CS 357: Algorithmic Game Theory

Lecture 6: Sponsored Search Auctions

Shikha Singh

Announcements

- **HW 3** due Tues in class
  - Short examples to practice Myerson payment rule
- **Assignment 2** is due next Friday (March 7) at noon
  - Partner assignment: fill out google form
- Assignment 1 grading: in progress
- **Exam 1** will be held in class on **March 14**
  - Short-"ish" questions on topics covered until the week before
  - Composed of mostly HW style questions with 1/2 open-ended
- Results on the first-price class auction:
  - Coming soon! Maybe even today if we have time

Recap from Last Time

- So far focus has been in solving an optimization problem of **maximizing a global welfare guarantee**, when valuation is private information
- Resorted to **sealed-bid mechanisms**:
  - Ask agents to self-report their value (collect bids)
  - Treat bids as if they are truthful and allocate to maximize or approximately maximize social welfare
  - Charge payments that enforce truthful bidding as a DSE
- **Myerson payment:** Only applies to single-parameter setting but can be applied even when not doing exact welfare maximization (to approximations)
- **VCG payments:** Applies to general settings but only for exact welfare maximizations

Leftovers: VCG Challenges

- Suffers from **collusion**, same way as second-price auctions
- **Intractability** of welfare maximization
    - This is a challenge even when restricted to a single-parameter setting
- **Budget balance**: If an agent has a **negative value** (say a seller who has a **cost** involved with outcomes) then the mechanism may not generate enough revenue to compensate the seller
        - Positive payments may not equal negative payments
    - That is, the VCG mechanism may incur a **budget deficit**
- **Non- monotonity of revenue:** It may generate worse revenue when the competition increases!

Today's Plan

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

Benjamin Edelman
Harvard University
bedelman@fas.harvard.edu

Michael Ostrovsky
Stanford University
ostrovsky@gsb.stanford.edu

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

October 3, 2005

Abstract

We investigate the "generalized second price" auction (GSP), a new mechanism which is used by search engines to sell online advertising that most Internet users encounter daily. GSP is tailored to its unique environment, and neither the mechanism nor the environment have previously been studied in the mechanism design literature. Although GSP looks similar to the Vickrey-Clarke-Groves (VCG) mechanism, its properties are very different. In particular, unlike the VCG mechanism, GSP generally does not have an equilibrium in dominant strategies, and truth-telling is not an equilibrium of GSP. To analyze the properties of GSP in a dynamic environment, we describe the generalized English auction that corresponds to the GSP and show that it has a unique equilibrium. This is an ex post equilibrium that results in the same payoffs to all players as the dominant strategy equilibrium of VCG.

Today's Plan

- Discuss Edelman et al. sponsored search paper
- Discussion structure:
    - Check in with each group: how did understanding the paper and related questions go? Thoughts? Questions?
    - Take turns to present at least one question by each group
    - Discuss the rest together as a class
    - Conclude: discuss how sponsored search as since evolved

Why Study Ad Markets?

The New York Times

THE ON TECH NEWSLETTER

Google and Facebook's Ad Empires

The tech giants talk a lot about the "metaverse" and cloud computing. What really powers them is selling us socks.

