jansen, 2006). Instead of the standard one-shot one-item auction, sponsored search auctions are continuous real-time multi-item auctions. Instead of two main participants, there are three: the search engine, the advertiser, and the user. Academic research is just starting to evaluate auction mechanism tradeoffs in terms of properties such as maximizing revenue, consumer welfare, price stability, and robustness to fraud. In this chapter, we describe some current research results in this area and their implications for practitioners.

### ***Auction Mechanism Background***

Google and Yahoo! are currently the major players in the sponsored search business, with several smaller players such as AskJeeves and LookSmart, and the more recent addition of the Microsoft adCenter. We focus here on the Yahoo! and Google contributors to the present day sponsored search auction format. GoTo.com (renamed Overture in 2001, and acquired by Yahoo! in 2003) originated the sponsored search auction in 1998, while Google's first sponsored search auction followed in 2002 (Fain and Pedersen, 2005). The original GoTo.com auction used a first price payment mechanism. The bidder with highest bid received the top advertising slot and was charged a price equal to the bid amount. The second highest bidder was assigned to the second highest slot, paying its bid amount, and so forth. This payment mechanism led to strategic bidding-war cycles that consumed advertiser time and diminished search engine profits (Edelman & Ostrovsky, 2005). For example, suppose advertiser *X* values the keyword "real estate" at \$2, while advertiser *Y* values it at \$1. If *X* starts the bidding by bidding the lowest possible price, say \$0.10, then *Y* bids \$0.11 to win the first ad slot. Advertiser *X* responds by bidding \$0.12, and so forth. Once *X* bids \$1, then *Y* will not bid \$1.01 since *Y* only values the keyword at \$1. To acquire the second ad slot, *Y* simply has to bid \$0.10. But now *X* only needs to bid \$0.11 to win the first slot, and so the cycle starts over.

Google AdWords added two significant twists to the auction mechanism. First, inspired by auction theory, they changed from a first price payment mechanism to a more robust second price one. In a second price auction, the highest bidder wins but only pays the second-highest bid price (plus some small delta). For example, suppose the bids are \$10, \$9, \$8 for a single item, then the \$10 bidder wins but only pays \$9. Under certain conditions (e.g., sealed-bid, one-shot), a second price auction is *incentive compatible*, meaning that the optimal strategy for bidders is to bid their true valuation (Vickrey, 1961). This eliminates time-consuming strategic game play and ensures that the item is sold to the bidder who values it the most. The Vickrey-Clark-Groves (VCG) auction is a multi-item extension to the single-item second price auction that retains the incentive compatibility property by charging each bidder the opportunity cost imposed upon the other bidders by being there (Krishna, 2002). Google's multi-item second price auction is known as generalized second price (GSP) auction, and is a straightforward extension of the single-item second price auction, where the  $k^{th}$  highest bidder pays the  $(k + 1)^{st}$  highest price plus some delta. Consider our previous example where advertiser  $X$  pays up to \$2 for the keyword "real estate", and advertiser  $Y$  will pay \$1. In a sealed-bid environment, where only the auctioneer knows the value of all bids, if  $X$  bids \$2, and  $Y$  bids \$1, then  $X$  pays  $\$1 + \$0.1$  (the minimum delta) while  $Y$  pays the minimum bid amount of \$0.10. If only one slot is available to advertisers, then the GSP and VCG auctions are equivalent. However, their properties differ when multiple slots are offered. Namely the GSP is not incentive compatible, but is easier to understand than the VCG, and can reduce the bidding war cycles of the first price mechanism. Also, while VCG maximizes consumer welfare, it may not maximize revenue or hold up well in the face of collusive and anti-social behavior by bidders (Ausubel & Milgrom, 2005; Brandt & Weiss, 2001).

The second main twist Google added to the original first-price auction mechanism is ranking its ads not by bid price alone, but by computing a quality score derived from the bid amount, the ad's click-through-rate, keyword relevancy along with landing-page and site quality. The *click-through-rate* (CTR) measures the rate that web searchers click through to the advertiser's web site when shown the ad. Google's other quality-based criteria serve to penalize sites that use deceptive practices or have exceptionally poor web sites, and is assessed through both human and automated methods. In this way, Google can protect the user's search experience, and also increase their profits, since users are more likely to click on relevant ads, and thereby generate Google revenue. These two auction mechanism changes have helped make Google's auction more stable and more profitable. Google's second price auction debuted in February 2002, and Overture (rebranded to Yahoo! Search Marketing in 2005) updated to a second-price auction in May 2002. However, Yahoo! has only recently decided to add a Quality Index (Yahoo!) to their ranking as of early 2007.

Sponsored search auctions demonstrate how markets can evolve to forms more suitable to their commercial environment. Researchers have also begun to formally describe both current properties of sponsored search auctions and consider potentially desirable modifications. However, it is difficult to formalize the full-fledged sponsored search auction environment where there are a multitude of continuously repeated auctions. In

game theory, the Folk Theorem states that any feasible and individually rational outcome of a repeated game can be sustained as an equilibrium. In other words, almost any reasonable outcome can be shown to happen across infinitely repeated games. Some researchers have simplified the problem by considering the auction mechanism for a one auction, one keyword setting. Under these conditions, the GSP auction has been shown theoretically to have a local equilibrium where advertisers find a bid amount such that they neither want to raise their bid to move up to the next highest slot, nor lower it to move to the next lower slot. This equilibrium concept was discovered independently as a *locally envy-free equilibrium* (Edelman, Ostrovsky, & Schwarz, 2005) or equivalently a *symmetric Nash equilibrium* (Varian, 2006).

Researchers have also stepped back to look at the entire competitive landscape by considering the problem facing an advertiser with multiple advertising campaigns, each consisting of multiple keywords. Optimizing bidding strategy for a single auction, may not be the best strategy across all auctions of interest (Balcan, Chen, Devanur, & Kumar, 2006). Most have focused on using the budget to evaluate the overall return on investment (ROI). One proposed bidding heuristic is that advertisers equalize their ROI across all keywords (Borgs et al., 2007). Others have incorporated minimum ROI into their bidding strategies, or applied budget constraints to single auctions (Abrams, 2006; Borgs et al., 2007; Lahaie, 2006; Szymanski & Lee, 2006).

