

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

Benjamin Edelman
Harvard University
bedelman@fas.harvard.edu

Michael Ostrovsky
Stanford University
ostrovsky@gsb.stanford.edu

Michael Schwarz*
UC Berkeley and NBER
mschwarz@berkeley.edu

August 12, 2005.

Abstract

We investigate the “generalized second price” auction (GSP), a new auction mechanism, which is used by Internet search engines to sell online advertising. GSP is tailored to its unique environment, and neither the mechanism nor the environment have previously been studied in the auction literature. Although GSP looks similar to the Vickrey auction, its properties are very different. In particular, unlike the Vickrey auction, GSP generally does not have an equilibrium in dominant strategies, and truth-telling is not an equilibrium of GSP.

We consider a subset of the equilibria of GSP that we refer to as “locally envy-free.” We show that the seller’s revenue in the dominant strategy equilibrium of VCG is equal to its revenue in its least preferred locally envy-free equilibrium. We also describe the generalized English auction that corresponds to the generalized second-price auction and show that it has a unique equilibrium. This is an ex post equilibrium that results in the same payoffs to all players as the dominant strategy equilibrium of VCG. The generalized English auction is a robust mechanism in the sense that the outcome depends only on the values of the participants and not on their beliefs about each others’ values.

*We are grateful to David Pennock and Yahoo! for the data and overall help. We also thank Drew Fudenberg, Louis Kaplow, Paul Milgrom, Muriel Niederle, Ariel Pakes, and Al Roth for helpful discussions.

1 Introduction

This paper investigates a new auction mechanism, which we call the “generalized second price” auction, or GSP. This mechanism was first implemented by Google in February 2002, and its variations are now used to sell billions of dollars of Internet advertising every year. GSP is tailored to the unique environment of the market for online ads, and neither the environment nor the mechanism have previously been studied in the auction literature.

Let us briefly describe how these auctions work. When an Internet user enters a keyword (i.e., a search term) 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., ads submitted by advertisers. 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. Facing scarcity in advertisement positions, search engines need a system for allocating the positions to advertisers, and auctions are a natural choice. Currently, the auction mechanisms most widely used by search engines are different variants of the generalized second price auction.

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, the ad with the highest bid at the top, the ad with the next highest bid in the second position, and so on. If a user subsequently clicks on an ad in position k , 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 $k + 1$. If a search engine offered only one advertisement per result page, this auction mechanism would be equivalent to the standard second price, or Vickrey, auction. With multiple positions available, the GSP generalizes the second price auction. (Hence the name.) Here, a winner pays the next highest bidder’s bid. But as we will explain later, the multi-unit GSP auction is no longer equivalent to the Vickrey auction, for multi-unit GSP lacks some of Vickrey’s desirable properties. In particular, unlike the Vickrey auction, GSP generally does not have an equilibrium in dominant strategies, and truth-telling is not an equilibrium of GSP.

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: over 98% of Google’s revenues, and, it is believed, over 50% of Yahoo!’s revenues are derived from sales via GSP auctions. To appreciate the size of the market dominated by GSP auctions, note that the combined market capitalization

of Google and Yahoo! is over \$125 billion. In comparison, the combined market capitalization of all US airlines is about \$20 billion.

A combination of several features makes the market for Internet advertising unique. First, bidding takes place continuously. For example, the advertiser with the second highest bid on a given keyword at some instant will be listed as the second sponsored link at that instant. But any other advertiser can revise his bid at any time, and the order of sponsored links and prices will change accordingly. These changes can be very rapid because advertisers can employ automated robots, including commercially available software, in responding to others' bids.

Second, the 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, much like in electricity markets. Of course, this environment is very different from a market for electricity: e.g., the marginal utility of advertisers from an additional click can be considered constant and hence can be reasonably well approximated by a single number, while the marginal benefit from electricity is rapidly diminishing, requiring bidders to submit entire demand curves; the market for electricity is two-sided, whereas sponsored search auctions are one-sided; etc.

