

American Economic Association

Internet Advertising and the Generalized Second-Price Auction: Selling Billions of Dollars Worth of Keywords

Author(s): Benjamin Edelman, Michael Ostrovsky, Michael Schwarz

Source: *The American Economic Review*, Vol. 97, No. 1 (Mar., 2007), pp. 242-259

Published by: [American Economic Association](#)

Stable URL: <http://www.jstor.org/stable/30034393>

Accessed: 31/08/2011 11:23

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



American Economic Association is collaborating with JSTOR to digitize, preserve and extend access to *The American Economic Review*.

<http://www.jstor.org>

Internet Advertising and the Generalized Second-Price Auction: Selling Billions of Dollars Worth of Keywords

By BENJAMIN EDELMAN, MICHAEL OSTROVSKY, AND MICHAEL SCHWARZ*

We investigate the “generalized second-price” (GSP) auction, a new mechanism used by search engines to sell online advertising. Although GSP looks similar to the Vickrey-Clarke-Groves (VCG) mechanism, its properties are very different. Unlike the VCG mechanism, GSP generally does not have an equilibrium in dominant strategies, and truth-telling is not an equilibrium of GSP. To analyze the properties of GSP, we describe the generalized English auction that corresponds to GSP and show that it has a unique equilibrium. This is an ex post equilibrium, with the same payoffs to all players as the dominant strategy equilibrium of VCG. (JEL D44, L81, M37)

This paper investigates a new auction mechanism, which we call the “generalized second-price” auction, or GSP. GSP is tailored to the unique environment of the market for online ads, and neither the environment nor the mechanism has previously been studied in the mechanism design literature. While studying the properties of a novel mechanism is often fascinating in itself, our interest is also motivated by the spectacular commercial success of GSP. It is the dominant transaction mechanism in a large and rapidly growing industry. For example, Google’s total revenue in 2005 was \$6.14 billion. Over 98 percent of its revenue came from GSP auctions. Yahoo!’s total revenue in 2005 was \$5.26 billion. A large share of Yahoo!’s revenue is derived from sales via GSP auctions. It is believed that over half of Yahoo!’s revenue is derived from sales via GSP auctions. As of May 2006, the combined market capitalization of these companies exceeded \$150 billion.

Let us briefly describe how these auctions work. When an Internet user enters a search

term (“query”) into a search engine, he gets back a page with results, containing both the links most relevant to the query and the sponsored links, i.e., paid advertisements. The ads are clearly distinguishable from the actual search results, and different searches yield different sponsored links: advertisers target their ads based on search keywords. For instance, if a travel agent buys the word “Hawaii,” then each time a user performs a search on this word, a link to the travel agent will appear on the search results page. When a user clicks on the sponsored link, he is sent to the advertiser’s Web page. The advertiser then pays the search engine for sending the user to its Web page, hence the name—“pay-per-click” pricing.

The number of ads that the search engine can show to a user is limited, and different positions on the search results page have different desirabilities for advertisers: an ad shown at the top of a page is more likely to be clicked than an ad shown at the bottom. Hence, search engines need a system for allocating the positions to advertisers, and auctions are a natural choice. Currently, the mechanisms most widely used by search engines are based on GSP.

In the simplest GSP auction, for a specific keyword, advertisers submit bids stating their maximum willingness to pay for a click. When a user enters a keyword, he receives search results along with sponsored links, the latter shown in decreasing order of bids. In particular,

* Edelman: Department of Economics, Harvard University, Cambridge, MA 02138 (e-mail: bedelman@fas.harvard.edu); Ostrovsky: Graduate School of Business, Stanford University, Stanford, CA 94305 (e-mail: ostrovsky@gsb.stanford.edu); Schwarz: Yahoo! Research, 1950 University Ave., Suite 200, Berkeley, CA 94704 (e-mail: mschwarz@yahoo-inc.com). We thank Drew Fudenberg, Louis Kaplow, Robin Lee, David McAdams, Paul Milgrom, Muriel Niederle, Ariel Pakes, David Pennock, and Al Roth for helpful discussions.

the ad with the highest bid is displayed at the top, the ad with the next highest bid is displayed in the second position, and so on. If a user subsequently clicks on an ad in position i , that advertiser is charged by the search engine an amount equal to the next highest bid, i.e., the bid of an advertiser in position $(i + 1)$. If a search engine offered only one advertisement per result page, this mechanism would be equivalent to the standard second-price auction, coinciding with the Vickrey-Clarke-Groves (VCG) mechanism (William Vickrey 1961; Edward H. Clarke 1971; Theodore Groves 1973), auction. With multiple positions available, GSP generalizes the second-price auction (hence the name). Here, each advertiser pays the next highest advertiser's bid. But as we will demonstrate, the multi-unit GSP auction is no longer equivalent to the VCG auction and lacks some of VCG's desirable properties. In particular, unlike the VCG mechanism, GSP generally does not have an equilibrium in dominant strategies, and truth-telling is not an equilibrium of GSP.

In Section I, we describe the evolution of the market for Internet advertisements and the unique features of the environment in this market. In Section II, we introduce a model of sponsored search auctions, and we begin our analysis of the model in Section III. Since advertisers can change their bids frequently, sponsored search auctions can be modeled as a continuous or an infinitely repeated game. By the folk theorem, however, such a game will have an extremely large set of equilibria, and so we focus instead on the one-shot, simultaneous-move, complete information stage game, introducing restrictions on advertisers' behavior suggested by the market's dynamic structure. We call the equilibria satisfying these restrictions "locally envy-free."

We then proceed to show that the set of locally envy-free equilibria contains an equilibrium in which the payoffs of the players are the same as in the dominant-strategy equilibrium of the VCG auction, even though both the bids of the players and the payment rules in the mechanisms are very different. Moreover, this equilibrium is the worst locally envy-free equilibrium for the search engine and the best locally envy-free equilibrium for the advertisers. Consequently, in *any* locally envy-free equilibrium of GSP, the total ex-

pected revenue to the seller is at least as high as in the dominant-strategy equilibrium of the VCG auction.

In Section IV, we present our main result. We introduce the generalized English auction with independent private values, which corresponds to the generalized second-price auction and is meant to capture the convergence of bidding behavior to the static equilibrium, in the same spirit as tâtonnement processes in the theory of general equilibrium and the deferred-acceptance salary adjustment process in the theory of matching in labor markets. The generalized English auction has several notable features. Although it is not dominant-strategy solvable, it has a unique, perfect Bayesian equilibrium in continuous strategies. In this equilibrium, all players receive VCG payoffs. Moreover, this equilibrium is *ex post*, i.e., even if a particular player learned the values of other players before the game, he would not want to change his strategy. This, in turn, implies that the equilibrium is robust, i.e., it does not depend on the underlying distribution of values: the profile of strategies that we identify is an *ex post* Bayesian Nash equilibrium for any set of distributions of advertisers' private values.

There are several recent theoretical and empirical papers related to sponsored search auctions. Gagan Aggarwal and Jason D. Hartline (2005), Aranyak Mehta et al. (2005), and Mohammad Mahdian, Hamid Nazerzadeh, and Amin Saberi (2006) propose computationally fast, near-optimal mechanisms for pricing and allocating slots to advertisers in the presence of budget constraints and random shocks. Christopher Meek, David M. Chickering, and David B. Wilson (2005) describe incentive-compatible auctions with stochastic allocation rules, generalizing Vickrey auctions, and argue that such auctions can be useful for selling Internet advertising despite being inefficient. Note that, in contrast to these papers, we study the mechanisms actually used by the search engines.

Xiaoquan Zhang (2005), Kursad Asdemir (2006), and Edelman and Ostrovsky (forthcoming) present empirical evidence of bid and ranking fluctuations in both generalized first-price and generalized second-price auctions. They argue that history-dependent strategies can give rise to such fluctuations. However, Hal R. Varian (forthcoming) empirically analyzes GSP

auction data from Google and reports that locally envy-free Nash equilibria “describe the basic properties of the prices observed in Google’s ad auction reasonably accurately.”¹

I. The Structure and Evolution of Sponsored Search Auctions

A. Notable Features of the Market for Internet Advertising

A combination of features makes the market for Internet advertising unique. First, bids can be changed at any time. An advertiser’s bid for a particular keyword will apply every time that keyword is entered by a search engine user, until the advertiser changes or withdraws the bid. For example, the advertiser with the second highest bid on a given keyword at some instant will be shown as the second sponsored link to a user searching for that keyword at that instant. The order of the ads may be different next time a user searches for that keyword, because the bids could have changed in the meantime.²

Second, search engines effectively sell flows of perishable advertising services rather than storable objects: if there are no ads for a particular search term during some period of time, the “capacity” is wasted.

Finally, unlike other centralized markets, where it is usually clear how to measure what is being sold, there is no “unit” of Internet advertising that is natural from the points of view of all involved parties. From the advertiser’s perspective, the relevant unit is the cost of attracting a customer who makes a purchase. This corresponds most directly to a pricing model in which an advertiser pays only when a customer actually completes a transaction. From the search engine’s perspective, the relevant unit is what it collects in revenues every time a user performs a search for a particular keyword. This

corresponds to a pricing model in which an advertiser is charged every time its link is shown to a potential customer. “Pay-per-click” is a middle ground between the two models: the advertiser pays every time a user clicks on the link. All three payment models are widely used on the Internet.³ The specific sector of Internet advertising that we study, sponsored search auctions, has converged to pay-per-click pricing.

