A Lightweight MPC Bidding Framework for Brand Auction Ads

Yuanlong Chen*
ychenfd@gmail.com
ByteDance Inc.
San Jose, California, USA

Bowen Zhu
bowen.zhu@bytedance.com
ByteDance Inc.
San Jose, California, USA

Bing Xia
xiabing.07@bytedance.com
ByteDance Inc.
San Jose, California, USA

Yichuan Wang
yichuan.wang@bytedance.com
ByteDance Inc.
San Jose, California, USA

Abstract

Brand advertising is a major objective in digital advertising, yet most real-time bidding methods are designed for performance campaigns. We propose a lightweight model predictive control framework for brand auction ads that leverages the relatively dense and timely feedback of upper-funnel objectives. The method estimates monotone bid-to-spend and bid-to-result responses online using isotonic regression and uses these models to select bids under budget and cost constraints. The framework is fully online, computationally lightweight, and easy to deploy. Offline simulations show improvements over PID and dual online gradient descent, and online budget-split A/B test demonstrates improved budget utilization, cost efficiency, and bid stability over production baseline.

CCS Concepts

• **Information systems → Online advertising; Display advertising**.

Keywords

Brand Advertising, Model Predictive Control, Real-time Bidding, Online Optimization, Optimal Bidding

1 Introduction

Online advertising has become a cornerstone of modern commerce, enabling businesses to reach targeted audiences efficiently and at scale. Digital advertising is also a primary revenue driver for major online platforms, where advertisers pursue objectives ranging from brand awareness and user engagement to clicks, conversions, and sales.

Digital ad impressions are commonly sold through real-time bidding (RTB), where bidding systems evaluate each opportunity and submit bids within milliseconds. The platform then runs an auction, selects the winning ad, and charges the advertiser according to the campaign's pricing rule, such as CPM, CPC, or CPV. Although RTB has been extensively studied, most existing bidding methods are designed for performance ads, which optimize lower-funnel outcomes such as clicks, conversions, and sales. These campaigns often face delayed and sparse feedback. Brand ads, by contrast, target upper-funnel objectives such as impression and video view, where feedback is typically faster and denser. This difference creates an opportunity for bidding algorithms tailored to brand auction ads.

This paper makes three contributions. First, we identify a production-relevant brand-advertising regime in which dense, timely upper-funnel feedback enables short-window online response estimation. Second, we develop a serving-light one-step RHC controller that estimates monotone bid-to-spend and bid-to-result curves using isotonic regression and uses them to select bids under budget and cost constraints, without a separate bid-landscape model-serving dependency. Third, we validate this design through offline simulations and a production budget-split A/B test, demonstrating improvements in budget utilization, cost efficiency and bid stability.

2 Related Work

Early digital advertising systems relied on manual bidding and throttling-based pacing to avoid premature budget exhaustion or underdelivery [2, 26]. More recent RTB work studies campaign-level online optimization under budget and cost constraints, including controller-based methods such as PID and dual online gradient descent [5, 6, 9, 11, 13, 16, 17, 21, 23, 27, 28]. Learning-based approaches formulate constrained bidding as sequential decision problems [22, 25]. A separate line of work studies platform-level online allocation and auto-bidding mechanisms [1, 3, 12, 18, 20]. For brand ads, most prior work focuses on guaranteed delivery or exposure planning rather than auction-based bidding [10, 14, 15, 19, 24]. Our work focuses on auction-based brand campaigns and proposes a lightweight online MPC approach for budget pacing without training or serving a high-dimensional policy.

3 Lightweight MPC Framework

We present a lightweight MPC framework for brand campaigns under a simplified second-price auction model. Starting from the budget-constrained max-delivery problem, we derive the optimal bidding structure, use online isotonic regression to model the bid landscape, and extend the framework to multi-constraint settings such as cost caps. In this framework, MPC updates bids repeatedly by combining the remaining budget with recent delivery feedback and short-horizon spend targets.

