

CSC200: Lecture 26

Announcements and Outline

- Announcements

- ▶ I have added questions to and finalized Assignment 3.
- ▶ Upcoming due dates:
 - ① Quiz 6 is on February 5 (scope to be announced next Monday)
 - ② Assignment 3 is due on February 10
 - ③ Term test 2 is on February 12
 - ④ February 15-19 Reading Week

- Today's agenda

- ▶ Review [sponsored search](#) (Chapter 15) and sponsored search as a matching markets problem.
- ▶ VCG and GSP prices

Selling slots in sponsored search

Web Images Maps News Books Translate Gmail more ▼

Google

About 22,000,000 results (0.37 seconds)

Your Ad Appears Here
When potential customers search on Google.

Everything
Videos
More

The web
Pages from Ireland

Any time
Past 2 days
More search tools

Learn French
www.RosettaStone.ie With Rosetta Stone. It's easy and effective. 10% off, free shipping.

French Lessons Dublin 2
www.Languagewise.ie/French_Lessons Evening French Lessons Start June 14th. Only Native, Qualified Tutors
★ Show map of 40 Fitzwilliam Street Upper, Dublin 2, Co. Dublin

French Classes Dublin
www.Standford.academy.ie D 2&6 All levels commence in June. Native, Qualified Tutors, Enrol Now!

Free French Lessons - Learn French at About - Learn, Speak, Teach...
Learn French with free French lessons, tips, and tools. Thousands of pages related to learning and teaching the French language: lessons, quizzes, ...
Grammar - Verb Conjugator - Daily French - For Beginners
french.about.com - 10 minutes ago - Cached - Similar

French Lessons - Learn French Online
Learn French with the About French Language site - hundreds of French lessons on grammar, vocabulary, and pronunciation. Find French Lessons by level or ...
french.about.com/.../lessons/French_Lessons_Learn_French_Online.htm - Cached - Similar
★ Show more results from french.about.com

Learn French - Bonjour, Online Free French Lessons to Learn French
Bonjour! Learn French and speak online for Free at Bonjour.com. Listen to French audio words and repeat. Perfect resource for kids, students adult and ...
www.bonjour.com/ - United States - Cached - Similar

Videos for french lessons

Basic French Lesson 1: Learn basic French
7 min - 13 Aug 2008
www.youtube.com

French lessons on numbers
7 min 59 sec - 13 Jan 2008
www.youtube.com

skool.ie - Junior cycle
Welcome to skool.ie - skool.ie's new Junior Cycle French where we have exciting and ...
ive French lessons on Me, My Hobbies & Interests, ...
ool.ie/junior.asp?id=2562 - Cached - Similar

Ir Rua - French lessons beckon for Irish
lessons beckon for Irish scrimageusloademmetmoloney.jpg border0 atEmmet
Mooney writes for the a hrefhttpwww.farmersjournal.ieThe Irish Farmers ...
www.irishnews.com/story.asp?id=3059 - 14 hours ago - Cached

Search Results

The French & Surf Academy
Expert French Tutor & Surf Instruction From Irish Team Members
www.frenchandsurf.com

Learn French For Free
Learn French easily online flexible and completely for free!
www.Bonus.com

French Summer Courses
Residential 2nd Level French course in Ireland 98% Satisfaction Level
www.etc.ie

Learn French in France
French and Surf. Intensive French French Conversation and many others
www.ExperienceALanguage.ie

French Classes in Dublin
City centre, evening courses. Various levels, excellent rates.
www.EdenSchool.org
Dublin

French Distance Course
Learn From Home In Your Spare Time Complete Tutor Across Ireland
www.KilroysCollege.ie

French Language Tutor
Your own Personal French Tutor Online. Effective, Affordable
www.myp8.com

Learn to Speak French
Audio CDs or Download & Save 55% Try a Free Lesson Today!
PimsleurAudio.com/try-a-free-lesson

- Google, Bing and other search engines display several paid advertisements along with the unpaid algorithmic “organic” ranking of documents in response to any query.
- Altavista and Yahoo did initially take payments for improved rankings.

