Budget Management In GSP (2018)

Yahoo!

March 18, 2018

Today's Presentation:

- "Budget Management Strategies in Repeated auctions", Balseiro, Kim, and Mahdian, WWW2017
- "Learning and Trust in Auction Markets", Jalaly, Nekipelov, and Tardos, (2017)

"Budget Management Strategies in Repeated auctions"
Balseiro, Kim, and Mahdian, WWW2017

Budget Management

- Advertisers declare the maximum daily amount they are willing to pay, and the platform adjust allocations and payments
- Enforce budget constraints:
    - Stop serving an advertiser as soon as his budget is exhausted
    - Throttle advertiser throughout the day at a rate that ensures the advertiser exhausts his budget close to the end of the day
    - Achieve the same rate of spend by shading the bids of the advertiser

Budget Management

- Advertisers declare the maximum daily amount they are willing to pay, and the platform adjust allocations and payments
- Enforce budget constraints:
  - Stop serving an advertiser as soon as his budget is exhausted
  - Throttle advertiser throughout the day at a rate that ensures the advertiser exhausts his budget close to the end of the day
  - Achieve the same rate of spend by shading the bids of the advertiser
- **Objective:** compare the system of equilibria of different budget management strategies in terms of the seller's profit and buyer's utility
- Budget management strategies implemented by the platform:
  - Probabilistic throttling
  - Thresholding
  - Bid shading
  - Reserve Pricing
  - Multiplicative Boosting

Budget Management

- Mechanisms are derived with simple modifications from the second price auction with the reserve price equal to the opportunity cost
- Budget management strategies implemented by the platform:
  - Throttling: controls expenditure by excluding a buyer independently and at random with fixed probability, and then running a second-price auction with reserve c. (missing good opportunities from not participating when buyer's value is high)
  - Thresholding: buyer participate in an auction when his bid is above a fixed threshold. Buyer does not miss the items he values the most. Payment: max of the second highest bid and c
  - Reserve Pricing: each buyer has a reserve price r. Payment: max of the second highest bid and the reserve price r_i (higher payments than thresholding).
  - Bid shading: buyers participate in all auctions with bids shaded by a constant multiplicative factor.
  - Multiplicative Boosting: similar to bid shading but only the allocation rule is modified, not payments (higher payments than bid shading).

Budget Management

Mechanism	Allocation/Payment Rules
Bid Shading (S) $(\mu_i \geq 0, \forall i)$	Each buyer i has a parameter $\mu_i \geq 0$ Buyer i wins if $\frac{b_i}{1+\mu_i} \geq max_{j \neq i} \frac{x_j}{1+\mu_j} \vee c$ Buyer i pays $max_{j \neq i} \frac{b_j}{1+\mu_j} \vee c$
Multiplicative Boosting (MB) $(\delta_i \geq 1, \forall i)$	Each buyer i has a parameter $\delta_i \geq 1$ Buyer i wins if $b_i \geq \delta_i(max_{j \neq i} b_j \vee c)$ Buyer i pays $max_{j \neq i} b_j \vee c$
Reserve Pricing (R) $(r_i \geq c, \forall i)$	Each buyer i has a parameter $r_i \geq c$ Buyer i wins if $b_i \geq max_{j \neq i, b_j \geq r_j} b_j \vee r_i$ Buyer i pays $max_{j \neq i, b_j \geq r_j} b_j \vee r_i$
Thresholding (T) $(\tau_i \geq c, \forall i)$	Each buyer i has a parameter $\tau_i \geq c$ Buyer i wins if $b_i \geq max_{j \neq i, b_j \geq \tau_j} b_j \vee \tau_i$ Buyer i pays $max_{j \neq i, b_j \geq \tau_j} b_j \vee c$
Throttling (TO) $(\theta_i \in [0,1], \forall i)$	Buyer i has a parameter θ_i and $I_i = 1$ with proba$1 - \theta_i$ Buyer i wins if $I_i = 1$ and $b_i \geq max_{j \neq i, I_j = 1} b_j \vee c$ Buyer i pays $max_{j \neq i, I_j = 1} b_j \vee c$