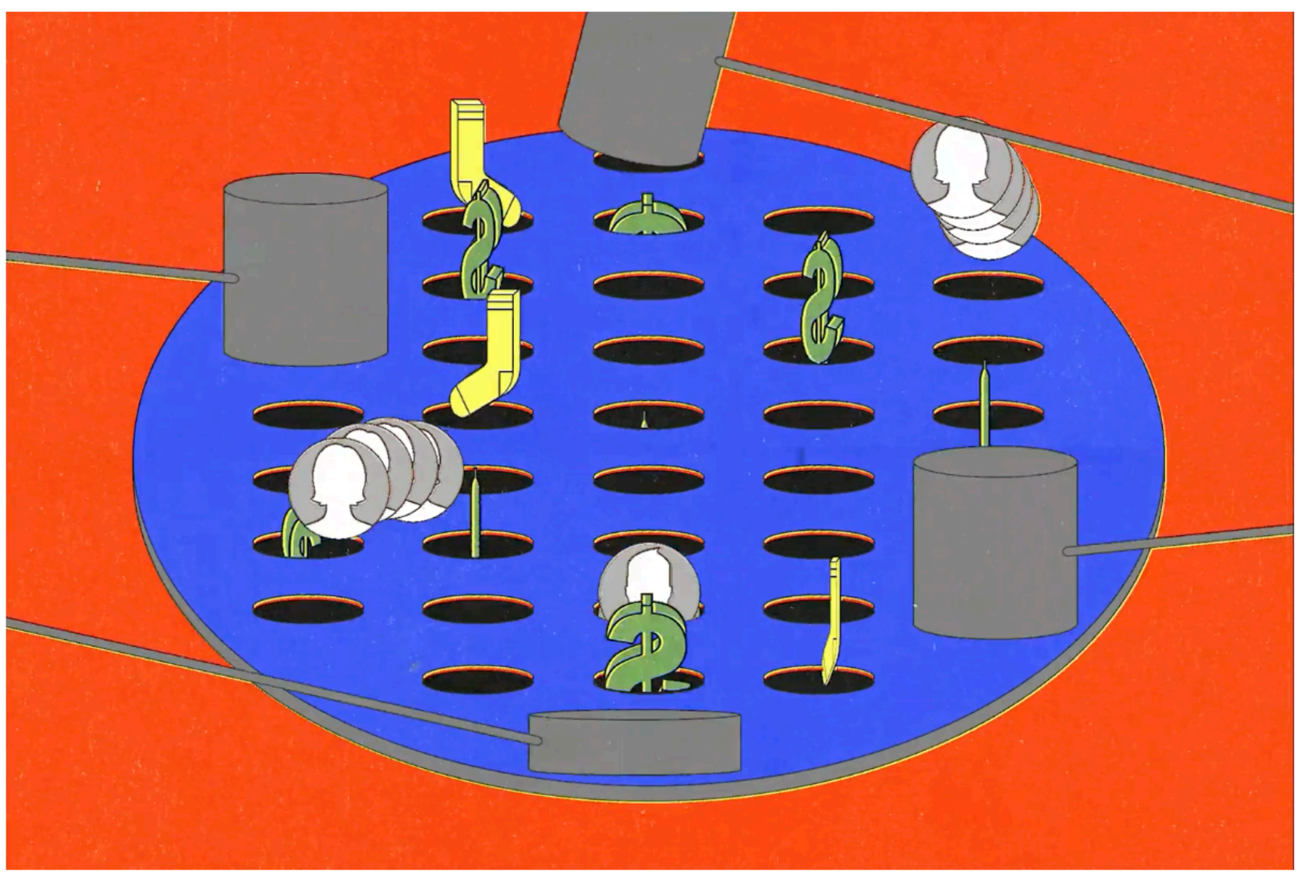

By Rapapawn

Google, Facebook, and Amazon to account for 64% of US digital ad spending this year

"AdSense counts more than ***2 million content publishers*** *as customers. Approved publishers can enter their Google code onto their sites or videos, and advertisers bid to show up* ***in those ad slots in auctions****. If a publisher's content displays an ad through AdSense, that publisher receives 68% of the revenue recognized by Google in connection with that service."*

https://www.cnbc.com/2021/05/18/how-does-google-make-money-advertising-business-breakdown-.html

Discussion

- History of how these systems evolved
- Why not use the "strategically-degenerate" truthful mechanism?
  - Re-engineering costs: search engines run campaigns on behalf of marketers and build considerable infrastructure
- What assumptions were made in the analysis? Do they seem reasonable?
- Any critiques?

GSP is not Dominant Strategyproof

- Truth telling is not a dominant strategy
- Only need to show **existence** of some **valuation and bid profile**, where utility of being not truthful is greater than truthful

