An Empirical Study of Reserve Price Optimisation in Display Advertising

Shuai Yuan, Bowei Chen, Jun Wang

Department of Computer Science

University College London

Peter Mason

Advance International Media

Sam Seljan

AppNexus

Reserve Price Optimisation

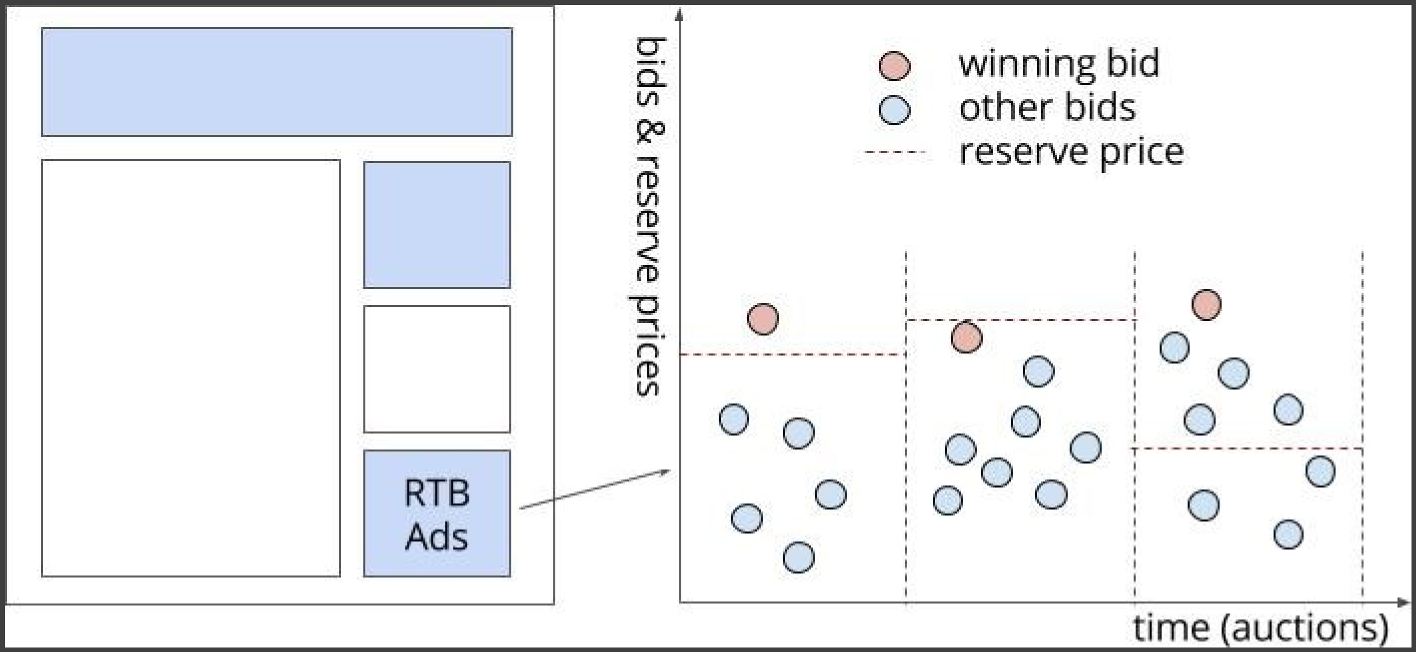

The task:

- To find the optimal reserve prices (hard floors)

Why

- Suppose it is second price auction
  - Normal case: $b_2 \geq \alpha$
  - Preferable case: $b_1 \geq \alpha > \ b_2$ (it increases the revenue)
  - Undesirable case: $\alpha > b_1$ (but there is risk)