- Rather than negotiating prices for various (say common) search terms or a promise of a certain number of impressions for particular keywords, there is an ongoing **automated auction** for the paid slots.



- What is the nature of this **auction**?
- How does Google **assign slots** and **charge prices** for a given query? Cannot know for sure, but fair to say that Google and others are believed to use variants of what we will be discussing.
- We will follow the text noting that this is a high level view focusing on the **algorithmic aspects of the auction**.
- Some important aspects are not considered and still the process is technically involved. It is an area of active research.

Matching markets and the adwords problem

- The adword problem (i.e. bidding for slots) is an example of matching markets as discussed in Chapter 10.
- In a matching market, we have **sellers for each item** (or one common seller for all the items) and **buyers desiring those items** where say buyer/agent j has some **intrinsic value $v(i, j)$** for item i .
- This is a generalization of a **single item auction** but it is a special case (i.e. we are assuming unit demand) of general combinatorial auctions.
- In the adwords problem, the items are slots, each slot i having a clickthrough rate $r(i)$ and an advertiser j having an expected value $v(j)$ for a click of the advertisement.
- We assume that each advertiser wants just one slot and we assume (the **separability assumption**) that the value of slot i for advertiser j is $v(i, j) = r(i) \cdot v(j)$.

The stated assumptions

Before preceding to the game theoretic aspect where advertisers are not necessarily making **truthful bids** as to their revenue per click values, we make the following assumptions:







- Advertisers and the **search engine know** the **click-through rates**. (This is not a worrisome assumption as statistics are kept.)
- Click-through rate is **independent** of the quality and relevance of ad.
 - ▶ This is not a valid assumption and the text goes on to show how to correct this.
 - ▶ Namely, we need to incorporate a **quality factor** $q(j)$ depending on the appropriateness of the advertiser j 's ad to the query.
 - ▶ This changes click-through to be $q(j)r(i)$ if advertiser j gets slot i . Equivalently, we can think of this as altering the valuation per click to be $v'(j) = q(j)v(j)$.
 - ▶ Hence the value $v(i, j)$ becomes $r(i)q(j)v(j) = r(i)v'(j)$.
- The click-through rate and valuations **do not depend on what is in the other slots**.
 - ▶ This is also not a valid assumption and is a complex issue. That is, we are assuming **no externalities**

The quality factor $q(j)$

- It is the search engine that determines the **quality factor $q(j)$**
⇒ essentially a search ranking problem using the proposed ads as the documents to be ranked.
- Overture (and then Yahoo which bought Overture) did not initially use a quality factor $q(j)$ for measuring the appropriateness of the advertisement for the query
- **It was Google that incorporated this critical factor.**
- Since a search engine makes its profit by the number of clicks, giving out slots just on the basis of an advertiser's valuation turned out to be a bad mistake.

Weighted bipartite graph of advertisers and slots

Without loss of generality we can assume that all edges exist and there are an equal number of advertisers and slots.

clickthrough rates	slots	advertisers	revenues per click
10			3
5			2
2			1

[Fig 15.2, E&K]

Setting $v(i, j) = r(i)v'(j)$

- Let $r(i)$ be the **click-through rate** for slot i
(= number of clicks per say hour)
- Let $v'(j)$ be the **(expected) revenue per click** for the j th advertiser.
- Usually one assumes that the click-through rates favour the top slots; more specifically we assume $r(i) > r(i + 1)$.

slots	advertisers	valuations
(a)	(x)	30, 15, 6
(b)	(y)	20, 10, 4
(c)	(z)	10, 5, 2


(a) Advertisers' valuations for the slots

[Fig 15.3 (a), E&K]

Maximum value matching

Fact (as discussed in matching markets)

- The problem of computing a **maximum value matching** in a bipartite graph (or any graph) can be computed optimally in time $O(|V|^{2.5})$.
- But as noted in the text, for the **separable case** that $v(i, j) = r(i)v(j)$ (or $r(i)v'(j) = r(i)q(j)v(j)$), the maximum matching problem is very simple; namely, assuming $r(i) > r(i + 1)$, greedily give out slots in order of the modified valuations $v'(j)$.



