Designing Online Advertising Markets

Susan Athey

Introduction

- Platform markets
  - Media markets, credit cards, dating, video games, operating systems…
  - Two groups of customers, externalities (typically indirect network effects)
    - Platform can help internalize externalities
- Auction-based platforms
  - Online advertising
  - Used cars
  - eBay
- Market design matters

Visit our eBay Store

Auction-Based Two-Sided Markets

- Like other two-sided markets, except…
  - Use of auction to match buyers and sellers, determine transaction price
  - It is a method for sorting as well as a method for platform to extract revenue
  - Limits discretion of market-maker
  - Still, rules and fees play an important role in determining the size and distribution of "the pie"

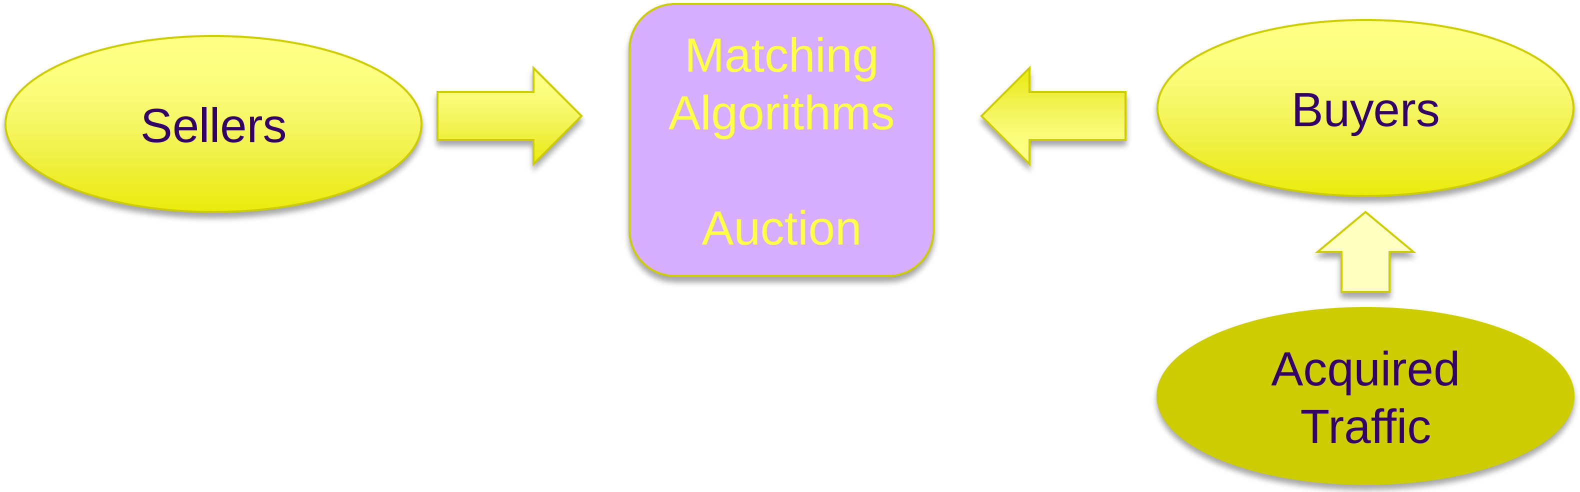

What Matters in Market Design for a Platform Market?

- Focus on efficiency and long-term participation
  - Participation crucial with competing platforms
- Extract rents to minimize distortions, keep more sensitive side of market engaged
- Auction design matters broadly
  - The design determines short term efficiency and distribution of rents
  - Determines participation and first-order issues of competition
    - See Athey, Levin and Seira (QJE, 2011) and Athey, Coey, and Levin (2011) for timber industry examples where impact of auction design on participation is quantitatively important:
      - First price v. open ascending auction format, small business set-asides

Market design v. Mechanism Design

- Mechanisms
  - Traditionally specify mapping from messages to actions, allocations and transfers, in a general setting with hidden information and/or actions
- Market Design includes mechanism design, plus, e.g.:
  - Defining the object for sale (e.g. impressions, clicks, conversions)
  - How and what information is:
    - Solicited from participants
    - Revealed to participants (accounting in search advertising, information about participants, transparency)
    - Presented (look of the website, search technology)
    - Created (upload technology, portability)
  - Set of standards, language for participants
  - Ex-post mechanisms: feedback, contract specification, remedies for bad outcomes