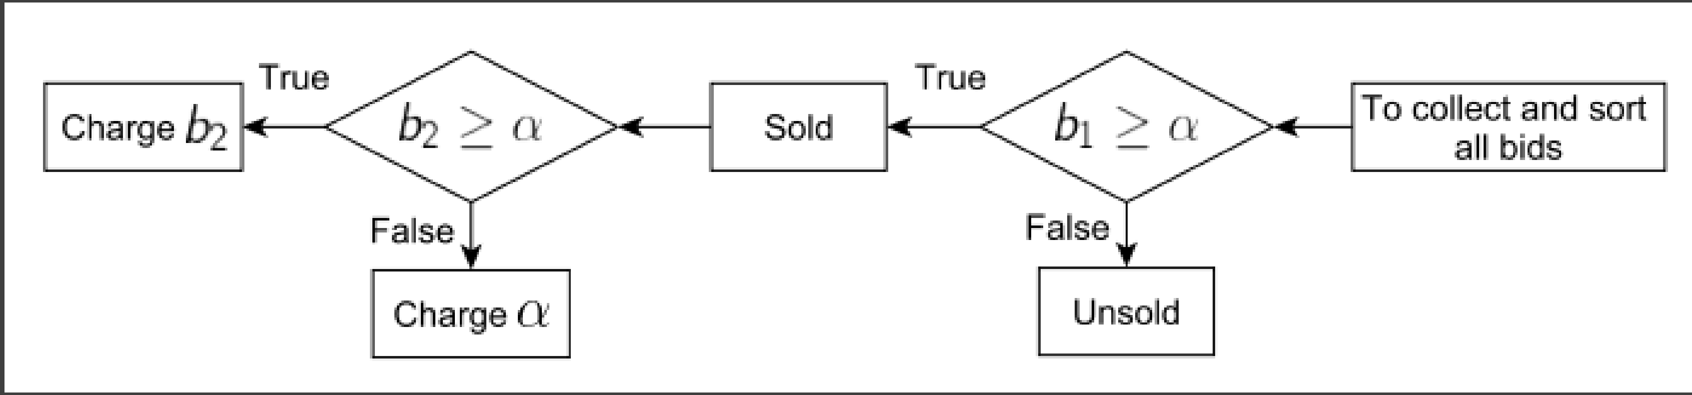

An example

- Suppose: two bidders, private values drawn from Uniform[0, 1]
- Without a reserve price (or $a = 0$), the payoff r is:

$$r = E[\min(b_1, b_2)] = 0.33$$

- With $a = 0.2$:

$$r = E[\min(b_1, b_2)\,|b_1 > 0.2, b_2 > 0.2] + 0.32 \times 0.2 = 0.36$$

- With $a = 0.5$:

$$r = E[\min(b_1, b_2)\,|b_1 > 0.5, b_2 > 0.5] + 0.5 \times 0.5 = 0.42$$

- With $a = 0.6$:

$$r = \underbrace{E[\min(b_1, b_2)\,|b_1 > 0.6, b_2 > 0.6]}_{\text{Paying the second highest price}} + \underbrace{(0.6 \times 0.4) \times 2 \times 0.6}_{\text{Paying the reserve price}} = 0.405$$

Reserve prices in internet advertising auctions: A field experiment, Ostrovsky and Schwarz, 2011

The optimal auction theory

- In the second price auctions, advertisers bid their private values $[b_1, \dots, b_K]$
- Private values -> Bids' distributions $F(\boldsymbol{b}) = F_1(b_1) \times \cdots \times F_K(b_K)$
  - Uniform
  - Log-normal
- The publisher also has a private value V_p
- The optimal reserve price is given by: $$\alpha - \frac{1 - F(\boldsymbol{b})}{F'(\boldsymbol{b})} - V_p = 0$$

Questions:

- How are advertisers bidding?
- Does Uniform/Log-normal fit well?

Optimal Reservation Prices in Auctions, Levin and Smith, 1996

Bidding is complicated

- They usually use a private regression model (No access to publishers)
- Perhaps they don't even know it! (Just try to maximise the ROI)

