Computational Analysis of Perfect-Information Position Auctions

Kevin Leyton-Brown

joint work with David Robert Martin Thompson

INFORMS: October 12, 2009

Operations Research Analogy

Consider mathematical programming:

- LP, MIP, QP (...) models of many interesting problems
- Many theoretical tools for analyzing these models
- General, computational solvers complement the theory

Operations Research Analogy

Consider mathematical programming:

- LP, MIP, QP (...) models of many interesting problems
- Many theoretical tools for analyzing these models
- General, computational solvers complement the theory

Now consider game theory, especially in the context of our focus today on sponsored search auctions:

- Expressive models
- Rich theoretical tools
- Few computational techniques

What if we could analyze auctions computationally?

Advantages:

- General valuation distribution
  - beyond e.g., strong monotonicity assumptions about value per click across slots

What if we could analyze auctions computationally?

Advantages:

- General valuation distribution
  - beyond e.g., strong monotonicity assumptions about value per click across slots
- General equilibrium concept
  - beyond e.g., locally envy-free; PSNE

What if we could analyze auctions computationally?

Advantages:

- General valuation distribution
  - beyond e.g., strong monotonicity assumptions about value per click across slots
- General equilibrium concept
  - beyond e.g., locally envy-free; PSNE
- Can handle reserve prices

What if we could analyze auctions computationally?

Advantages:

- General valuation distribution
  - beyond e.g., strong monotonicity assumptions about value per click across slots
- General equilibrium concept
  - beyond e.g., locally envy-free; PSNE
- Can handle reserve prices
- Can answer quantitative questions
  - e.g., what fraction of optimal social welfare?
  - e.g., which auction design achieves higher revenue?

What if we could analyze auctions computationally?

Advantages:

- General valuation distribution
  - beyond e.g., strong monotonicity assumptions about value per click across slots
- General equilibrium concept
  - beyond e.g., locally envy-free; PSNE
- Can handle reserve prices
- Can answer quantitative questions
  - e.g., what fraction of optimal social welfare?
  - e.g., which auction design achieves higher revenue?

(Potential) drawbacks:

- Results tied to specific valuation distributions

What if we could analyze auctions computationally?

Advantages:

- General valuation distribution
  - beyond e.g., strong monotonicity assumptions about value per click across slots
- General equilibrium concept
  - beyond e.g., locally envy-free; PSNE
- Can handle reserve prices
- Can answer quantitative questions
  - e.g., what fraction of optimal social welfare?
  - e.g., which auction design achieves higher revenue?

(Potential) drawbacks:

- Results tied to specific valuation distributions
- Discrete (rounding and tie-breaking)

Outline

Outline

Types of position auctions

- **GFP**: Yahoo! and Overture 1997–2002
- **uGSP**: Yahoo! 2002–2007
- **wGSP**: Google, MSN Live, Yahoo! 2007–present

Types of position auctions

- **GFP**: Yahoo! and Overture 1997–2002
- **uGSP**: Yahoo! 2002–2007
- **wGSP**: Google, MSN Live, Yahoo! 2007–present

Question

Is wGSP better than GFP and uGSP?

Types of position auctions

- **GFP**: Yahoo! and Overture 1997–2002
- **uGSP**: Yahoo! 2002–2007
- **wGSP**: Google, MSN Live, Yahoo! 2007–present

Question

Is wGSP better than GFP and uGSP?

- Better by what metric?
  - revenue
  - efficiency

What valuation model(s) should we consider?

Edelman, Ostrovsky & Schwarz (2007)

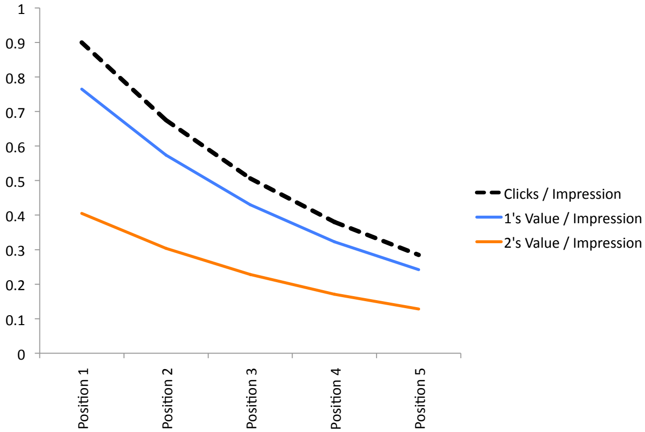

One click-through rate for everyone

Varian (2007)