Influencing Market Design

- Theoretical framework is key
  - Makes arguments coherent and precise
  - Identifies equilibrium effects
  - Advertiser and consumer choices incorporated
- What kinds of theories to focus on?
  - Theories tested and fruitfully applied before
  - Key parameters can be calibrated to resolve tradeoffs
  - Robustness to real-world departures from assumptions
  - Business/gov't focuses on expected outcomes within a range, avoiding disasters, or particular upsides
    - Mitigation could include changing the rules ex post

The Role of Data and Experimentation

- In online marketplaces, data plays a key role
- Data can inform what kinds of designs will work better or worse in range of environments similar to existing one
- Advocacy for design issues is much more effective with theory and data combined
- Experimentation crucial but also has limitations
  - Short-term experiments can't show long-term outcomes, feedback effects
- One part of empirical economics focuses structure on empirical analysis in order to learn model "primitives" and perform "counterfactuals"
  - Learn bidder values, predict equilibrium responses
  - Map between short-run user experience and long-term willingness to click

Search Advertising Auctions

- Advertiser submit bids for keywords
  - Offer a dollar payment *per click*.
  - Alternatives: price per impression, or per conversion.
- Separate auction for every query
  - Positions awarded in order of bid (more on this later).
  - Advertisers pay bid of the advertiser in the position below.
  - "Generalized second price" auction format.
- Some important features
  - Multiple positions, but advertisers submit only a single bid
    - Choice to sell clicks rather than different positions on the page key design choice, results in much thicker market
  - Search is highly targeted, and transaction oriented.

Current Auction Format

- Real-time
- Pay-per-click
- Click weighted
- Generalized second price auction

Generalized Second Price Auction with Click-Weighting

- Price for position *m* determined using *m* + 1st revenue per impression
- Bidder 1 pays $s_2 b_2 / s_1$ per click; this is the lowest price that would have put him in the first position.
- If *s* is the click-through rate and *R* is a per impression reserve price, these would be prices (but note that *s* may incorporate other factors):

Per-Click Bid	Estimated Revenue Bid	Price Per Click	Estimated Revenue (normalized to first position)
b_1	$s_1 b_1$	$s_2 b_2 / s_1$	$s_2 b_2$
b_2	$s_2 b_2$	$s_3 b_3 / s_2$	$s_3 b_3$
b_3	$s_3 b_3$	$s_4 b_4 / s_3$	$s_4 b_4$
b_4	$s_4 b_4$	R / s_4	R

Why a real-time auction?

- Real-time auction
  - Prices vary widely: $.05 to $30, $50 or even higher
  - Millions of search phrases
  - Hundreds of thousands of advertisers
  - Small advertisers join system all the time
  - Time sensitive items
  - Changing products and profits
- Costs of real-time auction
  - Firms need to monitor and fine-tune
  - Commitment to auction limits ability to price discriminate (but granular reserve prices available…)

Why a pay-per click (PPC) auction? (1)

- Alternatives
  - Pay per “impression” – this is what is being sold
  - Pay per action – pay for conversions

- Pay per impression
  - Advertiser bears risk for traffic quality
  - Search engine has private information about bot traffic, etc.
  - Traffic quality on partner network especially uncertain
  - In long term, advertisers monitor performance and adjust bids, but short term risk from fluctuations

Why a pay-per click (PPC) auction? (2)

- Pay per action
  - Use cookies to track conversions
  - Allows firms to enter profit on each item to advertising platform and let platform optimize for you
  - Events are rare, difficult to estimate probabilities
  - Advertiser controls (and can manipulate) probabilities
    - Agarwal, Athey and Yang (AER, 2009) show how multi-dimensional actions lead to manipulation in equilibrium
    - Successful pay per action models often have a close relationship with retailer, can monitor conversions
- Pay per click
  - Thicker market than pay per impression
  - Minimizes risk for advertisers relative to pay per impression, generalizes to broad match
  - Simpler for search engine, easy technology

Click-weighted auctions

- All three search engines use click-weighted auctions
  - Original design: rank firms by their bids
  - Yahoo! last to introduce in 2008 (now Bing provides Y! results)
- Unweighted per-click auctions can lead to lower revenue
- Example: search for Paris
  - Paris, France travel
    - Profit $.50/click, click-through rate 5% = $.025
  - Paris Hilton sex videos
    - Profit $5/click, click-through rate .25% = $.0125
  - Paris Hilton sex videos outbids Paris, France travel, revenue lower
- Advantages of unweighted or partially weighted auctions
  - Price discrimination (Lahaie and Pennock)
  - Do not require search engine to estimate click-through rates
  - Advertisers don’t want unnecessary clicks, write accurate ad text (see Athey and Ellison; note that in click-weighted auctions, more clicks leads to proportionally lower PPC so advertiser “does not pay” for the clicks in steady state)