Strategic Behavior

- When a DSE does not exist, harder to explain "outcome"
- Next best thing: Nash

Nash Equilibrium

- Lots of Nash equilibrium, some are inefficient and seem unlikely

High Level Overview

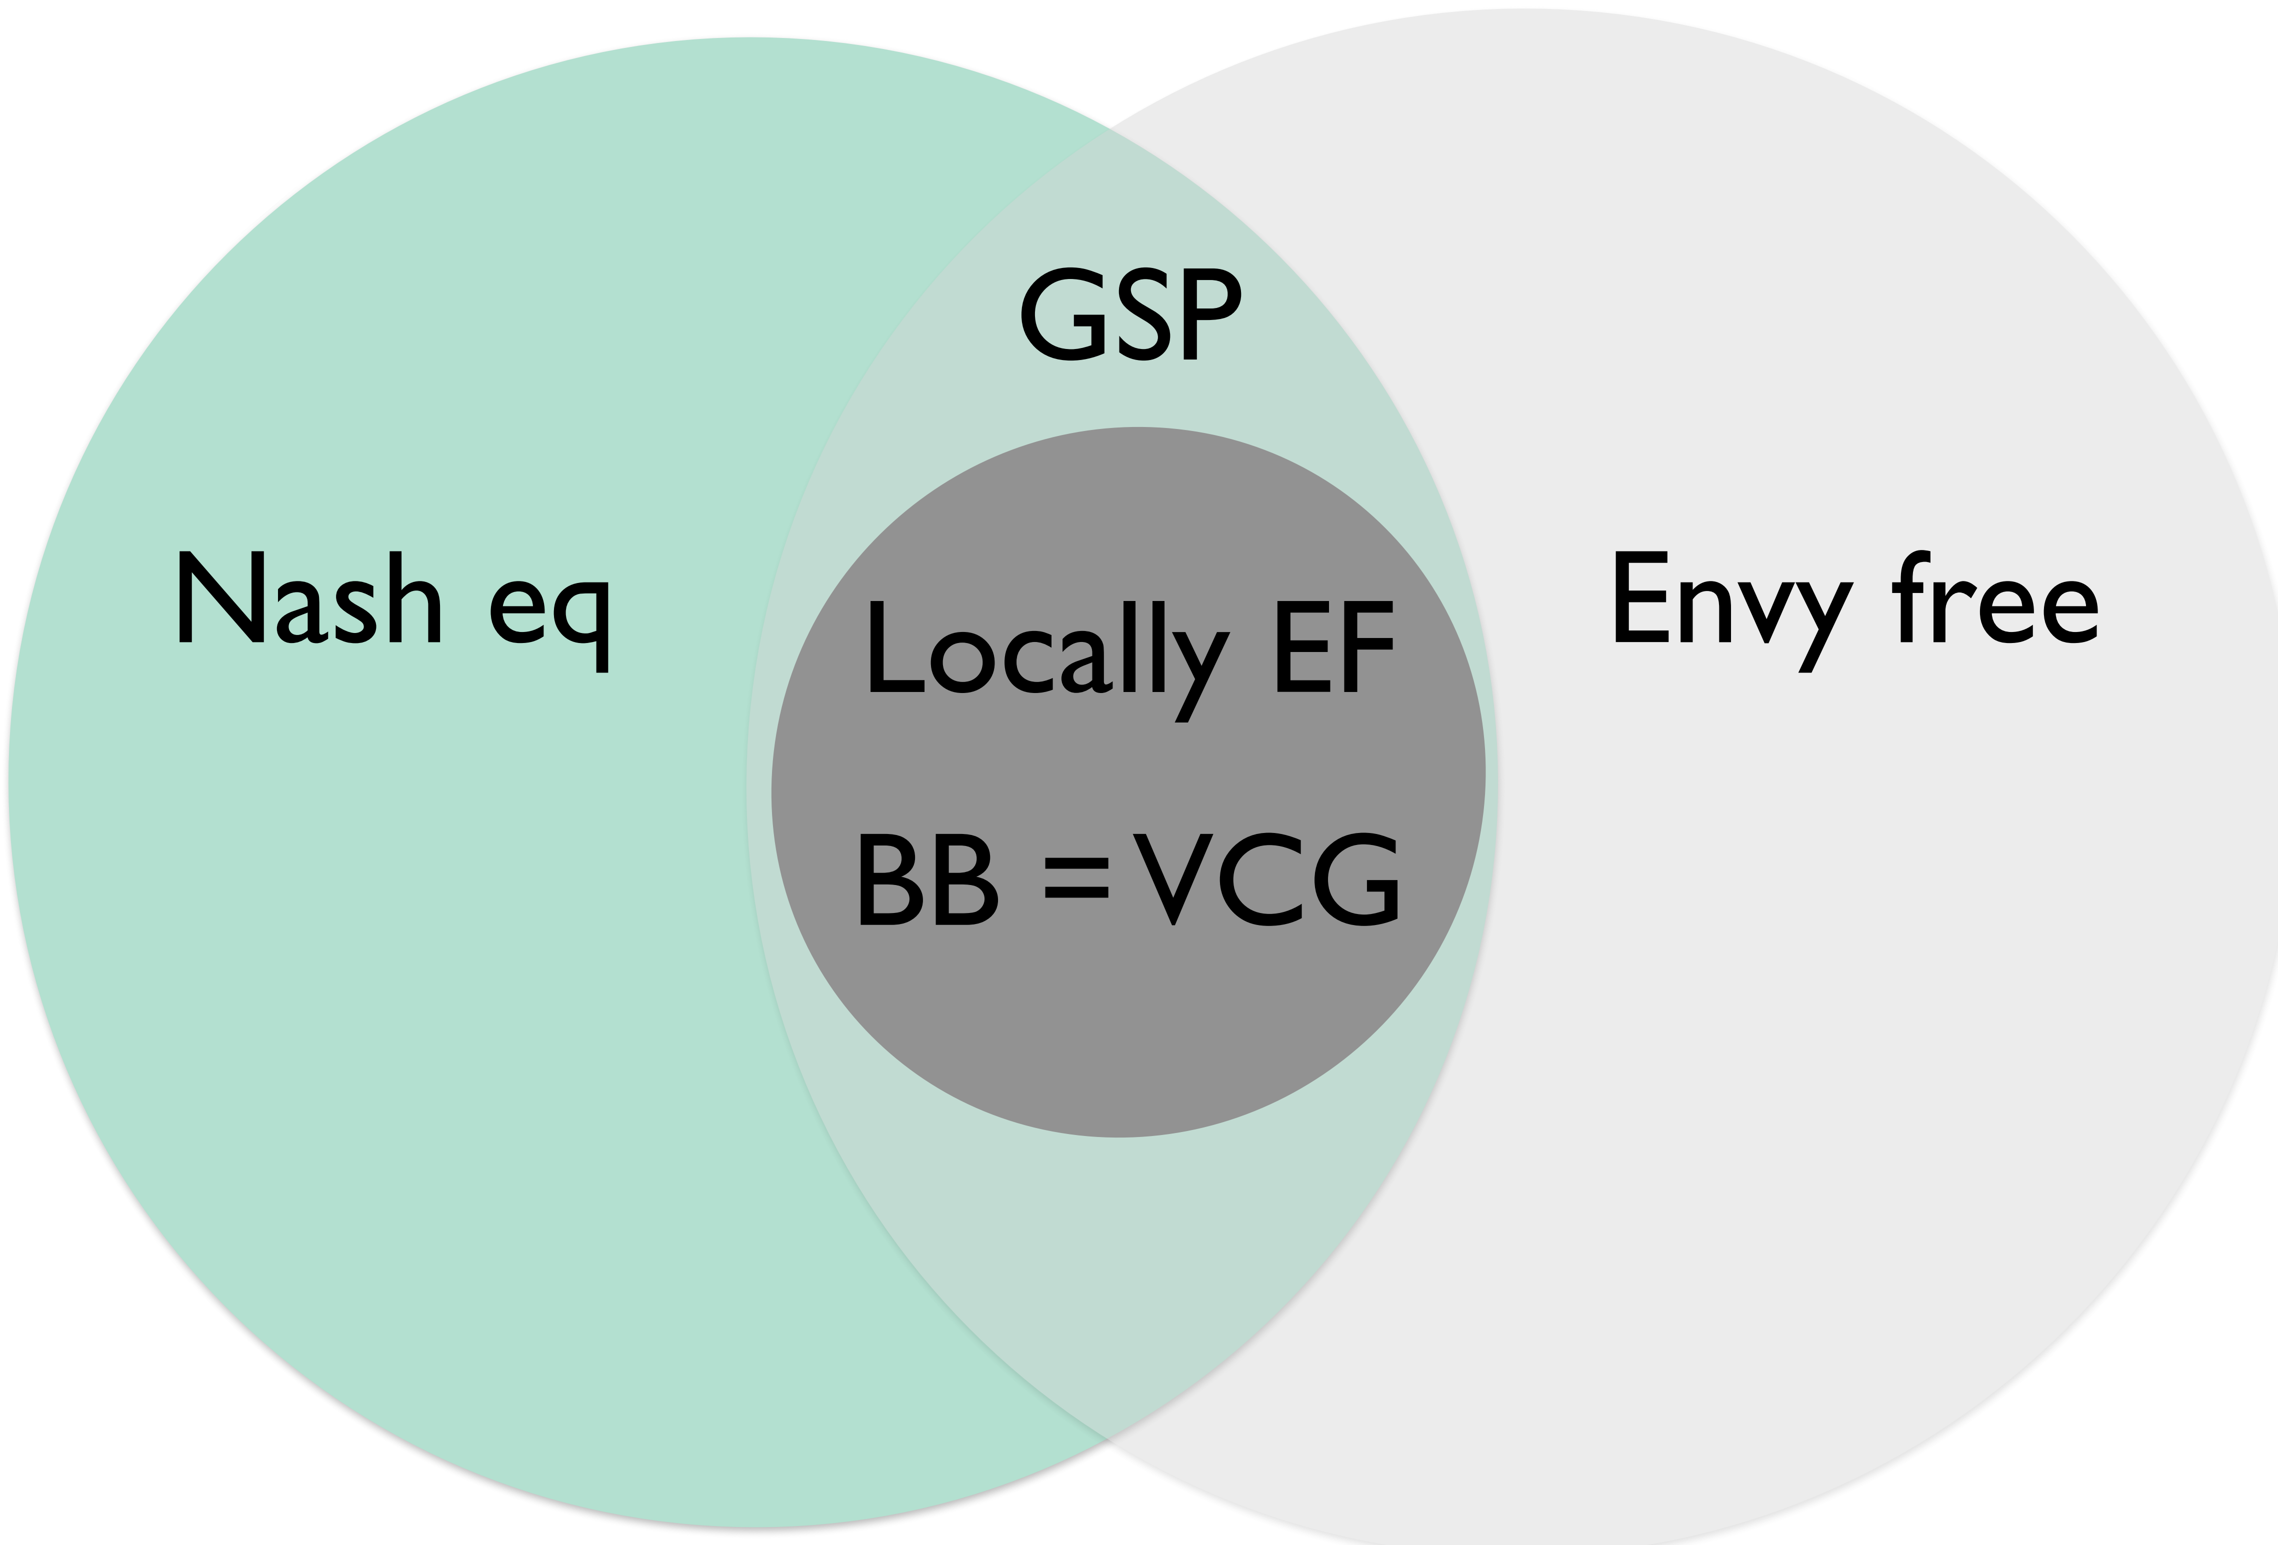

GSP's Bad Nash Equilibrium Example

- Bid profile is a Nash equilibrium, but does not maximize welfare! Economically inefficient outcome

$\alpha_1 = 0.2$	$v_1 = 4$	$b_1 = 4$	$p_1 = 2.1$
$\alpha_2 = 0.18$	$v_2 = 10$	$b_2 = 2.1$	$p_2 = 2$
$\alpha_3 = 0.1$	$v_3 = 2$	$b_3 = 2$	$p_3 = 1$
	$v_3 = 1$	$b_4 = 1$	

Nash Equilibrium

- Lots of Nash equilibrium, some are inefficient and seem unlikely
- Envy-free (solution concept) in GSP $\Longrightarrow$ Nash equilibrium in GSP

Envy Freeness: Solution Concept

- Envy-free outcome. We say that a bid profile $\mathbf{b} = (b_1, \ldots, b_n)$ where $b_1 \geq b_2 \geq \ldots \geq b_n$ is envy-free if for every bidder $\boldsymbol{i}$

$$\alpha_i(v_i - b_{i+1}) \geq \alpha_j(v_i - b_{j+1})$$