Budget Management

- Assumptions:
  - Advertisers bid truthfully and are only interested in their total expenditures fully meeting their budget constraints $\rightarrow$ Seller maximizes expected payment
  - All n buyers have the same distribution of values $F(\cdot)$ and budget B
  - Restrict attention to symmetric equilibria in which the platforms uses the same parameter for all buyers
  - Let G the expected expenditure of one buyer when the same parameter is used for all buyers
  - Let U be the expected utility of all buyers
  - Let I the expected number of items sold over the horizon
- Seller's profit: $P = nG - cI$
- Buyers' utility: $U = nV - nG$
- Total welfare: $W = nV - cI$

Budget Management

- Use simulated and real data to validate their results

- Theorem 4: The following dominance and non-dominance relations hold for the seller's profit:
    1. Reserve Pricing $\geq$ Thresholding $\geq$ Throttling $\geq_{IF}$ Bid Shading
    2. Reserve Pricing $\geq_{U}$ Mult. Boosting $\geq_{U}$ Bid Shading
    3. Mult. Boosting $||$ Thresholding
    4. Mult. Boosting $||$ Throttling

Budget Management

- Use simulated and real data to validate their results
- Theorem 4: The following dominance and non-dominance relations hold for the seller's profit:
  1. Reserve Pricing $\geq$ Thresholding $\geq$ Throttling $\geq_{IF}$ Bid Shading
  2. Reserve Pricing $\geq_U$ Mult. Boosting $\geq_U$ Bid Shading
  3. Mult. Boosting $||$ Thresholding
  4. Mult. Boosting $||$ Throttling
- Theorem 5: The following dominance and non-dominance relations hold for buyer's utility:
  1. Bid Shading $\geq$ Thresholding $\geq$ Reserve Pricing
  2. Bid Shading $\geq_U$ Mult. Boosting $\geq_U$ Reserve Pricing
  3. Thresholding $||$ Throttling $||$ Mult. Boosting (three-way comparison)

Budget Management: Seller's Objective

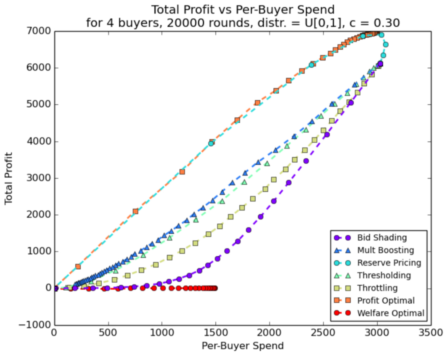

Budget Management: Buyer's Objective

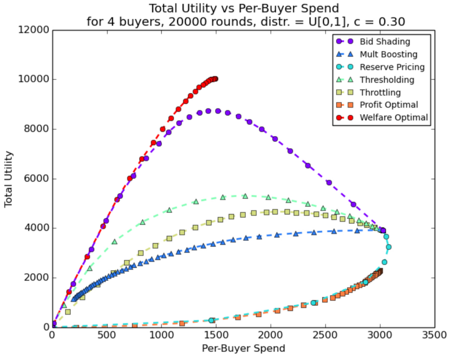

"Learning and Trust in Auction Markets"

Jalaly, Nekipelov, and Tardos, (2017)

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

- **Today's Objective:**

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?)
    - 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**)

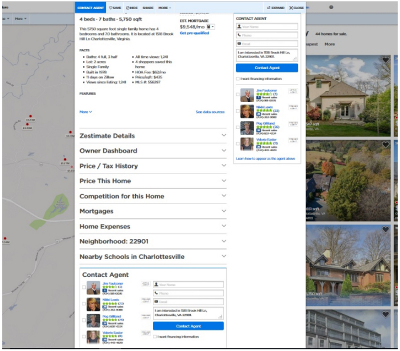

Auction Mechanism:

- GSP
- Agents have small budgets → 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