Finally, unlike other centralized markets, where it is usually clear how to measure what is being sold, there is no obvious definition of a "unit" of Internet advertisement. 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 consumer. "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.¹ The specific sector of Internet advertising that we study, sponsored search auctions, has converged to pay-per-click pricing after several years of evolution.

The history of this market is of interest as a case study of whether, how, and how quickly markets come to address their structural shortcomings (See Section 2). Many important auctions have recently been designed essentially from scratch, entirely replacing completely different historical allocation mechanisms: radio spectrum auctions (Milgrom, 2000; Binmore and Klemperer, 2002), electricity auctions (Wilson, 2002), and others. In contrast, reminiscent of the gradual evolution of medical residency match rules (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

¹A prominent example of "pay-per-transaction," and even "pay-per-dollar of revenue" is Amazon.com's Associates program, <http://www.amazon.com/gp/browse.html?&node=3435371> (accessed August 10, 2005). Under this program, a website that sends customers to Amazon.com receives a percentage of customers' purchases. "Pay-per-impression" advertising, in the form of banner ads, is very popular on major Internet portals, such as yahoo.com, msn.com, and aol.com.

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.

Since GSP evolved in the market for online advertising, its rules reflect the environment’s unique characteristics. GSP insists that for each keyword, bidders submit a single bid—even though several different items are for sale: position 1 is very different from position 5. The 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. Consequently, even though the environment is multi-object, buyer valuations can be adequately represented by one-dimensional types. However, one bid per keyword is probably not sufficiently expressive to fully convey the preferences of the bidders: e.g., it does not allow for the possibility that the 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, and we will likewise ignore these possibilities in our model, presented in Section 3. One important possibility that we abstract away from is that ads from different advertisers 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 CTRs, ranks the bidders purely in the order of decreasing bids, and charges the next-highest bidder’s bid. In contrast, Google reportedly multiplies each advertiser’s bid by its CTR to compute its expected revenue, ranks the ads by these expected revenues, and then charges each bidder the smallest amount sufficient to exceed the expected revenue of the next bidder’s bid times his CTR. In our analysis, we assume that all bidders have identical CTRs, in which case Google’s and Yahoo!’s mechanisms are identical. The analysis would remain largely the same if all advertisers’ CTRs in any given position x differed by a constant factor.²

We begin our analysis of the model in Section 4. Since advertisers can change their bids frequently, this environment can be modeled as a continuous or an infinitely repeated game. Our analysis begins with a static stage game, thinking about the restrictions that the dynamic structure imposes on the behavior of the players. We introduce these restrictions, and call the static equilibria satisfying them “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 Vickrey 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 bidders. Consequently, in *any* locally

²The analysis would have to change considerably if there are specific advertiser-position effects. The magnitude of these specific advertiser-position effects is ultimately an empirical question, and we do not have the kind of data that would allow us to answer it, but 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 average CTR across all positions), we believe it to be small.

envy-free equilibrium of GSP, the total expected revenue to the seller is at least as high as in the dominant-strategy equilibrium of the Vickrey auction.

In Section 5, we introduce the generalized English auction, which can be viewed as a dynamic game corresponding to the GSP auction. Although the generalized English auction is not dominant-strategy solvable, it has a unique perfect Bayesian equilibrium. In this equilibrium, all players receive the same payoffs as in the dominant strategy equilibrium of VCG. A remarkable feature of the generalized English auction is that it is a robust mechanism in the sense that the outcome depends only on the values of the participants and not on their beliefs about each others' values.

2 The Structure and History of Sponsored Search Auctions

We begin with a brief chronological review of the development of sponsored search auctions. Beginning in 1994, Internet advertisements were largely sold on a per-impression basis. Advertisers would pay flat fees to show their ads a fixed number of times (typically, one thousand showings, or “impressions”). Minimum contracts for advertising purchases were typically a few thousand dollars a month, and entry was slow.³

2.1 Generalized First-Price Auctions

In 1997, Overture (then GoTo; now part of Yahoo!) introduced a completely new model of selling 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 website, 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 thousand 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, 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 dynamic nature of the environment.