Budget can also serve as a proxy for risk management (e.g., against click fraud or unexpected user query surges). Auction payment mechanisms can also effect how risk is shared between the search engine and the advertiser. Pay-per-action (PPA) pricing allows advertisers to specify what action (e.g., selling a product, sales lead, etc.) a user would have to take on their website before they would have to pay. PPA advertising can help mitigate the risks of fraudulent clicks since the advertiser pays only if a customer takes a further action. Both pay-per-click and pay-per-action allow the search engine to bear more of the market variation risk than pay-per-impression. For example, if an ad gets shown more often at 5am, but users are less likely to buy then, an advertiser will not incur any budgetary loss. Another possibility is extending the auction mechanism to include futures – florists who want to lock up advertisements on and shortly before Valentine’s Day, might buy a keyword future for the word “flowers” (Balcan, Chen, Devanur, & Kumar, 2006). Finally, different auction mechanisms can also serve to apportion risk differently across search engine and advertiser (Liu, Chen, & Whinston, 2006).

### ***Click Fraud***

Click fraud occurs when an advertiser’s ad is clicked with the sole intent of generating a charge to the advertiser. It can be done automatically by computer scripts or directly by humans, and can occur for a variety of reasons, including a competitors desire to minimize the impact of an ad campaign, simple vandalism, or a desire by a publisher to increase their income (Jansen, 2006). To see this, consider Google’s Adwords and AdSense programs. In the Adwords program, Google displays advertiser’s ads on their search engine as a response to user queries. Clicking on an advertisement costs the

advertiser and runs down their ad budget. In the AdSense program, Google serves ads to a website whenever a user visits that website. Clicking on an advertisement costs the advertiser, but also provides revenue to the website owner. Google and Yahoo! search for suspicious patterns of clicking, and automatically do not charge advertisers for any clicks deemed to be an instance of click fraud. However, the process of assessing click fraud is complex and not provided to the public. Since the advertising networks benefit from the increased traffic, albeit unintentionally, both advertisers and website owners have concerns that fraudulent clicks be classified correctly.

The class action lawsuit, the Lane's Gifts v. Google, charged Google with not doing enough to detect click fraud. As part of the settlement, Dr. Alexander Tuzhilin from NYU was commissioned to do an independent study of Google's handling of click fraud (Tuzhilin, 2006). He was given access to Google's engineers and inside documents. His main finding was that Google's efforts to detect and combat click fraud seem to be "reasonable", with Google efforts becoming more effective by 2005. However, given the lack of full conversion data, there was way to give hard scientific evidence as to exactly how well the measures were working. Indeed, for some clicks, especially human ones, it may be impossible to determine whether it was click fraud simply from the data. Given this, Tuzhilin lends support to the notion that using CPA pricing may be necessary to reduce or eliminate click fraud. He also confirmed that AdSense is more prone to fraud than AdWords. Other involve designing click fraud resistant algorithms (Goodman, 2005; Immorlica, Jain, Mahdian, & Talwar, 2005).

### ***Recommendations***

#### *Identify business objectives.*

Online advertising is here to stay. In 2005, Proctor & Gamble was the nation's biggest advertiser, spending \$33.5 million of its \$4.6 billion ad budget on online ads. Similarly the second-biggest advertiser, General Motors, spent 2.5 percent of its total ad budget of \$4.35 billion online (Story, 2006). However, these large established companies tend to be more interested in furthering their brand recognition online, and concentrating on banner ads and other display advertisements rather than search ads. Clearly before starting a sponsored-search advertising campaign, a company must identify their business objectives and see if sponsored search is a good match. Some standard business objectives include selling products online, gathering contact information, generating traffic on your website, or providing information to drive visitors to local physical business locations. While most sponsored search pricing is based on CPC, it may be more straightforward for a business to evaluate its objective using different metrics. Cost-per-mille (CPM) measures the number of times (in thousands) that users have seen the ad, or the user impressions. It is most valuable when brand recognition is the goal, and comes directly from off-line advertising. Cost-per-action (CPA) measures how many of the clicks on an advertisement result in a sale or some other valuable "action". This is known as the *conversion rate*. For example, if the objective is to generate customer leads, then the action being measured might be how many users fill out a form with their name and email for further contact. Pricing based on CPM and CPA may be offered

directly by the search engine, or advertisers can buy advertisements using CPC pricing and then use the CTR, CPM or CPA rates supplied by the search engine data to evaluate how well their ad campaign outcome aligns with their business goals.

*Devote time to discovering the best keywords.*

Discovering the right keywords for a sponsored search ad campaign can make a big difference in the CTR and conversion rate. For example, consider the keywords “real estate” vs. “buy real estate”. Users who clicked on ads for “buy real estate” were much more likely to generate actions that led to sale leads, at least in 2003 (Ellam et. al. 2003). While there might be fewer searches for “buy real estate” vs. “real estate”, users who include the keyword “buy” in their search had a higher conversion rate and thus were more valuable. Advertisers bidding on “real estate” would need to bid higher to rank highly in a “buy real estate” auction. However, by bidding higher, they then show up as highly ranked when more casual users searched for “real estate” and end up paying a higher price for users with a lower conversion rate. To help address this issue, a common bidding strategy for advertisers was to estimate their conversion rate for the different keyword phrases, and then set their bid to maximize return on investment (ROI). So the lesson learned here is that these keyword markets are highly interrelated, and it may take some exploration to discover the market with the best ROI for a given business objective or product.

Similarly, although advertisers do not pay when an ad is shown and users do not click on it, bidding for a large number of general keywords is not an effective strategy as your advertisement and site may not have a high relevance for a general keyword, and your CTR will probably be low. Recall that in quality-based bidding, both of those factors can negatively impact the final advertisement ranking, requiring the advertiser to increase their bid price to keep the same final ranking.

Finding the right set of “golden-nugget” keywords requires knowing the market and demographics for a given business or product. In addition, competitors’ web sites may help give insight into useful keywords. Determining the right set of keywords can also be used to upgrade a business’ web site thereby improving both its bidding quality score as well as its unpaid search placement results. Google (Google) and Yahoo! (Yahoo!) provide advice on keyword selection, guidelines for increasing the quality score of an ad or site, along with tools that can help track hits for individual keywords. For example, when advertising for a particular product, Google recommends sending the user directly to that product’s web page rather than to the site’s home page (Google). Negative keywords can also be used to further target the audience. If you are a chocolate company specializing in truffles, you may want to exclude searches on “mushroom truffles” (i.e., “-mushroom”).

