Bid Shading by Win-Rate Estimation and Surplus Maximization*

Shengjun Pan
Verizon Media
AdKDD ’20

* Work by: S Pan, B Kitts, T Zhou, H He, B Shetty, A Flores, D Gligorijevic, J Pan, T Mao, S Gultekin and J Zhang

What is bid shading?

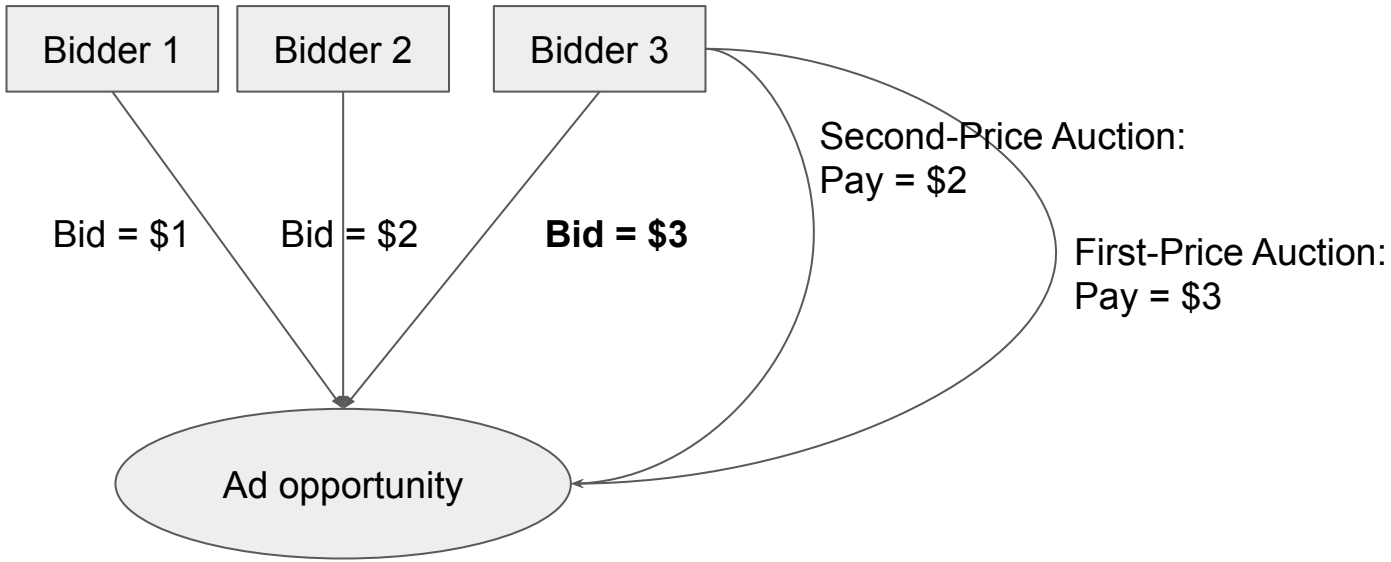

Key difference: Bidding strategy for First-Price Auction replies on competing bidders.
Bid Shading: $3 → $2.01

Surplus

Tradeoff between shading and winning:
more shading, more saving, less likely to win

$$surplus = (V - b)\mathbf{I}(b > mb2w) = \begin{cases} V - b, & \text{if } b > mb2w, \\ 0, & \text{otherwise.} \end{cases}$$

V = value of current ad opportunity
mb2w = minimum bid to win (highest competing bid price)