Example. Suppose there are two slots on a page and three bidders. An ad in the first slot receives 200 clicks per hour, while the second slot gets 100. Bidders 1, 2, and 3 have values per click of \$10, \$4, and \$2, respectively. Suppose bidder 2 bids \$2.01, to guarantee that he gets a slot. Then bidder 1 will not want to bid more than \$2.02—he does not need to pay more than that to

³See http://www.worlddata.com/wdnet8/articles/the_history_of_Internet_Advertising.htm and <http://www.zakon.org/robert/internet/timeline> (both accessed August 10, 2005).

get the top spot. But then bidder 2 will want to revise his bid to \$2.03 to get the top spot, bidder 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 bidders best respond to each other, they will want to revise their bids as often as possible. Figure 1 shows an example of this behavior on Overture.

2.2 Generalized Second-Price Auctions

Under the generalized first-price auction, the bidder who could react to its competitors' moves fastest had a substantial advantage. The mechanism therefore encouraged inefficient investments in gaming the system. It also created volatile prices that in turn caused allocative inefficiencies. Google recognized these problems when it introduced its own pay-per-click system, AdWords Select, in February 2002. Google also recognized that a bidder in position i will never want to pay more than one bid increment above the bid of the advertiser in position $(i + 1)$, and Google 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.⁴ Every agent submits a bid. Advertisers are arranged on the page in descending order of their bids. The agent in the first position pays a price per click that equals the bid of the second agent plus an increment; the second agent pays the price offered by the third agent plus an increment; and so forth.⁵

Example (continued). Let us now consider the payments in the environment of the previous example under GSP mechanism. If all advertisers bid truthfully, then bids are \$10, \$4, \$2. Payments in GSP will be \$4.01 and \$2.01. Truth-telling is indeed an equilibrium in this example, because no bidder can benefit by changing his bid. Note that total payments of bidders one and two are \$800 and \$200, respectively.

2.3 Generalized Second-Price and Vickrey Auctions

GSP shares an important property with the Vickrey-Clarke-Groves mechanism: both mechanisms set each agent's payment only based on the allocation and the bids of other players, not based on that agent's own bid. This is reflected in Google's marketing materials, which mention Vickrey and state that Google's "unique auction model uses Nobel Prize-winning economic theory to eliminate

⁴We focus on Overture's implementation, because Google's system is somewhat more complex. However, it is straightforward to generalize our analysis to Google's mechanism.

⁵For this version of GSP to be efficient, it is necessary that the number of clicks received by an advertisement in a given position depends on the ad's position and not on the advertiser's identity. Recognizing that this assumption may not hold if some ads attract more clicks than others, Google adjusts effective bids based on ads' click-through rates. But under the assumption that all ads have the same click-through rates conditional on position, Google's and Yahoo's versions of GSP are identical. With modification for Google's adjustment procedure, our results for the Yahoo version of GSP also hold for Google adjusted GSP.

... that feeling that you've paid too much".⁶ But GSP is not VCG. In particular, unlike the Vickrey 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.)

With only one slot, VCG and GSP would be identical. With several slots, the mechanisms are different: while GSP essentially charges bidder i the *bid* of bidder $i + 1$, VCG charges bidder i the *externality* that he imposes on others, i.e., the decrease in the value of clicks received by other agents because of i 's presence.

Example (continued). Let us compute VCG payments for the example considered above. The second agent's payment is \$200, as before. However, the payment of the first advertiser is now \$600: \$200 for the externality that he imposes on agent 3 (by forcing him out of position 2) and \$400 for the externality that he imposes on agent 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, if advertisers were to bid their true values under both mechanisms, revenues would always be higher under GSP.

2.4 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 a second-price mechanism, 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, because VCG revenues are lower than GSP revenues for the same bids and bidders 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 (Kagel, Harstad, and Levin, 1987). And of course, simply implementing and testing a new system may be costly. Consequently, new entrants such as Ask Jeeves and Microsoft Network⁷ have a comparative advantage over the established players in implementing VCG.

3 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

⁶See <https://www.google.com/adsense/afs.pdf> (accessed August 10, 2005).

