Sponsored Search Markets

What is Sponsored Search Advertising (or Keyword Auction)

- Advertisers submit bids for some keywords to a search engine, together with sponsored links related to the keywords. When a keyword matches a query of a search engine user, the search engine will show, along with the normal unpaid search result (called *organic* result), a limited number of matching sponsored links. The order of listed sponsored links is in general based on the rankings of the bids for the keyword. If the user clicks on a link, the advertiser will pay to the search engine a fee corresponding to the advertiser's bid for the keyword.

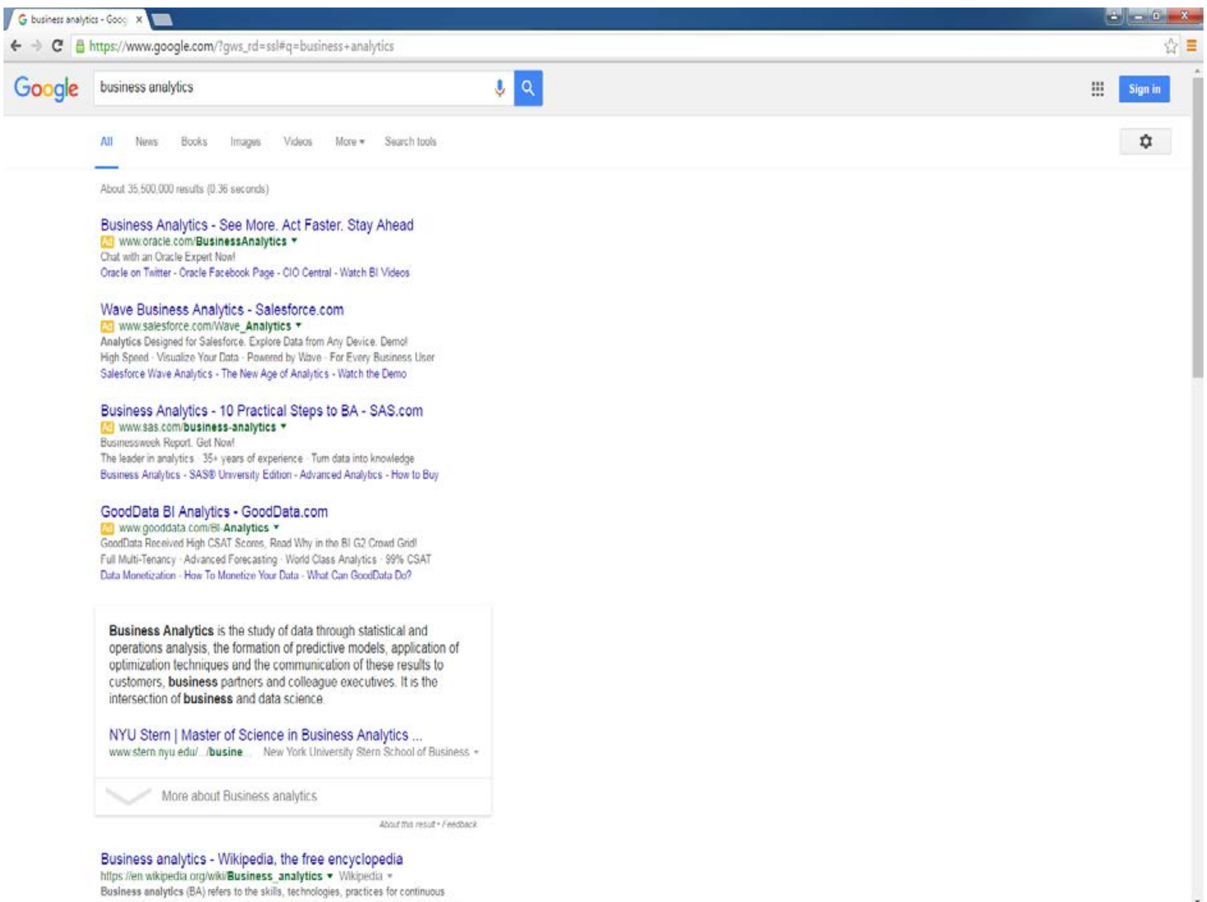

- Google's total revenue in fiscal year 2015 was $74.54 billion (http://www.statista.com/statistics/266206/googles-annual-global-revenue/). Advertising revenue amounted to $67.39 billion