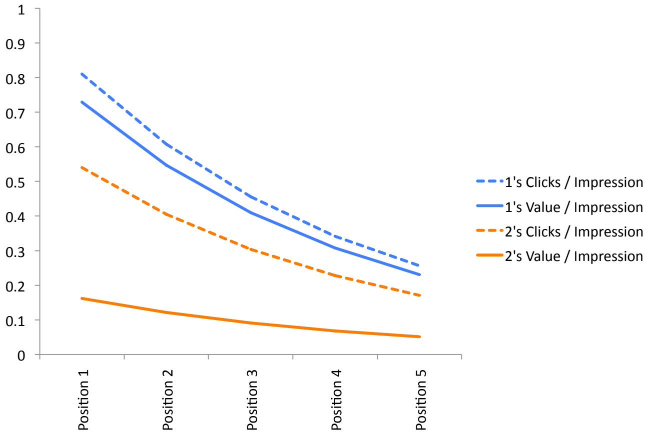

Click-through rates for different bidders are proportional

Blumrosen, Hartline & Nong (2008)

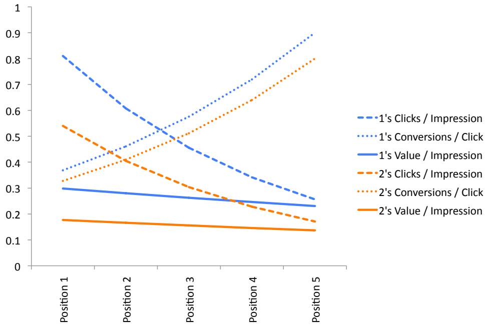

- Proportional, per-bidder click-through rates
- Proportional, per-bidder conversion rates
- Fewer clicks, higher conversion rate in lower slots

Benisch, Sadeh & Sandholm (2008)

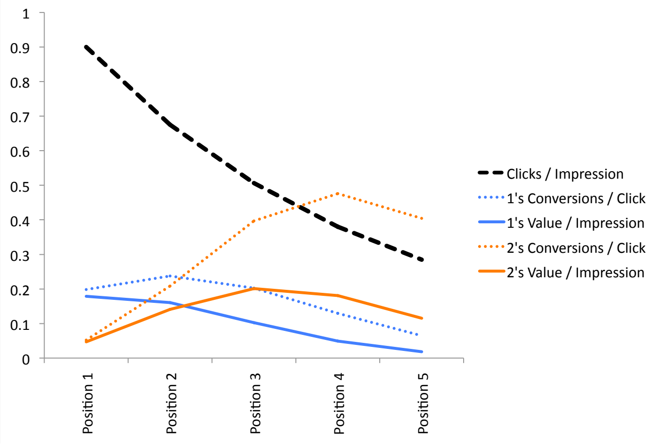

- One click-through rate for everyone
- Conversion rates are single-peaked, not proportional

Outline

Analyzing Position Auctions as Games

- Most existing literature analyzes position auctions as unrepeated, perfect-information interactions
  - unrepeated: probability one user will click on an ad is independent of the probability for the next user
  - perfect info: bidders can probe each others' values
- Given a valuation model for each advertiser and a fixed number of bid increments, we have a big normal-form game.

Analyzing Position Auctions as Games

- Most existing literature analyzes position auctions as unrepeated, perfect-information interactions
  - unrepeated: probability one user will click on an ad is independent of the probability for the next user
  - perfect info: bidders can probe each others' values
- Given a valuation model for each advertiser and a fixed number of bid increments, we have a big normal-form game.
- Problem: it's a really big normal-form game:
  - e.g., 10 bidders, 8 slots, bids in $\{0, 1, \ldots, 40\}$: ~700,000TB

Action Graph Games [Bhat, L-B, 2004; Jiang, L-B, 2006]

- A compact representation for perfect-information, simultaneous-move games
  - Like Bayes nets or graphical games: big table $\longrightarrow$ directed graph and small tables
  - Nodes correspond to actions. Table gives utility for playing a given action based on number of agents playing each neighboring action.