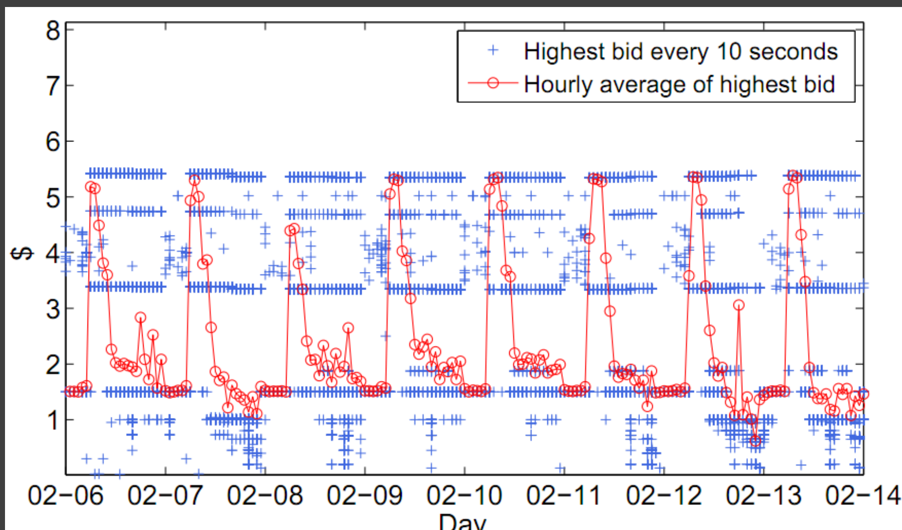

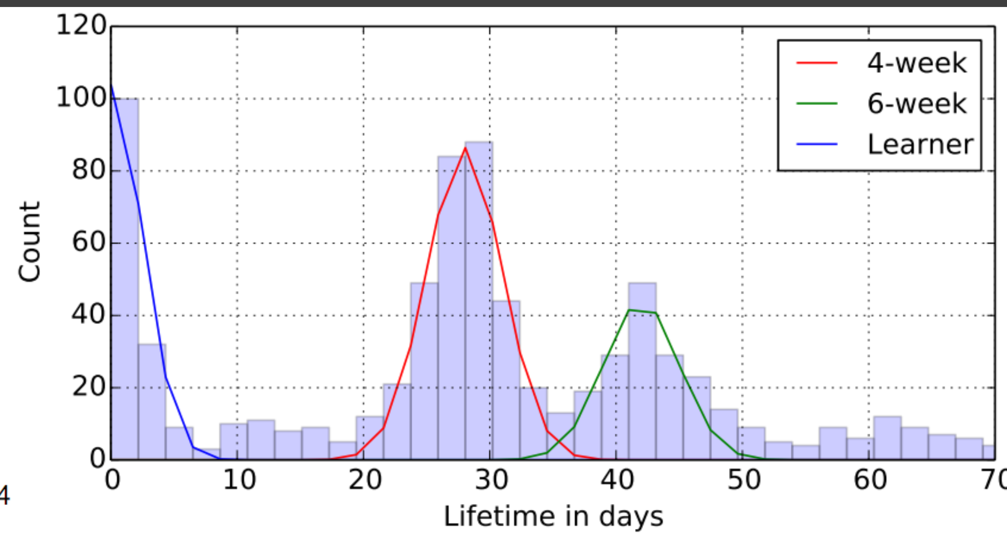

Many advertisers bid at fixed values
(Think about a decision tree)

And they come and go
(with different lifetime span)