3.1 Problem Formulation

We formulate the max-delivery problem for a brand advertising campaign (e.g., an awareness or video ad) charged on a per-impression basis. Let T be the predicted number of auction opportunities in a day. For the t-th auction, let r_t denote the utility, such as $r_t = 1$ for impression ads or the video play rate (VPR) for video ads. Let c_t be

*Corresponding author, work done while at ByteDance Inc. Current affiliation: Netflix Inc.

the auction cost per impression. In a simplified second-price eCPM auction, if the highest competing bid is c_t^{eCPM}, then $c_t = c_t^{\mathrm{eCPM}}/1000$. Let $x_t \in \{0,1\}$ indicate whether the campaign wins auction t. The budget-constrained max delivery problem is

$$\begin{aligned} \max_{x_t \in \{0,1\}} \quad & \sum_{t=1}^{T} x_t \cdot r_t \\ \text{s.t.} \quad & \sum_{t=1}^{T} x_t \cdot c_t \le B. \end{aligned} \tag{1}$$

Under standard regularity conditions, such as i.i.d. c_t and r_t, the optimal bid per impression takes the form [5, 17]

$$b^*_{t,imp} = \frac{r_t}{\lambda^*} = b^* \cdot r_t, \tag{2}$$

where λ^* is the Lagrangian dual variable for the budget constraint and $b^* = 1/\lambda^*$ is the optimal campaign-level bid per result. For impression ads, each impression can be treated as a result. When sufficient supply exists and the budget constraint is binding, the optimal bid is chosen to exhaust the budget in expectation by campaign end, motivating a target spend rate based on the remaining budget and predicted auction opportunities. PID controllers and dual online gradient descent methods use this idea by adjusting bids or dual variables according to the gap between actual and target spend rates.

In practice, bids are usually updated in batches rather than at every auction. The bid multiplier, or equivalently the dual variable λ, is updated every Δt, and the bid remains fixed within each pacing cycle. The interval Δt is typically chosen from a few seconds to several minutes to balance update frequency and variance of realized signals such as spend-rate.

3.2 General RHC Formulation and One-Step MPC

Suppose the campaign is divided into K pacing cycles. Let $\mathcal{A}_j$ denote the auction opportunities in cycle j. If bid multiplier b_j is held fixed during that cycle, its realized spend and results are

$$S_j(b_j) = \sum_{t \in \mathcal{A}_j} c_t \mathbf{1}\{b_j r_t \ge c_t\}, \tag{3}$$

$$V_j(b_j) = \sum_{t \in \mathcal{A}_j} r_t \mathbf{1}\{b_j r_t \ge c_t\}. \tag{4}$$

These responses are unknown in advance and must be estimated online.

At cycle k, an H-step RHC controller plans bid multipliers $\{b_{k+\ell|k}\}_{\ell=0}^{H_k-1}$, where $b_{k+\ell|k}$ denotes the bid for cycle $k+\ell$ computed using information available at cycle k, and $H_k = \min\{H, K-k\}$. Let $\widehat{S}_{k+\ell|k}(b)$ and $\widehat{V}_{k+\ell|k}(b)$ denote the corresponding predicted spend and result responses. The controller solves

$$\begin{aligned} \max_{\{b_{k+\ell|k}\}_{\ell=0}^{H_k-1}} \quad & \sum_{\ell=0}^{H_k-1} \widehat{V}_{k+\ell|k}(b_{k+\ell|k}) \\ \text{s.t.} \quad & \sum_{\ell=0}^{H_k-1} \widehat{S}_{k+\ell|k}(b_{k+\ell|k}) \le B^{\mathrm{tar}}_{k,H}, \\ & 0 \le b_{k+\ell|k} \le b_u, \quad \ell = 0, \dots, H_k - 1, \end{aligned} \tag{5}$$

where b_u is the bid upper bound and

$$B^{\mathrm{tar}}_{k,H} = B_k \frac{\sum_{\ell=0}^{H_k-1} \widehat{N}_{k+\ell|k}}{\sum_{j=k}^{K-1} \widehat{N}_{j|k}} \tag{6}$$

allocates the remaining budget B_k to the horizon in proportion to predicted auction supply. Because predicted spend is nonnegative and $B^{\mathrm{tar}}_{k,H} \le B_k$, the budget-state recursion and nonnegativity constraints are implicit in (5). Only $b^*_{k|k}$ is applied; the controller then observes delivery, updates its predictions, and re-solves at cycle $k+1$.