Since GSP evolved in the market for online advertising, its rules reflect the environment’s unique characteristics. GSP insists that for each keyword, advertisers submit a single bid—even though several different items are for sale (different advertising positions). GSP’s unusual one-bid requirement makes sense in this setting: the value of being in each position is proportional to the number of clicks associated with that position; the benefit of placing an ad in a higher position is that the ad is clicked more, but the users who click on ads in different positions are assumed to have the same values to advertisers (e.g., the same purchase probabilities). Consequently, even though the GSP environment is multi-object, buyer valuations can be adequately represented by one-dimensional types. For some advertisers, one bid per keyword may not be sufficiently expressive to fully convey preferences. For example, a single bid ignores the possibility that users who click on position 5 are somehow different from those who click on position 2; it does not allow for the possibility that advertisers care about the allocation of other positions, and so on. Nonetheless, these limitations are apparently not large enough to justify added complexity in the bidding language. Nico Brooks (2004) finds only moderate differences in purchase probabilities when ads are shown in different positions. Following search engines’ approaches and Brooks’ empirical findings, we likewise assume the value of a click is the same in all positions.

¹ Varian discovered envy-free Nash equilibria independently and called them “Symmetric Nash Equilibria” in his paper.

² For manual bidding through online advertiser centers, both Google and Yahoo! allow advertisers to make unlimited changes. In contrast, the search engines impose restrictions on the behavior of software bidding agents: e.g., Yahoo! limits the number of times an advertiser can change his bid in a given period of time.

³ A prominent example of “pay-per-transaction,” and even “pay-per-dollar of revenue” (“revenue sharing”), is Amazon.com’s Associates Program, www.amazon.com/gp/browse.html?&node=3435371 (accessed June 10, 2006). Under this program, a Web site that sends customers to Amazon.com receives a percentage of customers’ purchases. “Pay-per-impression” advertising, in the form of banner ads, remains popular on major Internet portals, such as yahoo.com, msn.com, and aol.com.

One important possibility that we abstract away from is that advertisers differ along dimensions other than per-click value, i.e., have different probabilities of being clicked when placed in the same position. (These probabilities are known in the industry as “click-through rates,” or CTRs.) Different search engines treat this possibility differently. Yahoo! ignores the differences, ranks the advertisers purely in decreasing order of bids, and charges the next-highest advertiser’s bid.⁴ Google multiplies each advertiser’s bid by its “quality score,” which is based on CTR and other factors, to compute its “rank number,” ranks the ads by rank numbers, and then charges each advertiser the smallest amount sufficient to exceed the rank number of the next advertiser.⁵ In our analysis, we assume that all advertisers are identical along dimensions other than per-click value, which eliminates this difference between the mechanisms used at Google and Yahoo!. As we discuss at the end of Section III, the analysis would remain largely the same if there were advertiser-specific differences in CTRs and “quality scores,” although the equilibria under Google and Yahoo! mechanisms would not be identical.^{6,7}

⁴ See <http://help.yahoo.com/help/us/performance/customer/dtc/bidding/>, link to “How do I figure out my cost?” and searchmarketing.yahoo.com/srch/index.php, link to “How Sponsored Search Works” (accessed June 10, 2006).

⁵ See www.google.com/adwords/learningcenter/#section1, links to “Ad Ranking” and “Cost Control” (accessed June 10, 2006). Initially, Google’s pricing mechanism was more transparent: quality score was equal to the estimated click-through rate.

⁶ The analysis would have to change considerably if there were specific advertiser-position effects. The magnitude of these advertiser-position effects is ultimately an empirical question, and we do not have the kind of data that would allow us to answer it; however, judging from the fact that the two major search engines effectively ignore it in their mechanisms (Yahoo! ignores CTRs altogether; Google computes an advertiser’s estimated CTR conditional on the advertiser attaining the first position), we believe it to be small.

⁷ Another important difference between the search engines’ implementations of GSP is the amount of information available to the advertisers. On Yahoo!, advertisers can directly observe the bids of their competitors (uv.bidtool.overture.com/d/search/tools/bidtool). On Google, they cannot. For any keyword and any bid amount, however, they can get an estimated average position and average cost-per-click they can expect (adwords.google.com/select/KeywordToolExternal, “Cost and ad position estimates”).

B. Evolution of Market Institutions

The history of sponsored search auctions is of interest as a case study of whether, how, and how quickly markets come to address their structural shortcomings. Many important mechanisms have recently been designed essentially from scratch, entirely replacing completely different historical allocation mechanisms: radio spectrum auctions (Paul Milgrom 2000; Ken Binmore and Paul Klemperer 2002), electricity auctions (Robert Wilson 2002), and others. In contrast, reminiscent of the gradual evolution of medical residency match rules (Alvin E. Roth 1984), sponsored search ad auctions have evolved in steps over time. In both medical residency and search advertising, flawed mechanisms were gradually replaced by increasingly superior designs. Notably, the Internet advertising market evolved much faster than the medical matching market. This may be due to the competitive pressures on mechanism designers present in the former but not in the latter, much lower costs of entry and experimentation, advances in the understanding of market mechanisms, and improved technology.

We proceed with a brief chronological review of the development of sponsored search mechanisms.

Early Internet Advertising.—Beginning in 1994, Internet advertisements were largely sold on a per-impression basis. Advertisers paid flat fees to show their ads a fixed number of times (typically, 1,000 showings or “impressions”). Contracts were negotiated on a case-by-case basis, minimum contracts for advertising purchases were large (typically, a few thousand dollars per month), and entry was slow.⁸

Generalized First-Price Auctions.—In 1997, Overture (then GoTo; now part of Yahoo!) introduced a completely new model of selling

From this information, they can back out estimates of their competitors’ rank numbers. Moreover, advertisers can experiment with changing their bids, which can also give them independent and relatively accurate estimates. Web sites accessed on June 10, 2006.

⁸ See http://www.worlddata.com/wdnet8/articles/the_history_of_Internet_Advertising.htm and www.zakon.org/robert/internet/timeline (both accessed June 10, 2006).

Internet advertising. In the original Overture auction design, each advertiser submitted a bid reporting the advertiser's willingness to pay on a per-click basis, for a particular keyword. The advertisers could now target their ads: instead of paying for a banner ad that would be shown to everyone visiting a Web site, advertisers could specify which keywords were relevant to their products and how much each of those keywords (or, more precisely, a user clicking on their ad after looking for that keyword) was worth to them. Also, advertising was no longer sold per 1,000 impressions; rather, it was sold one click at a time. Every time a consumer clicked on a sponsored link, an advertiser's account was automatically billed the amount of the advertiser's most recent bid. The links to advertisers were arranged in descending order of bids, making highest bids the most prominent. The ease of use, the very low entry costs, and the transparency of the mechanism quickly led to the success of Overture's paid search platform as the advertising provider for major search engines, including Yahoo! and MSN. However, the underlying auction mechanism itself was far from perfect. In particular, Overture and advertisers quickly learned that the mechanism was unstable due to the fact that bids could be changed very frequently.

Example. Suppose there are two slots on a page and three advertisers. An ad in the first slot receives 200 clicks per hour, while the second slot gets 100. Advertisers 1, 2, and 3 have values per click of \$10, \$4, and \$2, respectively. Suppose advertiser 2 bids \$2.01, to guarantee that he gets a slot. Then advertiser 1 will not want to bid more than \$2.02—he does not need to pay more than that to get the top spot. But then advertiser 2 will want to revise his bid to \$2.03 to get the top spot, advertiser 1 will in turn raise his bid to \$2.04, and so on. Clearly, there is no pure strategy equilibrium in the one-shot version of the game, and so if advertisers best respond to each other, they will want to revise their bids as often as possible.

Hence, advertisers will make socially inefficient investments into bidding robots, which can also be detrimental for the revenues of search engines. David McAdams and Schwarz (forthcoming) argue that in various settings, the costs that buyers incur while trying to “game” an auction mechanism are fully passed through to the

seller. Moreover, if the “speed” of the robots varies across advertisers, revenues can be very low even if advertisers' values are high. For instance, in the example above, suppose advertiser 1 has a robot that can adjust the bid very quickly, while advertisers 2 and 3 are humans and can change their bids at most once a day. In this case, as long as advertiser 3 does not bid more than his value, the revenues of a search engine are at most \$2.02 per click. Indeed, suppose advertiser 3 bids \$2.00. If advertiser 2 bids \$2.01, he will be in the second position paying \$2.01. If he bids any amount greater than that but lower than his value, he will remain in the second position and will pay more per click, because the robot of advertiser 1 will quickly outbid him. The revenue would not change even if the values of advertisers 1 and 2 were much higher.

Generalized Second-Price Auctions.—Under the generalized first-price auction, the advertiser who could react to competitors' moves fastest had a substantial advantage. The mechanism therefore encouraged inefficient investments in gaming the system, causing volatile prices and allocative inefficiencies. Google addressed these problems when it introduced its own pay-per-click system, AdWords Select, in February 2002. Google also recognized that an advertiser in position i will never want to pay more than one bid increment above the bid of the advertiser in position $(i + 1)$, and adopted this principle in its newly designed generalized second-price auction mechanism. In the simplest GSP auction, an advertiser in position i pays a price per click equal to the bid of an advertiser in position $(i + 1)$ plus a minimum increment (typically \$0.01). This second-price structure makes the market more user friendly and less susceptible to gaming.

Recognizing these advantages, Yahoo!/Overture also switched to GSP. Let us describe the version of GSP that it implemented.⁹ Every

⁹ We focus on Overture's implementation, because Google's system is somewhat more complex. Google adjusts effective bids based on ads' click-through rates and other factors, such as “relevance.” But under the assumption that all ads have the same relevance and click-through rates conditional on position, Google's and Yahoo!'s versions of GSP are identical. As we show in Section III, it is straightforward to generalize our analysis to Google's mechanism.

advertiser submits a bid. Advertisers are arranged on the page in descending order of their bids. The advertiser in the first position pays a price per click that equals the bid of the second advertiser plus an increment; the second advertiser pays the bid of the third advertiser plus an increment; and so forth.