Uniform/Log-normal distributions do NOT fit well

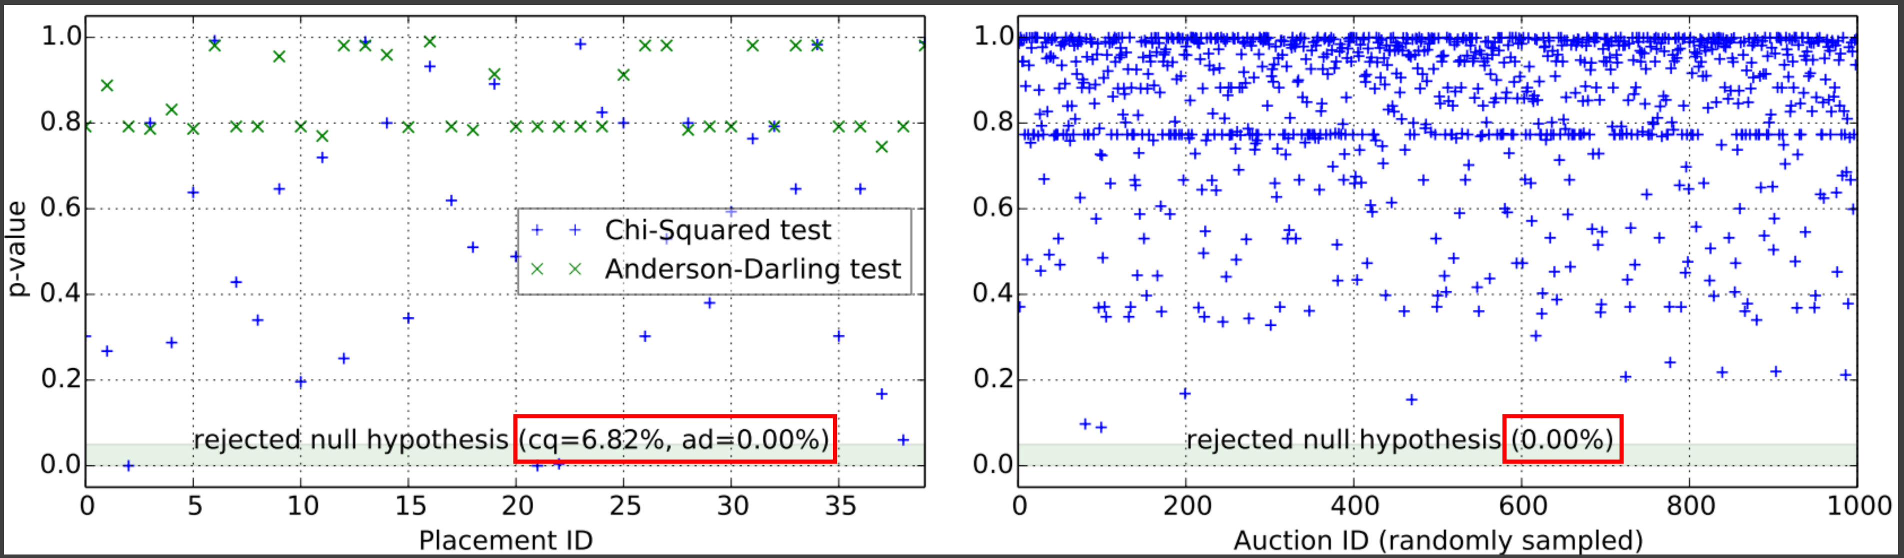

Test at the placement level
(because we usually set reserve prices
on placements)

Test at the auction level

- Chi-squared test for Uniformity
- Anderson-Darling test for Normality

Results from a field experiment

- On Yahoo! Sponsored search
- Using the Optimal Auction Theory

Table 7: Restricted sample (optimal reserve price < 20¢)

Variable	Value	t-statistic	p-value
Number of keywords (T – treatment group)	222,249		
Number of keywords (C – control group)	11,615		
(Mean change in depth in T)−(mean change in depth in C)	−0.8612	−60.29	< 0.0001
(Mean change in revenue in T)−(mean change in revenue in C)	−11.88%	−2.45	0.0144
Estimated impact of reserve prices on revenues	−9.19%	−11.1	< 0.0001

Mixed results

Table 8: Restricted sample (optimal reserve price ≥ 20¢)

Variable	Value	t-statistic	p-value
Number of keywords (T – treatment group)	216,383		
Number of keywords (C – control group)	11,401		
(Mean change in depth in T)−(mean change in depth in C)	−0.9664	−55.09	< 0.0001
(Mean change in revenue in T)−(mean change in revenue in C)	14.59%	1.79	0.0736
Estimated impact of reserve prices on revenues	3.80%	5.41	< 0.0001

Reserve prices in internet advertising auctions: A field experiment, Ostrovsky and Schwarz, 2011

Our solution

- A dynamic and one-shot game between the winner (w) and the publisher (p)
- Extension form representation
  - Information nodes:
    - I_1: Auction succeeded: the winning bid b_1 is higher
    - I_2: Auction failed: the reserve price α is higher
  - Actions:
    - a_{w1}: to increase b_1 so that $b_1 \geq \alpha$
    - a_{w2}: to increase b_1 so that $b_1 < \alpha$
    - a_{w3}: to decrease b_1 so that $b_1 \geq \alpha$
    - a_{w4}: to decrease b_1 so that $b_1 < \alpha$
    - a_{p1}: to increase α so that $\alpha \geq b_1$
    - a_{p2}: to increase α so that $\alpha < b_1$
    - a_{p3}: to decrease α so that $\alpha \geq b_1$
    - a_{p4}: to decrease α so that $\alpha < b_1$

1) Expected payoff of advertiser, publisher

2) Payoff for the advertiser could be negative if one has been bidding the max price
(a_{w1}: to increase b_1 so that $b_1 \geq \alpha$)