- Interpretation: (current price-per-click of slot $\boldsymbol{j}$ is $p_j = b_{j+1}$)
    - *each bidder $\boldsymbol{i}$ is as happy getting its current slot at its current price as it would be getting any other slot at that slot's current price*

Nash Equilibrium

- Lots of Nash equilibrium, some are inefficient and seem unlikely
- **Envy-free in GSP $\Longrightarrow$ Nash equilibrium in GSP**
- Lots of envy-free Nash still! Which ones are likely to be played?
    - Ones that emerge out of a reasonable "best response dynamics"

Many EF Equilibrium

- Can you come up with another bid profile that is envy free?
  - Of course 1 can increase his bid without affecting anything
  - Lets assume wlog $b_1 = v_1$

$\alpha_1 = 0.2$	$v_1 = 10$	$b_1 = 4$	$p_1 = 2$	$u_1 = 1.6$
$\alpha_2 = 0.18$	$v_2 = 4$	$b_2 = 2$	$p_2 = 1.5$	$u_2 = 0.449$
$\alpha_3 = 0.1$	$v_3 = 2$	$b_3 = 1.5$	$p_3 = 1$	$u_3 = 0.1$
	$v_3 = 1$	$b_4 = 1$		

Many EF Equilibrium

- Now how about bidder 2? Can they bid higher? Is that envy-free?