⁷See http://www.irconnect.com/askjinc/pages/news_releases.html?d=83045 and <http://advertising.msn.com/searchadv/> (both accessed August 10, 2005).

bidders (advertisers). The (expected) number of clicks per period received by the bidder whose ad was placed in position i is α_i . The value per click to bidder k is s_k . Bidders are risk-neutral, and bidder 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 a descending order: for any j and k such that $j < k$, we have $\alpha_j > \alpha_k$.

We model the Generalized Second Price Auction (GSP) 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 time 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 bidder, respectively. If several advertisers submit the same bid, they are ordered randomly.⁸ The auction then allocates the top position to the agent with the highest bid, $g(1)$, the second position to $g(2)$, and so on, down to position $\min\{N, K\}$. Note that each bidder gets at most one object. If a user clicks on a bidder's link, the bidder's payment per click is equal to the next bidder's bid. So bidder $g(i)$'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 $\alpha_i (s^{(i)} - b^{(i+1)})$. If there are at least as many positions as advertisers ($N \geq K$), then the last bidder's payment $p^{(K)}$ is equal to zero.

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—the higher the bid, the better the position—but the payments are different. Each player'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 bidder 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 static game of complete information resembling a sealed-bid second-price auction and as a dynamic game of incomplete information resembling an 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. Alternatively, take any agent j who is allocated object k , and suppose that in his absence, object k would be allocated to some other agent j' . Agent j 's VCG payment is the drop in the aggregate utility of other players caused by his presence, *given that they can re-optimize the allocation of positions assigned to them*, whereas his GSP payment is the value of

⁸The actual practice at Overture is to show equal bids according to the order in which the bidders placed their bids.

player j' for object k , and is therefore (since all other players remain indifferent) equal to the drop in the aggregate utility of other players caused by his presence, *given that they cannot re-optimize*. Of course, the former is at most as large as the latter.

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 2. There are still three bidders, 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 bidder 1's payoff is equal to $(\$10 - \$4) * 200 = \$1200$. 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 = \$1592 > \1200 .

4 Static GSP and 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 agents 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. However, the strategies required to support such equilibria are usually quite complex, requiring precise knowledge of the environment and careful implementation. It may not be reasonable to expect the advertisers to be able to execute such strategies: they often manage thousands of keywords, and implementing sophisticated dynamic strategies for many keywords is likely to be prohibitively expensive and complex. In theory, advertisers could implement such strategies via automated robots, but at present they cannot do so: any bidding software has to be authorized by the search engines, and they are unlikely to allow sophisticated strategies that would allow bidders to collude and reduce revenues. For example, we have talked to several suppliers of bidding software for sponsored search auctions, and none of their products makes it possible to condition one's bid on the prior behavior of a specific competitor.

We therefore focus on simple strategies that bidders can reasonably execute, and we 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, bidders are likely to learn each other's values. Second, since bids can be changed at any time, stable bids must be static best responses to each other—otherwise an agent 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 static one-shot game of complete information. Third, what are the

simple strategies that a bidder 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 agent k bids b_k and is assigned to position $i + 1$ and agent k' bids $b_{k'} > b_k$ and is assigned to position i . Note that if k raises his bid, his own payoff does not change, but the payoff of the player above him decreases. Of course, player k' can retaliate, and the most he can do is to slightly underbid bidder k , effectively swapping places with him. If bidder 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 k should not want to “exchange” positions with the advertiser in position $k - 1$. We call such vectors of bids “locally envy-free.”

Definition 4 *An equilibrium of the static game Γ imposed by GSP is locally envy-free if a player cannot improve his payoff by exchanging bids with the player ranked one position above him. More formally, in a locally envy-free equilibrium, for any $i \leq \min\{N, K\}$, we have $\alpha_i s_{g(i)} - p^{(i)} \geq \alpha_{i-1} s_{g(i)} - p^{(i-1)}$.*

Of course, it is possible that in the dynamic game bids change over time, depending on the player's strategies and information structure. However, as long as the restrictions are satisfied, if the dynamic game ever converges to a static vector of bids, that static equilibrium should correspond to a “locally envy-free” equilibrium of the static game Γ induced by the 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 Crawford and Knoer (1981), Kelso and Crawford (1982), and Demange, Gale, and Sotomayor (1986), and has been studied in detail in a recent paper by Hatfield and Milgrom (2005). Note, however, that in our case the non-standard auction is very different from those in the above papers.