3) One won't do that,
so discounted publisher's payoff

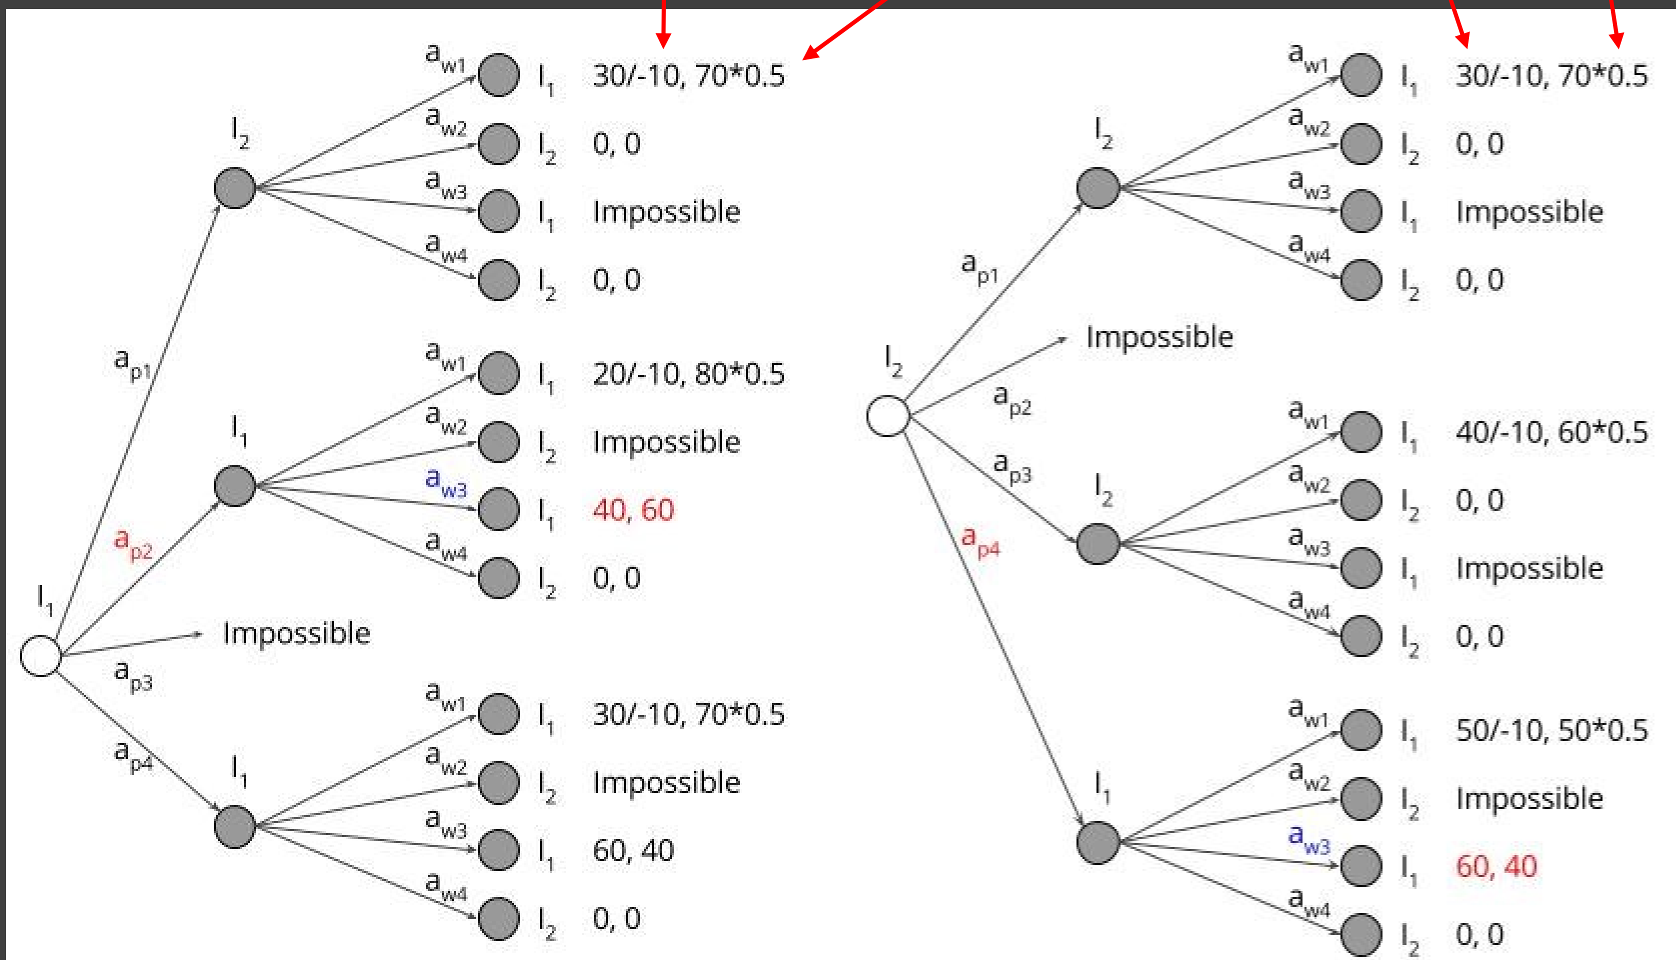

Heuristics and Modification

- If the reserve price is too high, lower it
- If too low, higher it
- If in the preferable range ($b_1 \geq \alpha \geq b_2$), slightly higher it
- A parameter λ allowing to converge over time

$$\Delta\alpha(t) = \begin{cases} \lambda^t h\alpha(t) & \alpha > b_1 \\ \lambda^t l\alpha(t) & b_2 > \alpha \\ \lambda^t p\alpha(t) & b_1 \geq \alpha \geq b_2 \end{cases}$$

Dataset (it’s online experiment)

- Observed 130m impressions from Dec 2012 to Feb 2013
  (Subsampled 10% due to computing power restraint)
- From 39 placements, 19 websites of different categories
- From a production Supply Side Platform in UK

Competing Algorithms

- Round-robin scheduling
  - Zero (the base line)
  - Weighted average (linear)
  - Optimal auction theory
  - Heuristics (OneShot)
  - Bayesian (univariant and bivariant)