Estimator tools that can be used without signing up for an advertising campaign include Google’s Adword Keyword (Google) and Traffic Estimator tools (Google). The Adword Keyword tool helps select keywords by displaying alternate potential keyword phrases, their cost and ad position estimates for a given bid, potential negative keywords, and

other related information. The Traffic Estimator tool provides estimates of ad position, clicks per day, and budget per day for multiple keywords. Figure 2 shows the Traffic Estimator output for the “buy real estate” and “real estate” keyword phrases given a \$1 maximum CPC, a \$20 daily budget, for the United States. The keyword phrase “buy real estate” receives fewer clicks than “real estate” for the above parameters, but to determine the better ROI, we would need to know the conversion rate.

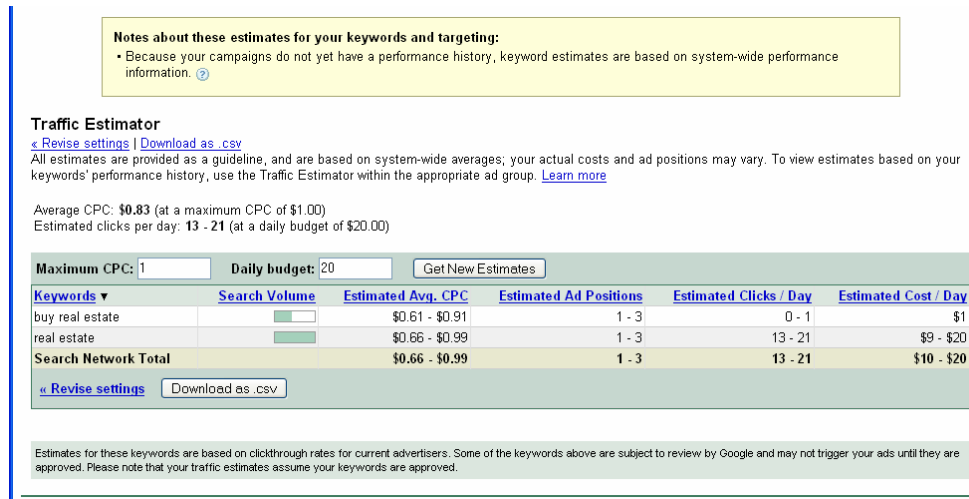


Figure 2: Google’s Adwords Traffic Estimator tool for “buy real estate” and “real estate”.

### *Monitor for click fraud*

Some recommendations from a short 10 day experiment by the Marketing Experiments Blog and Clicks2Customers for reducing click fraud include ("Click Fraud", 2005):

1. Monitor for variations in clicks and conversions in ad campaign performance.
2. Be conservative using ads placed on third-party websites (e.g., Google’s AdSense). Click fraud from search engines tends to be quite low, not so for third-party websites.
3. Only display ads in the geographic areas where products and services are sold.
4. Review server logs, including repeated IP addresses or IP ranges. If click fraud occurs mainly during certain times of day, then pause ad campaigns for those hours. Consider applications like KeywordMax, Urchin, Clicktracks, or companies like ClickDefense to track click fraud levels if the campaign size warrants it.
5. Do not simply bid to be number one, measure your results and try to maximize your ROI.
6. High traffic, expensive keywords tend to attract more fraud, try to find alternatives.

*Consider companies that setup, monitor and track ad campaigns.*

Depending on the size of an ad campaign and the in-house support available for creation and ongoing tracking of campaign success, a business may want to outsource their

marketing effort. The Buyer's Guide to Paid Search Ad (PPC) Agencies 2007 profiles 65 firms. Two major players include iProspect ([www.iprospect.com](http://www.iprospect.com)) and Clicks2Customers ([www.clicks2customers.com](http://www.clicks2customers.com)), and each provides substantial details on their respective websites about their services.

## **PREDICTION MARKETS AS DECISION SUPPORT TOOLS**

### ***Introduction***

Increasingly competitive market pressures, driven by globalization and enabled by information technology, have spawned new products and innovations at a rapid pace. Accurately assessing future demand for a new product, given non-existent or limited historical data, is challenging. Some companies are experimenting with a potential new addition to traditional forecasting tools called prediction markets. A *prediction market* is a financial market created for the sole purpose of producing forecasts, or predictions. Conventional wisdom holds that futures market prices can accurately forecast the likelihood of future events, in part as participants generally have a clear incentive to be honest and “put their money where their mouth is”. This insight has been backed up by more quantitative studies, including Roll's 1984 demonstration that the price of orange-juice futures is a better predictor of Florida weather than the National Weather Service (Roll, 1984). Companies such as Hewlett Packard, Google, and Lily are exploring the use of prediction markets as decision support tools to forecast sales, new product demand, or product release dates (King, 2006; Kiviat, 2004). By harnessing the distributed knowledge of their employees and other stakeholders, they hope to use market-driven consensus to arrive at an accurate forecast for questions that may not be easily addressed using other methodologies. Exactly when and where prediction markets can best be used as a decision support tool is still an open question.

### ***Existing markets***

Early prediction markets were set up to forecast events of interest to the general public. The Iowa Electronic Markets ([www.biz.uiowa.edu/iem/](http://www.biz.uiowa.edu/iem/)) is a limited real-money online futures market founded in 1988 that focuses on predictions of economic and political events, such as who will win presidential elections. Figure 3 shows the IEM daily midnight closing prices for the 2004 US Presidential Winner Takes All Market. Analysis of their trading data shows that prices in these markets generally predict the election outcomes better than polls (Berg, Forsythe, Nelson, & Rietz, 2001).




Figure 3: Kerry versus Bush Aggregate Probabilities from IEM

The Hollywood Stock Exchange ([www.hsx.com](http://www.hsx.com)) is a play money exchange that predicts how new films will perform at box office and trading data is sold to entertainment companies. Other examples include Yahoo! Tech Buzz ([buzz.research.yahoo.com](http://buzz.research.yahoo.com)), Tradesports ([www.tradesports.com](http://www.tradesports.com)) and NewsFutures ([us.newsfutures.com](http://us.newsfutures.com)).