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

Lemma 5 *The outcome of any locally envy-free equilibrium of auction Γ is a stable assignment.*

Lemma 6 *If the number of bidders is greater than the number of available positions, then any stable assignment is an outcome of a locally envy-free equilibrium of auction Γ .*

The proof of Lemma 5 shows first that in any locally envy-free equilibrium, each player does not envy not only the player assigned to the position right above him, and also that each player

does not envy *any* other player: he cannot improve his payoff by swapping positions with anyone else.

We will now construct a particular locally envy-free equilibrium of game Γ . This equilibrium has two important properties. First, in this equilibrium bidders' 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 it is the *best* locally envy-free equilibrium for the bidders. 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^* . For each advertiser $i \in \{2, \dots, \min\{N, K\}\}$, bid b_i^* is equal to $\frac{p^{V,(i-1)}}{\alpha_{i-1}}$, where $p^{V,(j)}$ is the payment of bidder j in the dominant-strategy equilibrium of VCG. Bid b_1^* is equal to s_1 .⁹

Theorem 7 *Strategy profile B^* is an envy-free equilibrium of game Γ . In this equilibrium, each bidder's position and payment is 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^* .*

5 GSP and the Generalized English Auction

In the model analyzed in the last section, we assume that bidders 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 show that, in fact, there are simple strategies that can quickly lead to such an outcome: each advertiser starts bidding at zero and keeps increasing his bid as long as he finds it profitable to do so. This simple myopic procedure leads to an equilibrium outcome in which bidders have no incentives to change their bids. Remarkably, the outcome is the same as the worst (for the search engine) locally-envy free equilibrium of the static GSP: all players get VCG payoffs.

We model this procedure through an analogue of the standard English auction. The generalized English Auction entails a clock showing the current price (that continuously increases over time). A player's 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 last remaining advertiser is placed in the best advertisement position, 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 its payment per click is equal to the third-highest advertiser's bid, and so on. In other words, the vector of bids obtained in the generalized English auction is used to allocate the objects and compute the prices according to the rules of GSP, just as the rules of the second price auction can be used to compute the payments in the English auction.¹⁰ The following theorem shows that this

⁹This bid does not affect any bidder's payment and can be set equal to any value greater than b_2^* .

¹⁰If several players drop out simultaneously, one of them is chosen randomly. Whenever a player drops out, the clock is stopped, and other players are also allowed to drop out; again, if several ones want to drop out, one is chosen

game has a unique perfect Bayesian equilibrium. All players' payoffs in this equilibrium are equal to the VCG payoffs.

For concreteness, assume that there are $N \geq 2$ slots and $K = N + 1$ advertisers. (Cases in which $N \neq K + 1$ require only minor modifications in the proof.) Click-through rates α_n are commonly known, with $\alpha_{N+1} \equiv 0$. Bidders' per-click valuations s_i 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 $p_i(k, h, s_i)$, where k is the number of bidders remaining (including the bidder himself), and $h = (b_{k+1}, \dots, b_K)$ is the history of prices at which bidders $K, K - 1, \dots, k + 1$ have dropped out. (As a result, the price that bidder i would have to pay per click if he dropped out next is equal to b_{k+1} , unless the history is empty, in which case we say that $b_{k+1} \equiv 0$.) s_i is the value per click of bidder i , and p_i is the price at which the bidder drops out.

Theorem 8 *In the unique perfect Bayesian equilibrium of the generalized English auction, an advertiser with value s_i drops out at price $p_i(k, h, s_i) = s_i - \frac{\alpha_k}{\alpha_{k-1}}(s_i - b_{k+1})$. 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.*

The intuition of the proof is as follows. First, with k players remaining and the next highest bid equal to b_{k+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 k and paying b_{k+1} per click and getting position $k - 1$ 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 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 only drop out at price levels at which he will find it unprofitable to compete with them for the higher positions.