Example (continued). Let us now consider the payments in the environment of the previous example under the GSP mechanism. If all advertisers bid truthfully, then bids are \$10, \$4, and \$2. Payments in GSP will be \$4 and \$2.¹⁰ Truth-telling is indeed an equilibrium in this example, because no advertiser can benefit by changing his bid. Note that total payments of advertisers 1 and 2 are \$800 and \$200, respectively.

Generalized Second-Price and VCG Auctions.—GSP looks similar to the VCG mechanism, because both mechanisms set each agent's payments based only on the allocation and bids of other players, not based on that agent's own bid. In fact, Google's advertising materials explicitly refer to Vickrey and state that Google's "unique auction model uses Nobel Prize-winning economic theory to eliminate ... that feeling that you've paid too much."¹¹ But GSP is not VCG.¹² In particular, unlike the VCG auction, GSP does not have an equilibrium in dominant strategies, and truth-telling is generally not an equilibrium strategy in GSP (see the example in Remark 3 in Section II). With only one slot, VCG and GSP would be identical. With several slots, the mechanisms are different. GSP charges the advertiser in position i the *bid* of the advertiser in position $i + 1$. In contrast, VCG charges the advertiser in position i the *externality* that he imposes on others by taking one of the slots away from them: the total payment of the advertiser in

position i is equal to the difference between the aggregate value of clicks that all other advertisers would have received if i were not present in the market and the aggregate value of clicks that all other advertisers receive when i is present. Note that an advertiser in position $j < i$ is not affected by i , and so the externality i imposes on her is zero, while an advertiser in position $j > i$ would have received position $(j - 1)$ in the absence of i , and so the externality i imposes on her is equal to her value per click multiplied by the difference in the number of clicks in positions j and $(j - 1)$.

Example (continued). Let us compute VCG payments for the example considered above. The second advertiser's payment is \$200, as before. However, the payment of the first advertiser is now \$600: \$200 for the externality that he imposes on advertiser 3 (by forcing him out of position 2) and \$400 for the externality that he imposes on advertiser 2 (by moving him from position 1 to position 2 and thus causing him to lose $(200 - 100) = 100$ clicks per hour). Note that in this example, revenues under VCG are lower than under GSP. As we will show later (Remark 1 in Section II), if advertisers were to bid their true values under both mechanisms, revenues would always be higher under GSP.

C. Assessing the Market's Development

The chronology above suggests three major stages in the development of the sponsored search advertising market. First, ads were sold manually, slowly, in large batches, and on a cost-per-impression basis. Second, Overture implemented keyword-targeted per-click sales and began to streamline advertisement sales with some self-serve bidding interfaces, but with a highly unstable first-price mechanism. Next, Google implemented the GSP auction, which was subsequently adopted by Overture (Yahoo!).

Interestingly, Google and Yahoo! still use GSP, rather than VCG, which would reduce incentives for strategizing and make life easier for advertisers. We see several possible reasons for this. First, VCG is hard to explain to typical advertising buyers. Second, switching to VCG may entail substantial transition costs: VCG

¹⁰ For convenience, we neglect the \$0.01 minimum increments.

¹¹ See <https://www.google.com/adsense/afs.pdf> (accessed June 10, 2006).

¹² Roth and Axel Ockenfels (2002) describe another example in which the architects of an auction may have tried to implement a mechanism strategically equivalent to the Vickrey auction, but did not get an important part of the mechanism right.

revenues are lower than GSP revenues for the same bids, and advertisers might be slow to stop shading their bids. Third, the revenue consequences of switching to VCG are uncertain: even the strategic equivalence of second-price and English auctions under private values fails to hold in experiments (John Kagel, Ronald M. Harstad, and Dan Levin 1987). And, of course, simply implementing and testing a new system may be costly—imposing switching costs on advertisers as well as on search engines.

II. The Rules of GSP

Let us now formally describe the rules of a sponsored search auction. For a given keyword, there are N objects (positions on the screen, where ads related to that keyword can be displayed) and K bidders (advertisers).¹³ The (expected) number of clicks per period received by the advertiser whose ad was placed in position i is α_i . The value per click to advertiser k is s_k . Advertisers are risk-neutral, and advertiser k 's payoff from being in position i is equal to $\alpha_i s_k$ minus his payments to the search engine. Note that these assumptions imply that the number of times a particular position is clicked does not depend on the ads in this and other positions, and also that an advertiser's value per click does not depend on the position in which its ad is displayed. Without loss of generality, positions are labeled in descending order: for any i and j such that $i < j$, we have $\alpha_i > \alpha_j$.

We model the GSP auction as follows. Suppose at some time t a search engine user enters a given keyword, and, for each k , advertiser k 's last bid submitted for this keyword prior to t was b_k ; if advertiser k did not submit a bid, we set $b_k = 0$. Let $b^{(j)}$ and $g(j)$ denote the bid and identity of the j -th highest advertiser, respectively. If several advertisers submit the same bid, they are ordered randomly.¹⁴ The mechanism then allocates the top position to the advertiser with the highest bid, $g(1)$, the second

position to $g(2)$, and so on, down to position $\min\{N, K\}$. Note that each advertiser gets at most one object. If a user clicks on an advertiser's link, the advertiser's payment per click is equal to the next advertiser's bid. So advertiser $g(i)$'s total payment $p^{(i)}$ is equal to $\alpha_i b^{(i+1)}$ for $i \in \{1, \dots, \min\{N, K\}\}$, and his payoff is equal to $\alpha_i (s_{g(i)} - b^{(i+1)})$. If there are at least as many positions as advertisers ($N \geq K$), then the last advertiser's payment $p^{(K)}$ is equal to zero.¹⁵

It is also useful to describe explicitly the rules that the VCG mechanism would impose in this setting. The rules for allocating positions are the same as under GSP: position i is assigned to advertiser $g(i)$ with the i -th highest bid $b^{(i)}$. The payments, however, are different. Each advertiser's payment is equal to the negative externality that he imposes on others, assuming that bids are equal to values. Thus, the payment of the last advertiser who gets allocated a spot is the same as under GSP: zero if $N \geq K$; $\alpha_N b^{(N+1)}$ otherwise. For all other $i < \min\{N, K\}$, payment p^V induced by VCG will be different from payment p induced by GSP. Namely, $p^{V(i)} = (\alpha_i - \alpha_{i+1})b^{(i+1)} + p^{V(i+1)}$.

In the following two sections, we will consider two alternative ways of completing the model: as a simultaneous-move game of complete information, resembling a sealed-bid second-price auction, and as an extensive-form game of incomplete information, resembling an ascending English auction. Before moving on to these models, let us make a few observations about GSP and VCG.

REMARK 1: *If all advertisers were to bid the same amounts under the two mechanisms, then each advertiser's payment would be at least as large under GSP as under VCG.*

This is easy to show by induction on advertisers' payments, starting with the last advertiser who gets assigned a position. For $i = \min\{K, N\}$, $p^{(i)} = p^{V(i)} = \alpha_i b^{(i+1)}$. For any $i < \min\{K,$

¹³ In actual sponsored search auctions at Google and Yahoo!, advertisers can also choose to place "broad match" bids that match searches that include a keyword along with additional search terms.

¹⁴ The actual practice at Overture is to show equal bids according to the order in which the advertisers placed their bids.

¹⁵ Although we set the reserve price to zero, search engines charge the last advertiser a positive reserve price. We also assume that a bid can be any nonnegative real number, while in practice bids can be specified only in \$.01 increments. Finally, we assume that advertiser $g(i)$ is charged the amount $b^{(i+1)}$ per click, while search engines typically charge one cent more, $(b^{(i+1)} + \$.01)$.

$$N\}, p^{V,(i)} - p^{V,(i+1)} = (\alpha_i - \alpha_{i+1})b^{(i+1)} \leq \alpha_i b^{(i+1)} - \alpha_{i+1} b^{(i+2)} = p^{(i)} - p^{(i+1)}.$$

REMARK 2: *Truth-telling is a dominant strategy under VCG.*

This is a well-known property of the VCG mechanism.

REMARK 3: *Truth-telling is not a dominant strategy under GSP.*

For instance, consider a slight modification of the example from Section I. There are still three advertisers, with values per click of \$10, \$4, and \$2, and two positions. However, the click-through rates of these positions are now almost the same: the first position receives 200 clicks per hour, and the second one gets 199. If all players bid truthfully, then advertiser 1's payoff is equal to $(\$10 - \$4) * 200 = \$1,200$. If, instead, he shades his bid and bids only \$3 per click, he will get the second position, and his payoff will be equal to $(\$10 - \$2) * 199 = \$1,592 > \$1,200$.

III. GSP and Locally Envy-Free Equilibria

Advertisers bidding on Yahoo! and Google can change their bids very frequently. We therefore think of these sponsored search auctions as continuous time or infinitely repeated games in which advertisers originally have private information about their types, gradually learn the values of others, and can adjust their bids repeatedly. In principle, the sets of equilibria in such repeated games can be very large, with players potentially punishing each other for deviations. The strategies required to support such equilibria are usually quite complex, however, requiring precise knowledge of the environment and careful implementation. In theory, advertisers could implement such strategies via automated robots, but in practice they may not be able to: bidding software must first be authorized by the search engines,¹⁶ and search engines are unlikely to permit strategies that would allow advertisers to collude and substantially reduce revenues.