Action Graph Games [Bhat, L-B, 2004; Jiang, L-B, 2006]

- A compact representation for perfect-information, simultaneous-move games
  - Like Bayes nets or graphical games: big table $\rightarrow$ directed graph and small tables
  - Nodes correspond to actions. Table gives utility for playing a given action based on number of agents playing each neighboring action.
- Representational savings:
  - Exponentially smaller
  - Even smaller using function nodes (e.g. sum, max)
- Computational savings:
  - Exponential speedup in expected utility calculations
  - Implies exponential speedup in
    - simpdiv [Scarf, 1967];
    - gnm [Govindan, Wilson, 2005]
    - both are implemented in Gambit [McKevley et al, 2006]

Representing Position Auctions as AGGs

- n bidders, m bid increments
  - nm actions

Representing Position Auctions as AGGs

- n bidders, m bid increments
  - nm actions
- Position depends on number of higher/equal bids
  - add 2 sum nodes per action

Representing Position Auctions as AGGs

- n bidders, m bid increments
  - nm actions
- Position depends on number of higher/equal bids
  - add 2 sum nodes per action
- GSP price depends on next highest bid
  - add 1 max node per action

Representing Position Auctions as AGGs

- n bidders, m bid increments
  - nm actions
- Position depends on number of higher/equal bids
  - add 2 sum nodes per action
- GSP price depends on next highest bid
  - add 1 max node per action
- utility tables for each action:
  - GFP: $O(n^2)$ ($\#$ possible tuples from sum nodes)
  - wGSP: $O(n^3m)$ (also includes values of max node, which depends on both per-bidder weight and amount)
- Overall: AGGs are $O(n^4m^2)$, vs NFGs $O(nm^n)$

Representing Position Auctions as AGGs

- n bidders, m bid increments
  - nm actions
- Position depends on number of higher/equal bids
  - add 2 sum nodes per action
- GSP price depends on next highest bid
  - add 1 max node per action
- utility tables for each action:
  - GFP: $O(n^2)$ ($\#$ possible tuples from sum nodes)
  - wGSP: $O(n^3 m)$ (also includes values of max node, which depends on both per-bidder weight and amount)
- Overall: AGGs are $O(n^4 m^2)$, vs NFGs $O(nm^n)$
- 10 bidders, 8 slots, bids in $\{0, 1, \ldots, 40\}$
  - NFGs: ∼700,000TB, vs. AGGs: <80MB

Outline

Specifying details

- Game size: 10 bidders, 8 slots, values in [0,40][1]
- Game instances: 100 draws from each model
  - assuming a uniform distribution on all free model parameters
  - normalizing the highest value to be equal to the highest bid amount, so that all increments are potentially useful
- Discretization: ties broken randomly, prices rounded up, 1 increment reserve price
- Multiple runs: 10 runs each of simpdiv and gnm, randomized starting points

[1] We also considered three other sizes in our paper.

Equilibrium selection

We need to decide which equilibria to report.

- Why?
  - Our solvers return arbitrary equilibria; many exist.
  - GSP best response set is interval (sets price for bidder above)

Equilibrium selection

We need to decide which equilibria to report.

- Why?
  - Our solvers return arbitrary equilibria; many exist.
  - GSP best response set is interval (sets price for bidder above)
- How?
  - Remove bids above value (always dominated)
    - Thus we restrict to *conservative Nash equilibria* [Paes Leme and Tardos, 2009]
  - Multiple runs
  - SLS through equilibrium space
    - maximize/minimize revenue/welfare

Statistical methods

- Goal: Quantitative, comparisons across mechanisms
  - Is A better than B?
- Problem: Possibly insignificant conclusions.
- Solution: A conservative, nonparametric statistical test, with multiple testing correction.
  - ** denotes significance at or above $p = 0.01$

Outline

Efficiency: what is known theoretically?

Theorem (Edelman, Ostrovsky & Schwarz, 2007; Varian, 2007)

In EOS and V models, wGSP is efficient in every envy-free Nash equilibrium.[2]

[2] Caveat: these results apply to continuous case without reserve price.

Efficiency: what is known theoretically?

Theorem (Edelman, Ostrovsky & Schwarz, 2007; Varian, 2007)

In EOS and V models, wGSP is efficient in every envy-free Nash equilibrium.[2]

Theorem (Paes Leme & Tardos, 2009)

In EOS and V models, wGSP is 1.62-efficient in every conservative Nash equilibrium.[2]

[2] Caveat: these results apply to continuous case without reserve price.

