Risk-Aware Bid Optimization for Online Display Advertisement

Rui Fan
Department of Decision Sciences
HEC MONTRÉAL

(joint work with Erick Delage)

CIKM '22

GERAD

Outline

CONTEXT

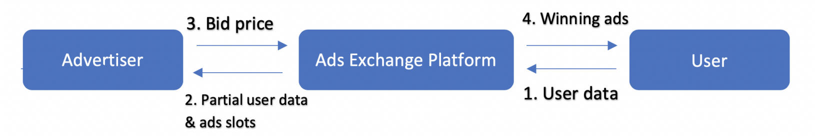

Figure: Ad exchange mechanism

PROBLEM DESCRIPTION

- An advertiser or agent (decision maker)
- Data: website users, ad slots format
- Find the optimal bidding policy given the predetermined budget for a certain period of time
- Map the bidding prices to each ad opportunity at once:
  Maximize the profit;
  Control the risk of violating the budget constraint.

Related Work

- Linear related to Click-through rate (CTR) [11] , or value of click [10, 2] (truthful bidding)
- Non-linear with estimations of CTR, winning probabilities, etc. [16, 5, 12]
- Multi-stage: RL-based models [8, 15, 17, 4]

Risk Management on Profit (RMP) Model [7]

- Static problem
- Risk of generated profit
- Risk from CTR estimation error (Bayesian logistic regression)
- Winning price is an independent variable

METHODOLOGY

- A stochastic model that mixes both empirical and parametric distributions
- Expected utility theory [1, 9] with entropic risk measure [13]
- Optimize a bid policy over a batch of M opportunities

Assumption

Assumption 1

The winning price W, realized click C and net value of the customer V are mutually independent given X.

MODELING CONDITIONAL CTR

- We assume that the CTR depends on the opportunity's features X, and formally denote [1]:

$$\theta(X) := \mathbb{P}(C|X)$$

[1] DeepFM model [6] is used to estimate $\theta(X)$

Modeling Conditional Winning Price Distribution

- The winning price W is conditional on observing X
- W follows the normal distribution $W \sim N(\hat{w}(X), \sigma(X))$
- Parametrized probability distribution function of the winning price W modeled as follows [2]:

$$f_{W|X}(w) = \frac{1}{\sigma(X)\sqrt{2\pi}} e^{-\frac{1}{2}\left(\frac{w-\hat{w}(X)}{\sigma(X)}\right)^2}.$$

[2] DeepFM model [6] is used to estimate $\hat{w}(X), \sigma(X)$

Modeling Conditional Winning Probability

- Expense only happens when the advertiser wins the bids [10]
- We model the probability of winning the bid depends on the bidding price b and the winning price W:

$$s(b, W) := 1\{b \geq W\} = \begin{cases} 1 & b \geq W \\ 0 & \text{otherwise} \end{cases}$$

MODELING CONDITIONAL VALUE OF CUSTOMER

- We model this variable conditional on the given opportunity X by [3] :

$$\hat{V}(X) := \mathbb{E}[V|X]$$

[3] In our experiments, V will be considered independent of X and known for simplicity.

Risk Neutral Model

- Considering a random batch of M i.i.d. opportunities denoted by $\{(X_i, W_i, C_i, V_i)\}_{i=1}^{M}$, with each V_i, C_i, and W_i mutually independent given X_i (as per Assumption 1)
- Maximizes the expected profit generated over the batch while satisfying the budget constraint:

$$
\begin{aligned}
b^{\text{rnp}}(\cdot) := \underset{b:\mathcal{X}\to\mathbb{R}^+}{\operatorname{argmax}} \quad & \mathbb{E}[\text{Batch profit}] \\
\text{s. t.} \quad & \mathbb{E}[\text{Batch expense}] \leq BM,
\end{aligned}
$$

MATHEMATICAL FORM

- Based on the linearity of expectation, batch expressions can be simplified to expected instantaneous value format:

$$
\begin{aligned}
b^{\mathrm{rnp}}(\cdot) = \underset{b:\mathcal{X}\to\mathbb{R}^+}{\operatorname{argmax}} \quad & \mathbb{E}[VCs(b(X), W)] - \mathbb{E}[Ws(b(X), W)] \\
\text{s. t.} \quad & \mathbb{E}[Ws(b(X), W)] \leq B. \qquad (1)
\end{aligned}
$$

Model the Risk

- Exponential utility function to model the risk aversion
- We replace the expected expense constraint with:

$$\mathbb{E}[u_\alpha((1/M)\text{Batch expense})] \geq u_\alpha^{-1}(B),$$

where $u_\alpha(y) := -\exp(\alpha y)$ is a concave utility function that allows the decision maker to control risk exposure using the parameter α.