Optimal bid price

Competing bid prices unknown → Estimate distribution instead

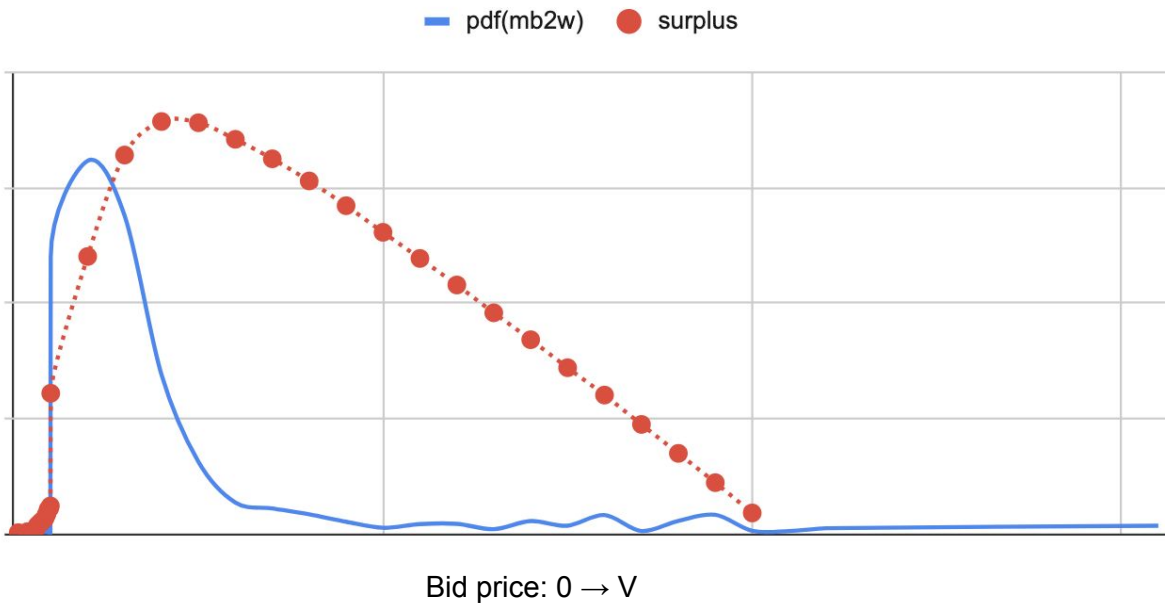

V = value of current ad opportunity
mb2w = minimum bid to win (highest competing bid price)

Previous approaches

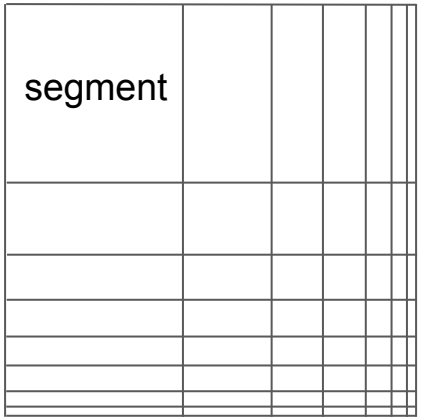

Publisher + user dimensions

	Segment	**Machine-Learning**
Method	Distribution of mb2w* → Maximize surplus	Predict bid shading factor → Multiplying value
Models	One model per segment	Single global model
Cons	● Cross-segment info unused ● Too many small segments; ● System burden;	● Fixed bid for a segment; ● Mb2w must be given.