Hewlett-Packard was one of the first to use internal prediction markets for enterprise forecasting, setting up markets to forecast computer workstation sales from 1996-1999 (K-Y. Chen & Plott, 2002). The traders were mostly product and finance managers, and their market forecasts came closer than official forecasts 6 out of 8 times. As the discipline matures, more companies such as Eli Lilly, Intel, and Microsoft are starting to use prediction markets for internal consumption (Kiviat, 2004). Google reports having run 143 markets so far, on questions from product launch dates to new office openings (Cowgill, 2005). Eli Lilly has experimented with drug development markets to predict the most successful drugs (Kiviat, 2004). Prediction markets have also been used as a means to provide early information on the expected benefits of a given aid project (Hahn & Tetlock, 2005) and to predict influenza outbreaks (Polgreen, Nelson, & Neumann, 2006). Third parties, such as NewsFutures and ConsensusPoint, currently provide commercial market platforms and support to companies wanting to run prediction markets.

Given ethical and legal constraints against real-money gambling, play money markets are often used for internal markets, which begs the question of whether market accuracy

results still apply. Several studies have shown that results are just as accurate (Feng, Bhargava, & Pennock, Forthcoming) and that the prices of securities in these play-money markets gave as accurate or more accurate predictions than expert opinions (Y. Chen, Chu, Mullen, & Pennock, 2005; Feng, Bhargava, & Pennock, Forthcoming; Pennock, Lawrence, Nielsen, & Giles, 2001). The key is to attract a knowledgeable and motivated pool of traders, and money is not the only means of doing that (Servan-Schreiber, Wolfers, Pennock, & Galebach, 2004). Play money markets typically offer participants small prizes along with public acknowledgement of their trading skills.

### *Simple enterprise market example*

Setting up a prediction market to capture this distributed knowledge requires first determining the appropriate market securities. A *security* pays some amount of money (real or play) contingent on future event(s). Consider the security “\$1 if sunny tomorrow”. A trader who purchases this security at thirty cents gets a profit of seventy cents if it is sunny tomorrow or else loses thirty cents if not. In the Iowa Electronic Markets model, traders initially purchase “unit portfolios”. For each \$1 a trader receives a basket of related securities. For example in the 2008 presidential election they might receive three contracts: “\$1 if Republican candidate wins”, “\$1 if Democratic candidate wins”, and “\$1 if another 3<sup>rd</sup> party wins”.

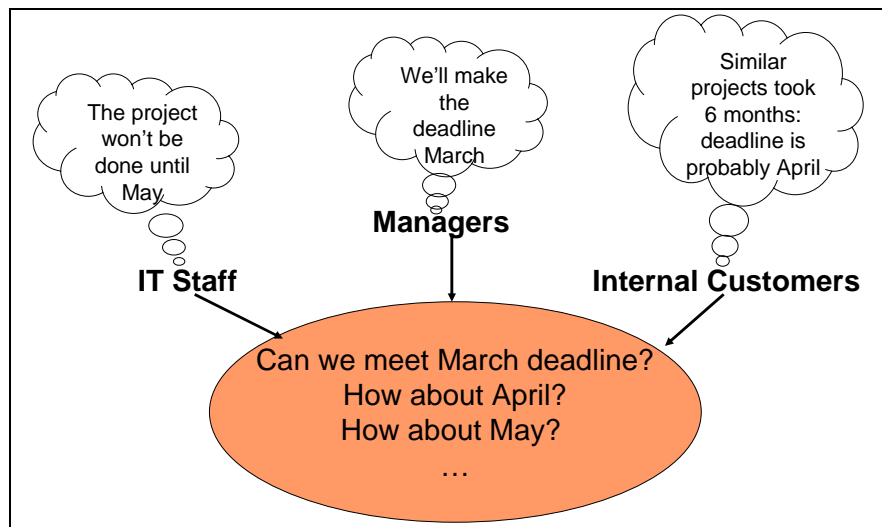


Figure 4: Harnessing distributed stakeholders' explicit and tacit knowledge

Moving to the enterprise domain, Figure 3 shows how different stakeholders' might think about an internal product release deadline. An example of a market security here might be “\$1 if product meets March deadline”, where “meets the deadline” would be specified with measurable criteria such as “implemented at least 90% of the features in the final requirements document”. However, there is more relevant information here to discover than just whether the product will meet the deadline. We might also want to know if the project will finish early (when has that ever happened?) or late. If it will finish late, then how late will it finish. To cover the entire space of possibilities, these securities can be sold in *bundles* or *baskets*: “\$1 if meet March deadline”, “\$1 if meet April deadline”, “\$1

if not done until May or after”. Each security would trade in a separate market, and following the IEM approach, each user would initially receive a bundle containing one of each security. A user who thought that the product would meet the April deadline with 100% probability would keep the “\$1 if meet April deadline” security and sell the other securities to achieve the greatest expected profit. A user who thought that the April deadline had a 75% probability of being met would keep the “\$1 if meet April deadline” security as long as the price in the April market was seventy-five cents or below and sell otherwise, and similarly for the other markets.

### ***Theory***

Theoretical and empirical research by economists shows that under certain conditions, markets aggregate less-than-perfect, diverse information across many individual traders, so that security prices summarize all the relevant information across traders into a *rational expectations equilibrium*, which is the assumption underlying the strong form of the *efficient markets hypothesis* in finance (Y. Chen, Mullen, & Chu, 2006; Y. Chen, Mullen, & Chu, 2004; Feigenbaum, Fortnow, Pennock, & Sami, 2003; Grossman, 1981; Jackwerth & Rubinstein, 1996; Plott & Sunder, 1988). Prediction markets are virtual security markets to aggregate information from a diverse set of traders. Depending on the context they are used, prediction markets are also called information markets, forecasting markets, decision markets, or game markets. The process of creating a prediction market starts with identifying a random variable of interest  $X$  (e.g., Product meets deadline?), then turning it into a financial instrument where the payoff is the realized value of variable (e.g., “Pays \$1 if Task completed by May 31<sup>st</sup>.”). If the security price incorporates all available information from market traders, it can be viewed as a consensus prediction about the value of the security reflecting not only individual trader beliefs, but also the effects of traders updating their beliefs in reaction to the market price and feeding this update back into the market, so that ideally, at time  $t$ ,  $price(t) \simeq E_t[X]$ . (Feng, Bhargava, & Pennock, Forthcoming). Recent preliminary theoretical work shows that a general interpretation of prediction market prices depends the joint distribution of trader’s beliefs, budgets, and risk preferences (Szymanski & Lee, 2006; Wolfers & Zitzewitz, 2005).

