Redacted

REDACTED FOR PUBLIC FILING & ABRIDGED

Ex. No.
UPX0889
1:20-cv-03010-APM

Outline

- Part 1: Background
  - Brief recap of our auction
  - What is a “pricing” launch?
  - How do we evaluate a pricing launch?
- Part 2: Recent Innovations
  - Deep dive into Robust Fine Grained Squashing (Potiron)
  - A peek into mechanisms under development
- Questions

Mention we will not cover format pricing and format pricing in this talk.

Part 1: Background

Auction recap

We run a generalized second price auction (GSP)

- We auction off each position separately
- Each candidate gets a score, called ***Long Term Value*** or ***LTV***
- For each position, we rank ads according to LTV (LTV must be > 0)
- We then pick the top to show in that position
- We move on to the next position and repeat, until we either run out of space or candidates

Long Term Value: The basic “vanilla” flavor

$$\mathrm{LTV} = \underbrace{bid \cdot pctr}_{\text{ECPM}} - \underbrace{\beta}_{\text{CPM Cost}}$$

ECPM
(Expected Cost Per Mille)

CPM Cost
(“Blindness” cost)

Redacted

Second Pricing

Winner pays minimum price needed to beat runner up

1. Equate LTV of winner to LTV of runner up
2. Solve for bid. Result is the cost-per-click (CPC)

In the equations:

- w: winner
- ru: Runner Up

$$\mathrm{LTV}_w = \mathrm{LTV}_{ru}$$

$$bid \cdot pctr_w - \beta_w = bid_{ru} \cdot pctr_{ru} - \beta_{ru}$$

$$bid = \frac{bid_{ru} \cdot pctr_{ru} - \beta_{ru} + \beta_w}{pctr_w}$$

This becomes the CPC

Reserve Pricing

Where there is no competition (i.e. no runner up) winner pays the "blindness reserve"

1. Equate LTV of winner to 0
2. Solve for bid. Result is the reserve cost-per-click (CPC)

$$\mathrm{LTV}_w = 0$$

$$bid \cdot pctr_w - \beta_w = 0$$

In the equations:

- w: winner
- ru: Runner Up

This becomes the CPC → $$bid = \frac{\beta_w}{pctr_w}$$

Efficient Auctions

The auction has three functions:

The outcome of these two steps is referred to as an ***allocation***

1. Select the ads to show
2. Rank the ads that are selected

This is ***pricing***

3. Price them

An ***allocation*** is ***efficient*** if it delivers clicks to advertisers who value them the most

Redacted

Allocation, one more time

PLA Module??

Ad3

Ad1

Ad2

Efficiency Loss

Today in our auction, ***selection***, ***ranking*** and ***pricing*** are all tied together

➔ Changing the mechanism to affect one, has ramifications on the others

In particular, ***pricing*** changes often lead to ***worse allocations***

➔ This is referred to as ***efficiency loss***

Auction Pricing Mechanisms

What are auction pricing mechanisms designed to do?

- Specifically designed to extract ***value*** from advertisers
  - Advertisers derive value from the clicks we deliver to them
  - Are we pricing them adequately for the value they receive?
- Designed to minimize ***efficiency loss***
  - A simple proxy to efficiency loss is click loss
  - Since we're paid for clicks, we don't want to disproportionately lose them in the process
- Designed to minimize adverse response from advertisers
  - Is pricing too aggressive compared to value?
  - Are certain advertisers at risk of withdrawing from the auction?

Auction Pricing Mechanisms

What are they **not** designed to do?

- Not designed to increase clicks
  - We're actually happy when we ***minimize*** the click loss
- Not designed to focus on the user
  - We resort to ***allocation*** mechanisms for this, e.g. Kumamon (go/kumamon-design)

Source of Efficiency Loss

Pricing mechanisms often have side effects which lead to efficiency loss. Some example are:

- Click Loss. Can happen one of two ways:
  - **Reranking**: Higher pCTR ads are forced down the rank
  - **Impression Loss**: You can't get clicks on ads that don't show!
- Adverse Advertiser Response
  - Advertiser lower bids, change targeting, or downright leave the auction

Redacte

But we have an auction designed to set prices… why do we need more?

When Vanilla pricing may no be sufficient

Second Pricing works great most of the time, but there are failure scenarios

- **Weak or lack of Auction Pressure**
  - When no competition is present, or when competition is of inferior quality
- **Reserve pricing**
  - Reserve prices are generally lower than their second price counterparts

We need a way to extract value more directly

➔ We need ***pricing mechanisms*** with ***pricing knobs***

Canonical example: Squashing

How it works?