Upper-funnel brand signals are generally denser and faster than deep-funnel conversions, motivating a locally stationary approximation over a recent fitting window. We therefore set $H = 1$ in production to reduce serving complexity while retaining the receding-horizon update: the controller predicts delivery for the next interval, applies the resulting bid, and re-optimizes after observing new feedback. When $H = 1$, the RHC problem reduces to choosing a single bid multiplier for the next pacing interval. The target spend for the next pacing interval is

$$TS_k = B_k \cdot \frac{\widehat{N}_{k|k}}{\sum_{j=k}^{K-1} \widehat{N}_{j|k}}. \tag{7}$$

Although the action horizon is one interval, TS_k is lifetime-aware because it depends on the remaining budget and predicted supply over the remaining campaign. The one-step RHC problem becomes

$$\begin{aligned} \max_{0 \le b \le b_u} \quad & \widehat{V}_{k|k}(b) \\ \text{s.t.} \quad & \widehat{S}_{k|k}(b) \le TS_k. \end{aligned} \tag{8}$$

Since both expected spend and expected result are monotone nondecreasing in the bid, the solution to (8) is the largest bid whose predicted spend does not exceed the target spend:

$$b^*_k = \sup\left\{ b \in [0, b_u] : \widehat{S}_{k|k}(b) \le TS_k \right\}. \tag{9}$$

In the rest of this paper, we implement the one-step rule in (9). The next subsection describes how the predicted bid-to-spend response $\widehat{S}_{k|k}(b)$ is estimated online using lightweight isotonic regression.

3.3 Lightweight Bid-to-X Modeling

We construct one-step bid-to-X response models using isotonic regression [7], where X may denote spend, impressions, or video plays over a fixed pacing interval. We use spend to illustrate the method; the other signals are handled analogously.

At pacing cycle k, we collect up to the N most recent completed bid-spend observations $\{(b_i, s_i)\}$, where b_i is fixed throughout interval i and $s_i = S_i(b_i)$ is its realized spend. During warm-up, when fewer than N observations are available, we use all available data points. Before fitting, observations with identical bids are collapsed by averaging their realized spends. We assume that auction supply and market composition remain sufficiently stable over this recent window and the next interval. If supply varies materially, the response can instead be estimated using spend per auction opportunity and rescaled by the predicted next-interval supply. Under this local-stationarity assumption, the latent conditional mean

$$\overline{S}_k(b) = \mathbb{E}[s_i \mid b_i = b]$$

Algorithm 1 Pool Adjacent Violators Algorithm (PAVA)

Input: $\{(b_k, s_k)\}_{k=1}^{n}$, sorted by bid
Output: $\{(b_k, \tilde{s}_k)\}_{k=1}^{n}$, where $\tilde{s}_1 \le \cdots \le \tilde{s}_n$
Initialize each observation as a block $(L_k, R_k, w_k, \mu_k) = (k, k, 1, s_k)$
Scan adjacent blocks from left to right
while two adjacent blocks satisfy $\mu_i > \mu_{i+1}$ **do**
  Merge them and set
$$w \leftarrow w_i + w_{i+1}, \qquad \mu \leftarrow \frac{w_i\mu_i + w_{i+1}\mu_{i+1}}{w}$$
  Recheck the preceding block
end while
for each final block (L, R, w, μ) **do**
  Set $\tilde{s}_k \leftarrow \mu$ for all $k = L, \ldots, R$
end for
return $\{(b_k, \tilde{s}_k)\}_{k=1}^{n}$

is nondecreasing in b: for a fixed auction population, increasing the bid can only expand the set of won auctions. Individual interval observations are nevertheless affected by stochastic auction arrivals, competing bids, and opportunity composition, and therefore are not expected to be monotone. Isotonic regression estimates this latent monotone response from the noisy observations; we denote the resulting next-interval prediction by $\widehat{S}_{k|k}(b)$.