Clickthrough Rates and Revenues per Click

- A slot *i* has a **clickthrough rate** r_i (e.g., the number of clicks per hour). We assume that:
  - Bidders and search engine know the clickthrough rates. This is reasonable because the clickthrough rates can be obtained from historical data.
  - The clickthrough rate depends only on the slot itself and not on the quality or suitability of the ad that is placed there. This might not be true since a high-quality ad will receive more clicks than an off-topic ad. The issue will be discussed later.
  - The clickthrough rate of a slot does not depend on the ads that are in other slots. This is not a very realistic assumption but it is a complex issue that is still not well understood even within the search industry.
  - Clickthrough rates are listed from top down; i.e., $r_1 > r_2 > r_3 > ...$
- A bidder (advertiser) *j* has a **revenue per click** v_j, the expected amount of revenue the advertiser receives per user who clicks on the ad, which is assumed to be independent of what was being shown on the page the user clicked on.

Matching Slots to Bidders with Market-Clearing Prices

- The problem of matching slots to bidders can be viewed as a problem of finding market-clearing prices and can be solved accordingly.
- Bidder *j*'s valuation for slot *i* is $v_{ij} = r_i v_j$.

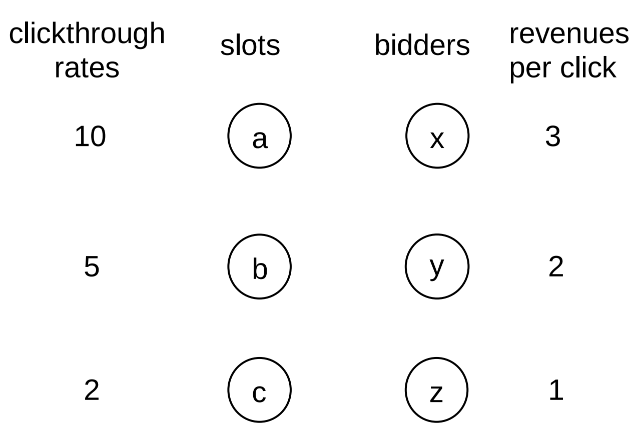

Figure 15.2

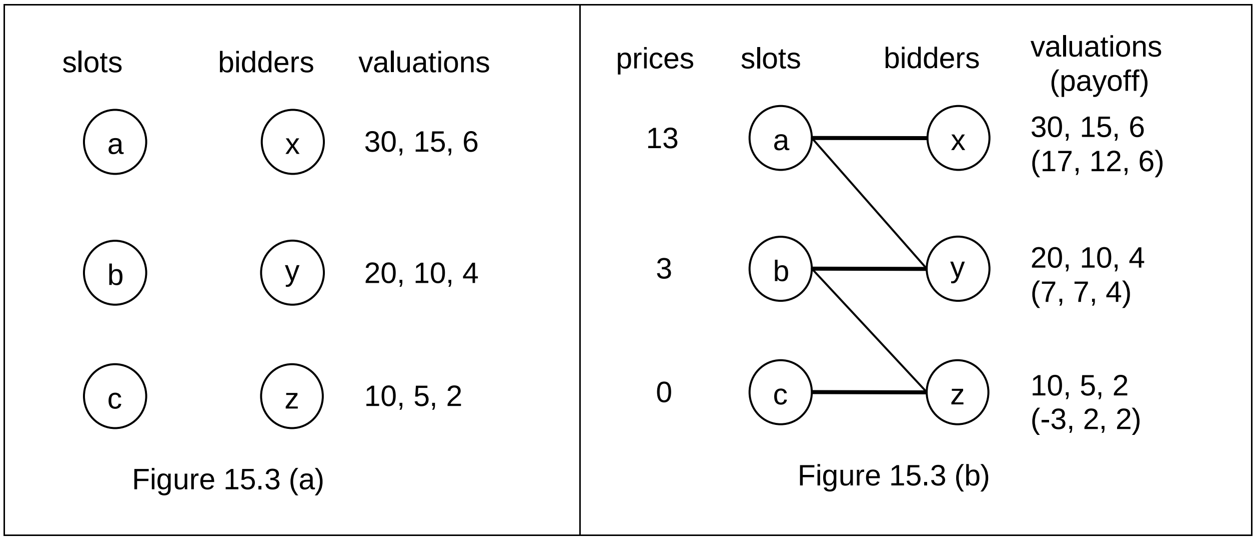

Figure 15.3 (a) Figure 15.3 (b)