- Compress dynamic range of pCTRs in a given auction
- Achieved by moving all candidate pCTRs in the direction of the max pCTR in that auction
- Effectively simulates auction pressure

$$\mathrm{LTV}_w = \mathrm{LTV}_{ru}$$

$$\sigma_w \cdot \mathrm{LTV}_w = \sigma_{ru} \cdot \mathrm{LTV}_{ru}$$

Squashing multipliers

$$\sigma_c = \frac{\lambda \cdot pctr_{max} + (1 - \lambda) \cdot pctr_c}{pctr_c}$$

Pricing Example under squashing

Scenario 1: Winner has the maximum pCTR in the auction

Strictly > 1.0 ⇒ LTV of runner up "improves" ⇒ Auction pressure increase

$$\mathrm{LTV}_w = \frac{\sigma_{ru}}{\sigma_w} \cdot \mathrm{LTV}_{ru}$$

Reduces to 1.0 since winner has max pCTR

$$\sigma_c = \frac{\lambda \cdot pctr_{max} + (1-\lambda) \cdot pctr_c}{pctr_c}$$

Pricing Example under squashing

Scenario 2: Runner up has the max pCTR in the auction

Reduces to 1.0 since runner up has max pCTR

$$\mathrm{LTV}_w = \frac{\sigma_{ru}}{\sigma_w} \cdot \mathrm{LTV}_{ru}$$

Strictly > 1.0 ⇒ LTV of runner up "worsens" ⇒ Auction pressure decreases

$$\sigma_c = \frac{\lambda \cdot pctr_{max} + (1-\lambda) \cdot pctr_c}{pctr_c}$$

Logs Analysis

- **Tail Impact**: What fraction of cost (or gains) come from adgroups whose CPCs are larger than a certain fraction
  - Weighted by **cost**: Provides a measure of ***risk***
  - Weighted by **gains**: Provides a measure of ***stickiness***
- **Top Division Impact**: For our top divisions, what does the CPC impact look like?
- **MH-CPC**: controls for advertiser mix

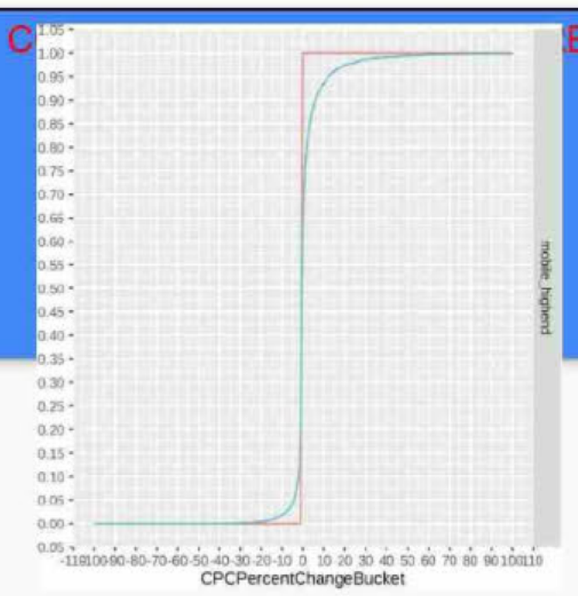

Division	Desktop FirstPredCPC	Desktop SecondPredCPC	Mobile FirstPredCPC	Mobile SecondPredCPC
Booking.com	11.15	13.10	9.37	9.17
Amazon.com 01 US	-0.01	0.59	+0.40	-0.35
Trivago	2.44	3.09	1.41	1.98
Expedia - AMER 100	0.75	0.90	1.09	2.30
Airbnb	0.67	2.07	3.02	2.54
Gucci	33.27	-13.13	-3.28	-8.71
rentalcars.com	11.73	13.15	9.62	6.66
Walmart Labs	1.36	2.96	-0.78	-0.32
Hotels.com - AMER 65/EMEA 45	2.23	2.56	6.20	3.39
TripAdvisor.com	-2.09	-0.39	-0.31	0.81
Wayfair.com	-0.24	-0.00	0.52	1.66
Amazon.com 01 DE	0.29	1.01	-0.69	0.29
Indeed	9.71	4.80	-0.83	1.59

Comparison of Tail Impacts of various mechanisms

Risk: % of Ad Groups with CPC change > Threshold (spend weighted)

Threshold	Potiron	Kabocha (full)	Momiji	GammaYellow (full)	Sapporo (full)
12.5%	4%	3.5%	13.4%	4%	1.54%
25.0%	2%	1.2%	0.57%	2%	0.92%

Stickiness: % of Ad Groups with CPC change > Threshold (gains weighted)

Threshold	Potiron	Kabocha (full)	Momiji	GammaYellow (full)	Sapporo (full)
12.5%	31%	20%	21%	45%	90%
25.0%	18%	5%	1%	31%	85%

Understanding response: **Advertiser Experiments**

1. Partition query space so as to maximize advertiser's interactions (i.e. discover micro markets)
2. Randomly partition the space into treatment and control
3. Apply treatment for several weeks
4. Run inference models to predict, as a function of dose, what the response under a launch would look like