Auction Format

- Pay-your-bid
  - Early auction designs used a pay-your-bid format
  - Outcome was cycling
  - Firms incrementally outbid one another prices so high that it is unprofitable to win; then drop bid dramatically
  - Leads to inefficient outcomes (Edelman & Ostrovsky)
- Generalized second price
  - Bing and Google now use a variant of a second-price auction, where advertisers pay minimum required to maintain position
  - Nice stability properties
  - The search engines continue to update their rules
- Vickrey
  - In simplified model, outcome-equivalent to the lowest-revenue envy-free Nash equilibrium of GSP
  - Athey-Nekipelov show that GSP is inefficient in more realistic environments

Choices in Auction Design in Baseline Model

- Baseline equilibrium analysis: Varian, Edelman/Ostrovsky/Schwarz
  - Model can be used to analyze market design
- How many slots to sell?
  - Revenue v. Efficiency
- Setting a reserve price?
  - Optimal reserve prices raise more revenue than restricting slots
  - Myerson's approach can be adapted and optimal reserve price is independent of the number of slots or competitors (Ostrovsky & Schwarz) assuming a priori symmetric bidders
- Clickability, "squashing", and quality scores for price discrimination
  - Downweighting clickability in rankings can increase revenue at the expense of efficiency (Lahaie and Pennock)
  - A form of price discrimination

Market Design Matters in Practice

- Historical benchmarking and various industry studies assert:
  - Ads per page, ads in top positions above algorithmic results
    - Google < Microsoft < Yahoo!
  - Google's ads more "relevant"
    - Better advertiser base, more exact match bidding, more data to predict user response to ads
    - More selective policies (filtering, matching, reserve prices)
  - Google has much higher RPS, higher cost per acquisition

Extending Theory to Incorporate Platform Economics

- Market design incorporating consumers:
  - Endogenize clicking behavior
  - Incorporate search costs and welfare
- Athey-Ellison (2008, forthcoming QJE) setup
  - Ads vary in quality, this is unknown to consumers
  - Distribution of consumer search costs in population
  - Consumer propensity to search depends on anticipated quality
  - Analyze the impact of market design choices on welfare

Athey-Ellison Results

- Eqm bidding → ads ranked by quality
- Propensity to click depends on perceived quality
  - Implies platform must consider impact of policies on quality
- Market design considerations altered to economize on consumer search and convey information to consumers
- Total welfare and consumer welfare are proportional
  - Consumer welfare is producer welfare LESS search costs--divergence
  - Consumers search in proportion to perceived quality of ads, which is proportional to value created for advertisers before fees
  - “Tuning dials” for consumer welfare also maximizes social welfare
- Platform profits in conflict with advertiser profits
- Optimal reserve prices derived
- Reduced use of click-through weighting incentives more accurate ad text, economizes on consumer search
- Asymmetries lead to non-existence of efficient equilibria

Consumer, Advertiser Participation Exogenous (Monopoly Platform)

- Reserve prices
  - Reduce supply of advertisements (size of "pie")
  - Redistribute surplus between search ad platform and advertisers
- Consumers are helped by low to moderate reserve prices (conserving search costs), but harmed by high reserve prices because fewer ads are displayed
- Monopoly search ad platforms raise reserve prices beyond the social optimum
  - At social optimum, further increase in reserve raises search ad platform profits faster than it hurts total surplus
  - Search ad platform sets reserve prices inefficiently high