To maximize the total valuation, we can always assign the slot with highest clickthrough rate to the bidder with maximum revenue per click, the slot with second highest rate to the bidder with second highest revenue per click, and so on. However, this assignment can be carried out by a search engine only if it actually knows the valuations of the bidders. In order to set prices when the search engine does not know these valuations, it is necessary to have a mechanism to encourage bidders to report these valuations truthfully.

Vickrey-Clarke-Groves (VCG) Principle

In a **Vickrey-Clarke-Groves** (**VCG**) auction, the price charged for an ad is equal to the total lose in revenue to the other bidders due to the presence of this bid in the auction. This value is equal to the total increase in valuation the other bidders would get if the optimal matching is computed without this bidder.

Apply the VCG principle to the example above:

- If *x* is not present, *y* gets *a* and *y*'s valuation is increased by 20 − 10 = 10; and *z* gets *b* and *z*'s valuation is increased by 5 − 2 = 3. The total loss to *y* and *z* caused by *x* is thus 10 + 3 = 13, and so this is the price charged to *x*.

- If *y* is not present, *x* still gets *a* (unaffected), while *z* gets *b*, for an improved valuation of 5 − 2 = 3. The total loss caused by *y* is 0 + 3 = 3, and so this is the price charged to *y*.

- Buyer *z*'s presence does not cause any loss to *x* and *y*, so the price charged to *z* is 0.

VCG Price

- S is a set of sellers (slots, items), numbered $1,...,|S|$; and B is a set of buyers (bidders), numbered $1,...,|B|$. Without loss of generality, we can assume $|S| = |B| = n$.
- V_B^S is the maximum total valuation over all possible perfect matchings of sellers and buyers (the value of the socially optimal outcome with all buyers and sellers present).
- V_{B-j}^{S-i} is the best total valuation the rest of the buyers can get if item *i* is assigned to buyer *j*, where
  - $S - i$ is the set of sellers with seller *i* removed;
  - $B - j$ is the set of buyers with buyer *j* removed.
- V_{B-j}^{S} is the best total valuation the rest of the buyers can get if buyer *j* is not present, but item *i* is still available.
- The **VCG price** p_{ij} that buyer *j* pays for item *i* is the total loss in revenue caused by buyer *j* to the rest of the buyers, which is the difference between the total valuation the rest of the buyers can get without *j* present and the total valuation with *j* present:

$$p_{ij} = V_{B-j}^{S} - V_{B-j}^{S-i} \tag{1}$$

- The second-price sealed-bid auction is a special case of the VCG auction where the number of items is 1.

VCG Mechanism: Truth-Telling as a Dominant Strategy

Claim (truthful or incentive-compatible bidding): *With the VCG mechanism, bidding with truthful valuation is a dominant strategy for each buyer, and the resulting assignment maximizes the total valuation of any perfect matching of slots and buyers.*

Intuition (informal proof):

a. Optimal matching without me does not depend on my bid.

b. Optimal matching without me and my assigned item does not depend on my bid.

c. Price paid by me for my item does not depend on my bid.

Therefore, I have no incentive to lie:

- I should not bid more than my valuation because I might pay too much; and
- I should not bid less than my valuation because I might not get the item I want.

Proof: The proof for the second part of this claim (that the total valuation is maximized) is straightforward because by definition the assignment of items is designed to maximize the total valuation, if buyers report their valuations truthfully.