- Reserve prices are set for each `(placement, hour_of_day)`

Findings

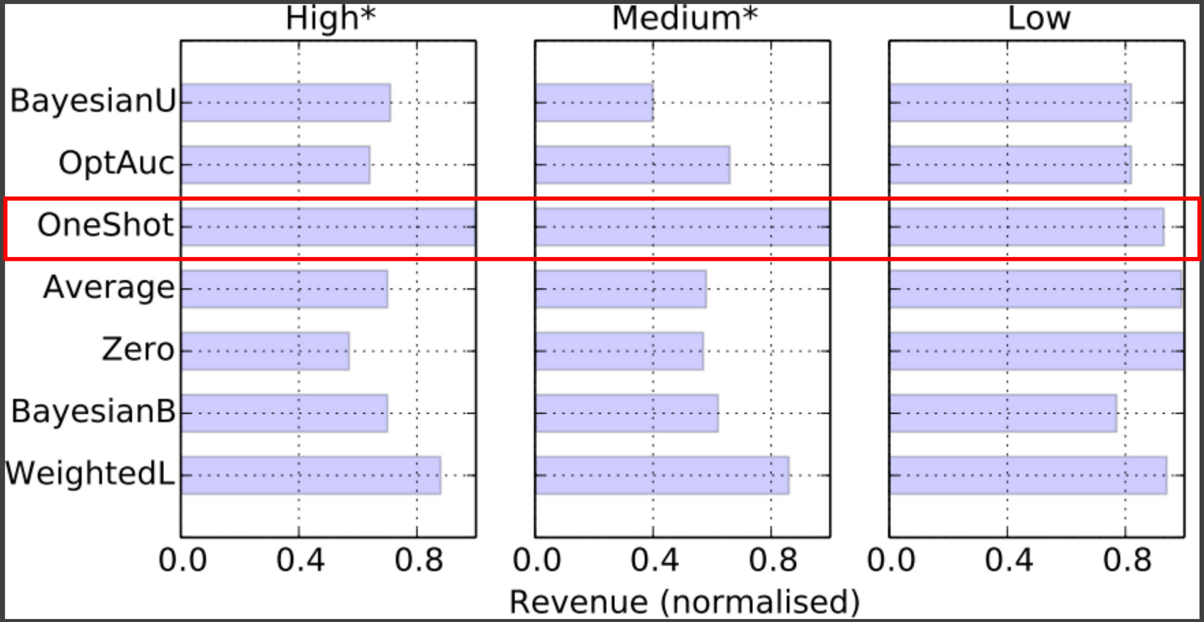

12.3% better than the 2nd best
28.5% better than the optimal auction theory