* mb2w = minimum bid to win

A unified approach: joint distribution estimation

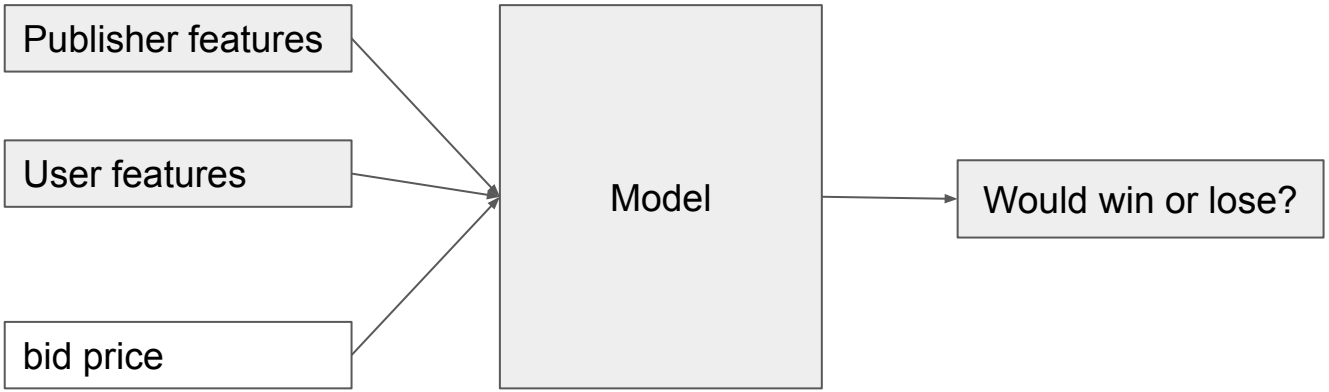

Bid price:

- known at training;
- to be searched for at serving

Essentially estimate win-rate CDFs for all segments simultaneously.

A unified approach: formulation, bisection search

$$\Pr(\text{win}) = \left(1 + e^{-(w_0 + \sum_{i=1}^{k} w_i x_i + \beta \log b)}\right)^{-1}$$

$$\begin{aligned} b^* &= \arg\max_{b>0} (V - b) \operatorname{logistic}\left(w_0 + \sum_{i=1}^{k} w_i x_i + \beta \log b\right) \\ &= \arg\max_{b>0} (V - b) \left(1 + e^{-w_0 - \sum_{i=1}^{k} w_i x_i - \beta \log b}\right)^{-1} \\ &= \arg\max_{b>0} \frac{V - b}{1 + e^{-\alpha} b^{-\beta}}, \end{aligned}$$

where $\alpha = w_0 + \sum_{i=1}^{k} w_i x_i$.

$$f(b) = \frac{V - b}{1 + e^{-\alpha} b^{-\beta}}$$

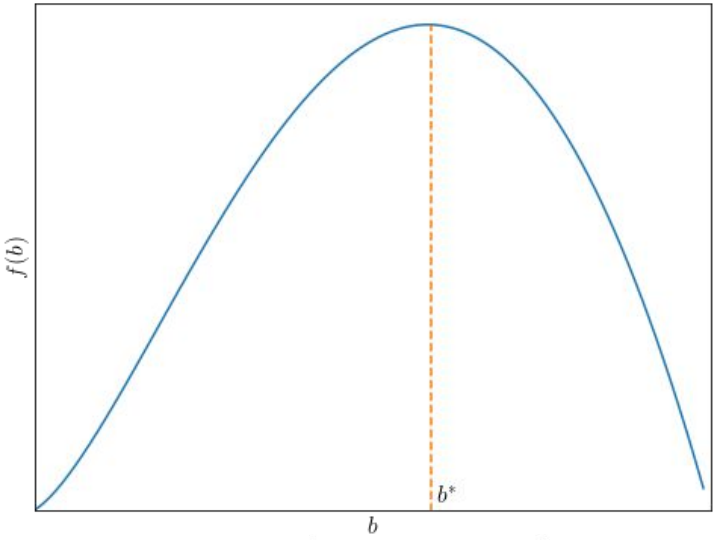

$$\frac{\beta}{\beta + 1 + e^{\alpha} V^{\beta}} V \le b^* < \frac{\beta}{\beta + 1} V.$$

A/B testing

Budget controller:
Shaded → Surplus (saved cost) → Reinvested → Better eCPX

A/B Testing	Spend	Surplus	eCPM	eCPC	eCPA
WR v.s. LR	+1.3%	+1.4%	-7.4%	-4.5%	-2.7%
WR v.s. SEG	+1.2%	+2.5%	-5.4%	-5.5%	-3.9%

Table 3: Improvements on Business Metrics

LR: logistic regression, predicting optimal bidding factor
SEG: per-segment distribution estimation
WR: our method, win-rate based unified approach

Challenges

- It is not obvious “what’s the right thing to do”
  - Different choices for formulation, e.g., what distribution to use
  - Different choices for metrics
- Convoluted with budget controllers
  - Offline experiments are not sufficient
  - Dependency on controllers
- Speed requirement at serving time
  - Inference or maximization must be fast at serving

Thank you!

Questions?