A Lightweight MPC Bidding Framework for Brand Auction Ads

Yuanlong Chen*, Bowen Zhu, Bing Xia, Yichuan Wang

ByteDance Inc. | *Current affiliation: Netflix Inc.

Why brand auction ads?

Performance ads

- Clicks and conversions are often sparse.
- Feedback can be delayed.
- Heavy bid-landscape models are common.

Brand auction ads

- Impressions and video views arrive quickly.
- Short-window response fitting is practical.
- Serving can stay low-dimensional.

Design implication: use recent pacing intervals to estimate the next-interval response online.

Control objective

Max delivery under budget

$$\max_{x_t \in \{0,1\}} \sum_t x_t r_t \quad \text{s.t.} \quad \sum_t x_t c_t \leq B.$$

r_t: brand result value; c_t: auction cost; B: campaign budget.

Optimal bid structure

$$b^*_{t,\mathrm{imp}} = b^* \cdot r_t$$

- Serving controls only the campaign multiplier b.
- The multiplier is updated once per pacing cycle.

Target spend → **Bid response** → **Next bid**

The key control question becomes: what multiplier should be used in the next interval?

One-step MPC

Four-step cycle

1. Remaining budget + supply
2. Compute TS_k
3. Fit $\widehat{S}_k(b)$
4. Choose b_k^*

Lifetime-aware target

$$TS_k = B_k \frac{\widehat{N}_k}{\sum_{j \geq k} \widehat{N}_j}$$

One-step rule

$$b_k^* = \sup\{b : \widehat{S}_k(b) \leq TS_k\}.$$

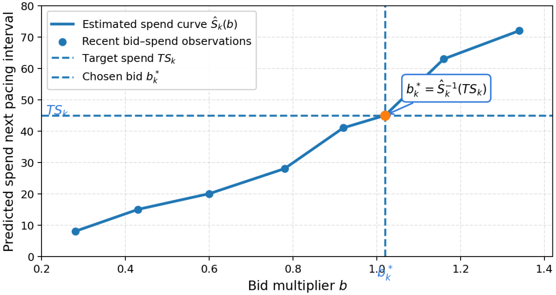

PAVA: lightweight bid-response estimation

Input

- Recent (b_i, s_i) pairs.
- Noisy interval spend.

PAVA output

- Monotone $\widehat{S}_k(b)$.
- Largest feasible bid b_k^*.

Online, one-dimensional, and easy to debug.

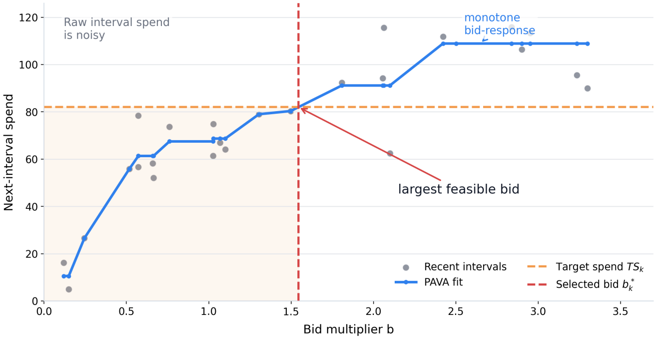

Cost cap extension

Add a result response

$$\widehat{S}_k(b) \quad \text{and} \quad \widehat{G}_k(b)$$

- $\widehat{S}_k(b)$: predicted spend.
- $\widehat{G}_k(b)$: predicted results.
- Both are monotone and fitted with PAVA.

One-step constrained search

$$\begin{aligned} \max_{0 \le b \le b_u} \quad & \widehat{G}_k(b) \\ \text{s.t.} \quad & \widehat{S}_k(b) \le TS_k, \\ & \widehat{S}_k(b)/\widehat{G}_k(b) \le TC_k. \end{aligned}$$

The ratio need not be monotone, so production uses a bid-grid feasibility check.

The controller remains low-dimensional: fit curves, test feasibility, apply one bid.

Offline simulations

100 max-delivery video-campaign episodes; baselines are tuned PID and dual online gradient descent.

CPV approaches oracle

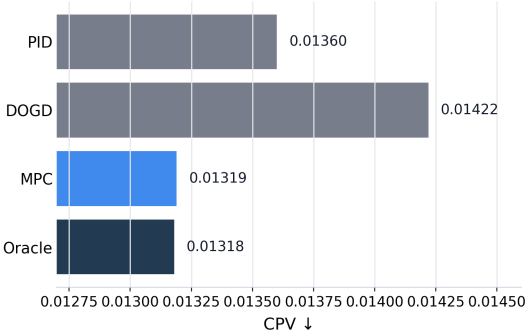

Bid variance is lower

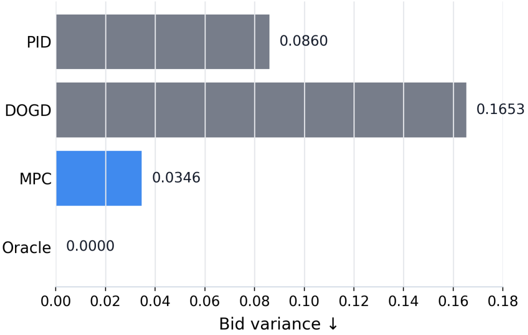

Model-based inversion reduces oscillation relative to error-only feedback.

Online A/B test and takeaway

Seven-day within-campaign budget-split test; only the bidding controller changed.

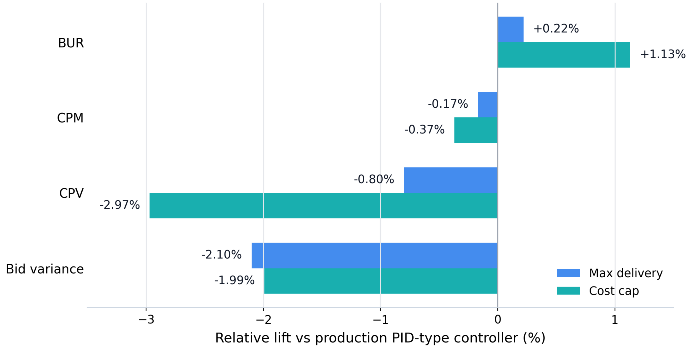

Main takeaways

- Higher budget utilization.
- Lower CPM and CPV.
- Lower bid variance.
- No extra model-serving dependency.

Best fit

Upper-funnel brand objectives with dense, fast feedback.