(b) Market-clearing prices for slots

[Fig 15.3 (b), E&K]

Some unstated assumptions/issues

- Advertisers usually state a **bid per click** $v(j)$ and also a **maximum budget** $\text{Budget}(j)$ which cannot be exceeded.
- The **conversion rate** (when click becomes purchase) can often be higher for slots not quite at the top.
- Advertisements can vary in the **amount of space** they occupy.
- Advertisers might want to vary the space and valuation depending on which slot they receive.
- The search engine can (and usually will) have a **reserved price** for slots or charge calculated price + some nominal amount (say \$.01 per click).
- This is a **repeated game** where advertisers are constantly monitoring and changing their bids.

The quality factor revisited

- As noted, it is the search engine that **determines the quality factor** $q(j)$ and this is essentially a search ranking problem using the proposed ads as the documents to be ranked.
- In practice the quality factor may depend on much more than the advertisement. There are additional factors that influence the quality factor such as:
 - ① As noted, the auction for a given adword is a repeated game and the search engine (and advertiser) may want to spread out ones ad placements for an adword by not exhausting ones daily budget quickly.
 - ② The search engine can keep track of a given user's queries and that might change the quality factor for a particular ad.
 - ★ For example, previous queries might indicate that the user is very interested in a car purchase and hence an auto ad might be appropriate even if the given query (i.e. adword) isn't so suggestive of an auto purchase.
 - ③ The search engine might come to understand that the ad itself is misleading and will then bar the advertiser from the auction

Should advertisers be truthful about their valuation per click?

- Just as in **single item auctions**, advertisers (as agents) might not want to reveal their true valuations for various slots.
- But unlike more general combinatorial auctions:
 - ▶ We can efficiently compute an allocation in matching markets that is socially optimal which allows us to efficiently utilize the **Vickrey-Clarke-Groves (VCG) auction**. And for the basic way we are setting $v(i, j)$ (say ignoring the quality factor) we can simply assign by sorting the $v(j)$ values and assigning slots “greedily” to achieve the optimal social welfare.
 - ▶ VCG generalizes the Vickery second price single item auction.
 - ▶ VCG **computes an optimal allocation** (with respect to social welfare) and then **sets “VCG prices”**.
- Recall: VCG prices are based on the following idea:

Every agent is charged a price equal to the total amount of additional valuation (to other agents) if that agent were not present.

What this means for a matching market

- VCG prices are independent of ones own bid. For a matching market they have a very nice form; namely the price that agent j pays to get slot i is determined by the other bids as follows:

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

- Here p_{ij} is the **price paid** and V_{B-j}^S (resp. V_{B-j}^{S-i}) is the **optimum social welfare** with set S (resp. $S - \{i\}$) of sellers/items and set $B - \{j\}$ of buyers (with respect to any set of bids made by the other agents).
- The VCG mechanism is **truthful**; that is bidding truthfully is a dominant strategy for the agents and by definition provides a socially optimal allocation. Note that this claim regarding truthfulness holds generally whenever the global objective is the social welfare (i.e. sum of all agents valuations).

Why VCG is truthful and therefore optimal for social welfare? The argument for a matching market.

- Since the price p_{ij} is determined by the other agents bids (independent of agent j 's bid for item i), the only deviation from truthfulness will be if agent j can change the item received (say from item i to item h).
- Thus to show truthfulness, it suffices to show there is no gain in j 's utility in obtaining item h instead of item i . That is,

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

- Using the VCG prices this becomes the following:


$$v_{ij} - [V_{B-j}^S - V_{B-j}^{S-i}] \geq v_{hj} - [V_{B-j}^S - V_{B-j}^{S-h}]$$

- Which simplifies to

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

- We argue that the left hand side is V_B^S while the right hand side is at most V_B^S since it restricts solutions to assign slot h to agent j .

Figure illustrating why VCG is truthful



(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 .

[Fig 15.5, E&K]

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

VCG recap

- **Most important property:** the price that agent j pays for its allocation (in this case slot i) is determined by the other bids. The argument presented shows that **VCG is truthful** no matter what the other agents bid (i.e. whether truthful or not).
- But the VCG truthfulness depends on the optimality of the allocation. Even an algorithm which is very close to being optimal might still not be incentive compatible (i.e. have prices that will make truthfulness a dominant strategy).
- Note that VCG generalizes Vickery auctions for a single item since the “harm” caused to the other agents (in this case just one agent) is precisely the second highest price.

