2008

Generalized Second Price Auctions with Hierarchical Bidding

Kevin Bryan
Virginia Commonwealth University

GENERALIZED SECOND PRICE AUCTIONS WITH HIERARCHICAL BIDDING

A Thesis submitted in partial fulfillment of the requirements for the degree of Master of Science at Virginia Commonwealth University.

by

KEVIN BRYAN
M.A., Boston University, 2006
B.A., Boston University, 2006

Director: JILL HARDIN
ASSOCIATE PROFESSOR, DEPT. OF STATISTICAL SCIENCES AND OPERATIONS RESEARCH

Virginia Commonwealth University
Richmond, Virginia
August 2008

Abstract

GENERALIZED SECOND PRICE AUCTIONS WITH HIERARCHICAL BIDDING

By Kevin Bryan, M.S.

A Thesis submitted in partial fulfillment of the requirements for the degree of Master of Science at Virginia Commonwealth University.

Virginia Commonwealth University, 2008

Major Director: Jill Hardin
Associate Professor, Statistical Sciences and Operations Research

The sale of text advertisements on search engines using an auction format called Generalized Second Price (GSP) has become increasingly common. GSP is unique in that it allows bidders to revise their bid if they are unhappy with the result of the auction, and because the auction sells multiple units of a related good simultaneously. We model this sale as a hierarchical game with complete information, allowing one potential bidder to bid in each stage. The hierarchical game has an entirely different set of equilibria from the simultaneous bid game studied in earlier research on this auction. Under hierarchical bidding, Vickrey-Clarke-Groves guarantees higher auctioneer revenue than any equilibrium in GSP.

1 Introduction

Text advertisements on search engines such as Google and Yahoo, responsible for over ten billion dollars of annual revenue at those two companies alone, are sold in a Generalized Second Price (GSP) auction. In this auction, firms who wish to display their advertisement when a given keyword is searched enter an auction for that keyword. This auction runs continuously, and a firm's bid reflects the amount that firm is willing to pay if its ad is clicked by a search engine user (a *clickthrough*). At any given time, the advertiser with the highest bid per clickthrough is given the ad slot at the top of the webpage, the advertiser with the second-highest bid per clickthrough is given the second ad slot, and so on. Each slot is associated with a clickthrough rate, which represents how often the slot attracts a clickthrough in a given time period. After each bid, bidders are told where their ad will be placed, and how much they will have to pay per clickthrough. If they are unsatisfied with their current position, bidders are allowed to revise their bid at any time. The firm in slot i pays the bid of the firm in slot $i + 1$ for each clickthrough, hence the name Generalized Second Price. Naively, GSP appears similar to standard, second-price Vickrey auctions, which give the auctioneer revenue at least as high as any other auction format (*revenue-optimality*), lead to the bidder with the highest willingness to

pay receiving the top ad slot (*efficiency*), and induce bidders to simply submit their willingness to pay as their bid (*truthfulness*); indeed, Google's own advertising materials were making these claims as of 2006.

Given the importance of GSP to internet firms, there has been substantial interest in recent years concerning properties of this auction. Previous research assumes that firms bid in stages, with the bids from each player at each stage being sent to the auctioneer *simultaneously*. Equilibria to a simultaneous game, in the Nash sense, are a vector of bids such that no bidder, in hindsight, wishes to unilaterally deviate from her strategy. This line of research has shown that, unlike second-price Vickrey auctions, Generalized Second Price auctions are not necessarily efficient - that is, GSP may lead to allocations that do not maximize the sum of payoffs to the auctioneer and the auction participants - and do not induce truthful bidding as a dominant strategy.

We propose an alternative structure, where bids are sent by each bidder in a hierarchy. An auction where, during any given momemt, one bidder is picked at random to change her bid if she desires, taking previous bids as given, fits this model. Since any firm can potentially be the first bidder in a sequence, we define equilibria in this auction as the intersection of equilibria of every possible perturbation of a hierarchical game. Though many sequential games have no equilibria in the intersec-

tion of every possible perturbation, the GSP auction has a computable analytic set of equilibria in this intersection. With this solution concept, equilibrium bidding is guaranteed to be efficient, unlike in the simultaneous bid solution. Further, the auctioneer revenue in the hierarchical game is shown to always be lower than Vickrey-Clarke-Groves (VCG) revenue, or the revenue to the auctioneer from distributing slots using the Vickrey-Clarke-Groves mechanism described in Section 2. This is the opposite result of the well-known paper by Edelman et al (2007), who solve for "locally envy-free" simultaneous bid equilibria, defined as Nash equilibria such that no bidder wants to switch bids with any other bidder. Locally envy-free equilibria are guaranteed to be *higher* than VCG revenue. That result has led many researchers to dismiss VCG as an allocative mechanism in the case of text advertisement sales. However, since VCG guarantees efficiency, is a truthful mechanism, and delivers higher revenue than any hierarchical GSP equilibrium, our results suggest that it may be preferable for search engines running auctions of this type to sell by VCG rather than GSP.

The most straightforward discussion of simultaneous bid GSP is Edelman et al (2007), which first presented many of the results described in Section 3 of this paper. Edelman and Ostrovsky (2007) discuss properties of the first-price auction used originally by Overture/Yahoo! to

sell keyword advertisements. Mehta et al (2007) and Borgs et al (2005) consider the GSP problem when bidders are budget constrained. In particular, when bidders are budget-constrained, it is optimal to bid as high as possible within your equilibrium interval in order to push your opponent out of the auction as quickly as possible. Zhou and Lukose (2007) consider a "vindictive equilibrium" where bidders not only maximize their own profit, but also try to minimize one or more opponent's profits. They show that with two vindictive bidders, equilibria exist and are efficient, but that Nash equilibria do not generally exist with three or more vindictive bidders. Aggarwal et al (2006) and Varian (2007) extend GSP by allowing each advertisement to have a different quality, and therefore a different clickthrough rate in any slot; this mimics the actual procedure used at Google. Borgers et al (2007) also extend GSP, but by giving payoffs that include an impression value, or a value that depends simply on having the ad shown rather than on having the ad clicked. Athey and Ellison (2008) allow consumers to vary their search habits depending on the quality of displayed advertisements, and discuss the usefulness of reserve prices in this context.

Cary et al (2008) is most similar to the spirit of this paper. They allow bids to be submitted sequentially, and discuss an algorithm called "Balanced Bidding" which converges to a Nash equilibrium of the static

game with the same ordering and payments as Vickrey-Clarke-Groves. But though the Balanced Bidding algorithm has asynchronous bidding, bidders do not take into account the nature in which their bid will affect the bids of future bidders. Because of this, the structure is not hierarchical, as in this paper, and the payments for every bidder are higher than in the hierarchical game.

This article proceeds as follows. In Section 2, we provide general background on concepts in game theory, auction theory and mechanism design that are used in later proofs. In Section 3, we describe the auction model and review the literature using the simultaneous solution. In Section 4, we describe the solution to the hierarchical game and its properties. Finally, in Section 5, we show that endogenizing the number of slots available is always in the best interest of the auctioneer, and that this leads hierarchical bid GSP auctions to be inefficient, unlike VCG.

2 Background

The proofs developed in this paper draw on definitions and theorems from three intellectual strands: game theory, auction theory and mechanism design. In this section, we briefly introduce concepts from those areas. In particular, in game theory we discuss equilibrium definitions, extensive form games, Stackelberg equilibria, and hierarchical games. In auction theory, we discuss the main results concerning first and second-

price single-good auctions and multiple-slot auctions, as well as the Revenue Equivalence Theorem. Finally, we introduce the Vickrey-Clarke-Groves mechanism.

2.1 Game Theory

A game is a situation of strategic interdependence. In many social science problems, one agent's optimization problem involves payoffs which depend on the action chosen by other agents. A canonical example of a game is the Prisoner's Dilemma, where two prisoners must independently decide whether to confess or lie about a crime, with the jail term for both depending on the actions taken by both. Though situations of strategic interdependence have surely been studied throughout history, and though specific games have been solved at least since Cournot (1838), modern game theory is considered to have begun with the work of Zermelo (1913), Borel (1921), and von Neumann and Morganstern (1944). A general history of game theory's development can be found in Dimand and Dimand (1996). Myerson (1997) provides a more rigorous introduction to the topics presented in this section.