Model the Risk

- For a batch of M opportunities, the constraint takes the form:

$$\mathbb{E}[u_\alpha(\frac{1}{M}\sum_{i=1}^{M} W_i s(b(X_i), W_i))] \geq u_\alpha(B). \tag{2}$$

Model the Risk

Lemma 1
Constraint (2)

$$\mathbb{E}[u_\alpha(\frac{1}{M}\sum_{i=1}^{M} W_i s(b(X_i), W_i))] \geq u_\alpha(B)$$

is equivalent to

$$\mathbb{E}[h(b(X), X)] \geq -1,$$

where

$$h(b, X) := -e^{\gamma_1(X)}\Phi(\frac{b - \hat{w}(X) - \alpha'\sigma(X)^2}{\sigma(X)}) - e^{\gamma_2} + e^{\gamma_2}\Phi(\frac{b - \hat{w}(X)}{\sigma(X)}) \tag{3}$$

with $\alpha' := \alpha/M$, $\gamma_1(X) := (1/2)(\alpha')^2\sigma(X)^2 + \alpha'\hat{w}(X) - \alpha' B$ *and* $\gamma_2 := -\alpha' B$.

Risk Averse Profit Maximization Model

- The risk-averse expected instantaneous profit maximization problem:

$$
\begin{aligned}
b^{\text{rap}}(\cdot) := \underset{b:\mathcal{X}\to\mathbb{R}^+}{\text{argmax}} \quad & \mathbb{E}[VCs(b(X), W)] - \mathbb{E}[Ws(b(X), W)] \\
\text{s. t.} \quad & \mathbb{E}[h(b(X), X)] \geq -1.
\end{aligned}
\tag{4}
$$

Lagrangian relaxation

$$\tilde{b}_{\lambda}^{\text{rap}}(\cdot) := \underset{b:\mathcal{X}\to\mathbb{R}^+}{\operatorname{argmax}} \quad \mathbb{E}[VCs(b(X),W)] - \mathbb{E}[Ws(b(X),W)] - \lambda(-1 - \mathbb{E}[h(b(X),X)])$$
$$= \underset{b:\mathcal{X}\to\mathbb{R}^+}{\operatorname{argmax}} \quad \mathbb{E}[\mathcal{G}_{\lambda}(b(X),X)],$$

Lagrangian Multiplier

- Indicate the strength of the budget constraint
- The relationship between λ and expected revenue/expense is monotonous
- The optimal λ can be found by the bisection method using the Training set that depends on the empirical distribution of X

Closed-form Solution of Lagrangian relaxation

Lemma 2

For any $\lambda \geq 0$, a maximizer of the Lagrangian relaxation takes the form:

$$\forall X \in \mathcal{X},\ \tilde{b}_{\lambda}^{rap}(X) := \arg \max_{b \in \{0,\ -\frac{\mathbf{W}(\lambda\alpha' e^{(\hat{V}(X)\theta(X)+\lambda e^{\gamma_2} - B)\alpha'})}{\alpha'} + \hat{V}(X)\theta(X) + \lambda e^{\gamma_2},\ \infty\}} \mathcal{G}_{\lambda}(b, X),$$

where $\mathbf{W}$ *is the Lambert W-function [3], i.e. the inverse of* $f(x) := xe^x$.

Evaluation Metrics

- Sharpe ratio [14]
- (Empirical) Early stop frequency
- Common KPIs: profit, expense, clicks, impression rate[4]

[4]The probability that advertiser successfully expose the ad to the customers, which is the realized winning rate

Risk Control

- Control on expense

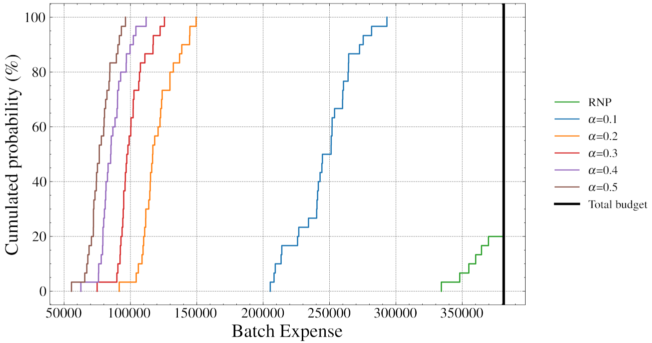

Figure: Empirical CDF of Batch expense under different risk level when B=$\bar{B}/2$

RISK CONTROL

- Early stop frequency

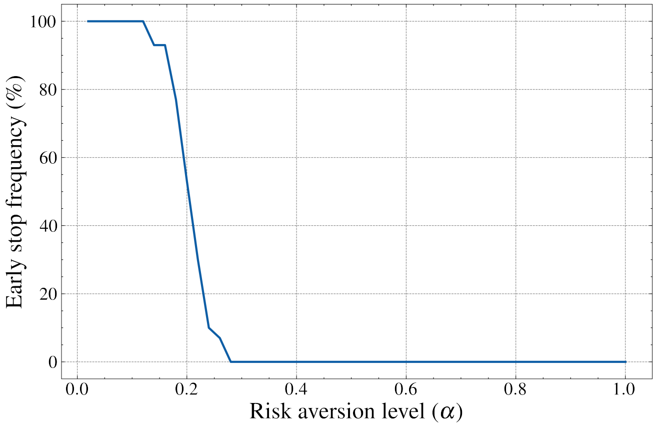

Figure: Empirical Early Stop Frequency under different risk level for the profit model with B=$\bar{B}/32$