To prove that truth-telling is a dominant strategy, we first recognize that if a deviation from truth-telling is to be beneficial for a buyer, the item received by the buyer must be different. Suppose, then, that buyer *j* lies about her valuations and gets item *h* instead of item *i*. To show that there is no incentive to lie and receive *h* instead of *i*, it suffices to show that

$$v_{ij} - p_{ij} \geq v_{hj} - p_{hj}\,.$$

Substituting (1) into the above inequality, we have

$$v_{ij} - (V_{B-j}^{S} - V_{B-j}^{S-i}) \geq v_{hj} - (V_{B-j}^{S} - V_{B-j}^{S-h})\,,$$

which simplifies to

$$v_{ij} + V_{B-j}^{S-i} \geq v_{hj} + V_{B-j}^{S-h}\,. \tag{2}$$

The left-hand side is V_B^S, which corresponds to the matching that achieves the maximum total valuation over all possible perfect matchings. In contrast, the matching on the right-hand side is constructed by assigning *j* with some other item *h*, and then optimally matching the remaining buyers and items. The two different matchings are shown in the figure below (E&K, Figure 15.5). The left-hand side of (2), which is the maximum valuation *with no restrictions* on who gets which slot, must be at least as large as the right-hand side, which is the maximum *with a restriction*. Therefore, inequality (2) holds.

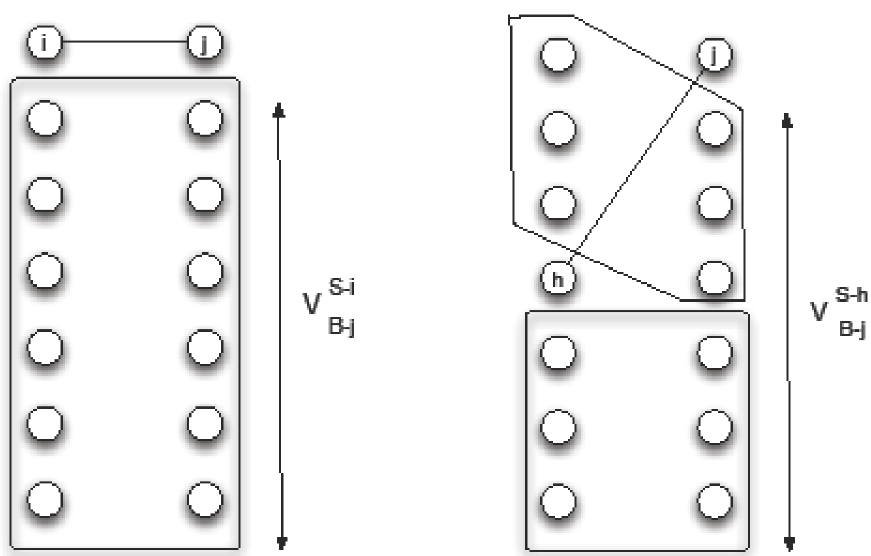

(a) $v_{ij} + V_{B-j}^{S-i}$ is the maximum valuation of any matching.

(b) $v_{hj} + V_{B-j}^{S-h}$ is the maximum valuation only over matchings constrained to assign *h* to *j*.

Figure 15.5: The heart of the proof that the VCG procedure encourages truthful bidding comes down to a comparison of the value of two matchings.

Generalized Second-Price (GSP) Auction

- In a generalized second-price (GSP) auction, each bidder submits a bid b_i. The highest bidder gets the first slot, the second-highest gets the second slot, and so on. But the highest bidder pays the price bid by the second-highest bidder, the second-highest pays the price bid by the third-highest, and so on.
- Formally, let's relabel the bidders and their bids so that $b_1 > b_2 > ... > b_n$. The GSP procedure assigns slot *i* to b_i, at a price per click equal to b_{i+1}, and collect a payment $r_i b_{i+1}$ (where r_i is the clickthrough rate for slot *i* described earlier).
- The second-price sealed-bid auction is a special case of the GSP auction when the number of items is 1. So, when only one item is listed for auction, VCG, GSP and the second-price sealed-bid auction are all the same.
- Google and search engine industry use GSP, not VCG. However, Google stated that its "unique auction model uses Nobel Prize-winning economic theory to eliminate ... that feeling that you've paid too much" [Ben Edelman, Michael Ostrovsky, Michael Schwarz (2007) "Internet Advertising and the Generalized Second Price Auction: Selling Billions of Dollars Worth of Keywords," *American Economic Review*, vol. 97, no. 1, pp.242-259].