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

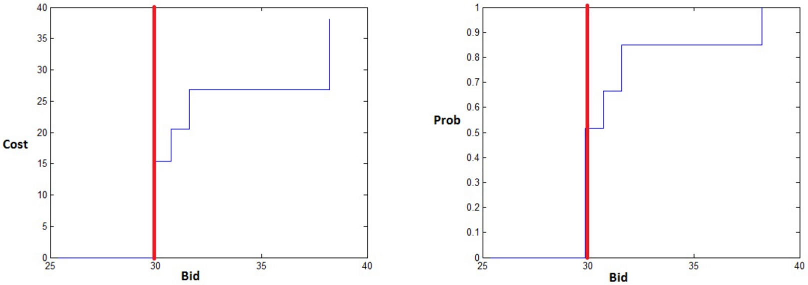

Bid Recommendation Tool

- Optimal Bid: Intersection eCPM and per Impression Budget curves

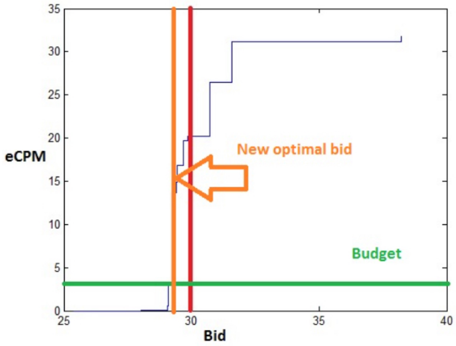

Model

- Expected Spent of bidder i (conditional on i not being filtered)

$$eCPM_i(b_i;\pi) = \sum_{N \in \mathcal{N}_i} \gamma_i^N \prod_{j \neq i} \pi_j^{n_j^N} (1-\pi_j)^{(1-n_j^N)} PRICE_i^N$$

where

- π_j probability of j being filtered
- $\mathcal{N}_i$ set of binary digits that indicates if bidder j is filtered or not conditional on i not being filtered
- N vector of filtered not filtered ads (specific row of $\mathcal{N}_i$)
- $n_i^N = 1$ and $n_j^N \in \{0,1\}$ $\forall j \neq i$ indicator functions telling if the ad has been filtered
- $PRICE_i^N$: GSP price given configuration N
- γ_i^N position weights. Probability that an ad ranked in one particular position is shown (only top four positions have $\gamma^N \neq 0$)

Model

- Expected Spent of bidder i (conditional on i not being filtered)

$$eCPM_i(b_i;\pi) = \sum_{N\in\mathcal{N}_i} \gamma_i^N \prod_{j\neq i} \pi_j^{n_j^N} (1-\pi_j)^{(1-n_j^N)} PRICE_i^N$$

where

  - π_j probability of j being filtered
  - $\mathcal{N}_i$ set of binary digits that indicates if bidder j is filtered or not conditional on i not being filtered
  - N vector of filtered not filtered ads (specific row of $\mathcal{N}_i$)
  - $n_i^N = 1$ and $n_j^N \in \{0,1\}\ \forall j \neq i$ indicator functions telling if the ad has been filtered
  - $PRICE_i^N$: GSP price given configuration N
  - γ_i^N position weights. Probability that an ad ranked in one particular position is shown (only top four positions have $\gamma^N \neq 0$)

- Expected impression share (probability of showing an i impression conditional on i not being filtered)

$$eQ_i(b_i;\pi) = \sum_{N\in\mathcal{N}_i} \gamma_i^N \prod_{j\neq i} \pi_j^{n_j^N} (1-\pi_j)^{(1-n_j^N)}$$

- Note that $eCPM_i(b_i;\pi)$ and $eQ_i(b_i;\pi)$ do not depend on π_i

Model

- Overall spent per impression of bidder i

$$\pi_i \times eCPM_i(b_i)$$

- Probability of i showing an impression

$$\pi_i \times eQ_i(b_i)$$

- Budget per impression $Budget_i$

$$Budget_i = \frac{Monthly\ \ Budget\ \ Bidder\ \ i}{Expected\ \ Inventory}$$

Budget-Smoothing Probabilities

- **Objective:** Given bids and budgets, what are the Budget-Smoothing Probabilities (π)?