The result of Theorem 8 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 depend on other player's bids.

The equilibrium described in Theorem 8 is an ex post equilibrium. Indeed as long as all agents other than agent i follow the equilibrium strategy described in Theorem 8, it is a best response for agent i to follow the equilibrium strategy regardless of realization of other agents' values. Thus the outcome implemented by this mechanism does not depend on the common knowledge assumption, the outcome depends only on the realization of bidders' values and does not depend on agents' beliefs about each others' types.

randomly. If several players end up dropping 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 next highest, and so on.

Obviously, any dominant strategy solvable game has ex post equilibrium. The Generalized English Auction is not dominant strategy solvable, it is ex-post equilibrium because it implements a social choice function that can be implemented in dominant strategies (by VCG). This combination of properties is quite striking: payoffs are uniquely determined, even though there are no dominant strategies *and* there is no need for players to know the distributions of other players' types and the equilibrium is unique.¹¹ 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."

References

- [1] Bergemann, D., and S. Morris (2004), "Robust Mechanism Design," *Econometrica*, forthcoming.
- [2] Binmore, K., and P. Klemperer (2002), "The Biggest Auction Ever: The Sale of the British 3G Telecom Licences," *Economic Journal*, 112, 74–96.
- [3] Crawford, V., and E. Knoer (1981), "Job Matching with Heterogeneous Firms and Workers," *Econometrica*, 49, 437–450.
- [4] Demange, G., D. Gale, and M. Sotomayor (1986), "Multi-Item Auctions," *Journal of Political Economy*, 94, 863–872.
- [5] Hatfield, J., and P. Milgrom (2005), "Auctions, Matching and the Law of Aggregate Demand," *American Economic Review*, forthcoming.
- [6] Kagel, J., R. Harstad, and D. Levin (1987), "Information Impact and Allocation Rules in Auctions with Affiliated Private Values: A Laboratory Study," *Econometrica*, 55, 1275–1304.
- [7] Kelso, A., and V. Crawford (1982), "Job Matching, Coalition Formation, and Gross Substitutes," *Econometrica*, 50, 1483–1504.
- [8] Milgrom, P., (2000), "Putting Auction Theory to Work: The Simultaneous Ascending Auction," *Journal of Political Economy*, 108, 245–272.
- [9] Roth, A., (1984), "The Evolution of the Labor Market for Medical Interns and Residents: A Case Study in Game Theory," *Journal of Political Economy*, 92, 991–1016.
- [10] Shapley, L., and M. Shubik (1972), "The Assignment Game I: The Core," *International Journal of Game Theory*, 1, 111–130.
- [11] Vickrey, W., (1961), "Counterspeculation, Auctions, and Competitive Sealed Tenders," *Journal of Finance*, 16, 8–27.
- [12] Wilson, R., (2002), "Architecture of Electric Power Markets," *Econometrica*, 70, 1299–1340.

¹¹In this environment a mechanism that has unique ex post equilibrium in dominant strategies is also available however this is not the mechanism used in practice.

Appendix

Proof of Lemma 5

By definition, in any locally envy-free equilibrium outcome, no player can profitably re-match with the position assigned to the player right above him. Also, no player (*a*) can profitably re-match with a position assigned to a player below him (*b*)—if such a profitable re-matching existed, player *a* would find it profitable to slightly undercut player *b* in game Γ and get *b*'s position and payment. But this would contradict the assumption that we are in equilibrium.¹²

Hence, we only need to show that no player can profitably re-match with the position assigned to a player more than one spot above him. First, note that in any envy-free equilibrium, the resulting matching must be assortative, i.e., for any i , the player assigned to position i has a higher per-click valuation than the player assigned to position $i + 1$, and therefore, player 1 must be assigned to the top position, player 2—to the second-highest position, and so on.