GSP: Not a Truth-Telling (Incentive-Compatible) Mechanism

Example: There are two slots, *a* and *b*, with clickthrough rates of 10 and 4 respectively. There are three bidders, *x*, *y*, and *z*, with values per click of 7, 6, and 1 respectively. We add a fake slot *c* to equalize the number of bidders and slots.

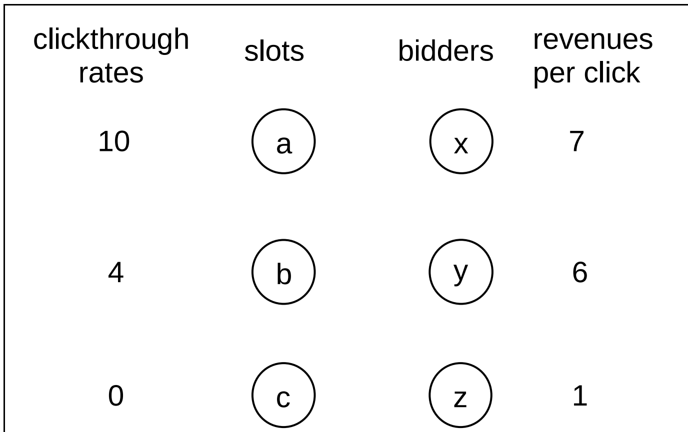

Figure 15.6

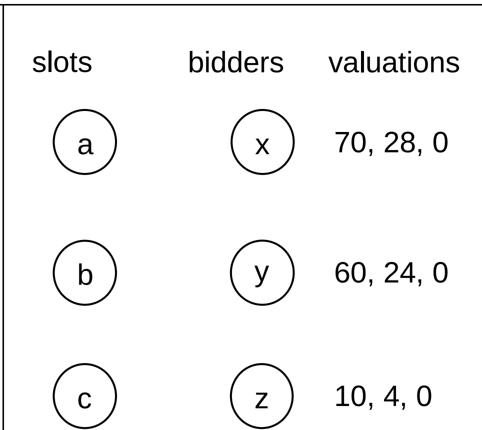

Figure 15.7

If each bidder bids truthfully, then *x* gets slot *a* at a price per click of 6; with a clickthrough rate of 10, *x* pays (6)(10) = 60 for slot *a*. The *x*'s valuation for *a* is (7)(10) = 70, so its payoff is 70 − 60 = 10. Now, if *x* were to lower its bid to 5, then it would get slot *b* for a price per click of 1, paying (1)(4) = 4 for slot *b*. The *x*'s valuation for *b* is (7)(4) = 28, so its payoff is 28 − 4 = 24, which is better than the payoff when bidding truthfully.

Nash Equilibrium

In game theory, **Nash Equilibrium** is a set of strategies (and the corresponding payoffs) having the property that no player can benefit by changing his/her strategy unilaterally while the other players keep their strategies unchanged.

- In order to analyze Nash equilibrium in the bidding game, it is typically assumed that *each bidder knows the values of all other bidders*. The argument for this assumption is that these bidders have been bidding against each other repeatedly and have learned each others' values from past experience.

Multiple and Non-Optimal Equilibria in GSP

1. For the example above, suppose that *x* bids 5, *y* bids 4, and *z* bids 2. We can confirm that this forms an equilibrium:

   - First, *x* does not want to lower its bid below 4 to get *b* instead of *a*, because $[(7-2)(4)=20]<[(7-4)(10)=30]$. Also, *x* does not want to lower its bid below 2 to get *c*, which is worse.

   - Second, *y* does not want to raise its bid above 5 to get *a* instead of *b*, because $[(6-5)(10)=10]<[(6-2)(4)=16]$. Similarly, *y* does not want to lower its bid below 2 to get *c*, which is worse.

   - Finally, *z* does not want to raise its bid above 4 to get *b* instead of *c*, because $[(1-4)(4)=-12]<0$. Raising the bid to get *a* is even worse.

   Since *x* gets *a*, *y* gets *b* and *z* gets *c*, this is an equilibrium that produces a socially optimal allocation of advertisers to slots.

2. Suppose now *x* bids 3, *y* bids 5, and *z* bids 1. We can confirm that this forms another equilibrium:

   - First, *x* does not want to raise its bid above 5 to get *a* instead of *b*, because $[(7-5)(10)=20]<[(7-1)(4)=24]$. Also, *x* does not want to lower its bid below 1 to get *c*, which is worse.