¹⁶ See, e.g., <http://help.yahoo.com/help/us/performance/customer/dtc/bidding>, questions 7 and 8 (accessed June 10, 2006).

We therefore focus on simple strategies and study the rest points of the bidding process: if the vector of bids stabilizes, at what bids can it stabilize? We impose several assumptions and restrictions. First, we assume that all values are common knowledge: over time, advertisers are likely to learn all relevant information about each other's values. Second, since bids can be changed at any time, stable bids must be best responses to each other—otherwise, an advertiser whose bid is not a best response would have an incentive to change it. Thus, we assume that the bids form an equilibrium in the simultaneous-move, one-shot game of complete information. Third, what are the simple strategies that an advertiser can use to increase his payoff, beyond simple best responses to the other players' bids?

One clear strategy is to try to force out the player who occupies the position immediately above. Suppose advertiser k bids b_k and is assigned to position i , and advertiser k' bids $b_{k'} > b_k$ and is assigned to position $(i - 1)$. Note that if k raises his bid slightly, his own payoff does not change, but the payoff of the player above him decreases. Of course, player k' can retaliate, and the most she can do is to slightly underbid advertiser k , effectively swapping places with him. If advertiser k is better off after such retaliation, he will indeed want to force player k' out, and the vector of bids will change. Thus, if the vector converges to a rest point, an advertiser in position i should not want to "exchange" positions with the advertiser in position $(i - 1)$. We call such vectors of bids "locally envy-free."¹⁷

DEFINITION 1: *An equilibrium of the simultaneous-move game induced by GSP is locally envy-free if a player cannot improve*

¹⁷ An alternative interpretation of this restriction is as follows. With only one slot, GSP coincides with the standard second-price auction, and the restriction of local envy-freeness simply says that the losing advertiser bids at least his own value, ruling out various implausible equilibria. Likewise, with multiple slots, the local envy-freeness restriction is equivalent to saying that the bid of the advertiser who gets position i and thus "loses" position $(i - 1)$ is such that his "marginal bid" (i.e., the difference between the highest amount that he could have paid if he had "won" position $(i - 1)$ and the amount he actually pays for position i) is at least as high as the marginal value for the extra clicks he would have received in position $(i - 1)$. To see this, simply rearrange equation (1) to get $p^{(i-1)} - p^{(i)} \geq s_{g(i)}(\alpha_{i-1} - \alpha_i)$.

his payoff by exchanging bids with the player ranked one position above him. More formally, in a locally envy-free equilibrium, for any $i \leq \min\{N + 1, K\}$,

$$(1) \quad \alpha_i s_{g(i)} - p^{(i)} \geq \alpha_{i-1} s_{g(i)} - p^{(i-1)}.$$

Of course, it is possible that bids change over time, depending on the players' strategies and information structure. However, if the behavior ever converges to a vector of bids, that vector should correspond to a locally envy-free equilibrium of the simultaneous-move game Γ induced by GSP. Consequently, we view a locally envy-free equilibrium Γ as a prediction regarding a rest point at which the vector of bids stabilizes. In this section, we study the set of locally envy-free equilibria.

We first show that the set of locally envy-free equilibria maps naturally to a set of stable assignments in a corresponding two-sided matching market. The idea that auctions and two-sided matching models are closely related is not new: it goes back to Vincent P. Crawford and Elsie M. Knoer (1981), Alexander S. Kelso and Crawford (1982), Herman B. Leonard (1983), and Gabrielle Demange, David Gale, and Marilda Sotomayor (1986), and has been studied in detail in a recent paper by John W. Hatfield and Milgrom (2005). Note, however, that in our case the nonstandard auction is very different from those in the papers noted above.

Our environment maps naturally into the most basic assignment model, studied first by Lloyd S. Shapley and Martin Shubik (1972). Consider each position as an agent who is looking for a match with an advertiser. The value of a position-advertiser pair (i, k) is equal to $\alpha_i s_k$. We call this assignment game A . The advertiser makes its payment p_{ik} for the position, and the advertiser is left with $\alpha_i s_k - p_{ik}$. The following pair of lemmas shows that there is a natural mapping from the set of locally envy-free equilibria of GSP to the set of stable assignments. All proofs are in the Appendix.

LEMMA 1: *The outcome of any locally envy-free equilibrium of auction Γ is a stable assignment.*

LEMMA 2: *If the number of advertisers is greater than the number of available positions,*

then any stable assignment is an outcome of a locally envy-free equilibrium of auction Γ .

We will now construct a particular locally envy-free equilibrium of game Γ . This equilibrium has two important properties. First, in this equilibrium, advertisers' payments coincide with their payments in the dominant-strategy equilibrium of VCG. Second, this equilibrium is the *worst* locally envy-free equilibrium for the search engine and the *best* locally envy-free equilibrium for the advertisers. Consequently, the revenues of a search engine are (weakly) higher in any locally envy-free equilibrium of GSP than in the dominant-strategy equilibrium of VCG.

Consider the following strategy profile B^* . Without loss of generality, assume that advertisers are labeled in decreasing order of their values, i.e., if $j < k$, then $s_j \geq s_k$. For each advertiser $j \in \{2, \dots, \min\{N + 1, K\}\}$, bid b_j^* is equal to $p^{V,j-1}/\alpha_{j-1}$, where $p^{V,j-1}$ is the payment of advertiser $j - 1$ in the dominant-strategy equilibrium of VCG where all advertisers bid truthfully. Bid b_1^* is equal to s_1 .¹⁸

THEOREM 1: *Strategy profile B^* is a locally envy-free equilibrium of game Γ . In this equilibrium, each advertiser's position and payment are equal to those in the dominant-strategy equilibrium of the game induced by VCG. In any other locally envy-free equilibrium of game Γ , the total revenue of the seller is at least as high as in B^* .*

To prove Theorem 1, we first note that payments under strategy profile B^* coincide with VCG payments and check that B^* is indeed a locally envy-free equilibrium. This follows from the fact that, by construction, each advertiser is indifferent between remaining in his position and swapping with the advertiser one position above him. Next, from Lemma 1 we know that every locally envy-free equilibrium corresponds to a stable assignment. The "core elongation" property of the set of stable assignments (Shapley and Shubik 1972; Crawford and Knoer 1981) implies that there exists an

¹⁸ This bid does not affect any advertiser's payment and can be set equal to any value greater than b_2^* .

“advertiser-optimal” assignment A in that set, such that in any other stable assignment, each advertiser pays at least as much to the search engine as he does in A . Moreover, Leonard (1983) and Demange, Gale, and Sotomayor (1986) show that in general assignment games, payoffs of “buyers” in the buyer-optimal stable assignment coincide with their VCG payoffs, which is sufficient to complete the proof. This is particularly easy to show in the specific environment that we consider, and so for completeness we include a short independent proof.

In the model, we assume that all advertisers are identical along dimensions other than per-click value, and in particular have identical click-through rates. The analysis remains largely the same if, instead, we assume that the CTRs of different advertisers are multiples of one another, i.e., if any advertiser k assigned to any position i receives $\alpha_i \beta_k$ clicks, where α_i is a position-specific factor and β_k is an advertiser-specific factor. In this case, the versions of GSP implemented by Yahoo! and Google differ.

Under Yahoo!’s system, advertisers are still ranked by bids, and each of them is charged the next-highest advertiser’s bid. Then, bids form a locally envy-free equilibrium if and only if, for any i and j , $\alpha_i \beta_{g(i)}(s_{g(i)} - b^{(i+1)}) \geq \alpha_j \beta_{g(i)}(s_{g(i)} - b^{(j+1)})$. Dividing both sides by the positive number $\beta_{g(i)}$, we get $\alpha_i(s_{g(i)} - b^{(i+1)}) \geq \alpha_j(s_{g(i)} - b^{(j+1)})$, i.e., the necessary and sufficient condition for a locally envy-free equilibrium in the case where all β_k are equal to one. Hence, under Yahoo!’s version of GSP, equilibria are not affected by changes in β s.

Under Google’s system, advertisers are arranged by “rank numbers.” Advertiser k ’s rank number is the product of his bid and “quality score” γ_k .¹⁹ Thus, under Google’s system, $g(1)$ is the advertiser with the highest rank number, $g(2)$ is the advertiser with the second highest rank number, and so on. Per-click payment of advertiser $g(i)$ is equal to the smallest amount $x^{(i)}$, such that $\gamma_{g(i+1)} x^{(i)}$ is greater than or equal to the next highest advertiser’s rank number, i.e., $x^{(i)} = \gamma_{g(i+1)} b_{g(i+1)} / \gamma_{g(i)}$. Then, bids form a

locally envy-free equilibrium if and only if, for any i and j , $\alpha_i \beta_{g(i)}(s_{g(i)} - \gamma_{g(i+1)} b_{g(i+1)} / \gamma_{g(i)}) \geq \alpha_j \beta_{g(i)}(s_{g(i)} - \gamma_{g(j+1)} b_{g(j+1)} / \gamma_{g(i)})$. Dividing both sides by $\beta_{g(i)}$ and multiplying by $\gamma_{g(i)}$, we get $\alpha_i(\gamma_{g(i)} s_{g(i)} - \gamma_{g(i+1)} b_{g(i+1)}) \geq \alpha_j(\gamma_{g(i)} s_{g(i)} - \gamma_{g(j+1)} b_{g(j+1)})$. Hence, the set of bids $\{b_k\}$ is a locally envy-free equilibrium under Google’s version of GSP with position-specific factors $\{\alpha_i\}$, advertiser-specific quality scores $\{\gamma_k\}$, and per-click values $\{s_k\}$, if and only if the set of bids $\{\gamma_k b_k\}$ is an equilibrium of our basic model with position-specific CTRs $\{\alpha_i\}$, per-click values $\{\gamma_k s_k\}$, and no quality scores or advertiser-specific factors in CTRs.