Findings

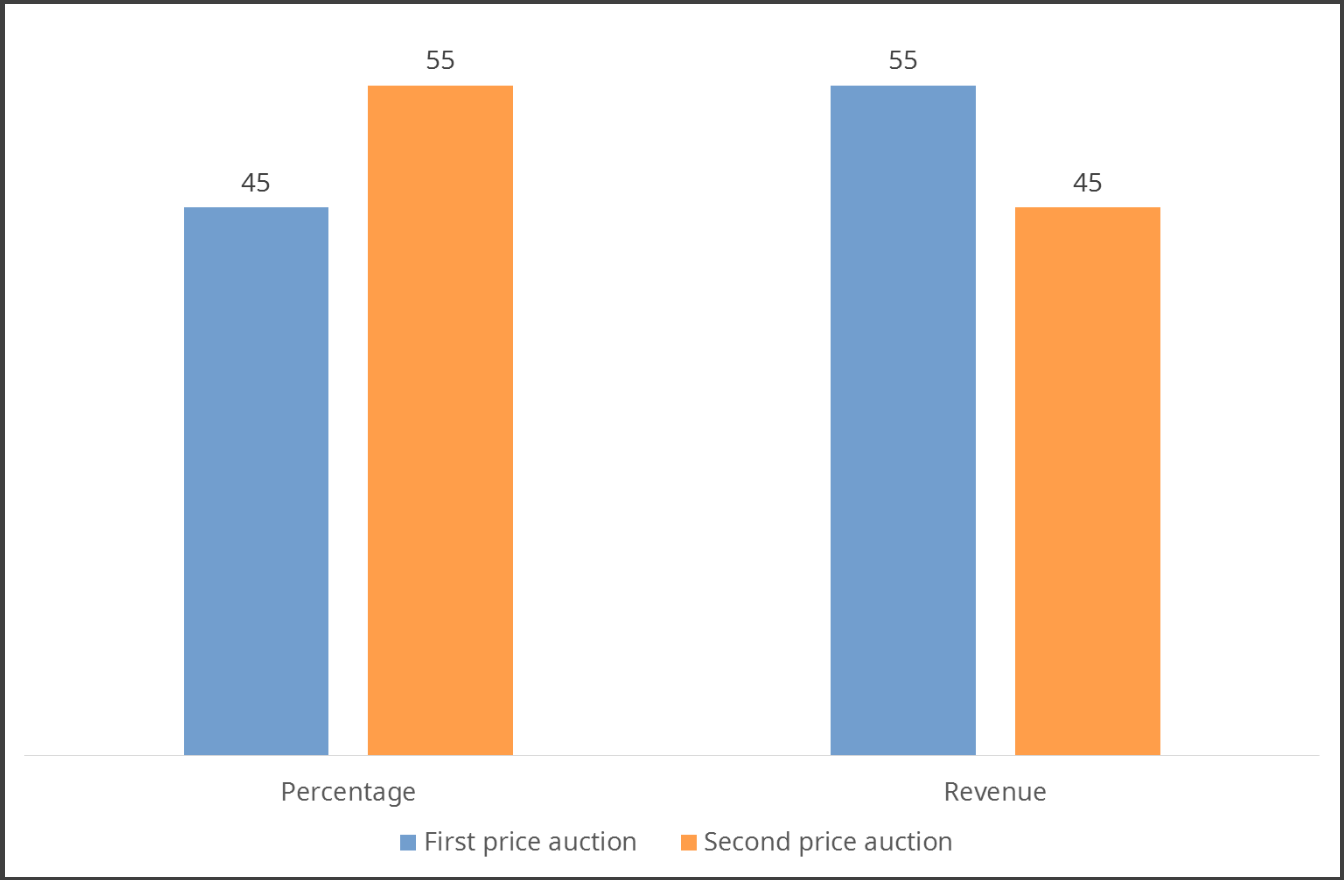

Advertisers are overpaying because of tricky set ups

(They don’t know it could be first price auction!)

(And seems they don’t care!)