$\alpha_1 = 0.2$	$v_1 = 10$	$b_1 = 10$	$p_1 = 2$	$u_1 = 1.6$
$\alpha_2 = 0.18$	$v_2 = 4$	$b_2 = 2$	$p_2 = 1.5$	$u_2 = 0.449$
$\alpha_3 = 0.1$	$v_3 = 2$	$b_3 = 1.5$	$p_3 = 1$	$u_3 = 0.1$
	$v_3 = 1$	$b_4 = 1$		

Many EF Equilibrium

- Why not bid $b_2 = \$9.99$ and increase the price for bidder 1 which can potentially drive that bidder out of future auctions?
- What can go wrong?

$\alpha_1 = 0.2$	$v_1 = 10$	$b_1 = 10$	$p_1 = 9.99$
$\alpha_2 = 0.18$	$v_2 = 4$	$b_2 = 9.99$	$p_2 = 1.5$
$\alpha_3 = 0.1$	$v_3 = 2$	$b_3 = 1.5$	$p_3 = 1$
	$v_3 = 1$	$b_4 = 1$	

Many EF Equilibrium

- **Potential concern.** Bidder 1 could retaliate and "jam" bidder 2 by bidding 9.98 which would put bidder 1 in slot 1 at price 9.98

$\alpha_1 = 0.2$ $\quad v_2 = 4$ $\quad b_2 = 9.99$ $\quad p_1 = 9.98$

$\alpha_2 = 0.18$ $\quad v_1 = 10$ $\quad b_1 = 9.98$ $\quad p_2 = 1.5$

$\alpha_3 = 0.1$ $\quad v_3 = 2$ $\quad b_3 = 1.5$ $\quad p_3 = 1$

$v_3 = 1$ $\quad b_4 = 1$

Intuition for LEF Equilibrium

- **Idea.** Bidders will prefer highest bids amongst those that achieve the same position as it drives up the price of their competitors

$\alpha_1 = 0.2$ $v_1 = 10$ $b_1 = 10$

$\alpha_2 = 0.18$ $v_2 = 4$ $b_2 = ?$

$\alpha_3 = 0.1$ $v_3 = 2$ $b_3 = ?$

$v_3 = 1$ $b_4 = 1$

Nash Equilibrium

- Lots of Nash equilibrium, some are inefficient and seem unlikely
- Envy-free in GSP $\Longrightarrow$ Nash equilibrium in GSP
- Lots of envy-free Nash still! Which ones are likely to be played?
    - Ones that emerge out of a reasonable "best response dynamics"
- **Locally envy free (no one wants to swap with one above)**
    - Still can be more than one LEF, need to refine more to get a unique-dominant strategy for everyone

Locally Envy Free

- Does such a bid b_i always exist?
    - As long as $b_{i+1} \leq v_i$ and $\alpha_i < \alpha_{i-1}$, then yes ($b_{i+1} < b_i \leq v_i$)

$$\underbrace{\alpha_i(v_i - b_{i+1})}_{\text{utility current position}} \geq \underbrace{\alpha_{i-1}(v_i - b_i)}_{\text{utility in case of retaliation}}$$