After ordering the observations by bid, isotonic regression solves

$$\min_{z_1 \le \cdots \le z_N} \sum_{i=1}^{N} (s_i - z_i)^2.$$

PAVA [4, 8] solves this problem in $O(N)$ time. In implementation, the moving window is maintained in bid order, with each insertion and deletion requiring $O(\log N)$ time, thus avoiding a complete re-sort at every cycle.

We linearly interpolate between the fitted points to construct $\widehat{S}_{k|k}(b)$. Below the observed range, we interpolate between the structural anchor $(0, 0)$ and the first fitted point, since a zero bid induces zero spend. Above the observed range, we extrapolate from the last fitted point using a predefined positive slope ρ_{ext}. This permits exploration beyond the observed bid range during cold start, with ρ_{ext} controlling the exploration aggressiveness. Algorithms 1 and 2 summarize the fitting and bid-update procedures.

3.4 Extension

We extend the framework to cost-cap bidding, which maximizes results subject to budget and average cost constraints [17]:

$$\begin{aligned} \max_{x_t \in \{0,1\}} \quad & \sum_{t=1}^{T} x_t r_t \\ \text{s.t.} \quad & \sum_{t=1}^{T} x_t c_t \le B, \\ & \sum_{t=1}^{T} x_t c_t \le C \left(\sum_{t=1}^{T} x_t r_t \right), \end{aligned} \tag{10}$$

where C is the advertiser-specified cost cap. Under standard regularity conditions, the optimal impression-level bid retains the

Algorithm 2 MPC-Based Max-Delivery Bid Update

Input: recent bid-spend observations $\{(b_i, s_i)\}_{i=1}^{N}$, sorted by bid; remaining budget B_k; predicted supply $\{\widehat{N}_{j|k}\}_{j=k}^{K-1}$; bid upper bound b_u; extrapolation slope $\rho_{\text{ext}} > 0$
Output: Bid multiplier b_k^*
if $B_k \le 0$ **or** $\sum_{j=k}^{K-1} \widehat{N}_{j|k} \le 0$ **then**
  return 0
end if
$$TS_k \leftarrow B_k \frac{\widehat{N}_{k|k}}{\sum_{j=k}^{K-1} \widehat{N}_{j|k}}$$
Apply Algorithm 1 to obtain $\{(b_i, \tilde{s}_i)\}_{i=1}^{N}$
if $TS_k < \tilde{s}_1$ **then** ▷ interpolate from $(0, 0)$
$$b_k^* \leftarrow b_1 \frac{TS_k}{\tilde{s}_1}$$
else if $TS_k \ge \tilde{s}_N$ **then** ▷ extrapolate above observed range
$$b_k^* \leftarrow \min\{b_u,\ b_N + \rho_{\text{ext}}\,(TS_k - \tilde{s}_N)\}$$
else ▷ interpolate within observed range
  Find the largest $i < N$ such that $\tilde{s}_i \le TS_k$
$$b_k^* \leftarrow b_i + \frac{TS_k - \tilde{s}_i}{\tilde{s}_{i+1} - \tilde{s}_i}(b_{i+1} - b_i)$$
end if
return b_k^*

multiplicative form

$$b_{t,\text{imp}}^* = \frac{1 + \mu^* C}{\lambda^* + \mu^*} r_t = b^* r_t,$$

where λ^* and μ^* are the dual variables associated with the budget and cost constraints.

At pacing cycle k, let NC_k denote the results accumulated so far. If the controller spends the remaining budget B_k, satisfying the lifetime cost cap requires

$$\frac{B}{NC_k + NC_k^{\text{rem}}} \le C,$$

where NC_k^{rem} denotes the additional results delivered over the remaining campaign. Thus,

$$NC_k^{\text{rem}} \ge \frac{B}{C} - NC_k.$$

Define the residual result requirement

$$D_k^C = \frac{B}{C} - NC_k.$$

When $D_k^C > 0$, the maximum affordable average cost per additional result is

$$TC_k = \frac{B_k}{D_k^C} = \frac{B_k}{B/C - NC_k}. \tag{11}$$

If $D_k^C \le 0$, sufficient results have already been accumulated to satisfy the lifetime cap after spending the remaining budget, so the cost constraint is inactive.