Does Google use VCG?

- So... since we have an optimal maximum value matching algorithm, we should be done?
- But Google (and other search companies) use something called **generalized second-price auction (GSP)** in determining its prices.
 - ▶ When agent j is allocated slot i , it pays the expected revenue for slot i that would be obtained by the bid that won slot $i + 1$.
 - ▶ This is essentially then the lowest bid that agent j can make and still obtain slot i .
- As the name suggests GSP is also a generalization of the Vickery second price auction and initially it was thought (by Google and by advertisers) that GSP was VCG or at least that it was truthful.
- **But** as the text points out subsequent analysis showed that VCG and GSP pricing is different and **GSP is not in general truthful**.

Why did the search industry use GSP?













- **Reasons to use GSP** Google adopted GSP when it was already in use by other search engines. Intuitively,
 - ▶ GSP is easier for advertisers (and us) to understand.
 - ▶ It is a natural generalization of the second price auction for a single item.
 - ▶ We will see that although not truthful, GSP always has pure **Nash equilibria (NE)** (i.e. bidding strategies) and we may want to think that over time advertisers are learning to bid according to such NE.
- The person responsible for choosing GSP thought it was truthful!
- Google's ultimate goal is not social welfare but **revenue maximization**.
- We will see that there can be many different NE and depending on the NE, the revenue can be more or less than the revenue of VCG.
- And there is always a NE for GSP that matches the revenue of VCG.

The good and the bad of GSP NEs

- **Bad:**
 - ▶ truth telling not necessarily a dominant strategy
 - ▶ multiple NE some of which do not maximize social welfare (= total of advertisers values for slots received)
- **Good:** there is a NE which does maximize social welfare
- **Bad:** Depending on which NE is chosen, revenue can be more or less than VCG
- The network in Figures 15.6 and 15.7 provides a nice example where GSP is not truthful, and where
 - ▶ there is an NE which is **socially optimal** and generates more revenue than VCG
 - ▶ An NE which is not socially optimal and generates less revenue than VCG

GSP equilibria examples

- In GSP, bidding truthfully (at 7 per click) x will win slot a and have utility $70 - 60 = 10$ whereas by lowering bid to 5 per click, x will win slot b obtaining utility $28 - 4 = 24$.
- The (x, y, z) bids $(5, 4, 2)$ achieve revenue $10 \times 4 + 4 \times 2 = 48$. This is a socially optimal NE. **What needs to be shown?**
- The (x, y, z) bids $(3, 5, 2)$ achieve revenue $10 \times 3 + 4 \times 2 = 38$. This is an NE that is **NOT** socially optimal
- VCG achieves revenue 44; $64 - 24 = 40$ from x and 4 from y .

clickthrough rates	slots	advertisers	revenues per click	slots	advertisers	valuations
10			7			70, 28, 0
4			6			60, 24, 0
0			1			10, 4, 0

[Fig 15.6, E&K]

[Fig 15.7, E&K]

Should a search engine use VCG or GSP?

As the text says, this is a complicated question.

- Since a search engines ultimate goal is revenue, should it care about social efficiency?

Should a search engine use VCG or GSP?

As the text says, this is a complicated question.

- Since a search engines ultimate goal is revenue, should it care about social efficiency? One can argue that over time, if there was a competing social engine that is more socially efficient, then advertisers might switch to use that search engine.

Should a search engine use VCG or GSP?

As the text says, this is a complicated question.

- Since a search engine's ultimate goal is revenue, should it care about social efficiency? One can argue that over time, if there was a competing search engine that is more socially efficient, then advertisers might switch to use that search engine.
- How would you set up a controlled experiment to gauge the comparative social efficiency and revenue that would be obtained **over time** by using VCG or GSP?
- For one recent laboratory experimental study, see <https://www.dropbox.com/s/flz75k4ebcibyh7/Online>