Indeed, let s_i and s_{i+1} denote the values of players assigned to positions i and $i + 1$. Equilibrium restrictions imply that $\alpha_i s_i - p_i \geq \alpha_{i+1} s_i - p_{i+1}$ (nobody wants to move one position down), and envy-freeness implies that $\alpha_{i+1} s_{i+1} - p_{i+1} \geq \alpha_i s_{i+1} - p_i$ (nobody wants to move one position up). Manipulating the above inequalities yields $\alpha_i s_i - \alpha_i s_{i+1} + \alpha_{i+1} s_{i+1} \geq \alpha_{i+1} s_i$, thus $(\alpha_i - \alpha_{i+1}) s_i \geq (\alpha_i - \alpha_{i+1}) s_{i+1}$. Since $\alpha_i > \alpha_{i+1}$, we have $s_i \geq s_{i+1}$, and hence the envy-free equilibrium outcome must be an assortative match.

Now, let us show that no player can profitably re-match with the position assigned to a player more than one spot above him. Suppose agent k is considering re-matching with position $m < k - 1$. Since the equilibrium is envy-free, we have

$$\begin{aligned} \alpha_k s_k - p_k &\geq \alpha_{k-1} s_k - p_{k-1} \\ \alpha_{k-1} s_{k-1} - p_{k-1} &\geq \alpha_{k-2} s_{k-1} - p_{k-2} \\ &\vdots \\ \alpha_{m+1} s_{m+1} - p_{m+1} &\geq \alpha_m s_{m+1} - p_m. \end{aligned}$$

Since $\alpha_j > \alpha_{j+1}$ for any j , and $s_j > s_k$ for any $j < k$, the above inequalities remain valid after replacing s_j with s_k . Doing that, then adding all inequalities up, and canceling out the redundant elements, we get $\alpha_k s_k - p_k \geq \alpha_m s_k - p_m$. But that implies that advertiser k cannot re-match profitably with position m , and we are done.

Proof of Lemma 6

Take a stable assignment. By a result of Shapley and Shubik (1972), this assignment must be efficient, and hence assortative, and so we can talk about advertiser i being matched with position

¹²This argument relies on the fact that in equilibrium, no two (or more) players bid the same amount, which is straightforward to prove: since all players' per-click values are different, and all ties are broken randomly with equal probabilities, at least one of such players would find it profitable to bid slightly higher or slightly lower.

i , with associated payment p_i .

Let us construct an envy-free equilibrium with the corresponding outcome. Let $b_1 = s_1$ and $b_i = \frac{p_{i-1}}{\alpha_{i-1}}$ for $i > 1$. Let us show that this set of strategies is an envy-free equilibrium. First, note that for any i , $b_i > b_{i+1}$ (because otherwise we would have, for some i , $\frac{p_{i-1}}{\alpha_{i-1}} \leq \frac{p_i}{\alpha_i} \Rightarrow s_i - \frac{p_{i-1}}{\alpha_{i-1}} \geq s_i - \frac{p_i}{\alpha_i} \Rightarrow \alpha_{i-1}(s_i - p_{i-1}) > \alpha_i(s_i - p_i)$, which would imply that player i could re-match 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 re-matching with the corresponding position in the assignment game. To see that this equilibrium is envy-free, note that the payoff from swapping with the bidder above is exactly equal to the payoff from re-matching with that player's position in the assignment game.

Proof of Theorem 7

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

$$\begin{aligned}
\frac{p^{V,(i-1)}}{\alpha_{i-1}} &> \frac{p^{V,(i)}}{\alpha_i} \\
&\Downarrow \\
\frac{(\alpha_{i-1} - \alpha_i)s_i + p^{V,(i)}}{\alpha_{i-1}} &> \frac{p^{V,(i)}}{\alpha_i} \\
&\Downarrow \\
\alpha_i(\alpha_{i-1} - \alpha_i)s_i &> (\alpha_{i-1} - \alpha_i)p^{V,(i)} \\
&\Downarrow \\
\alpha_i s_i &> p^{V,(i)}.
\end{aligned}$$