Efficiency: what is known theoretically?

Theorem (Edelman, Ostrovsky & Schwarz, 2007; Varian, 2007)

In EOS and V models, wGSP is efficient in every envy-free Nash equilibrium.[2]

Theorem (Paes Leme & Tardos, 2009)

In EOS and V models, wGSP is 1.62-efficient in every conservative Nash equilibrium.[2]

Theorem (Blumrosen, Hartline & Nong, 2008)

There are cases in the BHN model where wGSP is not efficient in any pure-strategy Nash equilibrium.

[2] Caveat: these results apply to continuous case without reserve price.

Efficiency: what is known theoretically?

Theorem (Edelman, Ostrovsky & Schwarz, 2007; Varian, 2007)

In EOS and V models, wGSP is efficient in every envy-free Nash equilibrium.[2]

Theorem (Paes Leme & Tardos, 2009)

In EOS and V models, wGSP is 1.62-efficient in every conservative Nash equilibrium.[2]

Theorem (Blumrosen, Hartline & Nong, 2008)

There are cases in the BHN model where wGSP is not efficient in any pure-strategy Nash equilibrium.

Theorem (Benisch, Sadeh & Sandholm, 2008)

There are cases in the BSS model where wGSP is not efficient in any pure-strategy Bayes-Nash equilibrium.[2]

[2] Caveat: these results apply to continuous case without reserve price.

Efficiency: Experimental Questions

Question

When we go beyond restricted equilibrium families (e.g., envy-free), what happens?

Question

How common are efficiency failures, and how severe are they?

Results: Efficiency

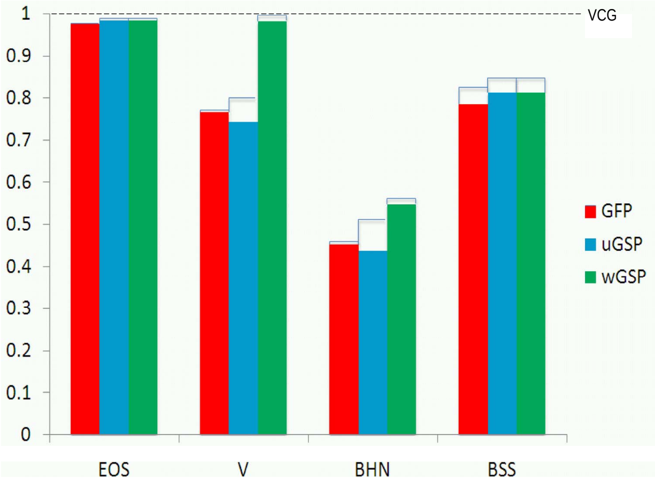

- Broad conclusion: {uGSP,GFP} $\leq^{**}$ wGSP $\leq^{**}$ VCG

Revenue: Theoretical Predictions and Questions

Theorem (Edelman, Ostrovsky & Schwarz, 2007; Varian, 2007)

In EOS and V models, wGSP generates more revenue than VCG in every "envy-free" Nash equilibrium.

Revenue: Theoretical Predictions and Questions

Theorem (Edelman, Ostrovsky & Schwarz, 2007; Varian, 2007)

In EOS and V models, wGSP generates more revenue than VCG in every "envy-free" Nash equilibrium.

Question

When we go beyond envy-free equilibria, does this result still hold?

Question

How do different auction designs compare in terms of revenue?

EOS: revenue range

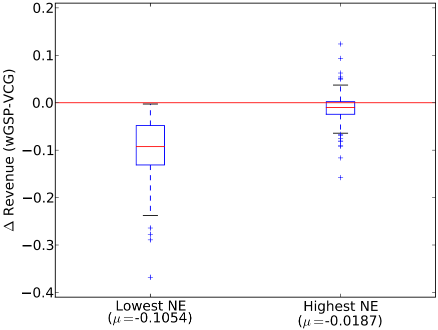

EOS: Without envy-free restriction but with restriction to conservative equilibria:

- expected worst wGSP revenue $<^{**}$ expected VCG revenue
- expected best wGSP revenue $<^{**}$ expected VCG revenue