A full description of a game involves knowing the players, the rules (including what actions and information are available to every player at all possible states of the game, represented by nodes), the outcomes

of every possible set of actions, and the payoffs to each player given those outcomes. This paper represents games in the extensive form Γ. Formally, an extensive form game consists of

- A set of nodes N, a set of actions A and a set of players P.
- A function p specifying a single node preceding each node except an initial node n_0, which has no predecessor.
- A function $a : N/n_0 \rightarrow A$ giving the unique action that leads to any noninitial node from its immediate predecessor.
- A collection of information sets X consisting of all possible sets of information known by the players, and a function $H : N \rightarrow X$ assigning each node to an information set. The set of actions available at any node with identical information must be identical.
- A collection of payoff functions U assigning payoffs to each player at each terminal node, where a terminal node is a node with no successor nodes.

That is, an extensive form game can be fully described as

$$\Gamma(N, A, P, p(.), a(.), X, H(.), U)$$

. A strategy for each player is a function $S : X \rightarrow A$ specifying the

Figure 1: The simultaneous game

actions taken by an agent at every possible information state, such that $S(X)$ is a possible action at that information state. More extensive details can be found in Mas-Colell, Whinston and Green (1995).

Example 1 (Myerson 1997) *Consider a game where two players choose left or right simultaneously. If they both choose left, they both receive payoff 2. If they both choose right, Player 1 receives 3 and Player 2 receives 1. If Player 1 chooses left and Player 2 chooses right, Player 1 receives 4 and Player 2 receives 0. Finally, if Player 1 chooses right and Player 2 left, then Player 1 receives 1 and Player 2 receives 0. Figure 1 displays this game.*

The two numbers above each node represent who is moving at each

node, and what information is available to that player. In this case, information set 1 implies that neither player has moved, and information set 2 implies that Player 1 has moved but that Player 2 does not know what Player 1 has done. Note that no matter what Player 2 does, Player 1 would be better off playing left. Player 2, realizing that Player 1 will always play left, will then also play left, and both players receive a payoff of 2.

The most well-known solution concept for a game is the pure strategy Nash equilibrium. A pure strategy Nash equilibrium is a set of strategies $S = \{S_1, S_{2....}\}$ such that, given payoff functions U_i and fixed bids from all other players $S_{\tilde{\imath}}$,

$$U_i(S_i, S_{\tilde{\imath}}) \geq U_i(S_i', S_{\tilde{\imath}}) \text{ for all possible strategies } S_i' \text{ and all players } i.$$

That is, taking other agents' actions as given, a strategy for a player is an equilibrium strategy if it is the best possible response to those other agent's actions. In the left-right game, (Left, Right) cannot be a Nash equilibrium because, if Player 2 knows that Player 1 will go Left, then Player 2 will also go Left. The only Nash equilibrium is (Left, Left).

Extensive form games can easily handle sequential moves. Consider the game in Figure 2, which is identical to that in Figure 1 except for

Figure 2: The sequential game

the information sets.

Information set 2 implies that Player 2 knows Player 1 has moved left. Information set 3 means that Player 2 knows Player 1 has moved right. That is, Player 1 now moves before Player 2, and Player 2 knows what Player 1 has done. This is called a game of perfect information - every player knows every other player's past actions, and therefore the information state at every node is unique. If Player 1 plays right, then Player 2 maximizes her payoff by moving right, with final payoffs (3,1). If Player 1 plays left, then Player 2 will also play left, and the payoffs will be (2,2). Knowing this, Player 1 will move right, and the final payoffs will be (3,1). Note that this is different from the simultaneous-move

equilibrium.

In general, moving first in a sequential game is advantageous, but this is not always the case. Consider a matching game, where both players choose Heads or Tails, Player 1 wins if there is no match, and Player 2 wins if there is a match. The player who moves first will always lose in a sequential game, but will win half the time playing optimally in the simultaneous game.

Equilibria in sequential games are slightly more complicated than equilibria in one-move games. In particular, game theorists are generally only interested in a set of equilibria called subgame-perfect equilibria. A subgame is a subset of a game containing only a single decision node and its successor nodes, such that if x is a node in the subgame, every other node x' with the same information state as x is also in the subgame. Note that an entire game is a subgame. A set of strategies S is a subgame-perfect Nash equilibrium (SPNE) if and only if the strategies are a Nash equilibrium in every subgame. The reason for this definition is to get rid of strategies based on non-credible threats. For instance, in the sequential left-right game, Player 2 may threaten to play left should Player 1 play right, giving Player 1 a payoff of 1 if he plays right, but a payoff of at least 2 if he plays left. This threat is not credible, however - once Player 1 plays right, a new subgame begins where the optimal play

for Player 2 is right, not left. It has been proven that every finite game of perfect information has a pure strategy SPNE.

A two-stage sequential game, with one player choosing her action in the first stage, and n other players choosing their actions simultaneously in the second stage, is commonly called the Stackelberg game. The Stackelberg game originated in article by Stackelberg (1934), in which the German economist considered production in an industry consisting of one large "leader" firm and n smaller "follower" firms. The Stackelberg leader is often, but not always, able to extract a greater payoff in the sequential game than in a simultaneous game, as seen in the sequential left-right example. A generalization of the Stackelberg game is the n-hierarchy game, where n firms act in sequence, with firm i acting as a follower for firms with rank $i' < i$ and as a leader of firms with rank $i' > i$. Hierarchical games have been used to examine, for instance, the amendment process in a legislature (Baron and Ferejohn 1989), sequential prisoners' dilemmas (Vukov and Szabo 2005) and economic market formation (Boyer and Moreaux 1986).

2.2 Auction Theory

There are many ways through which objects may be allocated to individuals - a common mechanism is called a *price*, where the first agent

who arrives willing to pay a set amount gets the object. Auctions are often a useful way to allocate objects when potential buyers' willingness to pay is unknown and when the cost of holding the auction, as opposed to simply setting a price, is low. An auction, formally, is a mechanism through which agents' willingness to pay for an item is elicited (with *bids*) given a set of rules for how an item is allocated after those bids. Auctions are *universal* in that an auction mechanism does not depend on the item being auctioned. Further, auctions are *anonymous* in that the identity of bidders plays no role in determining the allocation of the item and the resulting payments (see, for example, Krishna 2002). The study of bidder behavior in auctions is a subdiscipline of game theory.

Auctions have a long and colorful history. Herodotus, in his *History of the Persian Wars,* discusses the Babylonian custom of auctioning wives:

> Once a year in each village the maidens of age to marry were collected all together into one place; while the men stood round them in a circle. Then a herald called up the damsels one by one, and offered them for sale. He began with the most beautiful. When she was sold for no small sum of money, he offered for sale the one who came next to her in beauty. All of them were sold to be wives. The richest of

> the Babylonians who wished to wed bid against each other for the loveliest maidens, while the humbler wife-seekers, who were indifferent about beauty, took the more homely damsels with marriage-portions. (Herodotus, I.196)

Since that time, auctions have been used to sell Dutch tulips, fine art, slaves, wireless spectra, the belongings of deceased Chinese monks, government procurement contracts, and even, at one point, the entire Roman Empire (Klemperer 2004, Damianov 2005 and Shubik 1983).[1] In this section, we briefly introduce the four most common single-unit auction formats, the Revenue Equivalence Theorem, the concept of optimal auctions, and first and second-price multiple-unit auctions.

The four most common auction formats are the English Auction, the Dutch Auction, the sealed-bid first-price auction, and the sealed-price second-price auction.

In the English auction, believed to be the format used to auction Babylonian wives, an auctioneer calls out a low price and potential bidders indicate whether they are interested in buying at that price. The auctioneer then begins increasing the price by small increments as long

[1] The Praetorian Guard auctioned off the Roman Empire to whoever was willing to pay the highest price per man to the Guard. Didius Julianus offered 25,000 sesterces per man, and was duly appointed Emperor. Generals in the provinces, however, were enraged at the auction, and within two months, Septimius Severus reached Rome and had Didius Julianus beheaded.

as at least two bidders remain interested. The last bidder to remain wins the item and pays an amount equal to the price at which the second-to-last bidder dropped out.

A Dutch auction (so called because it was used to auction tulips in 17th century Holland) proceeds in reverse; the auctioneer begins by calling out a very high price, and lowers the price by small increments until one bidder enters. The first bidder to enter wins the item at the given price. Together, Dutch and English auctions are called open auctions, since each bid is public rather than sealed.