IV. Main Result: GSP and Generalized English Auction

In the model analyzed in the previous section, we assume that advertisers have converged to a long-run steady state, have learned each other’s values, and no longer have incentives to change their bids. But how do they converge to such a situation? In this section, we introduce the generalized English auction, an analogue of the standard English auction corresponding to GSP, to help us answer this question.

In the generalized English auction, there is a clock showing the current price, which continuously increases over time. Initially, the price on the clock is zero, and all advertisers are in the auction. An advertiser can drop out at any time, and his bid is the price on the clock at the time when he drops out. The auction is over when the next-to-last advertiser drops out. The ad of the last remaining advertiser is placed in the best position on the screen, and this advertiser’s payment per click is equal to the price at which the next-to-last advertiser dropped out. The ad of the next-to-last advertiser is placed second, and his payment per click is equal to the third-highest advertiser’s bid, and so on.²⁰ In other words, the vector of bids obtained in the generalized English auction is used to allocate

¹⁹ Initially, Google simply used click-through rates to determine quality scores, setting $\gamma_k = \beta_k$. Later, however, it switched to a less transparent system for determining quality scores, incorporating such factors as the relevance of an ad’s text and the quality of an advertiser’s Web page.

²⁰ If several advertisers drop out simultaneously, one of them is chosen randomly. Whenever an advertiser drops out, the clock is stopped, and other advertisers are also allowed to drop out; again, if several advertisers want to drop out, one of them is chosen randomly. If several advertisers end up dropping out at the same price, the first one to drop out is placed in the lowest position of the still available ones, the next one to the position right above that, and so on.

the objects and compute the prices according to the rules of GSP. With one object, the generalized English auction becomes a simple English auction.²¹

We view the generalized English auction in the same light as the tâtonnement processes in the theory of general equilibrium (see, e.g., Andreu Mas-Colell, Michael D. Whinston, and Jerry R. Green 1995, sect. 17.H) and the salary adjustment process in the theory of matching in labor markets with heterogeneous firms and workers (Crawford and Knoer 1981).²² While all these processes, taken literally, happen in "imaginary time," they are meant to resemble the underlying dynamics of the actual markets, to help us distinguish more plausible equilibria from less plausible ones, characterize their stability and other properties, and examine the significance of the underlying assumptions. As in the case of the tâtonnement and salary adjustment processes, there are many features of real markets not captured by the generalized English auction, but we believe that it provides a natural and useful approximation.

To define the game formally, assume that there are $N \geq 2$ slots and $K = N + 1$ advertisers. (Cases with $K \neq N + 1$ require only minor modifications in the proof.) Click-through rates α_i are commonly known, with $\alpha_{N+1} \equiv 0$. Advertisers' per-click valuations s_k are drawn from a continuous distribution $F(\cdot)$ on $[0; +\infty)$ with a continuous density function $f(\cdot)$ that is positive everywhere on $(0, +\infty)$. Each advertiser knows his valuation and the distribution of other advertisers' valuations.

The strategy of an advertiser assigns the choice of dropping out or not for any history of the game, given that the advertiser has not previously dropped out. In other words, the strategy can be represented as a function $p_k(i, h, s_k)$, where s_k is the value per click of advertiser k , p_k

is the price at which he drops out, i is the number of advertisers remaining (including advertiser k), and $h = (b_{i+1}, \dots, b_{N+1})$ is the history of prices at which previous advertisers have dropped out. (As a result, the price that advertiser k would have to pay per click if he dropped out next is equal to b_{i+1} , unless the history is empty, in which case we say that $b_{i+1} \equiv 0$.) The following theorem shows that this game has a unique perfect Bayesian equilibrium with strategies continuous in advertisers' valuations.²³ The payoffs of all advertisers in this equilibrium are equal to VCG payoffs.

THEOREM 2: *In the unique perfect Bayesian equilibrium of the generalized English auction with strategies continuous in s_k , an advertiser with value s_k drops out at price*

$$(2) \quad p_k(i, h, s_k) = s_k - \frac{\alpha_i}{\alpha_{i-1}} (s_k - b_{i+1}).$$

In this equilibrium, each advertiser's resulting position and payoff are the same as in the dominant-strategy equilibrium of the game induced by VCG. This equilibrium is ex post: the strategy of each advertiser is a best response to other advertisers' strategies regardless of their realized values.

The intuition of the proof is as follows. First, with i players remaining and the next highest bid equal to b_{i+1} , it is a dominated strategy for a player with value s to drop out before price p reaches the level at which he is indifferent between getting position i and paying b_{i+1} per click and getting position $i - 1$ and paying p per click. Next, if for some set of types it is not optimal to drop out at this "borderline" price level, we can consider the lowest such type, and

²¹ This version of the English auction is also known as the "Japanese" or "button" auction.

²² Of course, the generalized English auction is very different from the salary adjustment process of Crawford and Knoer (1981) and its application to multi-unit auctions (Demange et al. 1986). In Demange et al. (1986), bidding proceeds simultaneously for all items and the auctioneer keeps track of a vector of item-specific prices, while in the generalized English auction bidding proceeds, in essence, sequentially, and the auctioneer keeps track of only one price.

²³ Without this restriction, multiple equilibria exist, even in the simplest English auction with two bidders and one object. For example, suppose there is one object for sale and two bidders with independent private values for this object distributed exponentially on $[0, \infty)$. Consider the following pair of strategies. If a bidder's value is in the interval $[0, 1]$ or in the interval $[2, \infty)$, he drops out when the clock reaches his value. If bidder 1's value is in the interval $(1, 2)$, he drops out at 1, and if bidder 2's value is in the interval $(1, 2)$, he drops out at 2. This pair of strategies, together with appropriate beliefs, forms a perfect Bayesian equilibrium.

then once the clock reaches this price level, a player of this type will know that he has the lowest per-click value of the remaining players. But then he will also know that the other remaining players will drop out only at price levels at which he will find it unprofitable to compete with them for the higher positions.

The result of Theorem 2 resembles the classic result on the equivalence of the English auction and the second-price sealed-bid auction under private values (Vickrey 1961). Note, however, that the intuition is very different: Vickrey's result follows simply from the existence of equilibria in dominant strategies, whereas in our case such strategies do not exist, and bids do depend on other player's bids. Also, our result is very different from the revenue equivalence theorem: payoffs in the generalized English auction coincide with VCG payments for all realizations of values, not only in expectation, and the result does not hinge on the assumptions of symmetric bidders or common priors.

The equilibrium described in Theorem 2 is an *ex post* equilibrium. As long as all advertisers other than advertiser k follow the equilibrium strategy described in Theorem 2, it is a best response for advertiser k to follow his equilibrium strategy, for any realization of other advertisers' values. Thus, the outcome implemented by this mechanism depends only on the realization of advertisers' values and does not depend on advertisers' beliefs about each other's types.

Clearly, any dominant strategy solvable game has an *ex post* equilibrium. However, the generalized English auction is not dominant strategy solvable. This combination of properties is quite striking: the equilibrium is unique and efficient, and the strategy of each advertiser does not depend on the distribution of other advertisers' values, yet advertisers do not have dominant strategies.²⁴ The generalized English

auction is a particularly interesting example, because it can be viewed as a model of a mechanism that has "emerged in the wild."

V. Conclusion

We investigate a new mechanism that we call the generalized second-price auction. GSP is tailored to the unique features of the market for Internet advertisements. As far as we know, this mechanism was first used in 2002. As of May 2006, the annual revenues from GSP auctions were on the order of \$10 billion.

GSP looks similar to the VCG mechanism, because just like in the standard second-price auction, the payment of a bidder does not directly depend on his bid. Although GSP looks similar to VCG, its properties are very different, and equilibrium behavior is far from straightforward. In particular, unlike the VCG mechanism, GSP generally does not have an equilibrium in dominant strategies, and truth-telling is not an equilibrium of GSP. We show that the generalized English auction that corresponds to the generalized second-price auction has a unique equilibrium.

This equilibrium has some notable properties. The bid functions have explicit analytic formulas, which, combined with equilibrium uniqueness, make our results a useful starting point for empirical analysis. Moreover, these functions do not depend on bidders' beliefs about each other's types: the outcome of the auction depends only on the realizations of bidders' values. This is one of the very few mechanisms encountered in practice that are not dominant strategy solvable and nevertheless have this property. It is particularly interesting that a mechanism with such notable features in theory and such enormous popularity in practice developed as a result of evolution of inefficient market institutions, which were gradually replaced by increasingly superior designs.

²⁴ Dirk Bergemann and Stephen Morris (2005) show that an outcome implementable by robust mechanisms must be implementable in dominant strategies. Indeed, the outcome implemented by the generalized English auction can be implemented in dominant strategies by the VCG mechanism; however, VCG is not the mechanism that is used in practice. Philippe Jehiel and Benny Moldovanu (2001) and Jehiel et al. (2006) show that, generically, any efficient choice function is not Bayes-Nash implementable and any nontrivial choice function is not *ex post* implementable, if values are interdependent and signals are multidimensional.

Of course, in our model, values are private, and, crucially, signals are single-dimensional, even though multiple different objects are for sale. This makes efficient *ex post* implementation feasible. For other examples of mechanisms that allocate multiple different objects to bidders with single-dimensional types, see Moldovanu and Aner Seia (2001) and Thomas Kittsteiner and Moldovanu (2005).