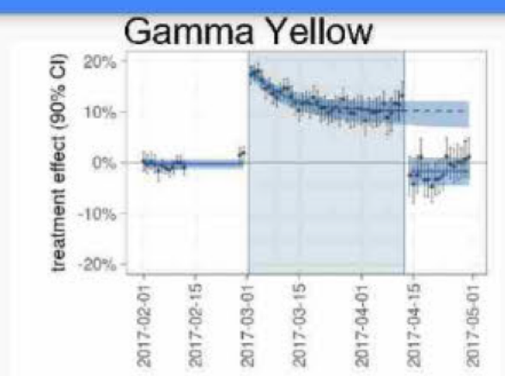

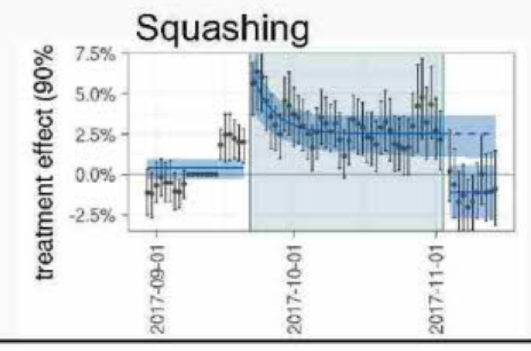

Handling contributions over time: **Holistic Pricing effort**

- Teams across AQ create value and move prices around
- Lower risk when we move prices along ***with*** value
- Developing tools to track, monitor the state of our system over time
  - Excess CPC rule of thumb
- Tune prices, safely, to
  - Ensure good value sharing between advertisers / google
  - Stay in touch with the additional value created over time
  - Limit risk

Redacted

Part 2: Recent Innovations

CONFIDENTIAL - DO NOT SHARE

Fine Grained Squashing

Project Potiron

Redacted

REDACTED FOR PUBLIC FILING & ABRIDGED

Motivation

Now that we know efficiency loss is a natural outcome of any pricing mechanism, how do we go about minimizing it?

- Can we identify, in a ***robust*** way, pockets of auctions that are more (or less) susceptible to efficiency loss?
  - Turn down the pricing knob for auctions that are ***more*** susceptible
  - Turn up the pricing knob for auctions that are ***less*** susceptible

Combatting Efficiency Loss: QSpace Squashing (aka Fine Grained)

QSpace is an AQ-wide service that clusters queries into ~23M clusters

- Could we fine tune squashing at that level?
  - Large fraction of spend lies in clusters that have a large fraction of auctions pCTR ranked
  - Opportunity for fine tuning it at that level

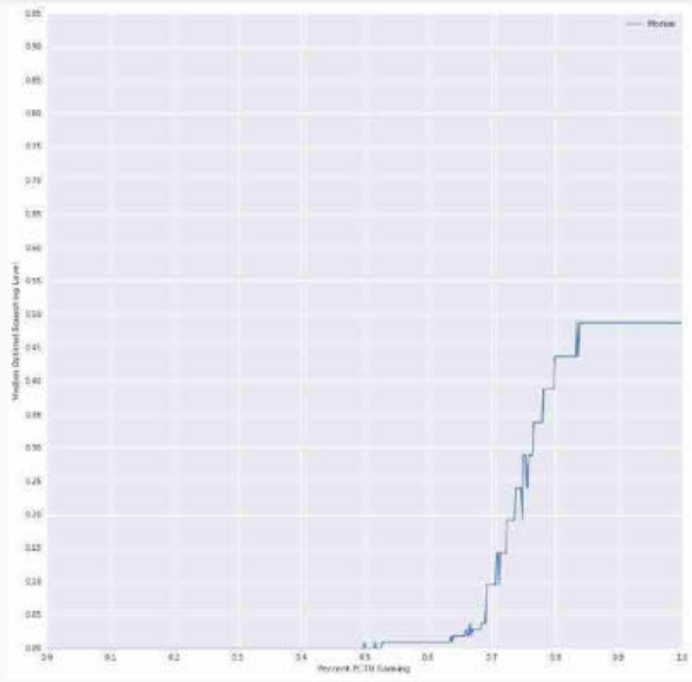

Redact

Re

A teaser of what's to come...

The team is actively developing several pricing mechanisms

- **Stateful Pricing**:
  - Borrow headroom from one auction to use in other auctions
  - Preliminary live experiments show a 8:1 Revenue-Efficiency Tradeoffs
  - Many infrastructure considerations
- **Probabilistic Click Pricing**
  - Randomly drop ads to achieve desired click-cost curve properties
  - A possible replacement for format pricing
- **Fractional Formats**
  - Probabilistically show incremental formats that are unsold due to low bids

Questions?