Compared with statistical forecasting methods, prediction markets can incorporate real-time information, which may not be contained in historical data. Conversely, price movements can be mined to indicate when new information has arrived into the system (Feng, Bhargava, & Pennock, Forthcoming). Compared with eliciting expert opinions, prediction markets are less constrained by space and time; they eliminate the effort of identifying experts and soliciting their participation, and they do not need to deal explicitly with conflicting opinions (French, 1985; Genest & Zidek, 1986). Most importantly, prediction markets make real-time predictions that take advantage of dispersed information, both explicit and tacit, which can sometimes be hard to capture using other forecasting methods. The market reaches its equilibrium only when all participants believe that the market price represents their predictions (otherwise they would trade).

Laboratory experiments have provided simplified environments for understanding how prediction market accuracy can be affected by a variety of factors such as market structure, aggregate uncertainty (i.e., even if traders pool their information freely, the outcome remains uncertain), and trader heterogeneity. In the case of market structure, negative effects can occur when the composition of available securities makes it difficult or too complex for trader to aggregate their information in the market (O'Brien & Srivastava, 1991; Plott & Sunder, 1988). Also, not surprisingly, outcomes with more aggregate uncertainty lead to less efficient information aggregation (Lundholm, 1991). However, the more heterogeneity of trader profiles, the more efficient the information aggregation (Forsythe & Lundholm, 1990). Sunder (Sunder, 1995) extensively summarized experimental work on information aggregation, concluding that identifying and understanding which factors facilitate or prevent information aggregation in these markets remains a key research issue.

---

### ***Recommendations***

Given the plethora of choices of market and asset design, we describe some of the strategic prediction market requirements that should drive subsequent tactical decisions. The first is to carefully select the events that the prediction market is to forecast. The second is to make the prediction markets simple and easy to use. The third is ensuring management buy-in and recruiting a diverse set of participants.

#### *Carefully select what to forecast and provide participants with proper incentives*

Determining the right data to collect to make a good forecast, and gathering it, can comprise a large portion of the forecasting effort. Prediction markets are no exception. The markets themselves require overhead to run and administer, while the participants must invest time to trade. The knowledge should be valuable enough to justify this time and effort. Hewlett Packard has modified their sales forecast markets to use a market-based method, rather than a prediction market per se, in part because a forecast accuracy increase of a few percentage points was not worth the effort of running markets continuously (Kay-Yut Chen, Fine, & Humberman, 2004; King, 2006).

Markets must also attract a reasonable number of participants with an interest in or knowledge of the forecasted events. At the simplest level, the more traders, the more market liquidity results in a more accurate market. In real-money markets, profit is a straightforward motivator. For play-money markets, while some will play for fun, most participants have to be provided with incentives to participate. The top few traders can be rewarded with prizes or their winnings can be converted to lottery tickets where the best traders have the best chance of winning lottery prizes (e.g., iPods, laptops, etc.). Another incentive is recognizing the top players through announcements or a "Top Five Traders" list.

Once the decision about what is worth forecasting is determined, then the appropriate market security must be designed. Our earlier example described a standard security of the form "\$1 if sunny tomorrow", but there are several variations. One possibility is that

if the early forecasting of events is critical, market creators can design a security with the appropriate weighted measures of performance, call *scoring rules*, that encourage early predictions (e.g., “\$1 if meet Feb deadline and predicted by Dec 31<sup>st</sup>, otherwise \$.90”).

Another example is conditional securities, which can be used for forecasts that involve assessing policy choices, such as whether to use development package *X* vs. *Y* (Berg & Rietz, 2003; Hanson, 1999). Thus given two conditional market securities: 1) “If using development package *X*, can meet product deadline”, and 2) “If using development package *Y*, can meet product deadline”, the relative value of the two asset prices reveals something about the confidence traders have in using development package *X* vs. *Y*. This can in turn be folded back into the enterprise decision making process. Of course, given that the results of these conditional markets may affect future decisions, such contacts will need to be set up to avoid self-defeating properties, so that they predict the outcome *given* the choices.

#### *Simplify user interfaces and interactions*

This requirement impacts both the user interface design and the market setup (e.g., markets should be easy to navigate, understand, and bid in). To keep the number of markets to a reasonable level, potential forecasting events should be prioritized. For example, not all deadlines should have prediction markets associated with them, only those with critical risk factors that can benefit from real-time monitoring.

Possible choices for prediction market mechanisms include continuous double auctions (CDAs), essentially stock market auctions, or pari-mutual markets similar to those used in horse betting. CDAs are used in exchanges such as Iowa Electronic Market, Hollywood Stock Exchange, Tradesports, NewsFutures, while dynamic pari-mutual markets can be seen at Yahoo! Tech Buzz. In a CDA, buyers and sellers continually place offers, and as soon as a buy offer is at least as great as a sell offer, a transaction occurs. A few exchanges use a market maker, or “virtual specialist”, who will always be willing to buy or sell, but it may lose money (although for fake money markets, this is not an issue). The Yahoo Tech Buzz pari-mutual auction also provides a market maker mechanism that ensures immediate trading, and does not lose money, but the interface interactions may not be as intuitive as the CDA. Some CDAs require traders to buy long or sell short to express their probabilities (i.e., if you don’t believe the product will meet the deadline, you don’t buy anything, you sell it short). However, unsophisticated and novice traders seem to feel uncomfortable with these concepts.