   - Second, *y* does not want to lower its bid below 3 to get *b* instead of *a*, because $[(6-1)(4)=20]<[(6-3)(10)=30]$. Similarly, *y* does not want to lower its bid below 1 to get *c*, which is worse.

   - Finally, *z* does not want to raise its bid above 3 to get *b* instead of *c*, because $[(1-3)(4)=-8]<0$. Raising the bid to get *a* is even worse.

   This equilibrium is not socially optimal because it assigns *y* to the highest slot and *x* to the second-highest.

Revenue of GSP and VCG

- For the first equilibrium above, the total revenue to the search engine is $(4)(10)+(2)(4)=48$. For the second equilibrium above, the total revenue to the search engine is $(1)(4)+(3)(10)=34$.

- Using VCG, the total valuation is maximized by assigning *a* to *x*, *b* to *y*, and *c* to *z*.
  - If *x* is not present, *y* gets *a* and *y*'s valuation is increased by 60 − 24 = 36; and *z* gets *b* and *z*'s valuation is increased by 4 − 0 = 4. So, the price charged to *x* is 36 + 4 = 40.
  - If *y* is not present, *x* still gets *a* (unaffected), while *z* gets *b*, for an improved valuation of 4 − 0 = 4. So, the price charged to *y* is 4.

  z's presence does not affect what *x* and *y* get, so the total revenue to the search engine is 40 + 4 = 44, which is in between the two revenues corresponding to the two GSP equilibria. Therefore, *the revenue to the search engine using GSP can be either higher or lower than that using VCG*.

Why Is Search Industry Using GSP?

- GSP is easier for advertisers to understand.
- Search engine's goal is to maximize revenue, not social welfare.
- Although not truth-telling, GSP has some nice properties (discussed below).

Equilibria of GSP Auction with Market-Clearing Prices: An Illustrative Example

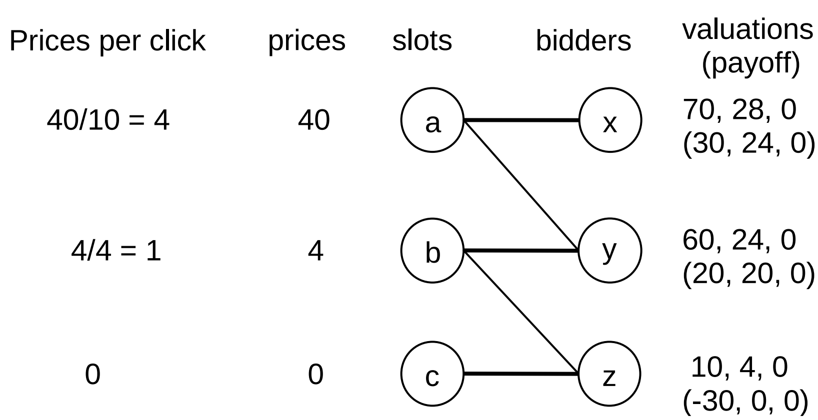

Figure 15.8

For the example above, with the market-clearing prices, the prices per click are 4 and 1 for the slots *a* and *b* respectively, so these should be the bids of *y* and *z* respectively. The *x*'s bid can be any value greater than 4. With these bids, *x* pays 4 per click for *a*, *y* pays 1 per click for *b*, and *z* pays 0 per click for *c* (the fake slot). The allocation is socially optimal.

We now use the market-clearing property to show that these bids form a Nash equilibrium. First, *x* has no incentive to lower its bid. Suppose *x* drops its bid below 4 but above 1. Then it gets *b* and pays 1 per click (the price that *y* is currently paying). This would result in a smaller payoff than that if *x* gets *a*. Similarly, *x* has no incentive to lower its bid to get *c*. For the same reason, *y* has no incentive to lower its bid.

Next, *y* has no incentive to raise its bid. In order for *y* to get *a*, it needs to pay a price per click above 4 (the price that *x* is currently paying). This would result in a smaller payoff than that if *y* gets *b*. Similarly, *z* has no incentive to raise its bid.