The next-interval target spend TS_k is computed as in (7). Using recent interval-level observations, we fit the predicted bid-to-spend and bid-to-result responses $\widehat{S}_{k|k}(b)$ and $\widehat{G}_{k|k}(b)$, respectively. When the cost constraint is active, the predicted cost per result is

$$\widehat{h}_{k|k}(b) = \frac{\widehat{S}_{k|k}(b)}{\widehat{G}_{k|k}(b)}, \qquad \widehat{G}_{k|k}(b) > 0.$$

The resulting one-step cost-cap problem is

$$\begin{aligned} \max_{0 \le b \le b_u} \quad & \widehat{G}_{k|k}(b) \\ \text{s.t.} \quad & \widehat{S}_{k|k}(b) \le TS_k, \\ & \frac{\widehat{S}_{k|k}(b)}{\widehat{G}_{k|k}(b)} \le TC_k, \quad \widehat{G}_{k|k}(b) > 0, \qquad \text{if } D_k^C > 0. \end{aligned} \tag{12}$$

The cost constraint is omitted when $D_k^C \le 0$. Although $\widehat{S}_{k|k}(b)$ and $\widehat{G}_{k|k}(b)$ are nondecreasing, their ratio need not be monotone, so we search over a bid grid. A bid with zero predicted results is infeasible when the cost constraint is active. If no positive bid is predicted feasible, the controller returns zero; in production, a configurable low fallback bid may instead be used under conservative spend safeguards to maintain limited delivery and refresh the response estimates.

4 Evaluation

We evaluate the proposed MPC framework through offline simulations against PID and dual online gradient descent baselines, followed by a large-scale online A/B test on TikTok brand auction ads.

4.1 Offline Simulations

We run simulations to quantitatively evaluate the performance of our proposed bidding algorithms.

4.1.1 Benchmarks. We compare our lightweight MPC against two industry-standard pacing strategies.

PID. Following [28], at each pacing interval starting at τ we compute the normalized error

$$e_\tau = 1 - \frac{AS_\tau}{TS_\tau},$$

where AS_τ is the actual spend and TS_τ is the target spend in that interval. The bid is then adjusted based on the PID controller update rule proposed in [28].

Dual Online Gradient Descent (DOGD). Following [17], at each pacing interval τ we update the dual variable λ via

$$\lambda \leftarrow \lambda \; - \; \epsilon_\tau \cdot \left(1 \; - \; \frac{AS_\tau/NR_\tau}{B/T}\right),$$

where AS_τ is the actual spend, NR_τ is the number of auctions, B is the total budget, T is the total predicted auction opportunities, and ϵ_τ is the learning rate. The bid per result is then set to $1/\lambda$ according to (2).

4.1.2 Experiment Set-up. We simulate a max-delivery video campaign with $B = 100$ and $T = 50{,}000$ auction opportunities, updating bids every 20 auctions. We sample $r_t = \min\{\text{LogNorm}(-1.4, 0.6), 0.8\}$ and $c_t = \max\{\text{LogNorm}(-4, 4), 10^{-5}\}$, where the log-normal parameters are fitted from internal production auction data. Bids are initialized at 10^{-4}, bounded by $b_u = 10$, and rate-limited to ±10% per update. MPC uses $N = 100$ intervals and upper extrapolation slope 0.2; PID uses gains $(0.1, 0.001, 0.05)$, and DOGD uses initial learning rate 0.3 with inverse-round decay. All simulations use random seed 314.

For each bidding algorithm, each simulation episode continues until either the campaign budget is fully depleted or all auction opportunities are exhausted. We run 100 episodes for each algorithm and report the average of metrics to assess performance.

4.1.3 Result Analysis. Several evaluation metrics are computed, including the *budget utilization rate (BUR)* to measure pacing efficiency, the *cost per video view (CPV)* to evaluate return on investment (ROI) for advertisers, and the *bid variance (BV)* to assess the stability of pacing dynamics.