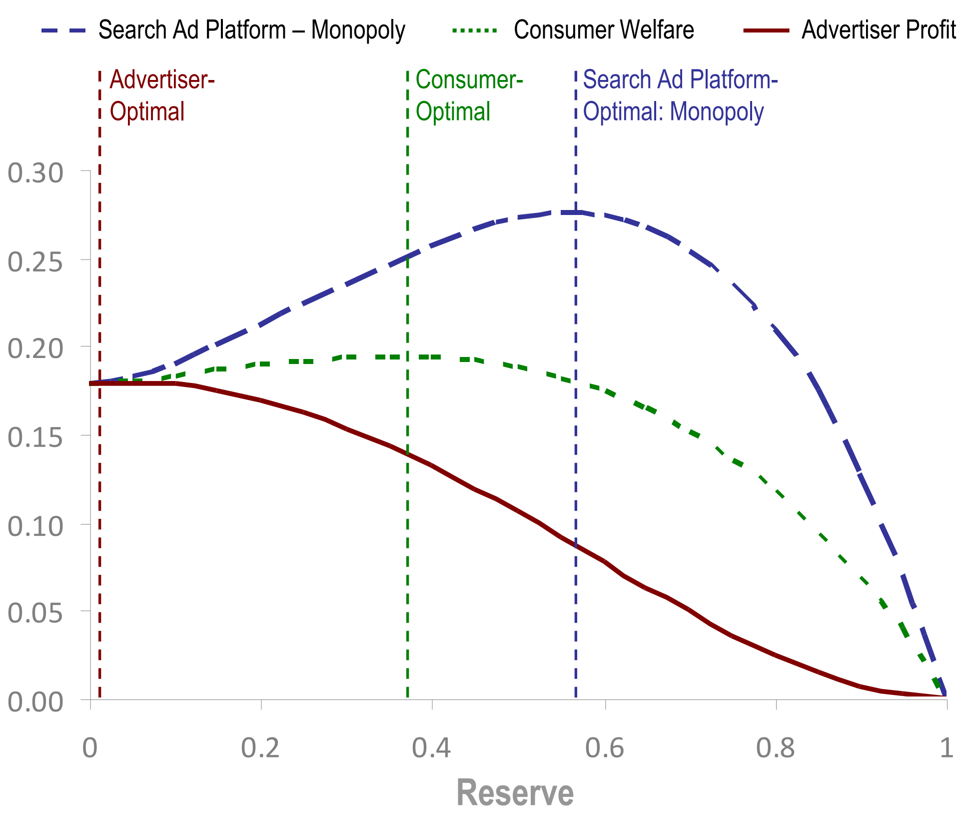

Competition Reduces Reserve Prices

- Competition pressures search ad platform to lower reserve price
  - Consumers and advertisers more responsive to reserve price in presence of a competing search ad platform
- Lower reserve prices attract more consumers and more advertising spend
- Advertiser profit is more sensitive to reserve prices than consumer welfare
- Consumer optimum is close to social optimum