#### *Get buy-in from management and ensure a diverse set of participants*

In addition to aggregating information, internal prediction markets allow workers to anonymously telegraph concerns to upper management. Thus it may be easier for workers to trade honestly and anonymously in the market (and get rewarded by the market for correct predictions) than to tell their manager directly that sales may go down next quarter, or that critical deadlines may not be met. However, for the same reason, it can be difficult to get support from upper management to implement such markets. They

may not want other employees to know that deadlines may not be met, fearing that it may impact morale. As cited by Business Week online (2006): "The risk of creating markets is that it's not fudgible, and if you ask the question, then you better want to know the answer," says Robin Hanson, a professor in the department of economics at George Mason University (King, 2006). Clearly this is an area where market designers need to be very sensitive to management concerns. There are also concerns regarding whether market results would create a disclosure dilemma. As Google's Cowgill (2006) puts it "Prediction markets could make everyone in the company an insider from the SEC [Securities — Exchange Commission] perspective, if you don't manage it well" (King, 2006). Thus companies generally steer away from forecasting revenue or other financial indicators.

Diversity of knowledge and skills across participants is also important since prediction markets aggregate information best when they combine the forecasts of a heterogeneous set of traders (Forsythe & Lundholm, 1990). Market designers should actively seek to recruit developers, project managers, financial and sales people and other types of stakeholders. No one person, indeed, no one department can independently track the emerging factors involved in forecasting complex events.

### ***Implications for Practitioners (post-manuscript)***

1. Title: Novel Business-Oriented Auctions
2. Short Description of the Manuscript:

With the advent of the Internet, traditional auction forms have evolved to fit into a plethora business niches, either integrating into traditional approaches or simply creating new opportunities. This chapter examines two novel uses for auctions in a business context, namely sponsored search auctions and prediction markets. Understanding the potential auction benefits and limitations can hopefully provide practitioners with a more informed and successful approach when employing these auction-based tools in their business.
3. Findings (in bullet format)
  - Sponsored search is big business making Google and Yahoo! \$6 and \$5 billion dollars respectively in 2005.
  - Sponsored search allows advertisers to target their audience and only pay when users click on their advertisement.
  - Sponsored search auctions have evolved from standard first-price auctions into variations of second-price auctions that are more suitable to their commercial environment.
  - Click fraud is a major concern in sponsored search.
  - Academic research is just starting to evaluate sponsored search auction mechanism tradeoffs in terms of properties such as maximizing revenue, consumer welfare, price stability, and robustness to fraud.
  - A prediction market is a futures market where prices forecast the likelihood of future events, in part since participants generally have a clear incentive to be honest and "put their money where their mouth is".
  - Companies such as Hewlett-Packard, Eli Lilly, Intel, Microsoft, and Google have used prediction markets internally for enterprise forecasting.

- Prediction markets make real-time predictions that take advantage of dispersed information, both explicit and tacit, which can sometimes be hard to capture using other forecasting methods.
  - Prediction markets they should be viewed as an integral part of the whole decision support system rather than a stand-alone tool.
4. Recommendations (in bullet format) or how the practitioners should put the findings to use and how they should interpret them
    - To use sponsored search auctions to best effect 1) identify business objectives, 2) devote time to discovering the best keywords, 3) monitor for click fraud, and 4) consider using companies that setup, monitor and track ad campaigns.
    - To use prediction markets to best effect 1) carefully select what to forecast and provide participants with proper incentives, 2) keep user interfaces and interactions simple, and 3) get buy-in from management and ensure a diverse set of participants.
  5. Caveats or limitations of using the results in practice:
 

Both sponsored search auctions and prediction markets are relatively new and practitioners should expect a certain amount of learning about how best to use them in their specific business environment. In the case of prediction markets, while they offer an alternative and potentially powerful forecasting methodology, it may not be appropriate to all business environments.

### ***Future Research Directions***

In both sponsored search auctions and prediction markets, market design requires the consideration complex tradeoffs to work effectively. In the case of sponsored search auctions, the mechanisms have evolved from simple first-price auctions to GSP auctions to GSP plus quality-based bidding. Driven by commercial pressures, the only constant in this arena is likely to be rapid innovation and change for the near future. A variety of pricing mechanisms, appropriate for different business objectives and for reducing click fraud, are either being offered or planned. More sophisticated bidding software agents could assist advertisers in achieving their campaign objectives by implementing and monitoring more complex bidding strategies. Yahoo! is exploring how to targeting mobile devices with sponsored search, especially within a given geographic region (Yahoo!, 2006). However, there are numerous potential personalization and privacy issues involved in this extension (Yuan and Zhang, 2003; Minch 2004).

Business prediction markets are starting to move from internal ones to industry-specific ones. For example, StorageMarkets ([www.storagemarkets.com](http://www.storagemarkets.com)) focuses on using prediction markets to gather competitive intelligence about the IT Storage industry. However, as prediction markets get more widely used, there is a danger that the financial payoffs involved in market outcomes may outweigh the market gains themselves (Rhode & Strumpf, 2007). While market methods have shown initial promise as enterprise forecasting tools, they are not a panacea, and come with their own set of pros and cons. Instead of viewing them as a stand-alone tool, they should be viewed as an integral part of the whole decision support system (Berg & Rietz, 2003). Further research and experience is needed to help determine when the market outperforms more

traditional forecasting methods, and where it can be used synergistically with other decision support tools.