In a sealed-bid first-price auction, bidders submit their willingness to pay, or *private value*, for an item in sealed envelopes; the highest bid submitted wins the item, paying the bid submitted. A sealed-bid second-price auction is similar except that the winning bidder pays the second-highest bid submitted. In the sealed-bid, second-price strategy, it is optimal for every bidder to bid their private value; if the winning bid is above a bidder's private value, then that bidder would have had to pay more than their private value to win the item, and therefore would not have wanted to win. If the winning bid is below a bidder's private value, the bidder wins the item and pays the second-highest private value. Since the winning bidder cannot control the second-highest bidder's bid, she can do no better than winning the item and

paying the second-highest private value.

An important feature of auctions is that each agent's private value is unknown to other bidders and to the auctioneer. If private values were known to the auctioneer, of course, it would be straightforward to simply allocate the item to the bidder with the highest private value and charge her exactly that amount. In this paper, we consider only private value auctions, where each bidder knows exactly her private value, rather than common value auctions, where the object's worth is unknown to the bidders at the start of the auction. An example of a common value auction might be the sale of an unexplored oil field; each firm may have some estimate of how much oil is in the field, but the bidders will also hope to improve their estimate based on the bidding behavior of other, knowledgeable firms.

Note that the Dutch auction and the sealed-bid first-price auction are equivalent. The only information learned during a Dutch auction is the value of the first bid submitted, which causes the auction to end. There is no strategic difference between offering a bid for a certain price in a sealed-bid first-price auction and planning to offer a bid when the Dutch auction reaches that price. Moreover, with independently-distributed private values, English auctions and sealed-bid second-price auctions have equivalent strategies. In the English auction, it is clearly optimal

to stay in the auction until the price passes a bidder's willingness to pay. The item will then be allocated to the bidder with the highest private value, and the final payment will be exactly the second highest private value. The final allocation and payments in the sealed-bid, second-price auction, then, will be exactly the same as the English auction. Note that the English auction and the second price auction have the property of *truthfulness* - every bidder reveals her true private value. Truthful bidding is induced because, regardless of what other bidders do, each bidder can maximize her profit by submitting her true private value. This is not true of the Dutch and the first price auctions, since the bidder with the highest private value has an incentive to bid under her true private value based on her expectation of the distribution of other bidder's private values. Indeed, the optimal bid is $E[Y|Y < x]$, where x is a bidder's private value and Y is the highest of the other independently drawn values; this result assumes that each bidder knows the distribution of other bidders' private values, but not the exact value.

Vickrey (1961) and more generally, Myerson (1981), proved that an auctioneer's expected revenue from any of the above auctions is the same. The Revenue Equivalence Theorem is of enormous importance in auction theory. It holds for all four auctions mentioned in this section as well as stranger auctions such as the all-pay auction, where only the winning

bidder gets the item but every bidder pays, as in political lobbying. The theorem is stated here, but the rather involved proof is omitted because of length; a version may be found in Krishna (2002).

Theorem 2 (Revenue Equivalence Theorem) *Assume there exist N risk-neutral bidders for an object, with each bidder holding a private value v_i drawn from a common, strictly increasing, atomless distribution on $[\underline{v}, \bar{v}]$. Then any auction where the object always goes to the bidder with the highest private value, and where a bidder with private value zero will pay zero, will provide the same revenue to the auctioneer and the same payments (in expectation) from the bidders (Klemperer 2004).*

An efficient auction is one in which the item is distributed to the bidder with the highest private value; the discussion above shows that the four most common auction formats are all, given the assumption of a common distribution of private values among bidders, efficient. An optimal auction is one in which the auctioneer's expected revenue is maximized. The Revenue Equivalence Theorem guarantees that all four auctions are optimal. Since each auction format is optimal, auctioneers should be indifferent, in theory, between selling goods with any of the formats.

Finally, consider auctions for which multiple items are sold. Multiple unit auctions have attracted enormous attention in recent years, due

to their use in stock IPOs, spectrum auctions, and Federal Reserve Treasury Auctions. For simplicity, we direct attention only to multiple unit auctions where each bidder desires at most one unit. Note that multiple units may be sold simultaneously, where each item is allocated after one auction is held according to some mechanism, or sequentially, where single-unit auctions are held until all of the units have been disposed of.

Two common simultaneous auctions for k units are the Vickrey auction and the discriminatory auction. In a discriminatory auction, the k highest bidders pay their bid and receive one unit. In the Vickrey auction, the k highest bidders all win one unit and each pay the amount of the highest non-winning bid. Given the assumption of single-unit demand, both auctions allocate efficiently, and a version of the Revenue Equivalence Theorem applies (see Krishna 2002 and Krishna and Perry 1998). It can be proven that sequential sealed-bid first-price and sealed-bid second-price auctions also allocate efficiently and are revenue equivalent to the two simultaneous auctions. However, the mathematics in the case of multi-unit demand, where bidders demand more than one unit according to some demand schedule, are considerably more difficult.

2.3 Mechanisms

While bidder behavior in auctions is a subset of game theory, the design of auctions themselves is a subset of mechanism design. Consider a planner who wishes to aggregate individual preferences and allocate resources accordingly. Individuals hold private information, such as their private valuation of a resource, and may not be willing to divulge true information if, given the allocation methods (or *mechanism*) of the planner, that individual may improve her welfare by lying. The information revelation problem constrains the mechanisms available to a planner; solving the problem of optimal mechanism design, therefore, has been an extraordinarily active field of social science research since Hurwicz's (1960) seminal paper on the topic. In this section, we briefly introduce two concepts from mechanism design which will be important to the study of GSP auctions: the revelation principle and the Vickrey-Clarke-Groves mechanism. An in-depth exploration of mechanism design can be found in Hurwicz and Reiter (2008).

Formally, a mechanism is a function $\Theta(S_1, ..., S_N, g(\cdot))$, where S_i represents the strategies available to agent i, and $g : S_1 \times ... \times S_N \rightarrow X$ is a function mapping the strategies of each player into an outcome X. A mechanism implements $f(\cdot)$ if there is a Nash equilibrium strategy profile R of the game induced by Θ such that $g(R) = f(\cdot)$. That

is, f is implementable only if there exists some mechanism inducing an equilibrium strategy that gives the allocation f. Let each agent have a piece of private information denoted λ_i. A mechanism for which $S_i = \lambda_i$ for each i is an equilibrium to the game induced by g is called a direct revelation mechanism, in that each agent simply tells $g(\cdot)$ her true type.

Example 3 *Consider a society wishing to choose its leader. One mechanism Θ might be that each individual gets to vote on the leader, and whichever leader gets the most votes will win. Assume that there are at least 3 candidates and at least 3 voters. Let each voter rank the candidates by preference, and let $f(\cdot)$ be a mapping from preferences to an outcome that is a) non-dictatorial (the election must account for the preferences of at least two voters), b) Pareto efficient (if all voters prefer A to B, then B must not be able to win), and c) disregards irrelevant alternatives (if the mechanism chooses A over B, then adding C into the election shouldn't cause the mechanism to rate B over A). Is $f(\cdot)$ implementable? That is, is there a voting mechanism such that self-interested voters, voting based on their preferences, will induce an election with the three properties listed above? There is not. This is the famous Arrow Impossibility Theorem (Arrow 1950).*

It may seem very difficult, then, to find what allocations are implementable, as this appears to involve searching over all possible mech-

anisms! However, the revelation principle says that, if an allocation can be implemented by some indirect mechanism, then it can be implemented by a direct mechanism (Gibbard 1973). Therefore, to find what allocations are implementable, only direct mechanisms need to be examined. The reasoning behind the revelation principle is straightforward. Assume that there exists a (perhaps very complicated) mechanism where agents input their private values v_i and play some complicated strategy s_i which, given the chosen mechanism, leads to outcome $g(s_i(v_i))$. Now create a outcome rule $f(\cdot)$ such that $f(v_i) = g(s_i(v_i))$. Since $s_i(v_i)$ is optimal under outcome rule $g(\cdot)$, v_i must be optimal under outcome rule $f(\cdot)$, and f as defined is a direct mechanism. In the context of auctions, the revelation principle means that if an auction outcome can be induced by non-truthful bidding given a certain auction rule, then there must exist some other auction rule such that the same outcome can be induced by bidders simply submitting their true private values.