APPENDIX: PROOFS

PROOF OF LEMMA 1:

By definition, in any locally envy-free equilibrium outcome, no advertiser can profitably rematch with the position assigned to the advertiser right above him. Also, no advertiser (*a*) can profitably rematch with a position assigned to an advertiser below him (*b*)—if such a profitable rematching existed, advertiser *a* would find it profitable to slightly undercut advertiser *b* in game Γ and get *b*'s position and payment. But this would contradict the assumption that we are in equilibrium.²⁵

Hence, we need only show that no advertiser can profitably rematch with the position assigned to an advertiser more than one spot above him. First, note that in any locally envy-free equilibrium, the resulting matching must be assortative, i.e., for any *i*, the advertiser assigned to position *i* has a higher per-click valuation than the advertiser assigned to position *i* + 1 and, therefore, the advertiser with the highest per-click value must be assigned to the top position, the advertiser with the second-highest per-click value to the second-highest position, and so on.

Indeed, suppose $s_{g(i)}$ and $s_{g(i+1)}$ are the values of advertisers assigned to positions *i* and *i* + 1. Equilibrium restrictions imply that $\alpha_i s_{g(i)} - p^{(i)} \geq \alpha_{i+1} s_{g(i)} - p^{(i+1)}$ (nobody wants to move one position down), and local envy-freeness implies that $\alpha_{i+1} s_{g(i+1)} - p^{(i+1)} \geq \alpha_i s_{g(i+1)} - p^{(i)}$ (nobody wants to move one position up). Manipulating the inequalities above yields $\alpha_i s_{g(i)} - \alpha_i s_{g(i+1)} + \alpha_{i+1} s_{g(i+1)} \geq \alpha_{i+1} s_{g(i)}$, thus $(\alpha_i - \alpha_{i+1}) s_{g(i)} \geq (\alpha_i - \alpha_{i+1}) s_{g(i+1)}$. Since $\alpha_i > \alpha_{i+1}$, we have $s_{g(i)} \geq s_{g(i+1)}$, and hence the locally envy-free equilibrium outcome must be an assortative match.

Now, let us show that no advertiser can profitably rematch with the position assigned to an advertiser more than one spot above

him. Suppose the advertiser assigned to position *i* is considering rematching with position $m < i - 1$. Since the equilibrium is locally envy-free, we have

$$\begin{aligned} \alpha_i s_{g(i)} - p^{(i)} &\geq \alpha_{i-1} s_{g(i)} - p^{(i-1)}, \\ \alpha_{i-1} s_{g(i-1)} - p^{(i-1)} &\geq \alpha_{i-2} s_{g(i-1)} - p^{(i-2)}, \\ &\vdots \\ \alpha_{m+1} s_{g(m+1)} - p^{(m+1)} &\geq \alpha_m s_{g(m+1)} - p^{(m)}. \end{aligned}$$

Since $\alpha_j > \alpha_{j+1}$ for any *j*, and $s_{g(i)} > s_{g(j)}$ for any $i < j$, the inequalities above remain valid after replacing $s_{g(i)}$ with $s_{g(j)}$. Doing that, then adding all inequalities up, and canceling out the redundant elements, we get $\alpha_i s_{g(i)} - p^{(i)} \geq \alpha_m s_{g(i)} - p^{(m)}$. But that implies that the advertiser assigned to position *i* cannot rematch profitably with position *m*, and we are done.

PROOF OF LEMMA 2:

Take a stable assignment. By a result of Shapley and Shubik (1972), this assignment must be efficient, and hence assortative, and so without loss of generality we can assume that advertisers are labeled in decreasing order of their bids (i.e., $s_j > s_k$ whenever $j < k$) and that advertiser *i* is matched with position *i*, with associated payment p_i .

Let us construct a locally envy-free equilibrium with the corresponding outcome. Let $b_1 = s_1$ and $b_i = p_{i-1}/\alpha_{i-1}$ for $i > 1$. Let us show that this set of strategies is a locally envy-free equilibrium. First, note that for any *i*, $b_i > b_{i+1}$ (because otherwise we would have, for some *i*, $p_{i-1}/\alpha_{i-1} \leq p_i/\alpha_i \Rightarrow s_i - p_{i-1}/\alpha_{i-1} \geq s_i - p_i/\alpha_i \Rightarrow \alpha_{i-1} s_i - p_{i-1} > \alpha_i s_i - p_i$, which would imply that player *i* could rematch profitably). Therefore, position allocations and payments resulting from this strategy profile will coincide with those in the original stable assignment. To see that this strategy profile is an equilibrium, note that deviating and moving to a different position in this strategy profile is at most as profitable for any player as rematching with the corresponding position in the assignment game. To see that this equilibrium is locally envy-free, note that the payoff from swapping with the bidder above is exactly equal

²⁵ This argument relies on the fact that in equilibrium, no two (or more) advertisers bid the same amount, which is straightforward to prove: since all advertisers' per-click values are different, and all ties are broken randomly with equal probabilities, at least one such advertiser would find it profitable to bid slightly higher or slightly lower.

to the payoff from rematching with that player's position in the assignment game.

PROOF OF THEOREM 1:

First, we need to check that the order of the bids is preserved, i.e., $b_j^* > b_{j+1}^*$ for any $j < \min\{N, K\}$. For $j \geq 2$, this is equivalent to

$$\begin{aligned}
 (3) \quad & \frac{p^{v,(j-1)}}{\alpha_{j-1}} > \frac{p^{v,(j)}}{\alpha_j} \\
 & \Downarrow \\
 & \frac{(\alpha_{j-1} - \alpha_j)s_j + p^{v,(j)}}{\alpha_{j-1}} > \frac{p^{v,(j)}}{\alpha_j} \\
 & \Downarrow \\
 & \alpha_j(\alpha_{j-1} - \alpha_j)s_j > (\alpha_{j-1} - \alpha_j)p^{v,(j)} \\
 & \Downarrow \\
 & \alpha_j s_j > p^{v,(j)}.
 \end{aligned}$$

For $j = 1$, $b_j^* > b_{j+1}^*$ is equivalent to

$$\begin{aligned}
 (4) \quad & s_1 > \frac{p^{v,(1)}}{\alpha_1} \\
 & \Downarrow \\
 & \alpha_1 s_1 > p^{v,(1)}.
 \end{aligned}$$

To see that for any j , $\alpha_j s_j > p^{v,(j)}$, note first that in the game induced by VCG, each player can guarantee himself the payoff of at least zero (by bidding zero), and hence in any equilibrium his payoff from clicks is at least as high as his payment. To prove that the inequality is strict, note that if player j 's value per click were slightly lower, e.g., $s_j - \Delta$ instead of s_j , $\Delta < s_j - s_{j+1}$, then his payment in the truth-telling equilibrium would still be the same (because it does not depend on his own bid, given the allocation of positions), and so $p^{v,(j)} \leq \alpha_j(s_j - \Delta) < \alpha_j s_j$. Thus, for any j , $b_j^* > b_{j+1}^*$, and therefore each bidder's position is the same as in the

truthful equilibrium of VCG. Therefore, by construction, payments are also the same.

Next, to see that no bidder j can benefit by bidding less than b_j^* , suppose that he bids an amount $b' < b_j^*$ that puts him in position $j' > j$. Then, by construction, his payment will be equal to the amount that he would need to pay to be in position j' under VCG, provided that other players bid truthfully. But truthful bidding is an equilibrium under VCG, and so such deviation cannot be profitable there—hence, it cannot be profitable in strategy profile B^* of game Γ either.

To see that no bidder j can benefit by bidding more than b_j^* , suppose that he bids an amount $b' > b_j^*$ that puts him in position $j' < j$. Then the net payoff from this deviation is equal to $(\alpha_{j'} - \alpha_j)s_j - (\alpha_j b_j^* - \alpha_j b_{j+1}) < (\alpha_{j'} - \alpha_j)s_j - (\alpha_j b_{j+1}^* - \alpha_j b_{j+1}) = \sum_{i=j'}^{j-1} (\alpha_i - \alpha_{i+1})s_j - \sum_{i=j'}^{j-1} (p^{v,(i)} - p^{v,(i+1)}) = \sum_{i=j'}^{j-1} (\alpha_i - \alpha_{i+1})s_j - \sum_{i=j'}^{j-1} (\alpha_i - \alpha_{i+1})s_{i+1}$. But since $s_j \leq s_{i+1}$ for any $i < j$, the last expression is less than or equal to zero, and hence the deviation is not profitable.

To check that this equilibrium is locally envy-free, note that if bidder j swapped his bids with bidder $j - 1$, his payoff would change by $(\alpha_{j-1} - \alpha_j)s_j - (\alpha_{j-1}b_j^* - \alpha_j b_{j+1}) = (\alpha_{j-1} - \alpha_j)s_j - (p^{v,(j-1)} - p^{v,(j)}) = (\alpha_{j-1} - \alpha_j)s_j - (p^{v,(j-1)} - p^{v,(j)}) = (\alpha_{j-1} - \alpha_j)s_j - (\alpha_{j-1} - \alpha_j)s_j = 0$. In other words, each bidder is indifferent between his actual payoff and his payoff after swapping bids with the bidder above, and hence the equilibrium is locally envy-free.

Let us now show that B^* is the best locally envy-free equilibrium for the bidders and the worst locally envy-free equilibrium for the search engine. The core-elongation property of the assignment game (Shapley and Shubik 1972; Crawford and Knoer 1981) implies that there exists an assignment that is the best stable assignment for all advertisers and the worst stable assignment for all positions. Suppose this assignment is characterized by a vector of payments $p = (p_1, \dots, p_K)$. Let $p^V = (p_1^V, \dots, p_K^V)$ be the set of dominant-strategy VCG payments, i.e., the set of payments in equilibrium B^* of game Γ .

