Learning and Trust in Auction Markets (2017)

Jalay, Nekipelov, Tardos

Yahoo!

March 18, 2018

Introduction:

- Study of behavior of bidders in an experimental launch of a new advertising auction platform by Zillow
- Zillow switched from negotiated contracts to auctions in several geographically isolated markets
- Local real estate agents bid on their own behalf, not using third-party intermediaries.
- Zillow also provides a recommendation tool that suggests the bid for each bidder
- **OBJECTIVE:** Paper focuses on the decisions of bidders whether or not to adopt the platform-provided bid recommendation

Introduction:

- Why agents may not be following the platform recommendation?
  - Do they use a different bidding strategy that improves their obtained utility?
  - Lack of trust?
- To answer the above questions, we need to infer the agents value for the impression (no-regret learning in repeated games vs Nash Equilibrium).

- Why is the problem interesting?
  - We are testing a new recommendation tool (Is it good?)
  - Understanding the advertisers value for queries (bid shade)
  - Compare vanilla GSP vs Line-print optimization algorithm (value change)
  - Budget smoothing mechanism
  - Budget and bid recommendations based on impression targets

Introduction:

- Zillow: Largest residential real state search platform in the US
- Platform monetized by showing ads of real estate agents offering services
- Negotiated contracts with real-estate agents for placing ads on the platform
- Experiment: GSP auction where agents pay for impressions
- Experiment: 1st agent is the listing agent of the property + 3 slots allocated via auctions (**Randomized order**)

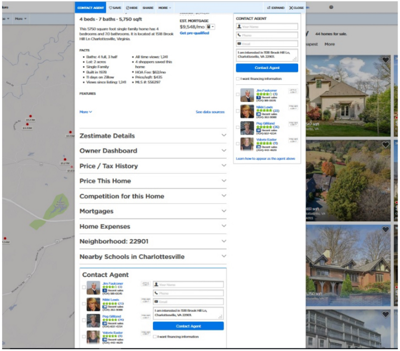

Auction Mechanism:

- GSP
- Agents have small budgets $\rightarrow$ budget-smoothing mechanism to have agents participate in auctions evenly across the time interval
- Sequence:
  1. Select eligible advertisers: advertisers bidding on the ZIP code of the property
  2. The system determines the filtering probabilities for budget smoothing. System estimates the expected spent of the agent given her bid and the filtering probabilities of other agents (fixed point computation)
  3. The remaining bidders are ranked by the order of their bids
  4. Three of the top four remaining bidders are displayed
  5. If the bidder is ranked j is shown, she plays the bid of the bidder ranked $j + 1$ (or reserve price) for the impression
  6. Top 3 bidders are randomly displayed

Market Environment

- 6 isolated markets (ZIP code)

Variable	Selected Regions			
	Mean	STD	25%	75%
Number of agents	19.33	2.29	18.0	20.75
Bids	23.94	14.14	17.3	19.31
Budgets (daily)	8.92	3.0	6.31	11.71
Active duration	85.97	10.38	78.03	91.5
Reserve price	11.65	7.03	7.99	10.74
Bid changes	0.73	0.26	0.54	0.85
Impression Volume	5.52	1.72	4.25	5.89

- Reserve Prices
- Small daily budgets relative to bids

Bid Recommendation Tool

- Bid recommendation based on bidder's monthly budget
- Tool is designed to set the bid that maximizes the expected number of impression that a given bidder gets given her budget
- Tool accounts for filtering probabilities

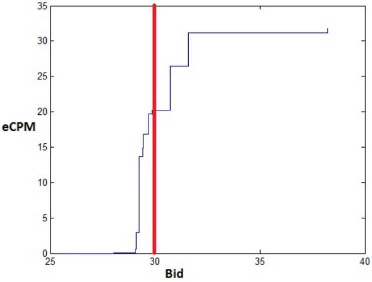

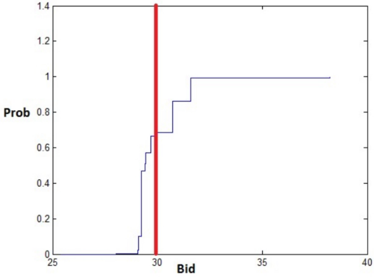

Bid Recommendation Tool

- Optimal Bid: Intersection eCPM and per Impression Budget curves

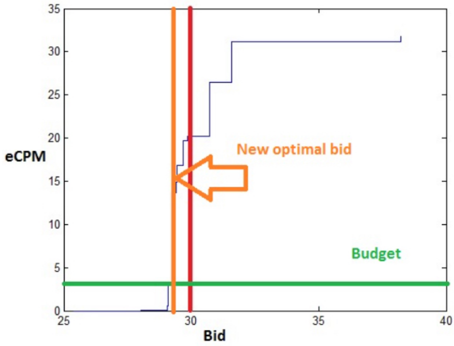

Adoption of bid tool

- Tool: allows bidders to submit their monthly budget for a market and the tool would provide the bid that maximizes the number of impressions that could be purchased within the month with the given budget.
- Tool utilization

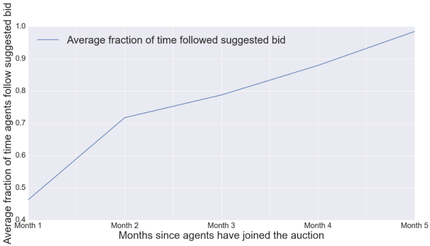

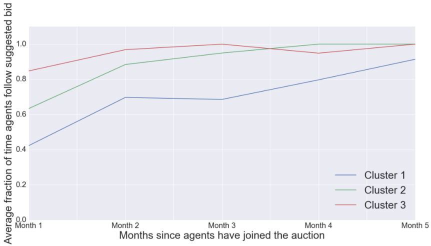

- Observation 1: at the beginning, bidders were willing to experiment with their bids by deviating from the recommendation
- Observation 2: bidders do not have full faith that the recommendations benefit them