Budget-Smoothing Probabilities

- **Objective:** Given bids and budgets, what are the Budget-Smoothing Probabilities (π)?

1. Sort bidders i by their bid b_i and assume bidders are numbered in this order
2. Construct array of 2^I binary I digit numbers
3. Take a subset of elements of $\mathcal{N}$ where $i-th$ digit equals 1 ($\mathcal{N}_i$)

Ad1	Ad2	$\cdots$	Ad i	$\cdots$	Ad I	
1	1	$\cdots$	1	$\cdots$	1	
0	1	$\cdots$	1	$\cdots$	1	$\leftarrow$ **N**
$\cdots$	$\cdots$	$\cdots$	1	$\cdots$	$\cdots$	
1	0	$\cdots$	1	$\cdots$	0	
0	1	$\cdots$	1	$\cdots$	0	

4. Compute $PRICE_i^N$ and $eCPM_i$

$$eCPM_i(b_i;\pi) = \sum_{N \in \mathcal{N}_i} \gamma_i^N \prod_{j \neq i} \pi_j^{n_j^N} (1-\pi_j)^{(1-n_j^N)} PRICE_i^N$$

5. Solve for $\pi_1, ..., \pi_I$ by solving a system of nonlinear equations

$$\pi_i = min\{1, \frac{Budget_i}{eCPM_i(b_i;\pi_1,...,\pi_I)}\}, i = 1, ..., I$$

Budget-Smoothing Probabilities

- Solve for $\pi_1, ..., \pi_I$ by solving a system of nonlinear equations

$$\pi_i = min\{1, \frac{Budget_i}{eCPM_i(b_i; \pi_1, ..., \pi_I)}\}, i = 1, ..., I$$

- For instance, we can find an approximate solution by minimizing the sum of squares using gradient descent or Newton's method

$$\sum_{i=1}^{I}(\pi_i - min\{1, \frac{Budget_i}{eCPM_i(b_i; \pi_1, ..., \pi_I)}\})^2$$

  with respect to $\pi_1, ..., \pi_I$
- Iterative algorithm for finding a fixed point. Stopping criteria $|\pi_s^{(k)} - \pi_s^{(k-1)}| \leq \epsilon$

Bid Recommendation Tool

- **Objective:** Tool designed to set the bid that maximizes the expected number of impressions that a given bidder gets given her budget

Bid Recommendation Tool

- **Objective:** Tool designed to set the bid that maximizes the expected number of impressions that a given bidder gets given her budget
- $eCPM_i(b_i)$ and $eQ_i(b_i)$ are monotone functions of the bid (b_i)
- If a bidder maximizes the probability of appearing in an impression as a function of the bid, the optimal bid (b_i^*) will be set such that:

$$eCPM_i(b_i^*) = Budget_i$$

- Intuition:

Bid Recommendation Tool

- **Objective:** Tool designed to set the bid that maximizes the expected number of impressions that a given bidder gets given her budget
- $eCPM_i(b_i)$ and $eQ_i(b_i)$ are monotone functions of the bid (b_i)
- If a bidder maximizes the probability of appearing in an impression as a function of the bid, the optimal bid (b_i^*) will be set such that:

$$eCPM_i(b_i^*) = Budget_i$$

- Intuition:

$$\underbrace{\pi_i \times eCPM_i(b_i)}_{spent(b_i)} \leq Budget_i \quad \pi_i = min\{1, \frac{Budget_i}{eCPM_i(b_i)}\} \quad Prob_i(b_i, Budget_i) = \pi_i \times eQ_i($$

1. When budget smoothing is not initiated $\pi_i = 1$ and probability of impression equals $eQ_i(b_i)$
2. Whenever budget smoothing is initiated $\pi_i < 1$, spent equals budget and

$$Prob_i(b_i, Budget_i) = \pi_i \times eQ_i(b_i) = \frac{Budget_i}{eCPM_i(b_i)} \times eQ_i(b_i)$$