In any stable assignment, p_K must be at least as high as $\alpha_K s_{K+1}$, since otherwise advertiser $K + 1$ would find it profitable to match with position K . On the other hand, $p_K^V = \alpha_K s_{K+1}$,

and hence in the advertiser-optimal stable assignment, $p_K = p_K^V$.

Next, in any stable assignment, it must be the case that $p_{K-1} - p_K \geq (\alpha_{K-1} - \alpha_K)s_K$ —otherwise, advertiser K would find it profitable to rematch with position $K - 1$. Hence, $p_{K-1} \geq (\alpha_{K-1} - \alpha_K)s_K + p_K \geq (\alpha_{K-1} - \alpha_K)s_K + p_K^V = p_{K-1}^V$, and so in the advertiser-optimal stable assignment, $p_{K-1} = p_{K-1}^V$.

Proceeding by induction, we get $p_j = p_j^V$ for any $j \leq K$ in the advertiser-optimal stable assignment, and so in any locally envy-free equilibrium of game Γ , the total revenue of the seller is at least as high as $\sum_{j=1}^K p_j^V$.

PROOF OF THEOREM 2:

First, note that in equilibrium, for any player k , any history h , and any number of remaining players i , the drop-out price $p_k(i, h, s_k)$ tends to infinity as s_k tends to infinity. (Otherwise, there would exist a player for whom it was optimal to deviate from his strategy and stay longer, for a sufficiently high value s .) Next, take any equilibrium of the generalized English auction. Note that if in this equilibrium $p_k(i, h, s_k) > p_k(i, h, s'_k)$ for some k, h, i , and types $s_k < s'_k$, then it has to be the case that both types s_k and s'_k are indifferent between dropping out at $p_k(i, h, s_k)$ and $p_k(i, h, s'_k)$. (Otherwise, one of them would be able to increase his payoff by mimicking the other.) Consequently, we can “swap” such players’ strategies, and therefore there exists an “observationally equivalent” equilibrium in which strategies are nondecreasing in types; also, they are still continuous in own values. Consider this equilibrium profile of strategies $p_k(i, h, s_k)$.

Let $q(i, b_{i+1}, s)$ be such a price that a player with value s is indifferent between getting position i at price b_{i+1} and position $i - 1$ at price $q(i, b_{i+1}, s)$. That is,

$$(5) \quad \alpha_{i-1}(s - q(i, b_{i+1}, s)) = \alpha_i(s - b_{i+1})$$

⇕

$$q(i, b_{i+1}, s) = s - \frac{\alpha_i}{\alpha_{i-1}}(s - b_{i+1}).$$

Slightly abusing notation, let $q(i, h, s) = q(i, b_{i+1}, s)$, where b_{i+1} is the last bid at which a player dropped out in history h . (This player

received position $i + 1$. If history h is empty, we set $b_{i+1} = 0$.) We will now show that for any i, k, h , and $s_k, p_k(i, h, s_k) = q(i, h, s_k)$.

Suppose that is not the case, and take the largest i for which there exist such history h (with the last player dropping out at b_{i+1}), player k , and type s_k (surviving with positive probability on the equilibrium path) that $p_k(i, h, s_k) \neq q(i, h, s_k)$. Since by assumption, all strategies up to this stage were $p_k(\cdot, \cdot, \cdot) = q(\cdot, \cdot, \cdot)$, we know that there exists a value $s_{min} \geq b_{i+1}$, such that all players with values less than s_{min} have dropped out, and all players with values greater than s_{min} are still in the auction.

Step 1: Suppose for some type $s \geq s_{min}$, $p_{max}(i, h, s) \equiv \max_k p_k(i, h, s) > q(i, h, s)$. Let s_0 be the smallest type, and let k be the corresponding player, such that $p_k(i, h, s_0) = p_{max}(i, h, s)$; clearly, $s_0 \leq s$. Without loss of generality, we can assume that there is a positive mass of types of other players dropping out at or before $p_k(i, h, s_0)$.²⁶

Step 1(a). Suppose first that there is a positive mass of types of other players dropping out at $p_k(i, h, s_0) = p_{max}(i, h, s)$. That implies that with positive probability, player k of type s_0 will remain in the subgame following the drop-out of some other player at $p_k(i, h, s_0)$ (since ties are broken randomly). Let us show that in this subgame, player k of type s_0 will be the first player to drop out with probability 1. Suppose that is not the case, and let $l < i - 1$ be the smallest number such that he gets position l with positive probability.

Consider any continuation of history h, h_{l+2} , such that the last player to drop out in that history gets position $l + 2$ and drops out at price b_{l+2} , player k of type s_0 is one of the remaining $l + 1$ players, there is a positive probability that player k gets position l in the continuation subgame following history h_{l+2} , and there is zero

²⁶ Otherwise, we have $\forall j \neq k, p_j(i, h, s_0) \leq p_j(i, h, s) \leq p_{max}(i, h, s) = p_k(i, h, s_0) \Rightarrow \forall j \neq k, p_j(i, h, s_0) = p_k(i, h, s_0)$ and $\forall s' > s_0, p_j(i, h, s') > p_k(i, h, s_0)$. But we also have $p_k(i, h, s_0) = p_{max}(i, h, s) > q(i, h, s) \geq q(i, h, s_0)$, and so for some $s' > s_0$ we have $p_{max}(i, h, s') > q(i, h, s')$ and $p_{max}(i, h, s') > p_{max}(i, h, s)$. We can then consider s'_0 in place of s_0 , where s'_0 is the smallest type, and k' is the corresponding player, such that $p_{k'}(i, h, s'_0) = p_{max}(i, h, s')$. There is a positive mass of types of other players dropping out before $p_{k'}(i, h, s'_0)$.

probability that player k gets position m for any $m < l$. Note that $s_0 \geq b_{l+2}$ —otherwise, it would have been optimal for player k to drop out earlier. Consider $p_k(l + 1, h_{l+2}, s_0)$. Since player k of type s_0 gets position l with positive probability in this subgame, there must be a positive mass of types of other players who drop out no later than $p_k(l + 1, h_{l+2}, s_0)$. Take the highest such type, s' , and the corresponding player j . It has to be the case that $s' > s_0 \geq b_{l+2}$. It also has to be the case that $q(l + 1, h_{l+2}, s')$ is less than or equal to $p_k(l + 1, h_{l+2}, s_0)$. (Otherwise, player j with value s' would be playing a strategy weakly dominated by dropping out at $q(l + 1, h_{l+2}, s')$, with a positive probability of earning strictly less than he would have earned if he waited until that price level.) Therefore, $p_k(l + 1, h_{l+2}, s_0) \geq q(l + 1, h_{l+2}, s') > q(l + 1, h_{l+2}, s_0) \geq b_{l+2}$. Let us show that it would be strictly better for player k with type s_0 to drop out at $q(l + 1, h_{l+2}, s_0)$ instead of waiting until $p_k(l + 1, h_{l+2}, s_0)$. Indeed, if nobody else drops out in between, or someone drops out before $q(l + 1, h_{l+2}, s_0)$, these strategies would result in identical payoffs. Otherwise, payoffs are different, and this happens with positive probability. Under the former strategy, player k earns

$$(6) \quad \alpha_{l+1}(s_0 - b_{l+2}).$$

Under the latter strategy, he earns

$$(7) \quad \alpha_l(s_0 - b_{l+1}),$$

where b_{l+1} is the price at which somebody else dropped out. (The probability of getting a spot $m < l$ is zero by construction.) With probability 1, $b_{l+1} \geq q(l + 1, h_{l+2}, s_0)$, and with positive probability, $b_{l+1} > q(l + 1, h_{l+2}, s_0)$, so the expected payoff from waiting until $p_k(l + 1, h_{l+2}, s_0)$ is strictly less than the expected payoff from dropping out at $q(l + 1, h_{l+2}, s_0)$: $E[\alpha_l(s_0 - b_{l+1})] < \alpha_l(s_0 - q(l + 1, h_{l+2}, s_0)) = \alpha_l(s_0 - (s_0 - (\alpha_{l+1}/\alpha_l)(s_0 - b_{l+2}))) = \alpha_{l+1}(s_0 - b_{l+2})$.

Therefore, in the subgame following the drop-out of some other player at $p_k(i, h, s_0)$, player k of type s_0 gets position $i - 1$ with probability 1, and therefore his payoff is

$\alpha_{i-1}(s_0 - p_k(i, h, s_0))$. Now suppose player k dropped out at a price $p_k(i, h, s_0) - \epsilon$ instead of waiting until $p_k(i, h, s_0)$. If somebody else drops out before $p_k(i, h, s_0) - \epsilon$ or after $p_k(i, h, s_0)$, or drops out at $p_k(i, h, s_0)$ but player k is chosen to drop out first, then these two strategies result in identical payoffs. The probability that somebody drops out in the interval $(p_k(i, h, s_0) - \epsilon, p_k(i, h, s_0))$ goes to zero as ϵ goes to zero, and the possible difference in the payoffs is finite, so the difference in payoffs due to this contingency goes to zero as ϵ goes to zero. Finally, there is a positive probability that somebody else drops out at $p_k(i, h, s_0)$ and is chosen to drop out first. If player k drops out before that, at $p_k(i, h, s_0) - \epsilon$, his payoff is $\alpha_i(s_0 - b_{i+1})$. If he waits until $p_k(i, h, s_0)$, we know that in the subsequent subgame his payoff is $\alpha_{i-1}(s_0 - p_k(i, h, s_0)) < \alpha_{i-1}(s_0 - q(i, h, s_0)) = \alpha_i(s_0 - b_{i+1})$. Therefore, for a sufficiently small ϵ , it is strictly better for player k with value s_0 to drop out at $p_k(i, h, s_0) - \epsilon$ instead of waiting until $p_k(i, h, s_0)$, which contradicts the assumption that $\{p_j(\cdot, \cdot, \cdot)\}$ is an equilibrium.