In auctions, a Vickrey-Clarke-Groves (VCG) auction is a direct mechanism such that the agent with the highest private value receives the object, and that each agent's payment is based on the indirect effect, or externality, that she imposes on other players by entering the auction. Consider a single-object auction with private values $\{v_a, v_b\} = \{10, 4\}$. First, allocate the object to the bidder with the highest value (in this

case, bidder a). Now note that, had a not been in the auction, b would have received value 4 from the object. Therefore, the VCG mechanism gives agent a the object, and charges a a price of 4. This result is simply the sealed-bid second-price auction; such auctions are sometimes called Vickrey auctions for this reason. The Vickrey auction for multiple units described in the Section 2.2 is simply the multiple unit analogue of the VCG mechanism. Proofs of why VCG induces truthful bidding can be found in Vickrey (1961), and generalizations can be found in Clarke (1971) and Groves (1973).

3 Simultaneous Bid Solution

The Generalized Second Price (GSP) auction is used principally to sell advertisements displayed alongside keyword searches at search engines. Before 1997, internet advertisers paid a fixed amount to have their advertisement displayed a set number of times; for instance, a cost-per-thousand (CPT) of $10.00 meant that a potential advertiser would pay ten dollars to have their advertisement displayed one thousand times. Beginning in 1997, Overture began selling advertisements in an auction, where bidders bid on a per-clickthrough basis on a particular search engine keyword. Every time a bidder's ad was clicked, Overture charged the bidder an amount equal to their most recent bid, and advertisements

Figure 3: Text ads from a Google search

with higher bids were displayed more prominently. Bidders were allowed to change their bid as frequently as they wanted (Edelman et al 2007 and Battelle 2005 discuss this history in greater depth). Figure 3 shows four text advertisements that were displayed during a recent search for "algebra."

The ability to change bids continuously in this auction leads to a lack of stability, however. Consider two bidders that value each clickthrough at v_a and v_b, such that $v_a > v_b$. Let there be two advertising slots available, with clickthrough rates c_1 and c_2, such that $c_1 > c_2$. Assume there exists an equilibrium such that player a bids b_a and player b bids b_b. If $b_a > b_b + \varepsilon$, then player a can improve her profit by simply lowering

her bid, and therefore her payment, by $\varepsilon/2$; a similar argument applies if $b_b > b_a + \varepsilon$. If $b_a = b_b$, then depending on how the auction breaks ties, the player assigned the second slot can improve her payoff by bidding ε higher, and getting $v_i(c_1 - c_2) - c_2\varepsilon$ more profit. These first-price auctions do not have any equilibria and bidders must update their bids over and over as fast as their bidding program is able to submit bids.

For this reason, both Google AdWords and Yahoo!/Overture today use second-price auctions where bidders are ranked by the amount of their bid, and the bidder in position i pays the bid of the bidder in position $i+1$. This auction may be described formally as follows:

Definition 4 *(Generalized Second Price Auctions) Consider an auction with n bidders and k slots, with $n \geq k$. Each slot is associated with a publicly-known clickthrough rate c_i. Each bidder has a linear private valuation of a clickthrough equal to v_i. Each bidder submits a publicly-known bid b_i. Given an n-by-1 vector of bids b, the bidders are ordered, and the top slot is given to the bidder with the highest bid, the second slot to the bidder with the second highest-bid, and so on. Once the auction ends, each bidder pays the bid of the player directly below them per clickthrough, and the profit of each bidder is said to be $(v_i - b_{i+1})c_i$. At any time, a bidder may revise her bid. The auction ends when no bidder wishes to further revise her bid.*

In practice, clickthrough rates and bids (other than b_{i+1}, since the bidder knows what she will pay should the auction end, and this amount is simply the bid of the player directly below her) are not necessarily known by all bidders. However, auctioneers such as Google have traditionally offered potential bidders an estimated clickthrough rate. Further, the bids of other players can be easily calculated by switching one's bid until one has been placed into every slot, and has therefore learned b_{i+1} for all i. In this paper, we follow the literature and assume bids and clickthrough rates are known to all bidders.

In order to solve for equilibria of this auction, the timing of the bids is critical. Consider a model where bids occur in stages, and every bidder submits a bid simultaneously in every stage. Given simultaneous bids, a Nash equilibrium to the auction means that no player can improve her profit by changing her bid, taking other players' bids as given. Note that, in order to move up one position, a player must bid higher than the bid of the player one spot above her, but in order to move down one position, a player must bid lower than the player one spot below her, and higher than the bid of the player *two* spots below her (Varian 2007). That is, a Nash equilibrium is

$$(v_i - b_{i+1})c_i \geq (v_i - b_j)c_j \text{ for } i > j$$

$$(v_i - b_{i+1})c_i \geq (v_i - b_{j+1})c_j \text{ for } i < j$$

for all i, j.

Since any change in a player's bid that does not change her position will not change her payments or profit, these inequalities are satisfied by intervals of bids, and thus Nash equilibria are not unique.

The Nash equilibria to simultaneous bid GSP have two unfortunate properties. First, there are equilibrium bids that are not efficient, meaning that bidders with lower valuation are assigned higher slots. Consider the following example from Borgers et al (2007):

Example 5 *Let there be three bidders with private values* $v = [16, 15, 14]$*, and three slots with clickthrough rates* $c = [3, 2, 1]$*. Consider the vector of bids* $[7, 9, 11]$*. Profits in this auction are* $[(16 - 0) * 1, (15 - 7) * 2, (14 - 9) * 3] = [16, 15, 15]$*. It is straightforward to check that the vector of bids is an equilibrium; no bidder can improve her payoff by unilaterally changing her bid. Note that the bidder with lowest private value receives the highest slot, and the bidder with the highest private value receives the lowest slot.*

Further, unlike in the second-price auctions discussed in Section 2, simultaneous bid GSP is not truthful. That is, the vector of bids equal to private values ($b = v$) is not necessarily an equilibrium. Aggarwal et

al (2006) offer the following example:

Example 6 *Let there be three bidders with private values* $v = [200, 180, 100]$ *and two slots with clickthrough rates* $c = [.5, .4]$. *Let* $b = [200, 180, 100]$. *Bidder 1 makes profit* $.5(200 - 180) = 10$ *with her current bid, but will make* $.4(200-100) = 40$ *if she bids below 180 but more than 100. Therefore* $b = [200, 180, 100]$ *is not an equilibrium.*

Edelman et al (2007) consider a simpler set of equilibria, called locally envy-free equilibria.

Definition 7 *(Locally envy-free) A set of bids is a locally envy-free equilibria if no player would prefer to switch bids with any other player. This is equivalent to requiring that no player wishes to switch bids with the player who bid directly above her. That is,*

$$c_i(v_i - b_{i+1}) \geq c_{i-1}(v_i - b_i)$$

for all i.

Locally envy-free equilibria are a subset of Nash equilibria, and are guaranteed to be efficient. Further, locally envy-free equilibria have a natural link to stable assignments in a two-sided matching game such as that studied in Shapley and Shubik (1972). However, given that

there are auctions without truthful Nash equilibria, auctions such as the one discussed in Example 6 also do not have a truthful locally envy-free equilibrium.

Since locally envy-free equilibria are efficient, the revelation principle guarantees the existence of a truthful mechanism which induces the same equilibrium payments and auctioneer revenue. Consider the Vickrey-Clarke-Groves analogue of the GSP auction. In VCG, bidders are ranked based on their bids, and each bidder pays the externality her entry into the auction imposed on other bidders by changing what slot they win.[2] In an auction with 5 bidders and 4 slots, let a bidder bid such that her bid is ranked second. Her entry into the auction places the formerly-second bidder in the third slot, places the formerly-third in the fourth slot, and forces the formerly-fourth to win no slot. Therefore, VCG would have the bidder ranked second pay $(c_2 - c_3)b_3 + (c_3 - c_4)b_4 + (c_4)b_5$. It is well-known that VCG mechanisms always have truthful bidding as a dominant strategy.