COMPARISON

Table: Metrics Results Compared with RMP model when B=$\bar{B}/2$

Metrics	RAP	RNP	RMP-A	RMP-N
Avg. batch clicks	5.600	**6.367**	3.867	3.700
Avg. batch profit	**192574**	169873	144471	142193
Avg. batch expense	292121	381178	190200	**178052**
Avg. impression rate	64.5%	**69.6%**	41.5%	38.8%
Sharpe ratio of profit	**1.083**	0.847	0.892	0.922
Early stop frequency	0	100%	0	0

COMPARISON

Table: Metrics Results Compared with RMP model when B=$\bar{B}/32$

Metrics	RAP	RNP	RMP-A	RMP-N
Avg. batch clicks	0.333	**1.267**	0.433	0.500
Avg. batch profit	20844	**85811**	30136	19454
Avg. batch expense	8006	23822	**7370**	23822
Avg. impression rate	**7.7%**	6.9%	5.0%	7.3%
Sharpe ratio of profit	0.450	**0.802**	0.487	0.363
Early stop frequency	0	100%	0	100%

Conclusion

- Effective risk control
- Competitive performance
- Interpretability & feasibility

BIBLIOGRAPHY I

[1] Daniel Bernoulli. Exposition of a new theory on the measurement of risk. *Econometrica*, 22(1):23–36, 1954.

[2] Ye Chen, Pavel Berkhin, Bo Anderson, and Nikhil R. Devanur. Real-time bidding algorithms for performance-based display ad allocation. In *Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '11, page 1307–1315, New York, NY, USA, 2011. Association for Computing Machinery.

[3] Robert Corless, Gaston Gonnet, D. Hare, David Jeffrey, and D. Knuth. On the lambert w function. *Advances in Computational Mathematics*, 5:329–359, 01 1996.

[4] Manxing Du, Redouane Sassioui, Georgios Varisteas, Radu State, Mats Brorsson, and Omar Cherkaoui. Improving real-time bidding using a constrained markov decision process, 10 2017.

[5] Joaquin Fernandez-Tapia, Olivier Guéant, and Jean-Michel Lasry. Optimal real-time bidding strategies, 2016.

[6] Huifeng Guo, Ruiming Tang, Yunming Ye, Zhenguo Li, and Xiuqiang He. Deepfm: A factorization-machine based neural network for ctr prediction, 2017.

[7] Zhang Haifeng, Zhang Weinan, Rong Yifei, Ren Kan, Li Wenxin, and Wang Jun. Managing risk of bidding in display advertising. *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining*, 02 2017.

[8] Junqi Jin, Chengru Song, Han Li, Kun Gai, Jun Wang, and Weinan Zhang. Real-time bidding with multi-agent reinforcement learning in display advertising. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*, CIKM '18, page 2193–2201, New York, NY, USA, 2018. Association for Computing Machinery.

[9] Oskar Morgenstern John von Neumann. *Theory of Games and Economic Behavior*. Princeton University Press, 1944.

[10] Michael Ostrovsky, Benjamin Edelman, and Michael Schwarz. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *American Economic Review*, 97:242–259, 2007.

BIBLIOGRAPHY II

[11] Claudia Perlich, Brian Dalessandro, Rod Hook, Ori Stitelman, Troy Raeder, and Foster Provost. Bid optimizing and inventory scoring in targeted online advertising. In *Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '12, page 804–812, New York, NY, USA, 2012. Association for Computing Machinery.

[12] Kan Ren, Weinan Zhang, Ke Chang, Yifei Rong, Yong Yu, and Jun Wang. Bidding machine: Learning to bid for directly optimizing profits in display advertising. *IEEE Transactions on Knowledge and Data Engineering*, 30(4):645–659, 2018.

[13] Birgit Rudloff and Ralf Wunderlich. Entropic risk constraints for utility, 2008.

[14] William F. Sharpe. The sharpe ratio. *The Journal of Portfolio Management*, 21(1):49–58, 1994.

[15] Di Wu, Xiujun Chen, Xun Yang, Hao Wang, Qing Tan, Xiaoxun Zhang, Jian Xu, and Kun Gai. Budget constrained bidding by model-free reinforcement learning in display advertising. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*, CIKM '18, page 1443–1451, New York, NY, USA, 2018. Association for Computing Machinery.

[16] Weinan Zhang, Shuai Yuan, and Jun Wang. Optimal real-time bidding for display advertising. In *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '14, page 1077–1086, New York, NY, USA, 2014. Association for Computing Machinery.

[17] Jun Zhao, Guang Qiu, Ziyu Guan, Wei Zhao, and Xiaofei He. Deep reinforcement learning for sponsored search real-time bidding. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, KDD '18, page 1021–1030, New York, NY, USA, 2018. Association for Computing Machinery.

QUESTIONS & COMMENTS

- Code:
  `https://github.com/ReneeRuiFAN/risk-aware_bid_optimization`
- Contact:
  rui.fan@hec.ca
  rui@rbcwmlab.com

Thank you!