The BUR is defined as the ratio of the delivered budget to the total allocated budget. To compute CPV, we approximate the total number of video views by summing the sampled VPR values across

---

Algorithm 3 MPC-Based Cost-Cap Bid Update

Input: total budget B; cost cap C; remaining budget B_k; accumulated results NC_k; predicted supply $\{\widehat{N}_{j|k}\}_{j=k}^{K-1}$; recent bid-spend and bid-result observations; bid upper bound b_u; grid step Δb
Output: bid multiplier b_k^*

$$D_k^N \leftarrow \sum_{j=k}^{K-1} \widehat{N}_{j|k}$$

if $B_k \le 0$ **or** $D_k^N \le 0$ **or** $\widehat{N}_{k|k} \le 0$ **then**
  return 0
end if

$$TS_k \leftarrow B_k \frac{\widehat{N}_{k|k}}{D_k^N}$$

$$D_k^C \leftarrow \frac{B}{C} - NC_k$$

if $D_k^C > 0$ **then**
  $TC_k \leftarrow B_k / D_k^C$
  CostActive $\leftarrow$ true
else
  $TC_k \leftarrow +\infty$
  CostActive $\leftarrow$ false
end if
Construct $\widehat{S}_{k|k}(b)$ and $\widehat{G}_{k|k}(b)$ using PAVA and the interpolation/extrapolation rules in Section 3.3
Initialize $b_k^* \leftarrow 0$
for $b \leftarrow \Delta b$ to b_u with step Δb **do**

$$s \leftarrow \widehat{S}_{k|k}(b), \qquad g \leftarrow \widehat{G}_{k|k}(b)$$

  if $s \le TS_k$ **and** $(\neg$CostActive **or** $(g > 0$ **and** $s/g \le TC_k))$ **then**
    $b_k^* \leftarrow b$
  end if
end for
return b_k^* ▷ 0 if no positive bid is feasible

Table 1: Performance Comparison of Max Delivery

Algorithm	BUR (%)	#Impressions	CPV	BV
PID	99.9	23,759	0.01360	0.0860
DOGD	99.9	22,754	0.01422	0.1653
MPC	99.8	24,582	0.01319	0.0346
Oracle	99.9	24,584	0.01318	0.00

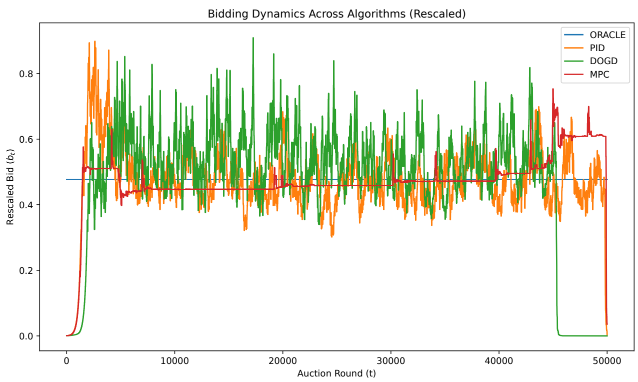

Figure 1: Illustrative bid trajectories from one simulation episode, with bids rescaled to $[0, 1]$.

all winning auction opportunities. The BV is defined as:

$$BV = \frac{1}{T}\sum_{t=1}^{T}\left(1 - \frac{b_t}{b_m}\right)^2,$$

where b_t is the bid at time t, and b_m is the average bid during the pacing day. A lower BV indicates more stable bidding behavior, leading to smoother budget pacing and delivery.

We first perform a hindsight grid search over constant bid multipliers using the realized auction sequence and report the best-performing value as the oracle constant-bid benchmark. The result is optimal within the searched constant-bid class, up to grid resolution. We then fine-tune the PID controller (by adjusting the control gains) and the DOGD algorithm (by tuning the learning rate) to identify their optimal configurations. The simulation results are summarized in Table 1. All tuned algorithms achieve near-full budget utilization. MPC approaches the oracle constant-bid result and attains lower CPV than the tested PID and DOGD baselines. This improvement is partially due to its reduction of bid variance, which helps stabilize the bid dynamics.