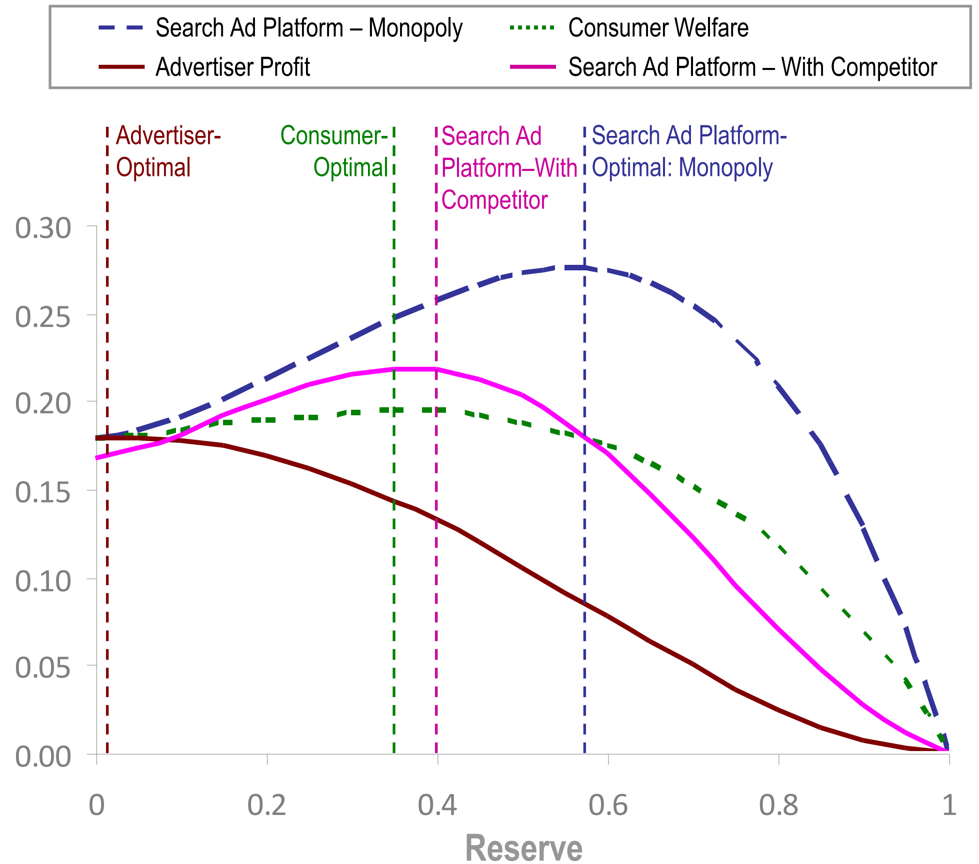

A Structural Model for Counterfactual Analysis

- Many auction design issues involve tradeoffs
  - Bidder-facing experiments expensive
  - With a calibrated model, you can quantify the magnitudes and make predictions about long-term bidder responses
- Athey-Nekipelov (2010)
  - The same bid applies to many user queries
  - Model the uncertainty faced by advertisers
  - Establish existence & uniqueness of equilibrium
  - Develop structural model, show identification, statistical properties
  - Estimate bidder values using historical data from Microsoft
  - Develop computational algorithm to compute equilibrium bids using homotopy method, apply to do counterfactual simulations
  - Shows that under uncertainty, GSP is inefficient, but revenue comparison with Vickrey is ambiguous

Advertiser Order Database

Delivery Engine

Scoring Algorithm

Bid for “mortgage calculator”: $X/click

Ad 1: **Mortgage Offers** - www.LendingTree.com
$400,000 for Only $1,910/Month or $200,000 for Only $955/Month!

Ad 2: **Mortgage - LendingTree®** - www.LendingTree.com
$200,000 for Only $955/Month. When Banks Compete, You Win.

1. User enters query
2. Delivery engine queries database to identify applicable bids
3. Scoring algorithm produces scores
4. Ads are selected, ranked and scored; no more than one ad per account on a page
5. User clicks on ads

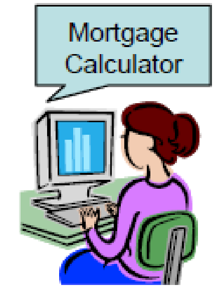

Process repeats for new user

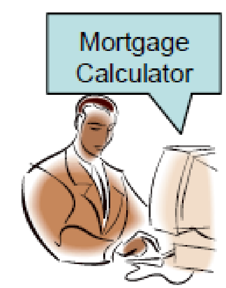

Reformulate Problem

$$\max_b v \cdot Q(b) - TE(b)$$

$$TC(q) = TE(Q^{-1}(q))$$

$$\max_q v \cdot q - TC(q) = q \cdot (v - AC(q))$$

$$\text{FOC: } v = MC(q) = AC(q) + q \cdot AC'(q)$$

- This is just classic monopsonist problem
- Can also relate it to more standard uniform-price auction objective function
- Can estimate these quantities from search engine data by simulating impact of hypothetical bid changes