For $i = 1$, $b_i^* > b_{i+1}^*$ is equivalent to

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

To see that for any i , $\alpha_i s_i > p^{V,(i)}$, 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 i 's value per click were slightly lower, e.g., $s_i - \Delta$ instead of s_i , $\Delta < s_i - s_{i+1}$, 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 $p^{V,(i)} \leq \alpha_i(s_i - \Delta) < \alpha_i s_i$. Thus, for any i , $b_i^* > b_{i+1}^*$, 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 i can benefit by bidding less than b_i^* , suppose that he bids an amount $b' < b_i^*$ that puts him in position $i' > i$. Then, by construction, his payment will be equal to the amount that he would need to pay to be in position i' 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 i can benefit by bidding more than b_i^* , suppose that he bids an amount $b' > b_i^*$ that puts him in position $i' < i$. Then the net payoff from this deviation is equal to $(\alpha_{i'} - \alpha_i)s_i - (\alpha_{i'}b_{i'}^* - \alpha_i b_{i+1}) < (\alpha_{i'} - \alpha_i)s_i - (\alpha_{i'}b_{i'+1}^* - \alpha_i b_{i+1}) = \sum_{j=i'}^{i-1} (\alpha_j - \alpha_{j+1})s_i - \sum_{j=i'}^{i-1} (p^{V,(j)} - p^{V,(j+1)}) = \sum_{j=i'}^{i-1} (\alpha_j - \alpha_{j+1})s_i - \sum_{j=i'}^{i-1} (\alpha_j - \alpha_{j+1})s_{j+1}$. But since $s_i \leq s_{j+1}$ for any $j < i$, 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 i swapped his bids with bidder $i-1$, his payoff would change by $(\alpha_{i-1} - \alpha_i)s_i - (\alpha_{i-1}b_i^* - \alpha_i b_{i+1}) = (\alpha_{i-1} - \alpha_i)s_i - (p^{V,(i-1)} - p^{V,(i)}) = (\alpha_{i-1} - \alpha_i)s_i - (p^{V,(i-1)} - p^{V,(i)}) = (\alpha_{i-1} - \alpha_i)s_i - (\alpha_{i-1} - \alpha_i)s_i = 0$. 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. A standard result of matching theory states that there exists an assignment that is the best stable assignment for all advertisers, and the worst stable assignment for all positions. A stable 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 buyer-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 re-match 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 buyer-optimal stable assignment, $p_{K-1} = p_{K-1}^V$.

Proceeding by induction, we get $p_k = p_k^V$ for any $k \leq K$ in the buyer-optimal (and therefore seller-pessimal) stable assignment, and so in any envy-free equilibrium of game Γ , the total revenue of the seller is at least as high as $\sum_{k=1}^K p_k^V$.

Proof of Theorem 8

TBA.



(a) 14 hours

(b) 1 week

Figure 1: “Sawtooth” pattern

Figure 1(a) presents the top bids, in dollars, for a specific keyword, every 15 minutes from 12:15 AM to 2:15 PM on July 18, 2002. There are two advertisers competing for the top spot. As indicated at point A, bidding starts with both of them below their maximum bids. The first advertiser (or, most likely, his robot) recognizes an opportunity to obtain the first listing position by raising the second bidder’s bid by \$0.01, and does so. But then the second advertiser sees that it has been outbid, and raises its bid in turn. This process continues until the bids reach one advertiser’s maximum bid (say, the first advertiser’s maximum), as shown at B. The first advertiser’s autobidder can no longer increase its bid to obtain the first place, so it instead looks to avoid overspending, which it does by lowering its bid to \$0.01 more than the third-place bidder, yielding the bid at C. But then the second advertiser sees that it can still obtain the first place by bidding \$0.01 more than the first advertiser’s newly-lowered bid. Bidding therefore begins to increase, yielding more and more “teeth” in the sawtooth pattern. Figure 1(b) shows the continuation of this pattern for the same keyword, for one week. We do not know what this keyword is, since in the data that we have received from Overture the keywords are replaced with identifying numbers. However, the sawtooth pattern appears to be common for most keywords.