To further validate our hypothesis, we plot a sampled path of bid(rescaled to $[0, 1]$) to illustrate the pacing dynamics of each algorithm in Figure 1. The proposed controller produces smoother bid updates. PID and DOGD update their bid or dual variable from realized pacing errors, whereas our controller maps a lifetime-aware target spend directly to a bid through the learned bid-to-spend model. This model-based update is consistent with the lower bid variance reported in Table 1.

We also test robustness to cold-start bid initialization by varying the initial bid from 0.0001 to 1.0. As shown in Figure 2, MPC remains close to the oracle constant-bid CPV across the tested initializations, whereas PID and DOGD are more sensitive to initialization.

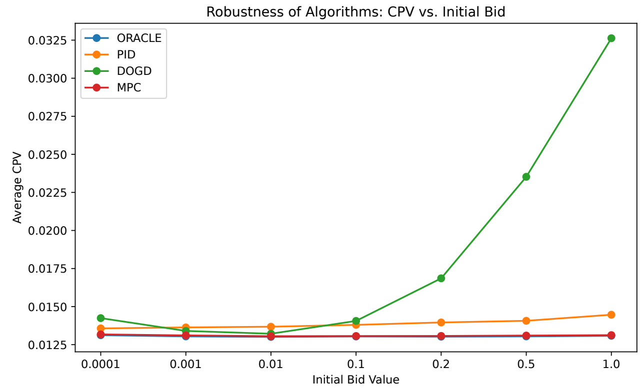

Figure 2: CPV versus initial bid value for each algorithm. MPC is robust to initialization, while PID and DOGD suffer as the cold-start bid deviates from oracle constant bid.

Table 2: Relative lift of the proposed controller over the production PID-type controller in the seven-day budget-split experiment. All entries are individually statistically significant at $p < 0.05$. Positive lift is better for BUR; negative lift is better for CPM, CPV and bid variance.

Metric	Max Delivery	Cost Cap
Budget Utilization Rate	+0.22%	+1.13%
CPM	-0.17%	-0.37%
CPV	-0.80%	-2.97%
Bid variance	-2.10%	-1.99%

4.2 Large-Scale Online Budget-Split A/B Test

We evaluated the proposed controller in a seven-day online experiment on TikTok brand auction campaigns using the platform's standard within-campaign budget-split framework. For each campaign, the budget was divided equally between treatment and control, forming two equal-budget demand cohorts. Eligible supply traffic was randomly assigned between the two cohorts with equal probability. All other production configurations, including targeting and ranking models, were held fixed; the only difference was the bidding controller. The treatment used the proposed MPC controller, whereas the control used the production PID-type controller. The experiment included max-delivery and cost-cap campaigns with awareness and video-view objectives. We omit the absolute campaign count for business-confidentiality reasons.

Table 2 reports the relative lift of treatment over control. Metrics are aggregated over the seven-day period for each campaign cohort, and significance is assessed using two-sided paired t-tests on the within-campaign treatment-control differences. All reported effects are significant at the 5% level, with 95% confidence intervals excluding zero.

5 Remarks

The unique attributes of brand auction ads make this lightweight framework particularly effective. The framework only requires recent bid-response pairs and fits low-dimensional monotone bid-to-spend and bid-to-result curves online. Compared with heavy ML-based landscape models or high-dimensional auction simulators, this design is easier to deploy, debug, and stabilize in production. It also avoids introducing a separate model-serving dependency into the bidding loop, which is important for real-time advertising systems.

In this paper, we focus mainly on upper-funnel brand objectives with relatively fast feedback. In practice, brand campaigns may optimize different billing or optimization events, such as impressions, 2-second views, 5-second views, or 15-second views. Longer video-view objectives introduce additional feedback delay, but in online experiments we find that such impact is relatively small on performance, even for 15-second video ads.

Extending the same framework to deep-funnel performance ads is more challenging. Clicks and conversions are delayed, sparse, and noisier than upper-funnel brand signals. As a result, the recent pacing-interval feedback used to estimate the bid-response models becomes less reliable. Applying this approach to performance campaigns would require additional components for prediction calibration, delayed-feedback correction, attribution, and variance reduction.