Step 1(b). Now, suppose there is mass zero of types of other players dropping out at $p_k(i, h, s_0) = p_{max}(i, h, s)$, but there is a positive mass dropping out before $p_k(i, h, s_0)$. Consider a sequence of small positive numbers (ϵ_n) converging to 0 as $n \rightarrow \infty$ and sequences (π_n^l) , $(\pi_n^2), \dots, (\pi_n^{i-1})$, where π_n^l is the probability that player k with value s_0 will end up in position l if another player drops out at price $(p_k(i, h, s_0) - \epsilon_n)$. Let $B = \alpha_i s_0$, i.e., the maximum payoff that a player with value s_0 can possibly get in the auction. Now, if (π_n^{i-1}) converges to 1 and (π_n^l) converges to zero for all $l < i - 1$, then, by an argument similar to the one at the end of Step 1(a), it is better for player k of type s_0 to drop out at some time $p_k(i, h, s_0) - \epsilon$.²⁷ If (π_n^{i-1}) does not converge to 1, take the smallest (i.e., best) l for which (π_n^l) does not converge to zero, and take a subsequence of ϵ_n along which (π_n^l) converges to some positive number ρ . Let s_1 be the value such that for a random draw of types of remaining players other than k , condi-

²⁷ If some other player drops out between $p_k(i, h, s_0) - \epsilon$ and $p_k(i, h, s_0)$, the benefit of staying longer tends to zero (it is at most $B(1 - \pi^{i-1})$), while the cost converges to a positive number (the difference between getting position i at price b_{i+1} and position $i - 1$ at price $p_k(i, h, s_0)$).

tional on each draw being greater than s_0 , the probability that at least one type is less than s_1 is equal to $\rho/2$ (i.e., $\prod_{j \neq k} [(1 - F_j(s_1))/(1 - F_j(s_0))] = 1 - \rho/2$). Clearly, $s_1 > s_0$. Take a small ε_n , and consider a subgame following some continuation of history h, h_{l+2} , where the $(l+2)^{\text{nd}}$ player drops out at b_{l+2} , $l+1$ players, including player k , remain, and player k gets position l with probability close to ρ (and any position better than l with probability close to zero). Consider $p_k(l+1, h_{l+2}, s_0)$. There must exist a player, j , such that $p_j(l+1, h_{l+2}, s_1) \leq p_k(l+1, h_{l+2}, s_0)$. (Otherwise, the probability of player k surviving until position l is less than or close to $\rho/2$, and thus cannot be close to ρ .) But then, by an argument similar to the one in Step 1(a), in this subgame it is strictly better for player k to drop out slightly earlier than $p_k(l+1, h_{l+2}, s_0)$: conditional on somebody else dropping out in between, the benefit is close to zero (the probability of getting a position better than l times the highest possible benefit B), while the cost is close to a positive number (the payoff from being in position $l+1$ at price b_{l+2} versus the payoff from being in position l at price at least $[s_1 - (\alpha_{l+1}/\alpha_l)(s_1 - b_{l+2})]$)—contradiction.

Step 2: In Step 1, we showed that $p_{\max}(i, h, s) \equiv \max_k p_k(i, h, s)$ cannot be greater than $q(i, h, s)$, and therefore for any player k and type $s \geq s_{\min}$, $p_k(i, h, s) \leq q(i, h, s)$. Take some value $s > s_{\min}$ and player k . Suppose $p_k(i, h, s) < q(i, h, s)$. Take some other player j . From Step 1, we have $p_j(i, h, s_{\min}) \leq q(i, h, s_{\min}) < q(i, h, s)$, and therefore if player k waited until $q(i, h, s)$ instead of dropping out at $p_k(i, h, s)$, the probability that someone dropped out in between would be positive, and hence the payoff would be strictly greater (by the definition of function $q(\cdot)$, player k with value s strictly prefers being in position $i-1$ or higher at any price less than $q(i, h, s)$ to being in position i at price b_{i+1}), which is impossible in equilibrium. Hence, $p_k(i, h, s) = q(i, h, s)$ for all $s > s_{\min}$. By continuity, we also have $p_k(i, h, s_{\min}) = q(i, h, s_{\min})$.

REFERENCES

- Aggarwal, Gagan, and Jason D. Hartline. 2005. "Knapsack Auctions." Paper presented at the First Workshop on Sponsored Search Auctions, Vancouver, BC.
- Asdemir, Kursad. 2006. "Bidding Patterns in Search Engine Auctions." Paper presented at the Second Workshop on Sponsored Search Auctions, Ann Arbor, MI.
- Bergemann, Dirk, and Stephen Morris. 2005. "Robust Mechanism Design." *Econometrica*, 73(6): 1771–1813.
- Binmore, Ken, and Paul Klemperer. 2002. "The Biggest Auction Ever: The Sale of the British 3G Telecom Licenses." *Economic Journal*, 112(478): C74–96.
- Brooks, Nico. 2004. *The Atlas Rank Report—Part II: How Search Engine Rank Impacts Conversion*. Seattle: Atlas Institute.
- Clarke, Edward H. 1971. "Multipart Pricing of Public Goods." *Public Choice*, 11(0): 17–33.
- Crawford, Vincent P., and Elsie M. Knoer. 1981. "Job Matching with Heterogeneous Firms and Workers." *Econometrica*, 49(2): 437–50.
- Demange, Gabrielle, David Gale, and Marilda Sotomayor. 1986. "Multi-Item Auctions." *Journal of Political Economy*, 94(4): 863–72.
- Edelman, Benjamin, and Michael Ostrovsky. Forthcoming. "Strategic Bidder Behavior in Sponsored Search Auctions." *Decision Support Systems*.
- Groves, Theodore. 1973. "Incentives in Teams." *Econometrica*, 41(4): 617–31.
- Hatfield, John William, and Paul R. Milgrom. 2005. "Matching with Contracts." *American Economic Review*, 95(4): 913–35.
- Jehiel, Philippe, and Benny Moldovanu. 2001. "Efficient Design with Interdependent Valuations." *Econometrica*, 69(5): 1237–59.
- Jehiel, Philippe, Moritz Meyer-ter-Vehn, Benny Moldovanu, and William R. Zame. 2006. "The Limits of Ex Post Implementation." *Econometrica*, 74(3): 585–610.
- Kagel, John H., Ronald M. Harstad, and Dan Levin. 1987. "Information Impact and Allocation Rules in Auctions with Affiliated Private Values: A Laboratory Study." *Econometrica*, 55(6): 1275–1304.
- Kelso, Alexander S., Jr., and Vincent P. Crawford. 1982. "Job Matching, Coalition Formation, and Gross Substitutes." *Econometrica*, 50(6): 1483–1504.
- Kittsteiner, Thomas, and Benny Moldovanu. 2005. *Management Science*, 51(2): 236–48.
- Leonard, Herman B. 1983. "Elicitation of Honest Preferences for the Assignment of Individuals to Positions." *Journal of Political Economy*, 91(3): 461–79.

- Mahdian, Mohammad, Hamid Nazerzadeh, and Amin Saberi.** 2006. "AdWords Allocation Problem with Unreliable Estimates." Unpublished.
- Mas-Colell, Andreu, Michael D. Whinston, and Jerry R. Green.** 1995. *Microeconomic Theory*. New York: Oxford University Press.
- McAdams, David, and Michael Schwarz.** Forthcoming. "Who Pays When Auction Rules are Bent?" *International Journal of Industrial Organization*.
- Meek, Christopher, David M. Chickering, and David B. Wilson.** 2005. "Stochastic and Contingent-Payment Auctions." Paper presented at the First Workshop on Sponsored Search Auctions, Vancouver, BC.
- Mehta, Aranyak, Amin Saberi, Umesh Vazirani, and Vijay Vazirani.** 2005. "AdWords and Generalized On-line Matching." In *Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science*, 264-73. Washington, DC: IEEE Computer Society.
- Milgrom, Paul.** 2000. "Putting Auction Theory to Work: The Simultaneous Ascending Auction." *Journal of Political Economy*, 108(2): 245-72.
- Moldovanu, Benny, and Aner Sela.** 2001. "The Optimal Allocation of Prizes in Contests." *American Economic Review*, 91(3): 542-58.
- Roth, Alvin E.** 1984. "The Evolution of the Labor Market for Medical Interns and Residents: A Case Study in Game Theory." *Journal of Political Economy*, 92(6): 991-1016.
- Roth, Alvin E., and Axel Ockenfels.** 2002. "Last-Minute Bidding and the Rules for Ending Second-Price Auctions: Evidence from eBay and Amazon Auctions on the Internet." *American Economic Review*, 92(4): 1093-1103.
- Shapley, Lloyd S., and Martin Shubik.** 1971. "The Assignment Game I: The Core." *International Journal of Game Theory*, 1(1): 111-30.
- Varian, Hal R.** Forthcoming. "Position Auctions." *International Journal of Industrial Organization*.
- Vickrey, William.** 1961. "Counterspeculation, Auctions, and Competitive Sealed Tenders." *Journal of Finance*, 16(1): 8-37.
- Wilson, Robert.** 2002. "Architecture of Power Markets." *Econometrica*, 70(4): 1299-1340
- Zhang, Xiaoquan.** 2005. "Finding Edgeworth Cycles in Online Advertising Auctions." Unpublished.