Estimates of AC(q), MC(q), and implied value for a high-value search phrase

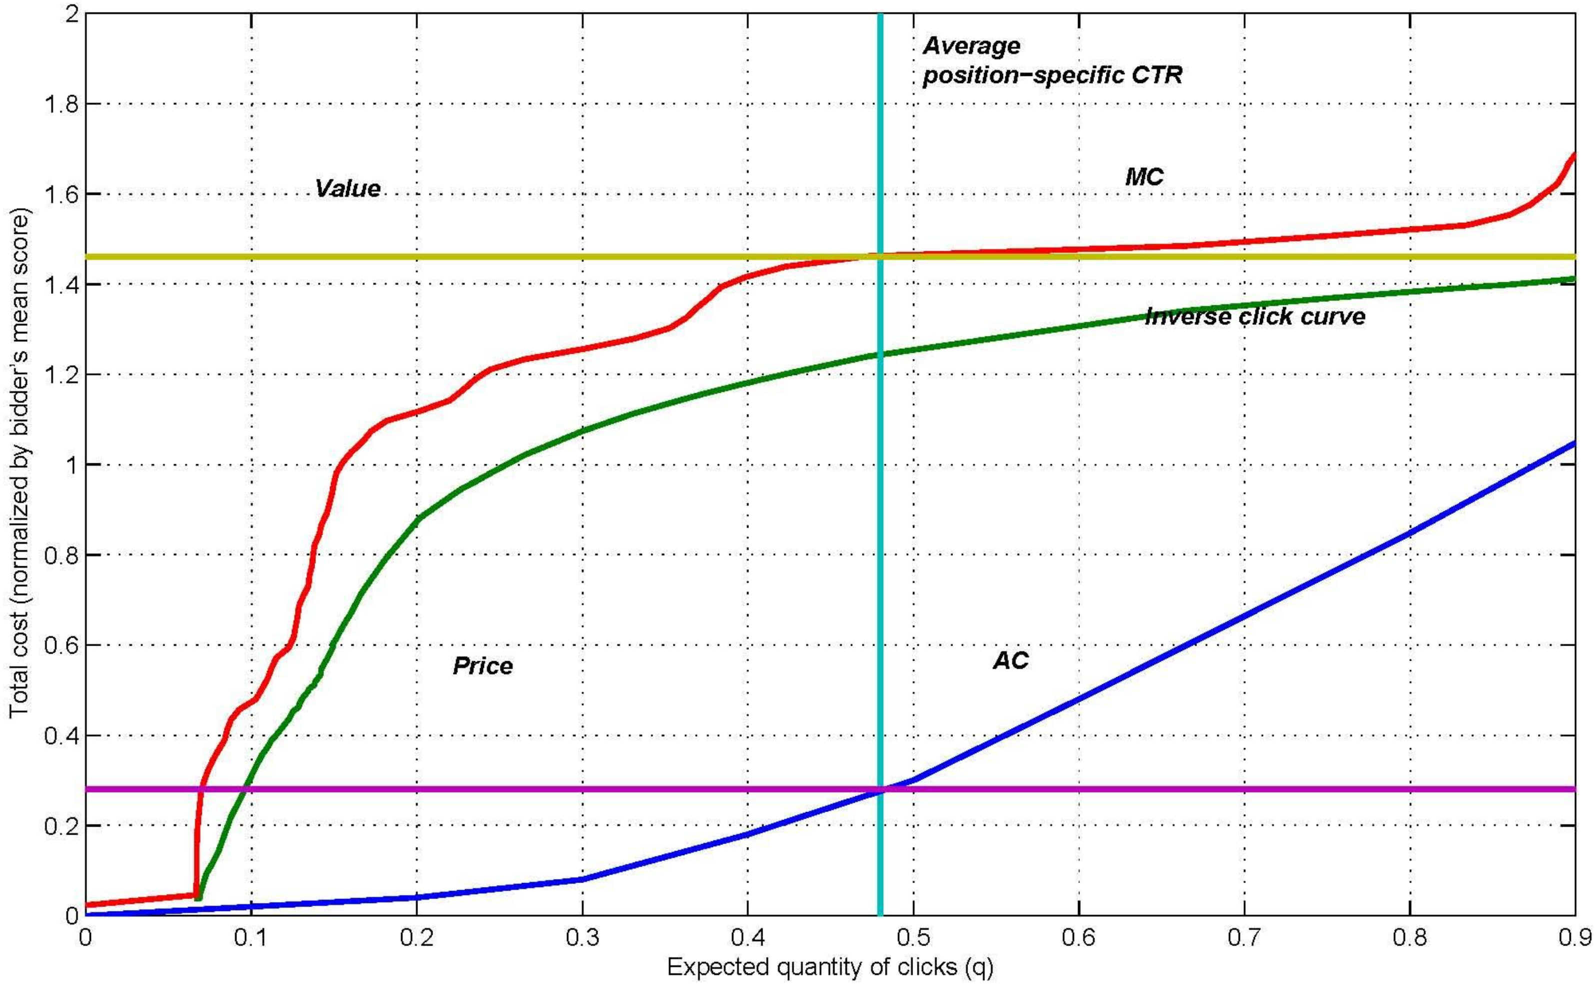

Model Predictions Out of Sample