Equilibria of GSP Auction with Market-Clearing Prices: General Claims

Let the bidders be labeled 1, 2, …, *n* in decreasing order of their valuations per click, and let the slots be labeled 1, 2, …, *n* in decreasing order of their clickthrough rates. Consider a set of market-clearing prices for the slots, denoted $p_1, p_2, ..., p_n$ (price per click multiplied by clickthrough rate). Since any set of market-clearing prices maximizes the total valuation of the bidder-slot matchings, slot *i* will be assigned to bidder *i*; that is, the bidder with the highest valuation gets the top slot, the bidder with next-highest valuation gets the second slot, and so on.

Claim: *Let the prices per click obtained from the market-clearing prices be* $p_i^* = p_i / r_i$. *Then,*

$$p_1^* \ge p_2^* \ge \cdots \ge p_n^*. \tag{3}$$

Proof: Consider any two slots *j* and *k*, where $j < k$ (i.e., *j* is ranked higher than *k*). Bidder *k*'s total payoff for slot *k* is $r_k(v_k - p_k^*)$, and for slot *j* it is $r_j(v_k - p_j^*)$ (where v_k is revenue per click for bidder *k*). Since prices are market-clearing, bidder *k* prefers slot *k* to slot *j*. So, $r_k(v_k - p_k^*) \ge r_j(v_k - p_j^*)$, or $(v_k - p_k^*)/(v_k - p_j^*) \ge r_j / r_k$. Because *j* is ranked higher than *k*, $r_j \ge r_k$. Therefore, $(v_k - p_k^*)/(v_k - p_j^*) \ge 1$, or $p_j^* \ge p_k^*$.

Claim: *For the GSP procedure, there always exists a set of bids in Nash equilibrium that produces socially optimal assignment of the bidders to the slots. This set of bids can be constructed by having bidder i place a bid of* p_{i-1}^* *for each* $i > 1$, *and having bidder* 1 *place any amount greater than* p_1^* (*with these bids, bidder i is assigned to slot i and is charged a price per click of* p_i^*).

Proof: The set of bids are socially optimal since they are constructed based on the market-clearing prices. We show these bids form a Nash equilibrium by establishing that no bidder has an incentive to lower or raise its bid.

Consider any bidder *j*, currently in slot *j*. If it were to lower its bid and get a lower slot *k*, $(k > j)$, it would need to bid just under bidder *k* and would pay an amount that bidder *k* is currently paying. But since the prices are market-clearing, *j* would get at most the same payoff from slot *k* as it gets from slot *j*. So there is no incentive for bidder *j* to lower its bid.

If bidder *j* were to raise its bid and get a higher slot i $(i < j)$, it would need to bid just above the current bidder *i*, pushing bidder *i* one slot down. Bidder *j* would pay the current bid of bidder *i*, which is actually larger than what bidder *i* is currently paying for slot *i* (which is the bid of bidder *i* + 1). So, bidder *j* would get slot *i* at a price higher than the current price. Since prices are market-clearing, bidder *j* either prefers slot *j* to slot *i* at current prices or is indifferent between getting any of the two slots; it certainly would not pay for slot *i* at a higher price. Therefore, there is no incentive for bidder *j* to raise its bid.

Ad Quality

- The assumption that clickthrough rate is independent of the quality and relevance of ad is in general not likely to be true. This is an important concern for search engines because their revenues are directly related to the number of clicks.

- In order to address this issue, Google introduced a quality factor q_j for advertiser *j*, which can be estimated by actually observing the clickthrough rate of the ad when shown on search results pages. Then, if advertiser *j* appears in slot *i*, the clickthrough rate is estimated to be $q_j r_i$, not just r_i. Rather than assigning advertisers to slots in decreasing orders of their bids b_j, Google assigns them in decreasing order of $q_j r_i$.

- In analyzing GSP auction with quality factors, we change the valuation of advertiser *j* for slot *i* from $v_{ij} = r_i v_j$ to $v_{ij} = q_j r_i v_j$. The analytical results corresponding to pure GSP still hold.

- Introduction of ad quality factors makes the keyword-based advertising market much more opaque to the advertisers. Since the ad quality factor is determined by the search engine, it increases the search engine's power to affect the actual ordering of the advertisers for a given set of bids.