Advertisers need to catch up
(at least from 1-year ago's point of view)

and consider cost in bidding algorithms

Weinan Zhang, Shuai Yuan, Jun Wang, Optimal Real-Time Bidding for Display Advertising, KDD 2014

An outlier
(Triggered by some random action)

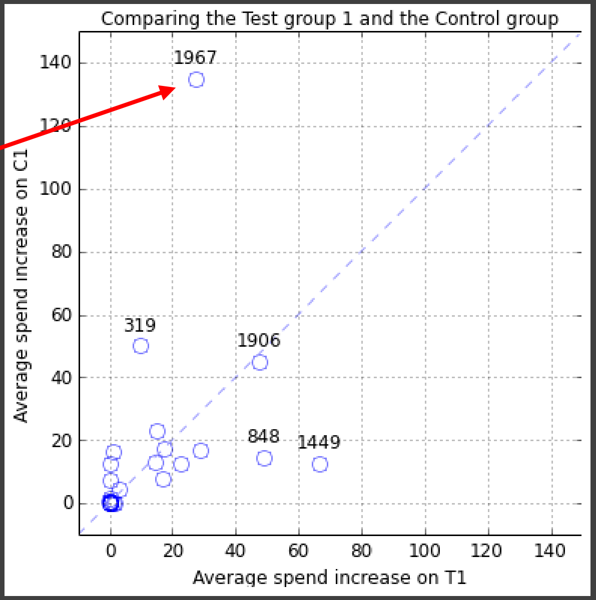

The unchanged budget allocation

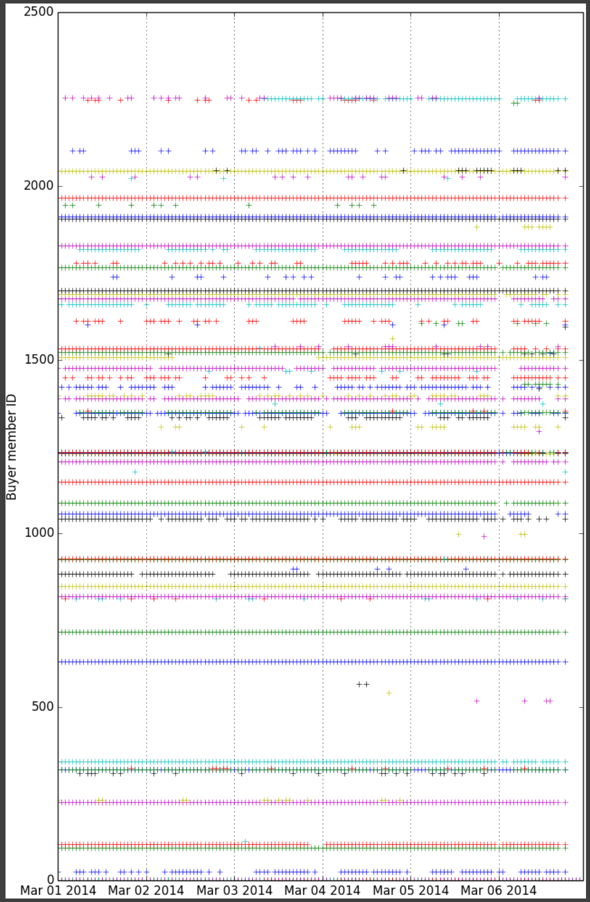

The continuous bidding activity

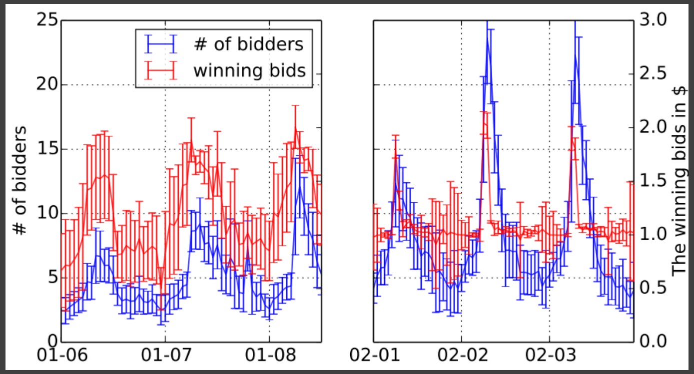

The unchanged bidding pattern

Future Work

- Reserve price optimisation
  - Audience data integration
    (Because the demand side is doing it!)
  - Finding better fitting distributions
- A unified supply side optimisation framework for big players
  - Enough volume for various online tests
  - Dynamic allocation of inventories
    (programmatic guarantee, private/public exchange, etc.)
    (Bowei Chen, Shuai Yuan, Jun Wang,
    A Dynamic Pricing Model for Unifying Programmatic Guarantee and Real-Time Bidding in Display Advertising,
    ADKDD 2014)
  - Joint optimisation

Q & A

- Thanks for your time!

- Shuai Yuan
- s.yuan@cs.ucl.ac.uk
- www.yuan-shuai.info