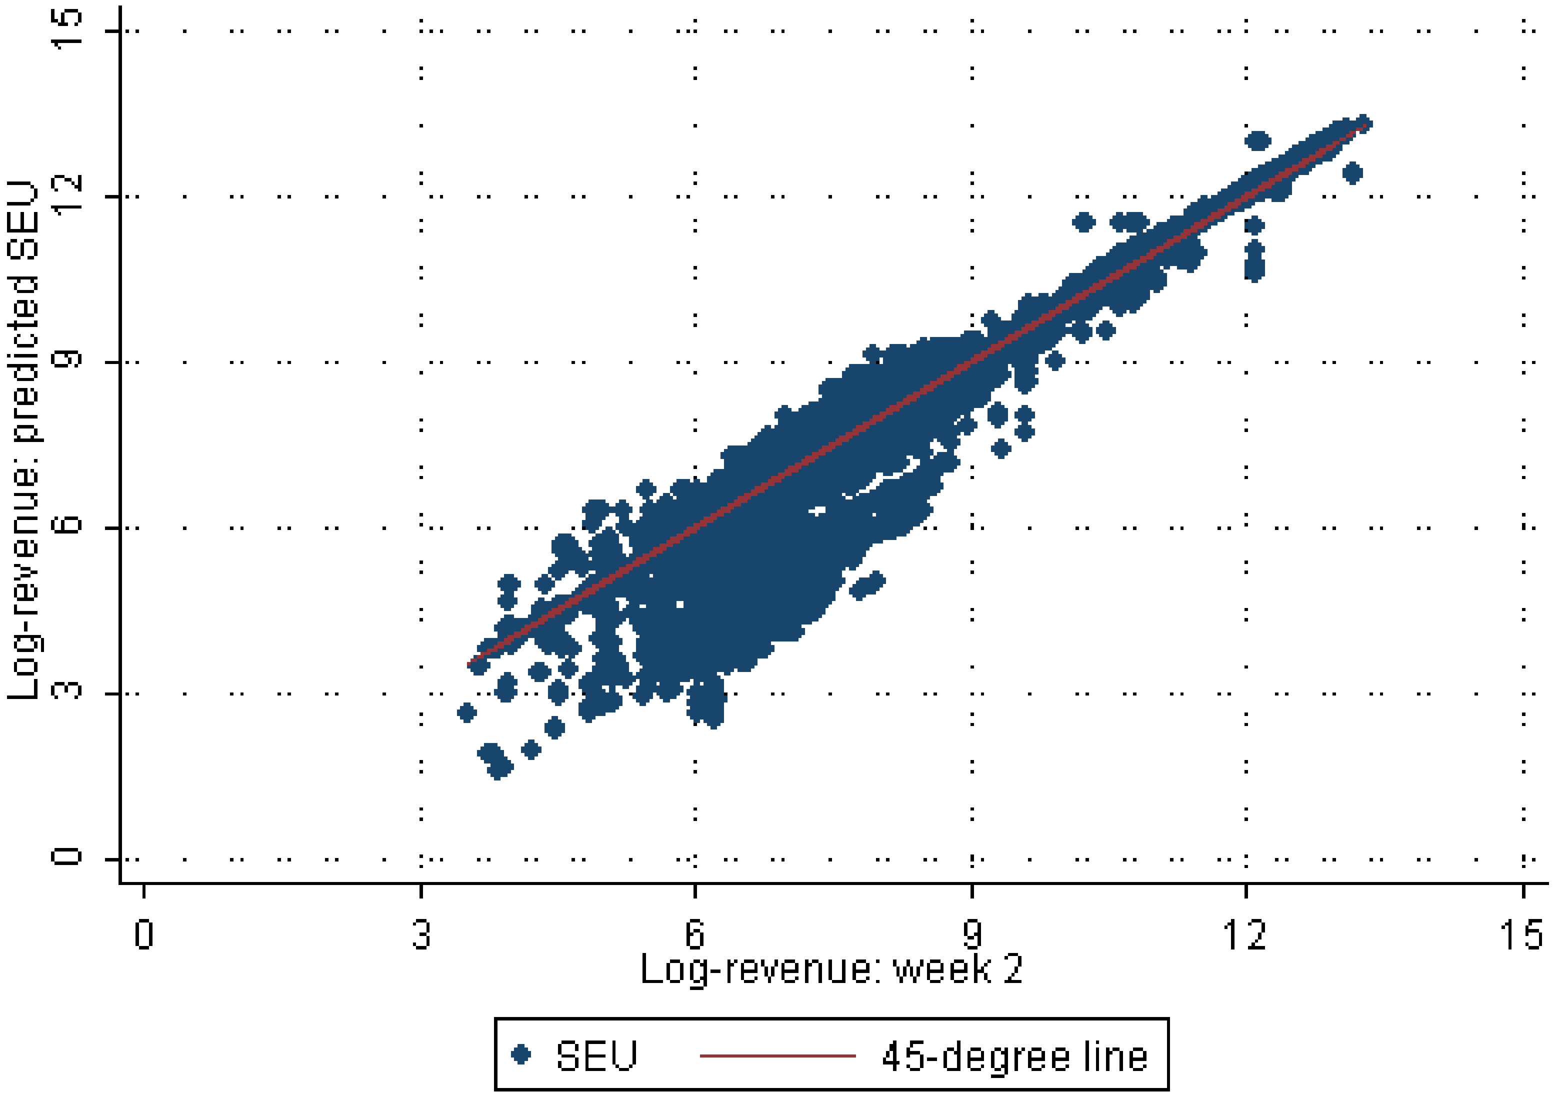

Bid Shading, Profits Per Click Vary by Position (implies inefficient allocation…)