## REFERENCES

- Abrams, Z. (2006). *Revenue maximization when bidders have budgets* Paper presented at the seventeenth annual ACM-SIAM symposium on Discrete algorithms, Miami, Florida.
- Ausubel, L. M., & Milgrom, P. (2005). The Lovely but Lonely Vickrey Auction. In P. Crampton, R. Steinberg & Y. Shoham (Eds.), *Combinatorial Auctions*: MIT Press.
- Balcan, M.-F., Chen, J., Devanur, N., & Kumar, A. (2006). Transcript of Panel discussion: Models for Sponsored Search: What are the right questions? [Electronic Version]. *Second Workshop on Sponsored Search Auctions* Retrieved April 5, 2007.
- Berg, J. E., Forsythe, R., Nelson, F. D., & Rietz, T. A. (2001). Results from a dozen years of election futures markets research. In C. A. Plott & V. Smith (Eds.), *Handbook of Experimental Economic Results*.
- Berg, J. E., & Rietz, T. A. (2003). Prediction Markets as Decision Support Systems. *Information Systems Frontiers*, 5(1), 79-93.
- Borgs, C., Chayes, J., Etesami, O., Immorlica, N., Jain, K., & Mahdian, M. (2007, May 8-12, 2006). *Dynamics of bid optimization in online advertisement auctions*. Paper presented at the 16th International World Wide Web Conference (WWW2007), Banff, Alberta, Canada.
- Brandt, F., & Weiss, G. (2001). Antisocial agents and Vickrey auctions. In J.-J. C. Meyer & M. Tambe (Eds.), *Intelligent Agents VIII* (Vol. 2333, pp. pp. 335-347): Springer-Verlag.
- Chen, K.-Y., Fine, L., & Humberman, B. A. (2004). Eliminating Public Knowledge Biases in Information-Aggregation Mechanisms. *Management Science*, 50(7), pp. 983-994.
- Chen, K.-Y., & Plott, C. (2002). *Information Aggregation Mechanisms: Concept, Design, and Implementation for a Sales Forecasting Problem* (No. Social Science Working Paper 1131): California Institute of Technology.
- Chen, Y., Chu, C.-H., Mullen, T., & Pennock, D. M. (2005, June 5-8, 2005). *Information Markets vs. Opinion Pools: An Empirical Comparison*. Paper presented at the ACM Conference on Electronic Commerce (EC-05), Vancouver, British Columbia, Canada.
- Chen, Y., Mullen, T., & Chu, C.-H. (2006). An In-Depth Analysis of Information Markets with Aggregate Uncertainty. *Electronic Commerce Research*, 6, 201-220.
- Chen, Y., Mullen, T., & Chu, C. H. (2004, June 2004). *Theoretical investigation of prediction markets with aggregate uncertainty*. Paper presented at the 7th International Conference on Electronic Commerce Research (ICECR-7), Dallas, TX.
- Click Fraud [Electronic (2005). Version]. *Marketing Experiments Journal*. Retrieved April 5, 2007 from <http://www.marketingexperiments.com/internet-online-advertising/click-fraud.html>.

- Cowgill, B. (2005, September 21, 2005). Putting crowd wisdom to work. Retrieved April 5, 2007, 2007, from <http://googleblog.blogspot.com/2005/09/putting-crowd-wisdom-to-work.html>
- Edelman, B., & Ostrovsky, M. (2005, June 5-8, 2005). *Strategic Bidder Behavior in Sponsored Search Auctions*. Paper presented at the ACM Electronic Commerce Conference, Vancouver, British Columbia, Canada.
- Edelman, B., Ostrovsky, M., & Schwarz, M. (2005). Internet Advertising and the Generalized Second Price Auction: Selling Billions of Dollars Worth of Keywords.
- Feigenbaum, J., Fortnow, L., Pennock, D. M., & Sami, R. (2003, June). *Computation in a distributed information market*. Paper presented at the 4th Annual ACM Conference on Electronic Commerce (EC'03), San Diego, CA.
- Feng, J., Bhargava, H. K., & Pennock, D. M. (Forthcoming). Implementing sponsored search in web search engines: Computational evaluation of alternative mechanisms. *Inform Journal on Computing*.
- Forsythe, R., & Lundholm, F. (1990). Information aggregation in an experimental market. *Econometrica*, 58, 309–347.
- French, S. (1985). Group consensus probability distributions: a critical survey. *Bayesian Statistics*, 2, 183–202.
- Genest, C., & Zidek, J. V. (1986). Combining probability distributions: A critique and an annotated bibliography. *Statistical Science*, 1(1), 114-148.
- Goodman, J. (2005). *Pay-per-percentage of impressions: an advertising method that is highly robust to fraud*. Paper presented at the Workshop on Sponsored Search Auctions.
- Google. Adwords Help Center. Retrieved April 5, 2007, 2007, from <https://adwords.google.com/support/>
- Google. Keyword Tool. Retrieved April 5, 2007, 2007, from <https://adwords.google.com/select/KeywordToolExternal>
- Google. Tips for success. Retrieved April 5, 2007, 2007, from <https://adwords.google.com/select/tips.html>
- Google. Traffic Estimator. Retrieved April 5, 2007, 2007, from <https://adwords.google.com/select/TrafficEstimatorSandbox>
- Grossman, S. J. (1981). An introduction to the theory of rational expectations under asymmetric information. *Review of Economic Studies*, 48(4), 541-559.
- Hahn, R. W., & Tetlock, P. C. (2005). Making Development Work. *Hoover Institution, Policy Review*, 132.
- Hanson, R. D. (1999). Decision Markets. *IEEE Intelligent Systems*, 16-19.
- Immorlica, N., Jain, K., Mahdian, M., & Talwar, K. (2005). Paper presented at the Workshop on Internet and Network Economics (WINE).
- Jackwerth, J. C., & Rubinstein, M. (1996). Recovering probability distribution from options prices. *Journal of Finance*, 51(5), 1611-1631.
- Jansen, B. (2006). Paid Search. *IEEE Computer*, 88-90.
- King, R. (2006). Workers, place your bets [Electronic Version]. *Business Week online*. Retrieved April 5, 2007 from [http://www.businessweek.com/technology/content/aug2006/tc20060803\\_012437.htm](http://www.businessweek.com/technology/content/aug2006/tc20060803_012437.htm).