Locally Envy Free Implies Envy Free

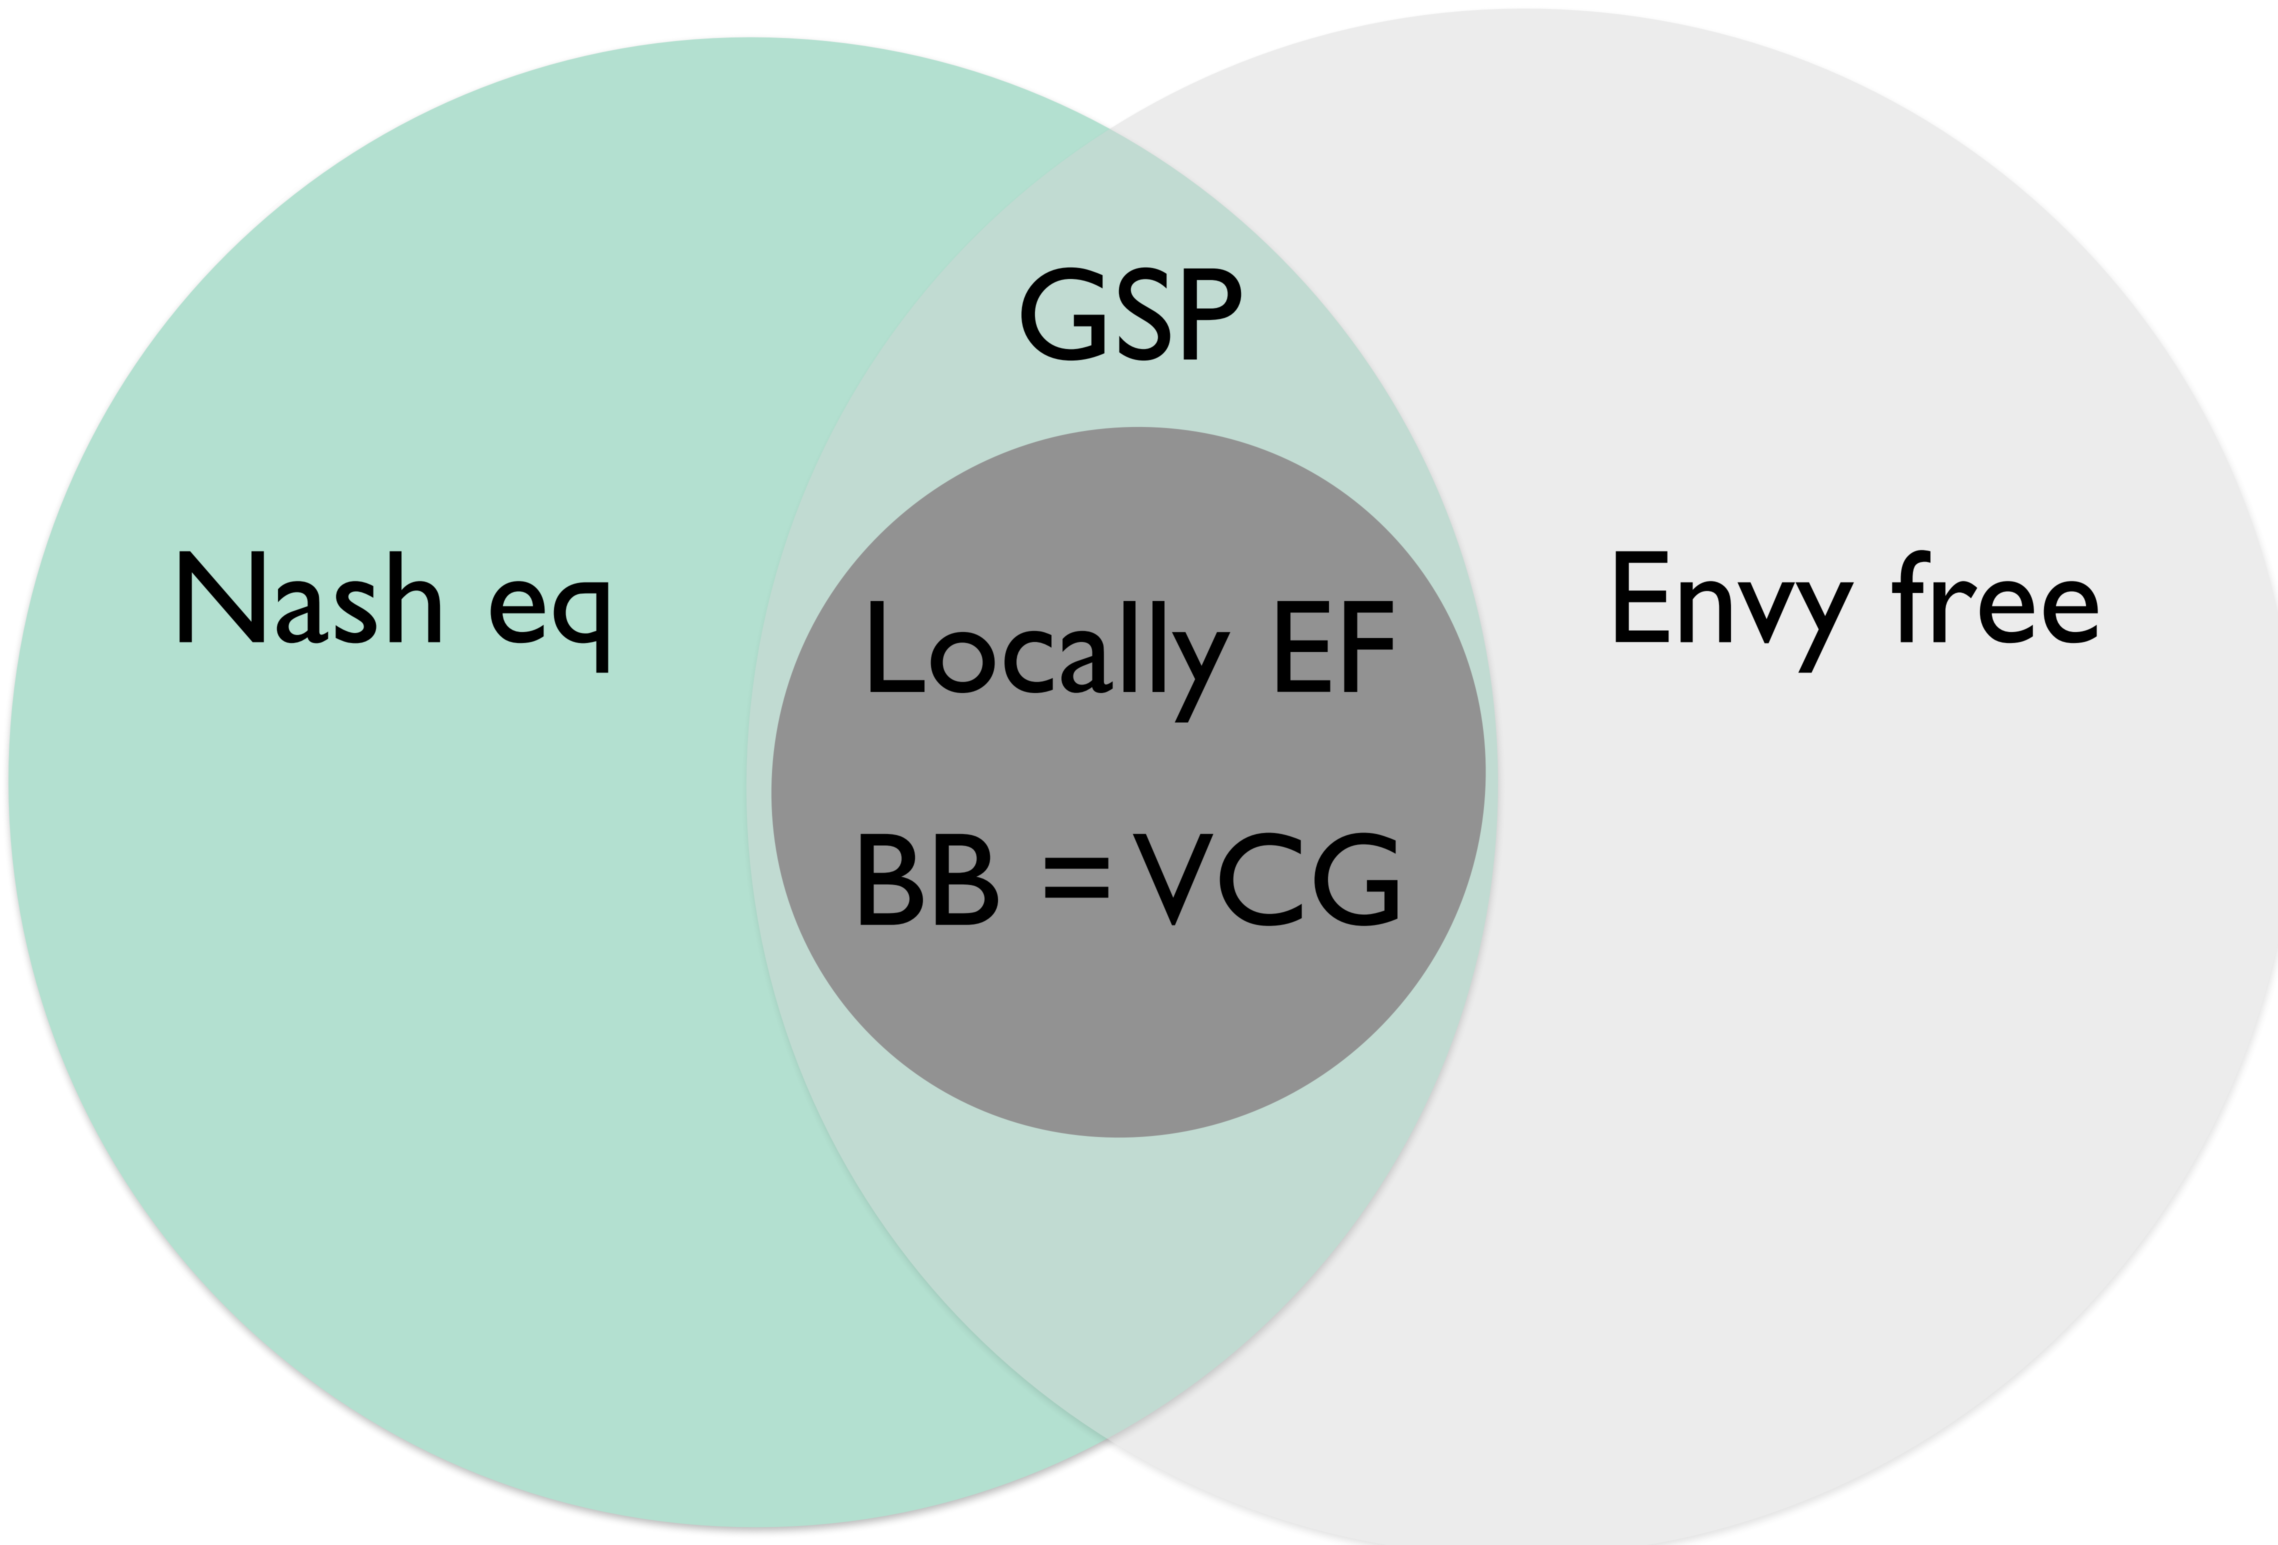

Balanced Bidding

- A bid profile $\mathbf{b} = (b_1, b_2, \ldots, b_n)$ satisfies the balanced bidding if
  - For bidder i for $2 \leq i \leq k$

$$\underbrace{\alpha_i(v_i - b_{i+1})}_{\text{utility current position}} = \underbrace{\alpha_{i-1}(v_i - b_i)}_{\text{utility in case of retaliation}}$$

  - Any unassigned bidder bids their true value
- For value ordered bids, the **balanced bidding requirement defines a unique bid profile** (up to the indifference of the top bidder)

Balanced Bidding

$$\underbrace{\alpha_i(v_i - b_{i+1})}_{\text{utility current position}} = \underbrace{\alpha_{i-1}(v_i - b_i)}_{\text{utility in case of retaliation}}$$

$\alpha_1 = 0.2$ $v_1 = 10$ $b_1 = 10$

$\alpha_2 = 0.18$ $v_2 = 4$ $b_2 = ?$