V: revenue range

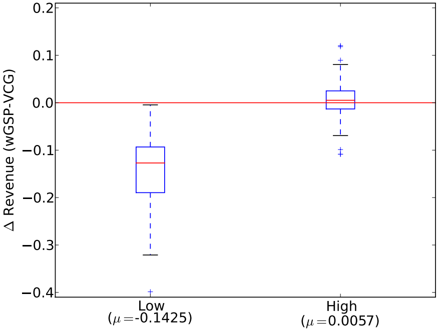

V: Without envy-free restriction but with restriction to conservative equilibria:

- expected worst wGSP revenue $<^{**}$ expected VCG revenue
- expected best wGSP revenue $>^{**}$ expected VCG revenue

V: best-case revenue

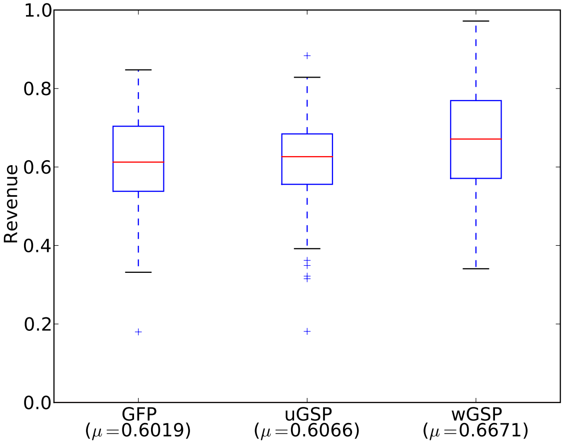

No significant revenue difference between the mechanisms.

V: best-case revenue

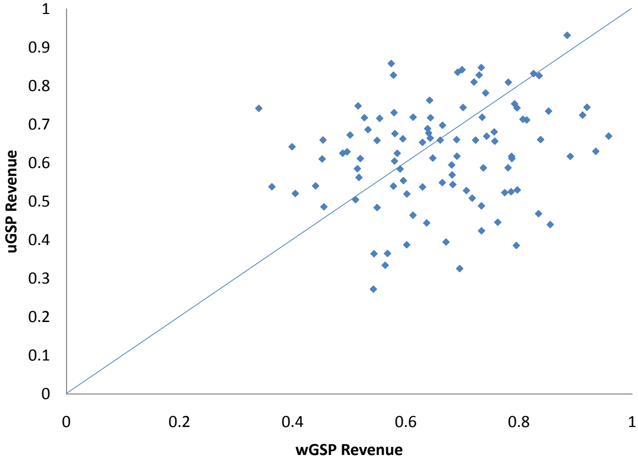

No significant revenue difference between the mechanisms.

BHN: revenue comparison

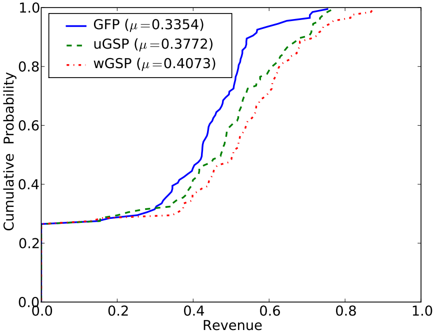

Expected wGSP revenue $>^{**}$ expected GFP/uGSP revenue

- not significant at all problem sizes we studied

BSS: revenue comparison

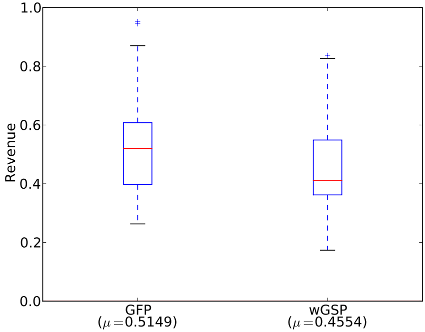

Expected GFP revenue $>^{**}$ expected uGSP/wGSP revenue

- not significant at all problem sizes we studied

Outline

Conclusion

- This approach is possible and yields real economic insights!
- Efficiency: wGSP is more efficient (even in difficult models) and very robust to equilibrium selection.
- Revenue: Ranking is unclear. Equilibrium selection and instance details have large impact.
- Code and data are available at: `http://www.cs.ubc.ca/research/position_auctions/`

This work was supported by Microsoft's Beyond Search program.

Future work

- Learning distributions from real-world data
- Generalize representation to other models (e.g. cascade)
- Better game solving techniques (e.g. provable bounds on revenue and welfare)