Example 8 *Consider an auction with private values* $v = [8, 5, 4]$ *and slots with* $c = [2, 1]$*. Truthful bidding is a Nash (and locally envy-free) equilibrium, with total payments for each bidder equal to* $[10, 4, 0]$ *and payoffs equal to* $[6, 1, 0]$*. Consider the VCG mechanism with true values*

[2] A derivation of VCG in this context is given in Appendix A.

submitted, however. If Player 1 were not in the auction, Player 2 would be in slot 1 and would receive an additional $(2-1)*5=5$*. Player 3 would be in slot 2 and would receive an additional* $(1-0)*4=4$*. Player 1, therefore, would pay* $5+4=9$ *under VCG. By a similar argument, Player 2 would pay* 4 *under VCG, and Player 3 would pay* 0*. VCG payments, then, are* $[9,4,0]$ *and payoffs are* $[7,1,0]$*.*

Note that VCG payoffs in this example are higher than under the locally envy-free equilibria examined. Edelman et al (2007) show that in any GSP auction, there exists a locally envy-free equilibrium with positions, payoffs and auctioneer revenue equal to VCG, but that this equilibrium has the lowest auctioneer revenue of any of the locally envy-free equilibria to that auction. That is, even though VCG is simple in that a bidder needs only to truthfully submit her private value, and even though VCG induces positions, payoffs and auctioneer revenue equal to a locally envy-free equilibrium, an auctioneer may not want to sell ads using VCG because any other locally envy-free equilibrium will give the auctioneer higher revenue. Even if all Nash equilibria, and not just locally envy-free equilibria, are considered, there are still many Nash equilibria with auctioneer revenue higher than under VCG.

4 The Hierarchical Game

In the simultaneous-bid GSP auction analyzed in the previous section, Nash equilibria can be inefficient and auctioneers may be hesitant to use VCG because GSP always induces bids that, in equilibrium, give the auctioneer higher revenue. In this section, we modify the timing of bids. Rather than having every bidder submit bids simultaneously each stage, we let one bidder bid in each stage, taking earlier bids from other bidders as given. That is, bidding will be hierarchical. Players who bid before a given bidder are called *leaders*, and players who bid after a given bidder are called *followers*. Every bidder will bid in best response to bids from leaders, and will bid assuming best responses from bidders who follow. In general, as in the Stackelberg Duopoly game, hierarchical bidding will not induce the same set of equilibria as simultaneous bidding.

Since bids are submitted in a hierarchy, the question remains as to what order the bidders should be allowed to submit bids. We remain agnostic on the ordering, but assume that an equilibrium set of bids is such that, no matter what order the bidders bid in, no bidder has an incentive to change her bid. Formally, an *ordered hierarchical equilibrium (OHE)* is an n-by-1 vector of bids b and resultant profits π for each bidder such that, for any ordering of bidders, the lowest equilibrium profit for every bidder is less than or equal to π.

Why is hierarchical bidding sensible in GSP auctions? Nothing inherent in the Google or Yahoo auctions requires bidders to submit bids simultaneously. In practice, they are always able to wait an extra second in order to view the latest bids submitted by other potential buyers. Further, as noted in Cary et al (2008), though bids may be updated as frequently as a bidder likes, there is a (very short) random delay between when a bid is submitted to the auctioneer and when that new bid is accepted. A model where, in each stage, one randomly selected bidder is allowed to change her bid is equivalent to hierarchical bidding.

The following example motivates why hierarchical bidding can lead to different equilibria.

Example 9 *Consider a vector of private values* $v = [10, 4, 2]$*, and click-through rates* $c = [2, 1, 0]$*. The vector of truthful bids* $b = [10, 4, 2]$ *is a locally envy-free equilibrium; no player wishes to switch bids with any other player. Because locally envy-free equilibria are a subset of Nash equilibria,* $[10, 4, 2]$ *is also a Nash equilibrium. Note that Bidder 2 is completely indifferent between bidding anywhere in the interval* $(2, 10)$*, but that Bidder 1 wants Bidder 2 to bid as low as possible. If bids are not simultaneous, then Bidder 1 can bid, for instance,* 3.50*, knowing that Bidder 2 will then improve her profit by then changing her bid to some-*

where in the interval $(2, 3.50)$. *Bidder 1 cannot bid below* 3*, however, because then Bidder 2 would prefer to bid above her. Let Bidder 1, then, bid* $3+\varepsilon$*, which induces Bidder 2 to then change her bid to somewhere in* $(2, 3)$. *No bidder can then improve their profit, and the auction ends with bids* $[3+\varepsilon, (2, 3), 2]$.

In general, it is difficult to solve for equilibria to hierarchical games. These games belong to a class of problems called multilevel programs, which are proven to be strongly NP-hard (Luo et al 1996). Intuitively, the bidder who bids first is solving an optimization problem taking followers' bids as best responses to her bid, the bidder who bids second solves an optimization problem taking the first bidder's bid as given and taking followers' bids as best responses to her own bid, and so on. In order to solve for the first bidder's bid, the last bidder's best response as a function of every other player's bids is substituted into the second-to-last player's optimization problem, and on up the chain until the second bidder's best response is substituted into the first player's optimization problem. Each player's best response function is often neither smooth nor connected.

In order to find equilibria, then, we first find equilibria to the hierarchical game where players bid in the order of their private values; this is called the *private-value ordering*. After showing that equilibria to this

ordering must be efficient, it is straightforward to solve for the equilibria algebraically. We then show that, given any other ordering, equilibria must be efficient. Finally, we show that, since equilibria in any ordering are efficient, no bidder can improve her profit from the profit attained in the private-value ordering. With equilibria in hand, we then show that OHE are a subset of Nash equilibria, and that the VCG solution gives auctioneer revenue at least as high as any OHE.

Lemma 10 *The equilibria in the private-value ordering hierarchical game are efficient.*

Proof. Let $b_i < b_j, v_i > v_j > v_k$. Let i bid before j. Assume that bids are in equilibrium. Bidder i receives profit $\pi_i = (v_i - b_k)c_i$ and bidder j receives profit $\pi_j = (v_j - b_i)c_j$. If bidder j bids directly below i instead, she will receive $(v_j - b_k)c_i$. Since bids are in equilibrium, $(v_j - b_i)c_j \geq (v_j - b_k)c_i$ and therefore $b_i \leq v_j - (v_j - b_k)\frac{c_i}{c_j}$.

Note that bidder i makes profit

$$(v_i - b_k)c_i = v_i c_i - b_k c_i \tag{1}$$

by letting $b_i \leq v_j - (v_j - b_k)\frac{c_i}{c_j}$, but makes profit $(v_i - b_j)c_j$ by letting $b_i > v_j - (v_j - b_k)\frac{c_i}{c_j}$ since j will then bid below i. In that case, since b_j will be less than b_i, and since b_i can be as low as $v_j - (v_j - b_k)\frac{c_i}{c_j}$, profit

for bidder i will be $(v_i - b_j)c_j \geq (v_i - (v_j - (v_j - b_k)\frac{c_i}{c_j}))c_j =$

$$v_i c_j - b_k c_i + v_j(c_i - c_j). \tag{2}$$

Note that profit in Eq. 2 is always greater than that of Eq. 1, since

$$\begin{aligned} v_i(c_j - c_i) &> v_j(c_j - c_i) \\ &\Rightarrow v_i(c_j - c_i) + b_k c_i > v_j(c_j - c_i) + b_k c_i \\ &\Rightarrow v_i(c_i - c_j) - b_k c_i < v_j(c_i - c_j) - b_k c_i \\ &\Rightarrow v_i c_i - b_k c_i < v_i c_j - b_k c_i + v_j(c_i - c_j). \end{aligned}$$

Therefore bidder i will always bid such that bidder j wishes to bid below her; this contradicts the assumption that $b_i < b_j$ can be an equilibrium. ■

If the private-value ordering equilibrium is efficient, three conditions must hold. First, $b_i < b_{i-1}$. Second, the bidder directly below i must not want to bid above b_i; that is, $(v_{i+1} - b_{i+2})c_{i+1} \geq (v_{i+1} - b_i)c_i$. Third, the bidder directly above bidder i must not want to bid below bidder i, so $(v_{i-1} - b_i)c_{i-1} \leq (v_{i-1} - b_{i+1})c_i$. Rearranging terms, $b_i \in (v_{i+1} - (v_{i+1} - b_{i+2})\frac{c_{i+1}}{c_i}, \min\{b_{i-1}, v_{i-1} - (v_{i-1} - b_{i+1})\frac{c_i}{c_{i-1}}\})$. In an auction with k slots, we assume that any bidder $j > k$ simply bids her private

value v_j; the bidder in that slot pays nothing and makes no profit. By bidding private value, she ensures that she will enter the auction if it becomes profitable for her to do so, such as if a bidder above her were to drop out. Also, for the bidder with the highest private value, the right hand side of the bid interval is ∞. Finally, for bidder k, the left hand side of the bid interval simplifies to v_{k+1}. Any efficient vector of bids, then, is constrained as follows, with slots $k = n - 1$ for simplicity.