Avg. Ranking	**(VPC-Bid)/ VPC**	**(VPC-Bid)/ CPC**	**(Bid-CPC)/ CPC**	**(VPC-CPC) / CPC**
Search Phrase #1				
[1,1.5)	0.61	3.61	1.18	4.79
[1.5,2.5)	0.22	0.54	0.86	1.39
[2.5,4)	0.21	0.36	0.31	0.68
[4,5.5)	0.20	0.34	0.31	0.65
[5.5,8)	0.20	0.34	0.28	0.62
Search Phrase #2				
[1,1.5)	0.32	0.75	0.44	1.19
[1.5,2.5)	0.36	0.81	0.40	1.21
[2.5,4)	0.43	1.12	0.43	1.55
[4,5.5)	0.35	0.79	0.34	1.13
[5.5,8)	0.28	0.52	0.30	0.82

GSP is inefficient; Revenue Comparison Ambiguous

Model	All	Pos. 1	Pos. 2-5	Pos. 6-8
Search Phrase #1				
Revenue GSP	2.35	1.51	0.80	0.03
Revenue Vickrey	1.94	1.05	0.87	0.03
Welfare GSP	10.80	8.85	1.89	0.05
Welfare Vickrey	10.92	8.92	1.95	0.05
Search Phrase #2				
Revenue GSP	3.18	1.84	1.23	0.10
Revenue Vickrey	3.25	1.90	1.24	0.11
Welfare GSP	6.92	3.58	3.16	0.18
Welfare Vickrey	6.97	3.70	3.10	0.17

“Squashing” Raises Revenue Substantially at Modest Efficiency Cost

Per-Click Bid	Estimated Revenue Bid	Price Per Click	Price Per Click with Squashing
b_1	s_1b_1	$b_2\ (s_2/s_1)$	$b_2\ (s_2/s_1)^a$
b_2	s_2b_2	R/s_2	$R/(s_2)^a$

Outcomes	All	Pos. 1	Pos. 2-5	Pos. 6-8
Advertiser Profit After Squashing	3.3	1.6	1.6	0.1
Advertiser Profit Before Squashing	4.0	2.1	1.7	0.1
Revenue After Squashing	3.6	1.8	1.7	0.1
Revenue Before Squashing	3.3	1.8	1.4	0.1
Welfare After Squashing	6.9	3.4	3.3	0.2
Welfare Before Squashing	7.3	3.9	3.2	0.2

Short-Term Effects Smaller than Long-Term
Limitation to Short-Term Experiments

Search Phrase #1	
Revenue After Squashing	3.29
Revenue After Squashing, Fixed Bids	3.15
Revenue Before Squashing	3.00
Search Phrase #2	
Revenue After Squashing	3.58
Revenue After Squashing, Fixed Bids	3.41
Revenue Before Squashing	3.32

This only incorporates advertiser long-term model.
Athey-Ellison model suggests additional user feedback: users click less, moderating effects.

Improving Click Prediction Accuracy Has Competing Effects: Reduced Welfare, Increased Revenue

Outcomes	All	Pos. 1	Pos. 2-5	Pos. 6-8
Search Phrase #1				
Profit Coarsened Scores	5.27	3.82	1.43	0.01
Profit Original	7.64	5.53	2.09	0.02
Revenue Coarsened Scores	5.05	3.34	1.68	0.04
Revenue Original	3.00	1.82	1.15	0.03
Welfare Coarsened Scores	10.32	7.16	3.11	0.06
Welfare Original	10.64	7.35	3.24	0.05
Search Phrase #2				
Profit Coarsened Scores	1.74	0.86	0.80	0.08
Profit Original	3.98	2.15	1.73	0.11
Revenue Coarsened Scores	4.57	2.05	2.35	0.17
Revenue Original	3.32	1.79	1.44	0.09
Welfare Coarsened Scores	6.31	2.90	3.15	0.26
Welfare Original	7.31	3.94	3.17	0.20

Revenue effects rely heavily on advertiser bidding response. With fixed bids, revenue increases are smaller and sometimes negative.

Conclusions

- Economic models can help guide design
- Platform considers welfare of all participants
  - Platform market considerations about sensitivity of various participants to decisions
  - Competition affects how participants are weighted
- Consumer search costs matter, and this affects decisions about reserve pricing and click weighting
- GSP creates incentives for “demand reduction” that are asymmetric across bidders
  - Inefficiency results, but relatively small