$\alpha_3 = 0.1$ $v_3 = 2$ $b_3 = 13/9$

$v_3 = 1$ $b_4 = 1$

Balanced Bidding

$$\underbrace{\alpha_i(v_i - b_{i+1})}_{\text{utility current position}} = \underbrace{\alpha_{i-1}(v_i - b_i)}_{\text{utility in case of retaliation}}$$

$\alpha_1 = 0.2$ $v_1 = 10$ $b_1 = 10$ $p_1 = 17/10$

$\alpha_2 = 0.18$ $v_2 = 4$ $b_2 = 17/10$ $p_2 = 13/9$

$\alpha_3 = 0.1$ $v_3 = 2$ $b_3 = 13/9$ $p_3 = 1$

$v_3 = 1$ $b_4 = 1$

Balanced Bidding

$$\underbrace{\alpha_i(v_i - b_{i+1})}_{\text{utility current position}} = \underbrace{\alpha_{i-1}(v_i - b_i)}_{\text{utility in case of retaliation}}$$

$\alpha_1 = 0.2$ $\quad v_1 = 10 \quad b_1 = 10 \quad p_1 = 17/10$

$\alpha_2 = 0.18$ $\quad v_2 = 4 \quad b_2 = 17/10 \quad p_2 = 13/9$

$\alpha_3 = 0.1$ $\quad v_3 = 2 \quad b_3 = 13/9 \quad p_3 = 1$

$v_3 = 1 \quad b_4 = 1$

These are exactly the VCG payments!!!