Lemma 11 *Every efficient vector of bids is constrained to the following form:*

$$\begin{aligned}
b_N &= v_N \qquad (3)\\
b_{N-1} &= (v_N, \min\{b_{N-2}, v_{N-2} - (v_{N-2} - b_N)\frac{c_{N-1}}{c_{N-2}}\})\\
&\ldots\\
b_i &= (v_{i+1} - (v_{i+1} - b_{i+2})\frac{c_{i+1}}{c_i}, \min\{b_{i-1}, v_{i-1} - (v_{i-1} - b_{i+1})\frac{c_i}{c_{i-1}}\})\\
&\ldots\\
b_2 &= (v_3 - (v_3 - b_4)\frac{c_3}{c_2}, \min\{b_1, v_1 - (v_1 - b_3)\frac{c_2}{c_1}\})\\
b_1 &= (v_2 - (v_2 - b_3)\frac{c_2}{c_1}, \infty)
\end{aligned}$$

Bidder 1's profit is equal to $(v_1 - b_2)\,c_1$, so profit is maximized by minimizing b_2. Since b_2 must be lower than $\min\{b_1, v_1 - (v_1 - b_3)\frac{c_2}{c_1}\}$, bidder 1 maximizes profit by choosing the lowest possible $b_1 = v_2 - (v_2 -$

$b_3)\frac{c_2}{c_1}$. Note that since $v_2 < v_1$, $b_1 = v_2 - (v_2 - b_3)\frac{c_2}{c_1} < v_1 - (v_1 - b_3)\frac{c_2}{c_1}$. Likewise, b_2 minimizes b_3, and therefore maximizes profit, by bidding $b_2 = v_3 - (v_3 - b_4)\frac{c_3}{c_2}$. Note that for any bid in $(v_N, \min\{b_{N-2}, v_{N-2} - (v_{N-2} - b_N)\frac{c_{N-1}}{c_N}\})$, bidder $N-1$'s profit is the same. This is because bidder N always bids v_N in an efficient auction.

Theorem 12 *The equilibrium bids of the private-value ordering hierarchical game are*

$$\begin{aligned} v_{i+1} - (v_{i+1} - b_{i+2})\frac{c_{i+1}}{c_i},\ i < k \\ (v_{i+1}, \min\{b_{i-1}, v_{i-1} - (v_{i-1} - b_{i+1})\frac{c_i}{c_{i-1}}\}),\ i = k \\ v_i,\ i > k. \end{aligned} \tag{4}$$

Given these bids, total payments from each bidder are $b_{i+1}c_i$.

Can any bidder, given some order of bidding, improve their profit from that received with the bids in Equation 4? They can not. First, the following lemma shows that under any perturbation of the order of bids, equilibria must still be efficient. The idea behind this proof is that, when a bidder with low private value tries to bid higher than a bidder with high private value, there is always an interval of optimal bids for the high private value bidder which will cause the low private value bidder to wish to switch her bid - that is, the vector of bids will

not be in equilibrium.

Since the bids must be efficient, every bidder knows that her profit is maximized by minimizing the bid of the bidder whose private value is directly below their own. In that case, if a bidder bids before the player with private value one spot lower than her, she will bid as in the private value OHE. If a bidder bids after the player one spot below her, she takes that player's bid as given, and bids somewhere in the interval that keeps her in her efficient slot.

Lemma 13 *Equilibria to perturbed order hierarchical game are efficient.*

Proof. Let $b_i < b_j, v_i > v_j > v_k$. Let j bid before i. Assume that bids are in equilibrium. Bidder i receives profit $\pi_i = (v_i - b_k)c_i$ and bidder j receives profit $\pi_j = (v_j - b_i)c_j$. If bidder j bids directly below i instead, she will receive $(v_j - b_k)c_i$. Since bids are in equilibrium, $(v_j - b_i)c_j \geq (v_j - b_k)c_i$ and therefore $b_i \leq v_j - (v_j - b_k)\frac{c_i}{c_j}$.

Note that bidder i makes profit

$$(v_i - b_k)c_i = v_i c_i - b_k c_i \tag{5}$$

by letting $b_i \leq v_j - (v_j - b_k)\frac{c_i}{c_j}$, but makes profit $(v_i - b_j)c_j$ by letting $b_i > v_j - (v_j - b_k)\frac{c_i}{c_j}$ since j will then bid below i. In that case, since b_j will be less than b_i, and since b_i can be as low as $v_j - (v_j - b_k)\frac{c_i}{c_j}$, profit

for bidder i will be $(v_i - b_j)c_j \geq (v_i - (v_j - (v_j - b_k)\frac{c_i}{c_j}))c_j =$

$$v_i c_j - b_k c_i + v_j(c_i - c_j). \tag{6}$$

As in the private-value ordering, note that profit in Eq. 6 is always greater than that of Eq. 5, since

$$\begin{aligned}
v_i(c_j - c_i) &> v_j(c_j - c_i) \\
&\Rightarrow v_i(c_j - c_i) + b_k c_i > v_j(c_j - c_i) + b_k c_i \\
&\Rightarrow v_i(c_i - c_j) - b_k c_i < v_j(c_i - c_j) - b_k c_i \\
&\Rightarrow v_i c_i - b_k c_i < v_i c_j - b_k c_i + v_j(c_i - c_j).
\end{aligned}$$

Finally, note that unless $b_j \geq v_i - (v_i - b_k)\frac{c_i}{c_j}$, bidder i will bid higher than bidder j, which would contradict the assumption that bids are in equilibrium. But since $v_i - (v_i - b_k)\frac{c_i}{c_j} > v_j - (v_j - b_k)\frac{c_i}{c_j}$, if $b_j \geq v_i - (v_i - b_k)\frac{c_i}{c_j}$, there must exist some $b_i \in (v_j - (v_j - b_k)\frac{c_i}{c_j}, v_i - (v_i - b_k)\frac{c_i}{c_j})$ such that bidder j would prefer to have bid below bidder i. This contradicts the assumption that $b_i < b_j$ can be an equilibrium. ■

Lemma 14 *Equilibria to the perturbed order hierarchical game give every bidder profit no larger than that of the private-value OHE.*

Proof. The previous lemma shows that equilibria to the perturbed order hierarchical game are efficient. Since the equilibria must be efficient, bids are constrained by the Nash intervals in Equation 3. That is, players must bid in the intervals

$$
\begin{aligned}
b_N &= v_N \\
b_{N-1} &= (v_N, \min\{b_{N-2}, v_{N-2} - (v_{N-2} - b_N)\frac{c_{N-1}}{c_{N-2}}\}) \\
&\dots \\
b_i &= (v_{i+1} - (v_{i+1} - b_{i+2})\frac{c_{i+1}}{c_i}, \min\{b_{i-1}, v_{i-1} - (v_{i-1} - b_{i+1})\frac{c_i}{c_{i-1}}\}) \\
&\dots \\
b_2 &= (v_3 - (v_3 - b_4)\frac{c_3}{c_2}, \min\{b_1, v_1 - (v_1 - b_3)\frac{c_2}{c_1}\}) \\
b_1 &= (v_2 - (v_2 - b_3)\frac{c_2}{c_1}, \infty)
\end{aligned}
$$

If a player bids before the player whose private value is directly below her, she wants to minimize that player's bid as in the private-value ordering, and will therefore bid exactly as in the private value OHE from Equation 4. If a player bids after the player whose private value is directly below her, she will bid anywhere in the Nash interval that makes the players above her and below her unwilling to outbid or underbid her, as in the case of bidder k in the private value OHE. Since the bids for bidders $i < k$ in the private value OHE are equal to the minimum of the Nash

interval, and the equilibrium bids for bidders $i \geq k$ are identical in the perturbed order case and in the private value OHE, bids for every player are at least as high in the perturbed order equilibrium as in the private value OHE. Therefore, the payments in the perturbed order equilibrium are at least as high as in the OHE, and the resultant profits are no higher, for all bidders. ■

Proposition 15 *Equilibria to the hierarchical game are equivalent to the private value OHE.*

Proof. The previous lemma establishes that every player has profits in a perturbed order auction that are no higher than their profits in the private value OHE. Therefore, no player has an incentive to change their bid once the private value OHE is reached, regardless of what order they are allowed to bid in. ■

OHE possess a few interesting properties. First, the constraints of hierarchical bidding in Eq. 3 are simply the set of Nash equilibria, so OHE are a subset of Nash equilibria. Second, consider the payments induced by the truthful VCG mechanism. VCG with truthful revelation induces payments of zero for bidders $i > k$, and total payments of $\sum_{a=i}^{N}(c_a - c_{a+1})v_{a+1}$ for each bidder $i \leq k$. That is, bidder k pays $c_k v_{k+1}$, bidder $k-1$ pays $(c_{k-1} - c_k)v_k + c_k v_{k+1}$, and so on. The following

proposition shows that VCG payments for every bidder are at least as big as the payments in any OHE, and therefore that auctioneer revenue is higher under VCG than under OHE.