Empirical analysis: trust of bid tool

- Why agents may not be following the platform recommendation? (own bidding strategy leading to improved outcomes vs lack of platform's trust)
- Infer impression **advertisers' value** using the concept of no-regret learning (Nekipelov, Syrgkanis, and Tardos (2015))
  - Experimenting with bids as off-equilibrium (market exploration)
  - Weaker assumption than Nash equilibrium
- No-regret learning: modeling players who have not reached a stable equilibrium, but rather use algorithmic learning, aiming to learn the best way to play from previous observations
- We measure **regret** as the difference between the time-averaged utility attained by bidder's bid sequence and the average utility attained by the best fixed bid in hindsight

No-regret learning

- We can express the expected utility of bidder i at t as

$$u_{it}(b_{it}, v_i) = v_i eQ_{it}(b_{it}) - eCPM_{it}(b_{it})$$

where at each time t bidder i with bid b_{it}, outstanding bids of other bidders $\vec{b}_{-i,t}$, and participation probability θ^t faces
  - probability allocation $eQ(b_{it}, \vec{b}_{-i,t} : \theta^t)$
  - expected cost $eCPM(b_{it}, \vec{b}_{-i,t} : \theta^t)$
- **Definition 1 (ϵ - Average Regret):** A sequence of play that we observe has ϵ_i average regret for bidder i if

$$\forall b' \in \mathcal{B} : \frac{1}{T}\sum_{t=1}^{T} u_{it}(b_{it}, v_i) \geq \frac{1}{T}\sum_{t=1}^{T} u_{it}(b', v_i) - \epsilon_i$$

where b' corresponds to a fixed bid from the bid sequence $\{b_{it}\}_{t=1}^{T}$
Nash Equilibrium Condition $\rightarrow \epsilon_i = 0$
- **Definition 2 (Rationalizable Set):** A pair (ϵ_i, v_i) of value v_i an error ϵ_i is a rationalizable pair for player i if it satisfies the above equation. We refer to the set of such pairs as the rationalizable set and denote it with $\mathcal{N}R$

Empirical analysis: trust of bid tool

- The $\epsilon-$regret condition can be rewritten as

$$\forall b' \in \mathcal{R}_+ : v_i \times \Delta eQ_i(b') \leq \Delta eCPM_i(b') + \epsilon_i$$

where $\Delta eCPM_i(b') = \frac{1}{T} \sum_{t=1}^{T} (eCPM_{it}(b') - eCPM_{it}(b_{it}))$, and
$\Delta eQ_i(b') = \frac{1}{T} \sum_{t=1}^{T} (eQ_{it}(b') - eQ_{it}(b_{it}))$

- Hence, the rationalizable set $\mathcal{N}R$ is an envelope of the family of half-spaces generated by inequalities and obtained by varying $b' \in \mathcal{R}_+$

- Rationalizable sets for 3 of the bidders most frequently changing their bids

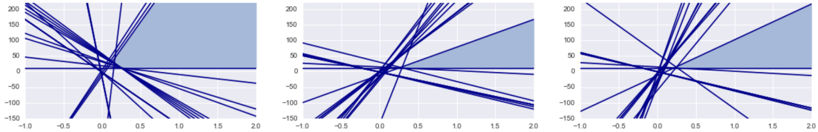

where vertical axis is the per **impression value** of the bidder and the horizontal axis is the **additive average regret** (in monetary units)

Empirical analysis: trust of bid tool

- Why agents may not be following the platform recommendation?

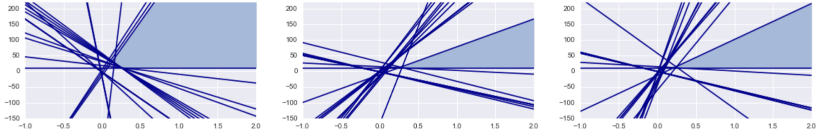

- Assume that the **impression value** for an advertiser corresponds to the pairs of value and regret in $\mathcal{N}R$ where the observed bid sequence has the smallest possible average regret.
- Compare bidder's utility: Recommended bid vs actual bid
- Agents regret by not using the platform bid recommendation.

Region	reg 1	reg 2	reg 3	reg 4	reg 5	reg 6	cl 1	cl 2	cl 3	all
worse	21.4%	20%	8.3%	0%	21.4%	50%	23.3%	25.9%	6.7%	20.8%
better	42.9%	20%	41.7%	28.6%	21.4%	20%	46.7%	22.2%	6.7%	29.2%
equal	35.7 %	60%	50%	71%	57.1%	30%	30%	51.8%	86.7%	50%

Empirical analysis: trust of bid tool

- Figure: difference between the regret of own bidding strategy and recommended bid across agents and the % of time agents use the recommended bid.

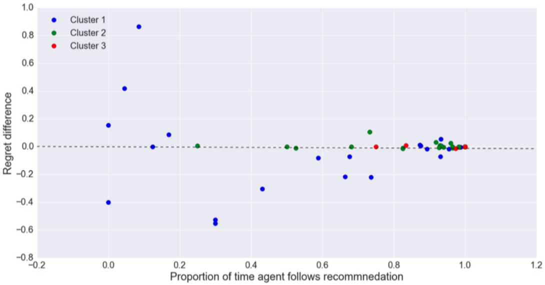

- On average, agents who use the recommended bid less do not show any improvements over the recommended bid measured by the average regret
- **Answer:** Results are explained by trust of agents in the platfom-provided bid recommendations