(GSP, Balanced Bidding) $\equiv$ (VCG with Truthful)

Balanced Bidding Proof

- **Theorem.** The outcome of the GSP auction in a locally envy-free Nash equilibrium bid profile $\mathbf{b}$ that satisfies balanced bidding is equal to the truthful outcome of the VCG auction.
- **Proof.** We can show $p_i[\text{VCG}] = p_i[\text{GSP}]$ using induction.
- Base case $i = k$. $p_k[\text{VCG}] = p_k[\text{GSP}] = \alpha_k b_{k+1}$
- For a slot $i < k$, we have: $p_i[\text{GSP}] = \alpha_i b_{i+1}$
    - Applying balanced bidding on bidder $i + 1$, we get $\alpha_{i+1}(v_{i+1} - b_{i+2}) = \alpha_i(v_{i+1} - b_{i+1})$
- That is, $p_i[\text{GSP}] = \alpha_i b_{i+1} = (\alpha_i - \alpha_{i+1})v_{i+1} + \alpha_{i+1}b_{i+2}$
    - $= (\alpha_i - \alpha_{i+1})v_{i+1} + p_{i+1}[\text{GSP}]$
    - $= (\alpha_i - \alpha_{i+1})v_{i+1} + p_{i+1}[\text{VCG}] = p_i[\text{VCG}]$ ∎

Revenue of Locally Envy Free

- **Theorem.** Balanced bidding equilibrium is the lowest-revenue among locally EF bid profiles.
- **Proof.** Similar inductive argument as before, but instead of using balanced bidding condition, use LEF inequality.
- Base case $i = k$. $p_k[\text{VCG}] = p_k[\text{GSP}] = \alpha_k b_{k+1}$
- For a slot $i < k$, we have: $p_i[\text{GSP}] = \alpha_i b_{i+1}$
  - Applying LEF condition on bidder $i+1$ not wanting slot i, we get $\alpha_{i+1}(v_{i+1} - b_{i+2}) \geq \alpha_i(v_{i+1} - b_{i+1})$
- That is, $p_i[\text{GSP}] \geq \alpha_i b_{i+1} = (\alpha_i - \alpha_{i+1})v_{i+1} + \alpha_{i+1} b_{i+2}$

$$\geq (\alpha_i - \alpha_{i+1})v_{i+1} + p_{i+1}[\text{GSP}]$$

$$\geq (\alpha_i - \alpha_{i+1})v_{i+1} + p_{i+1}[\text{VCG}] = p_i[\text{VCG}] \quad \blacksquare$$

Design Trade Offs

- VCG is used for contextual non real-time advertising, e.g. by X & Facebook
- Switch from GSP to VCG :
  - 2012, Google switched from GSP to VCG for its ad network AdSense
  - 2015: Yandex search engine
- **Reasons to prefer VCG over GSP?**
  - Truthful behavior: no need for bidders to strategize
  - Easier for sellers to estimate revenue
  - Enables faster experimentation: seeing how reserve prices effect revenue, etc.
  - Flexibility: VCG auction is highly configurable to different preferences and contexts

Back to First Price

- Ad exchanges moved from second-price sealed bid to first-price sealed-bid, with Google switching during 2019
- **Transparency.** Some businesses are both sellers and buyers
- Composability with different types of ads

	Non real-time	**Real-time (programmatic)**
Sponsored search	∘ Google and Bing **GSP** ∘ Yandex **VCG**	n/a
Contextual	Own inventory ∘ Twitter and Facebook **VCG** 3rd party inventory (ad networks) ∘ Google AdSense, FB Audience Network **VCG** ∘ Microsoft Audience Ads **GSP**	3rd party inventory (ad exchanges) ∘ AppNexus, Twitter, MoPub, and Google DoubleClick **First price** (was second price)

Why Do Competitive Markets Converge to First-Price Auctions?

Renato Paes Leme
renatoppl@google.com
Google Research
New York, NY

Balasubramanian Sivan
balusivan@google.com
Google Research
New York, NY

Yifeng Teng*
yifengt@cs.wisc.edu
University of Wisconsin-Madison
Madison, WI

ABSTRACT

We consider a setting in which bidders participate in multiple auctions run by different sellers, and optimize their bids for the *aggregate* auction. We analyze this setting by formulating a game between sellers, where a seller's strategy is to pick an auction to run. Our analysis aims to shed light on the recent change in the Display Ads market landscape: here, ad exchanges (sellers) were mostly running second-price auctions earlier and over time they switched to variants of the first-price auction, culminating in Google's Ad Exchange moving to a first-price auction in 2019. Our model and results offer an explanation for why the first-price auction occurs as a natural equilibrium in such competitive markets.

"Moving to a first-price auction puts Google at parity with other exchanges and SSPs (supply-side platforms) in the market, and will contribute to a much fairer transactional process across demand sources." : Scott Mulqueen

Digging Deeper

- Sponsored search markets have been a topic of extensive research in AGT
- If interested, lots of avenues to delve deeper (as part of Final Project)
  - Greedy Bidding Strategies for Keyword Auctions, Cary et al https://homes.cs.washington.edu/~karlin/papers/ecc.pdf
  - On Revenue in the Generalized Second Price Auction Lucier et al. https://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.298.4176&rep=rep1&type=pdf
  - Internet Advertising and the Generalized Second-Price Auction: Selling Billions of Dollars Worth of Keywords https://www.benedelman.org/publications/gsp-060801.pdf
- Roughgarden's course on Foundations of Sponsored Search: http://timroughgarden.org/f07/f07.html (a bit outdated but has useful resources)