Proposition 16 *VCG revenue is greater than or equal to OHE revenue.*

Proof. Consider the highest-revenue OHE equilibrium, where bidder k bids b_{i-1}. Since VCG and OHE are both efficient, revenue for all bidders is the same with both mechanisms. Under both VCG and OHE, payments for bidder k are $v_{k+1}c_k$. Under the highest-revenue OHE, payments for bidder $k-1$ are $b_{k-1}c_{k-1}$. For all other bidders, highest-revenue OHE payments for bidder i are $b_{i+1}c_i \leq b_i c_i = c_i(v_{i+1} - (v_{i+1} - b_{i+2})\frac{c_{i+1}}{c_i}) = v_{i+1}(c_i - c_{i+1}) + c_{i+1}b_{i+2} \leq v_{i+1}(c_i - c_{i+1}) + c_{i+1}b_{i+1} \leq \cdots \leq \sum_{a=i}^{N}(c_a - c_{a+1})v_{a+1} =$ VCG payments. ■

The following numerical example displays the above properties.

Example 17 *Consider a GSP auction with* $v = [11, 8, 7, 1]$ *and* $c = [10, 2, 1, 0]$*. Under VCG, bidder 4 pays 0, bidder 3 pays* $(1-0)*1 = 1$*, bidder 2 pays* $(2-1)*7+1 = 8$*, and bidder 1 pays* $(10-2)*8+8 = 72$*. Total auctioneer revenue is* 81*. OHE bids are* $[7.20, 4, (1, 4), 1]$*. The total payment vector, then, is* $[40, (2, 8), 1]$*, and auctioneer revenue is* $[43, 49] < 81$*.*

Edelman et al (2007) propose reasons why auctioneers might be using GSP rather than VCG to sell keywords, even though VCG requires bidders simply to submit their true private values and would therefore seem to be make bidding much simpler for potential advertisers. Under their locally envy-free equilibria, which assumed simultaneous bidding, VCG provides the lowest-possible auctioneer revenue, so auctioneers might rightfully be unwilling to switch to a mechanism that guarantees minimal revenue. However, if bids are hierarchical in nature, OHE guarantees that VCG has higher revenue than any GSP equilibria.

5 Extension: Endogenous Slots

In this section, we describe a straightforward extension of the hierarchical bidding model where the auctioneer endogenizes the number of slots available based on the current vector of bids. For mathematical simplicity, in Section 4 we followed the literature by letting the number of slots equal a constant $k \leq n$. Let $k = n$ be the greatest number of slots an auctioneer can offer. If an auctioneer can costlessly remove slots, and by removing them improve auctioneer revenue, then endogenous k allows a more robust formulation of auctioneer problem. It turns out to be the case that, for any vector of private values and clickthrough rates, the auctioneer can improve revenue by removing slots for some values of

k.

Consider an auction with n bidders and the maximum number of slots $k = n$. In this situation, the kth bidder pays nothing for her advertisements. More importantly, her equilibrium bid, as shown in Section 4, is in the interval $(v_{i+1}, \min\{b_{i-1}, v_{i-1}-(v_{i-1}-b_{i+1})\frac{c_i}{c_{i-1}}\})$, where $v_{i+1} = 0$. Bidders $k < n$ then bid their OHE value $v_{i+1}-(v_{i+1}-b_{i+2})\frac{c_{i+1}}{c_i}$. Bidder $k-1$ optimally bids only $(1-\frac{c_k}{c_{k-1}})v_k < v_k$. These two equations imply that both b_k and b_{k-1} are lower than v_k. In the case where $k = n-1$, the kth bidder will bid her private value, and bidder $k-1$ will be in some interval above that. That is, the bottom two bids are higher in the $k = n-1$ case than in the $k = n$ case. Since every other bid can be solved recursively as an increasing function of these two bids, since payments are based on these bids, and since the kth bidder pays zero to the auctioneer in both cases, it must be the case that removing the bottom slot from this auction increases auctioneer revenue.

Proposition 18 *For any vectors c and v, the auctioneer has higher OHE revenue with $n = k-1$ than with $n = k$.*

Proof. Let $n = k$. In equilibrium, the bidder with the lowest private value bids in the interval $(0, b_{k-1})$, the bidder with the second lowest private value bids $v_k - (v_k - 0)\frac{c_k}{c_{k-1}} = (1-\frac{c_k}{c_{k-1}})v_k < v_k$, and all other bidders bid $v_{i+1} - (v_{i+1} - b_{i+2})\frac{c_{i+1}}{c_i} = v_{i+1}(1-\frac{c_{i+1}}{c_i}) + b_{i+2}\frac{c_{i+1}}{c_i}$,

which is an increasing function of b_{i+2}. Revenue to the auctioneer is $(0,(1-\frac{c_k}{c_{k-1}})v_k)c_{k-1}+(1-\frac{c_{k-1}}{c_{k-2}})v_k c_{k-2}+\sum_{i=k-3}^{N} c_i(v_{i+1}(1-\frac{c_{i+1}}{c_i})+b_{i+2}\frac{c_{i+1}}{c_i})$. Now consider the case where $c_N=0$; that is, $k=n-1$. The bidder with lowest private value bids v_k, the bidder with the second lowest private value bids in the interval (v_k, b_{k-2}), and all other bidders bid $v_{i+1}-(v_{i+1}-b_{i+2})\frac{c_{i+1}}{c_i}=v_{i+1}(1-\frac{c_{i+1}}{c_i})+b_{i+2}\frac{c_{i+1}}{c_i}$. Revenue to the auctioneer, then, is $v_k c_{k-1}+(v_k,b_{k-2})c_{k-2}+\sum_{i=k-3}^{N} c_k(v_{i+1}(1-\frac{c_{i+1}}{c_i})+b_{i+2}\frac{c_{i+1}}{c_i})$. Since the payments of every bidder are strictly higher than in the case where $n=k$, auctioneer revenue improves by removing the final slot. ■

The following example gives revenue under two situations with endogenous slots.

Example 19 *Let* $v=[10,4,2]$*, and* $c=[3,2,1]$*. In the case where* $k=3$*, OHE bids are* $[2,1,(0,1)]$*. Revenue to the auctioneer is* $3+(0,2)+0=(3,5)$*. If the auctioneer auctions only two slots (*$c=[3,2,0]$*), OHE bids are* $[3,(2,3),2]$*. Revenue to the auctioneer is* $(6,9)+4=(10,13)$*. In this example, endogenizing slots can triple auctioneer revenue.*

The benefit to the auctioneer from removing slots other than the nth is more difficult to compute, and depend on the auctioneer's belief about the distribution of bidder private values.

Finally, note that this result is worrying for the efficiency properties of GSP even under hierarchical bidding. If slots are endogenous, GSP is

no longer efficient. Under the VCG mechanism, total auctioneer revenue is equal to $\sum_{i=1}^{N}\sum_{a=i}^{k}(c_a - c_{a+1})v_{a+1}$, which is increasing in k. There is no incentive under VCG to restrict the number of slots. Therefore, VCG might be preferable to GSP not only because auctioneer revenue is higher, but also because it guarantees efficiency.