   This function decreases as a function of bid, implying that the probability of getting an impression increases up to b_i^* and then decreases whenever the budget smoothing is initated

Bid Recommendation Tool

- Optimal Bid: Intersection eCPM and per Impression Budget curves

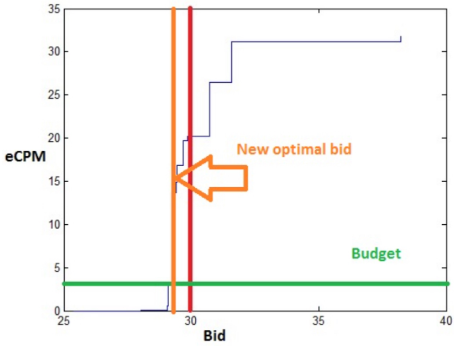

Bid Recommendation Tool

- Optimal Bid: spent equals per impression budget

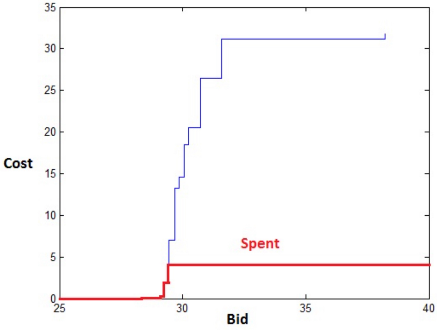

Budget and bid recommendations

- **Objective:** Tool designed to make recommendations for the monthly budget and the corresponding bid that meet a given impression target

Budget and bid recommendations

- **Objective:** Tool designed to make recommendations for the monthly budget and the corresponding bid that meet a given impression target
- Expected spent in a given impression

$$spent_i(b_i) = \pi_i \times eCPM_i(b_i) \leq Budget_i$$

- Expected probability of appearing in the impression

$$Prob_i(b_i, Budget_i) = \begin{cases} 1 \times eQ_i(b_i), & \text{if } eCPM_i(b_i) \leq Budget_i. \\ \frac{Budget_i}{eCPM_i(b_i)} \times eQ_i(b_i), & \text{if } eCPM_i(b_i) > Budget_i. \end{cases} \quad (1)$$

- Let *Inventory* be the total number of available impressions and $Goal_i$ the impression target for bidder i.

Budget and bid recommendations

- **Objective:** Tool designed to make recommendations for the monthly budget and the corresponding bid that meet a given impression target
- Expected spent in a given impression

$$spent_i(b_i) = \pi_i \times eCPM_i(b_i) \leq Budget_i$$

- Expected probability of appearing in the impression

$$Prob_i(b_i, Budget_i) = \begin{cases} 1 \times eQ_i(b_i), & \text{if } eCPM_i(b_i) \leq Budget_i. \\ \frac{Budget_i}{eCPM_i(b_i)} \times eQ_i(b_i), & \text{if } eCPM_i(b_i) > Budget_i. \end{cases} \tag{1}$$

- Let *Inventory* be the total number of available impressions and $Goal_i$ the impression target for bidder i.
- Then, the optimum bid for a given budget is

$$Prob_i(b_i, Budget_i) \leq \frac{Goal_i}{Inventory}$$

- The minimum budget per impression for which the impression goal is met

$$Budget_i = eCPM_i(b_i)$$

- Leading to condition

$$eQ_i(b_i) = \frac{Goal_i}{Inventory}$$

Budget and bid recommendations

- System of equations

$$\pi_j = min\left\{1, \frac{Budget_j}{eCPM_j(b_j)}\right\}, \quad j \neq i$$

$$\pi_i = 1,$$

$$eQ_i(b_i^*) = \frac{Goal_i}{Inventory}$$

with unknowns π_j and b_i^*

- The recommended bid is the solution b_i^* and the budget recommendation equals

$$Budget_i^* = eCPM_i(b_i^*)$$

- The algorithm also applies for multiple bidders and deal with corner cases