- Kiviat, B. (2004, July 06). The End Of Management? *Time*.
- Krishna, V. (2002). *Auction Theory*: Academic Press.
- Lahaie, S. (2006). *An analysis of alternative slot auction designs for sponsored search*. Paper presented at the Seventh ACM conference on Electronic commerce Ann Arbor, MI.
- Liu, D., Chen, J., & Whinston, A. B. (2006). *Weighted Unit-Price Auctions*. Paper presented at the Second Workshop on Sponsored Search Auctions, Ann Arbor, MI.
- Lundholm, R. (1991). What affects the efficiency of the market? Some answers from the laboratory. *The Accounting Review*, 66, 486–515.
- O'Brien, J., & Srivastava, S. (1991). Dynamic stock markets with multiple assets: An experimental analysis. *Journal of Finance*, 46, 1811–1838.
- Pennock, D. M., Lawrence, S., Nielsen, F. A., & Giles, C. L. (2001). *Extracting collective probabilistic forecasts from web games*. Paper presented at the 7th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Francisco, CA.
- Plott, C., & Sunder, S. (1988). Rational expectations and the aggregation of diverse information in laboratory security markets. *Econometrica*, 56, 1085–1118.
- Polgreen, P. M., Nelson, F. D., & Neumann, G. R. (2006). Using Prediction Markets to Forecast Trends in Infectious Diseases. *Microbe*, 1(10), pp. 459-465.
- Rhode, P. W., & Strumpf, K. S. (2007). *Manipulating Political Stock Markets: A Field Experiment and a Century of Observational Data* University of Arizona.
- Roll, R. (1984). Orange juice and weather. *American Economic Review*, 74, pp. 861-880.
- Servan-Schreiber, E., Wolfers, J., Pennock, D. M., & Galebach, B. (2004). Prediction Markets: Does money matter? *Electronic Markets*, 14(3).
- Story, L. (2006, October 30, 2006). Marketers Demanding Better Count of the Clicks. *New York Times*.
- Sunder, S. (1995). Experimental asset markets. In J. H. K. a. A. E. Roth (Ed.), *The Handbook of Experimental Economics* (pp. 445–500): Princeton University Press.
- Szymanski, B. K., & Lee, J.-S. (2006, June 11, 2006). *Impact of ROI on Bidding and Revenue in Sponsored Search Advertisement Auctions*. Paper presented at the Second Workshop on Sponsored Search Auctions, Ann Arbor, MI.
- Tuzhilin, A. (2006). *The Lane's Gifts v. Google Report*.
- Varian, H. R. (2006). *Position Auctions*: UC Berkeley.
- Vickrey, W. (1961). Counterspeculation, Auctions, and Competitive Sealed Tenders. *Journal of Finance*, 16, pp. 8–27.
- Wolfers, J., & Zitzewitz, E. (2005). *Interpreting Prediction Market Prices as Probabilities*: Wharton School, University of Pennsylvania.
- Yahoo! (April 5, 2007). Search Marketing Overview: Quality Index. Retrieved April 5, 2007, 2007, from [http://help.yahoo.com/l/us/yahoo/ysm/sps/start/overview\\_qualityindex.html](http://help.yahoo.com/l/us/yahoo/ysm/sps/start/overview_qualityindex.html)
- Yahoo! Search Marketing: Sponsored Search. Retrieved April 5, 2007, 2007, from <http://searchmarketing.yahoo.com/srch/index.php>
- Yahoo! (2006, August 16, 2006). Yahoo! and go2 Sign Mobile Search Advertising Distribution Agreement; Yahoo! Advertisers to Reach Consumers Through go2's

Mobile Local Content Channels. Retrieved April 5, 2007, 2007, from <http://yhoo.client.shareholder.com/press/ReleaseDetail.cfm?ReleaseID=207538>

## **ADDITIONAL READING**

Asdemir, K. (2006). *Bidding Patterns in Search Engine Auctions*. Paper presented at the Second Workshop on Sponsored Search Auctions, Ann Arbor, MI.

Camerer, C. (1998). Can Asset Markets be Manipulated? A Field Experiment with Racetrack Betting. *Journal of Political Economy*, 106, 457-82.

Even-Dar, E., Kearns, M., & Wortman, J. (2007) *Sponsored Search with Contexts*. Paper presented at the Third Workshop on Sponsored Search Auctions, Banff, Canada.

Giles, J. (2005). Wisdom of the crowd: decision makers, wrestling with thorny choices, are tapping into the collective foresight of ordinary people. *Nature*, 438, 281.

Hansen, J., Schmidt, C. & Strobel, M. (2004). Manipulation in Political Stock Markets—Preconditions and Evidence. *Applied Economics Letters*, 11, 459-463.

Hanson, R. (1999). Decision Markets. *IEEE Intelligent Systems*, 14, 16-19.

Hanson, R., Oprea, R. & Porter, D. (2006). Information Aggregation and Manipulation in an Experimental Market. *Journal of Economic Behavior and Organization*, 60(4), 449-459.

Kitts, B., Laxminar, P., LeBlanc, B. J., & Meech R. (2005). *A formal analysis of search auctions including predictions on click fraud and bidding tactics*. Paper presented at the Second Workshop on Sponsored Search Auctions.

Kitts, B. & LeBlanc, B. J.. (2004) *A trading agent and simulator for keyword auctions*. Paper presented at the Third International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2004), New York, NY.

Malone, T. W. (2004). *The Future of Work*. Harvard Business School Press.

Mehta, A., Saberi, A., Vazirani, U., & Vazirani V. (2005) *AdWords and Generalized Online Matching*. Paper presented at Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science (FOCS).

Milgrom, P. (2004). *Putting Auction Theory to Work*. Cambridge University Press.

- Oliven, K. & Rietz, T. (2004). Suckers are Born but Markets Are Made: Individual Rationality, Arbitrage, and Market Efficiency on an Electronic Futures Market. *Management Science*, 50, 336-351.
- Pennock, D., Lawrence, S., Giles, C. L., & Nielsen, F. A. (2001). The real power of artificial markets. *Science*, 291, 987-988.
- Pennock, D. (2004). *A Dynamic Pari-Mutuel Market for Hedging, Wagering, and Information Aggregation*. Paper presented at the 5th ACM conference on Electronic commerce, New York.
- Rhode, P. W. & Strumpf K. (2004). Historical Presidential Betting Markets. *Journal of Economic Perspectives*, 18(2), 127-142.
- Surowiecki, J. (2004). *The Wisdom of Crowds*. New York: Random House, Inc.
- Wolfers, J. & Zitzewitz, E. (2004). Prediction Markets. *Journal of Economic Perspective*, 18, 107-126.
- Zhang, X.M. and J. Feng. (2005) *Price Cycles in Online Advertising Auctions*. Paper presented at the 26th International Conference on Information Systems (ICIS), Las Vegas, NV.
- Zhou, Y. & Lukose, R. (2006) *Vindictive Bidding in Keyword Auctions*. Paper presented at the Second Workshop on Sponsored Search Auctions, Ann Arbor, MI.