6 Conclusion

Keyword auctions are now the most important source of revenue in the rapidly growing online advertisement market. The novelty of auctions with multiple slots and continuous updating of bids remains poorly understood, however. This paper presents an equilibrium for a generalized second-price auction where bids are submitted in a hierarchy, with every bidder acting as a Stackelberg follower when it comes to previous bids. The properties of these ordered hierarchical equilibria are very different from that of the locally envy-free and Nash equilibria examined in previous papers.

We show that ordered hierarchical equilibria always give lower auctioneer revenue than any locally envy-free equilibrium, and in particular, lower revenue than the Vickrey-Clarke-Groves mechanism. Further, when the number of slots sold is endogenized by the auctioneer, ordered hierarchical equilibria lose the efficiency property, which VCG maintains. Because of this, and contrary to earlier research, it may be optimal from

a revenue perspective for auctioneers to adopt VCG payments instead of a second-price auction, with the side benefit that bid calculations are significantly simpler for advertisers in VCG.

References

[1] Aggarwal, Gagan, Ashish Goel, and Rajeev Motwani (2006): "Truthful Auctions for Pricing Search Keywords," ACM Conference on Electronic Commerce.

[2] Arrow, Kenneth (1950): "A Difficulty in the Concept of Social Welfare," Journal of Political Economy 58(4), 328–346.

[3] Athey, Susan and Glenn Ellison (2008): "Position Auctions with Consumer Search," Mimeo.

[4] Baron, David and John Ferejohn (1989): "Bargaining in Legislatures," American Political Science Review 83(4), 1181-1206.

[5] Battelle, John (2005): The Search, Portfolio Hardcover, New York.

[6] Borel, Emile (1921): "La Theorie de Jeu et les Equations Integrales a Noyau Symetrique," Comptes Rendus de l'Academie des Sciences 173, 1304-1308.

[7] Borgers, Tilman, Ingemar Cox, Martin Pesendorfer, Vaclav Petricek (2006): "Equilibrium Bids in Sponsored Search Auctions: Theory and Evidence," Mimeo.

[8] Borgs, Christian, Jennifer Chayes, Nicole Immorlica, Mohammad Mahdian, and Amin Saberi (2006): "Multi-unit Auctions with Budget Constrained Bidders," Mimeo, Microsoft Research.

[9] Boyer, Marcel and Michel Moreaux (1986): "Perfect competition as

the limit of a hierarchical market game," Economics Letters 22(2-3), 115-118.

[10] Cary, Matthew, Aparna Das, Benjamin Edelman, Ioannis Giotis, Kurtis Heimerl, Anna Karlin, Claire Kenyon Mathieu and Michael Schwarz (2008): "On Best-Response Bidding in GSP Auctions," NBER Working Paper W13788.

[11] Clarke, Edward H. (1971): "Multipart Pricing of Public Goods," Public Choice 11(0), 17-33.

[12] Cournot, Augustin (1838): Recherches sur les Principes Mathematiques de la Theorie des Richesses. Paris: Hachette.

[13] Damianov, Damian (2005): "Auctions with Variable Supply and Competing Auctioneers," Dissertation, Heidelberg University.

[14] Dimand, Mary Ann and Robert W. Dimand (1996): A History of Game Theory, Volume 1: From the Beginnings to 1945, Routledge, London.

[15] Edelman, Benjamin and Michael Ostrovsky (2007): "Strategic Bidder Behavior in Sponsored Search Auctions," Decision Support Systems 43(1), 192-198.

[16] Edelman, Benjamin, Michael Ostrovsky, and Michael Schwarz (2007): "Internet Advertising and the Generalized Second-Price Auction: Selling Billions of Dollars Worth of Keywords," Ameri-

can Economic Review 97(1), 242-259.

[17] Gibbard, Allan (1973): “Manipulation of Voting Schemes: A General Result,” Econometrica 41, 587–601.

[18] Groves, Theodore (1973): “Incentives in Teams,” Econometrica 41(4), 617-631.

[19] Herodotus (c. 430 BC): History of the Persian Wars. Translation available at Ancient History Sourcebook, Fordham University, http://www.fordham.edu/halsall/ancient/greek-babylon.html#Herodotus.

[20] Hurwicz, Leonid (1960): “Optimality and Informational Efficiency in Resource Allocation Processes” in Mathematical Methods in the Social Sciences, Stanford University Press, Stanford, CA.

[21] Hurwicz, Leonid and Stanley Reiter (2008): Designing Economic Mechanisms, Cambridge University Press, Cambridge.

[22] Klemperer, Paul (2004): “Auctions: Theory and Practice,” Economics Papers 2004-W09, Economics Group, Nuffield College, University of Oxford.

[23] Krishna, Vijay and Motty Perry (1998): ”Efficient Mechanism Design,” mimeo.

[24] Krishna, Vijay (2002): Auction Theory, Elsevier, San Diego.

[25] Luo, Zhiquan, Jongshi Pang and Daniel Ralph (1996). Mathemat-

ical Programs with Equilibrium Constraints, Cambridge University Press, Cambridge.

[26] Mas-Colell, Andreu, Michael D. Whinston and Jerry R. Green (1995): Microeconomic Theory, Oxford University Press, New York.

[27] Mehta, Aranyak, Amin Saberi, Umesh Vazirani and Vijay Vazirani (2007): "AdWords and Generalized Online Matching," Journal of the ACM, Article 22.

[28] Myerson, Roger (1981): "Optimal Auction Design," Mathematics of Operations Research 6, 58-73.

[29] Myerson, Roger (1997): Game Theory: Analysis of Conflict, Harvard University Press, Cambridge, MA.

[30] Shapley, Lloyd S. and Martin Shubik (1971): "The Assignment Game I: The Core," International Journal of Game Theory 1(1), 111-130.

[31] Shubik, Martin (1983): "Auctions, Bidding, and Markets: A Historical Sketch," in Auctions, Bidding and Contracting: Uses and Theory, NYU Press, New York.

[32] Varian, Hal R. (2007): "Position Auctions," International Journal of Industrial Organization 25(6), 1163-1178.

[33] Vickrey, William (1961): "Counterspeculation Auctions, and Competitive Sealed Tenders." Journal of Finance 16(1), 8-37.

[34] von Neumann, John and Oskar Morgenstern (1944): Theory of Games and Economic Behavior, Princeton University Press, Princeton.

[35] von Stackelberg, Heinrich (1934): Marktform und Gleichgewicht, Springer Verlag, Vienna.

[36] Vukov, Jeremos and Gyorgy Szabo (2005): "Evolutionary Prisoner's Dilemma Game on Hierarchical Lattices," Physical Review E 71.

[37] Zermelo, Ernst (1913): "Uber eine Anwendung der Mengenlehre auf die Theorie des Schachspiels," Proc. Fifth Congress Mathematicians, (Cambridge 1912), Cambridge University Press 1913, 501-504.

[38] Zhou, Yunhong and Rajan Lukose (2007): "Vindictive Bidding in Keyword Auctions," Proceedings of the Ninth International Conference on Electronic Commerce, 141-146.

A Appendix: Vickrey-Clarke-Groves

The Vickrey-Clarke-Groves mechanism induces bidders to reveal true private values by aligning the incentives of both the bidders and a mechanism designer who desires truthful revelation. The section modifies a discussion in Varian (2007).

Consider a case with n agents. Let the mechanism designer choose some assignment of bidders to slots z, and let each agent have a true private value for each ordering $v_i(z)$ but signal to the mechanism designer a value $s_i(z)$. If the mechanism designer offers to maximize the sum of bidder reported values, $s_i(z)+\sum_{i\neq j} s_j(z)$, and bidder i truly cares about maximizing $v_i(z)+\sum_{i\neq j} s_j(z)$, then it is optimal for the bidder to simply submit her true value, or $s_i(z) = v_i(z)$.

In this case, the mechanism designer is paying each bidder in order to extract true private values. In order to minimize the sum of these payments, an amount that is independent from the reported value of the agent can be subtracted. Let the subtracted payment equal the maximized sum of reported values for each ordering if bidder i were excluded. The payment to bidder i is then

$$v_i(z) + \sum_{i\neq j} s_j(z) - \max_y \sum_{i\neq j} s_j(y).$$

Since every bidder now is best off by submitting her true private values, and since in the absence of bidder i, all bidders below bidder i will move up one slot, the above equation simplifies to

$$v_i c_i - \sum_{j>i} v_{j+1}(c_i - c_{i+1}).$$