6 Conclusion

This work presented a lightweight MPC-based bidding framework for brand auction ads. By learning monotone bid-response models online, the method provides an interpretable and low-overhead way to control budget pacing and cost efficiency. Offline simulations and large-scale online A/B testing validate its practical effectiveness, while extending the approach to delayed and sparse performance-ad feedback remains future work.

References

[1] Zoë Abrams, S Sathiya Keerthi, Ofer Mendelevitch, and John A Tomlin. 2008. Ad delivery with budgeted advertisers: A comprehensive LP approach. *Journal of Electronic Commerce Research* 9, 1 (2008).

[2] Deepak Agarwal, Souvik Ghosh, Kai Wei, and Siyu You. 2014. Budget pacing for targeted online advertisements at linkedin. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*. 1613–1619.

[3] Gagan Aggarwal, Ashwinkumar Badanidiyuru, and Aranyak Mehta. 2019. Autobidding with constraints. In *Web and Internet Economics: 15th International Conference, WINE 2019, New York, NY, USA, December 10–12, 2019, Proceedings 15*. Springer, 17–30.

[4] Miriam Ayer, H Daniel Brunk, George M Ewing, William T Reid, and Edward Silverman. 1955. An empirical distribution function for sampling with incomplete information. *The annals of mathematical statistics* (1955), 641–647.

[5] Santiago R Balseiro and Yonatan Gur. 2019. Learning in repeated auctions with budgets: Regret minimization and equilibrium. *Management Science* 65, 9 (2019), 3952–3968.

[6] Santiago R Balseiro, Haihao Lu, Vahab Mirrokni, and Balasubramanian Sivan. 2022. Analysis of dual-based pid controllers through convolutional mirror descent. *arXiv preprint arXiv:2202.06152* (2022).

[7] Richard E Barlow. 1972. Statistical inference under order restrictions: The theory and application of isotonic regression. *(No Title)* (1972).

[8] Frank MTA Busing. 2022. Monotone regression: A simple and fast O (n) PAVA implementation. *Journal of Statistical Software* 102 (2022), 1–25.

[9] Matteo Castiglioni, Andrea Celli, and Christian Kroer. 2022. Online learning with knapsacks: the best of both worlds. In *International Conference on Machine Learning*. PMLR, 2767–2783.

[10] Bowei Chen, Shuai Yuan, and Jun Wang. 2014. A dynamic pricing model for unifying programmatic guarantee and real-time bidding in display advertising. In *Proceedings of the Eighth International Workshop on Data Mining for Online Advertising*. 1–9.

[11] Yuanlong Chen. 2025. A Practical Guide to Budget Pacing Algorithms in Digital Advertising. *arXiv preprint arXiv:2503.06942* (2025).

[12] Ye Chen, Pavel Berkhin, Bo Anderson, and Nikhil R Devanur. 2011. 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*. 1307–1315.

[13] Vincent Conitzer, Christian Kroer, Eric Sodomka, and Nicolas E Stier-Moses. 2022. Multiplicative pacing equilibria in auction markets. *Operations Research* 70, 2 (2022), 963–989.

[14] Liang Dai, Kejie Lyu, Chengcheng Zhang, Guangming Zhao, Zhonglin Zu, Liang Wang, and Bo Zheng. 2024. Percentile risk-constrained budget pacing for guaranteed display advertising in online optimization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 38. 7987–7994.

[15] Zhen Fang, Yang Li, Chuanren Liu, Wenxiang Zhu, Yu Zheng, and Wenjun Zhou. 2019. Large-scale personalized delivery for guaranteed display advertising with real-time pacing. In *2019 IEEE International Conference on Data Mining (ICDM)*. IEEE, 190–199.

[16] Joaquin Fernandez-Tapia, Olivier Guéant, and Jean-Michel Lasry. 2017. Optimal real-time bidding strategies. *Applied mathematics research express* 2017, 1 (2017), 142–183.

[17] Yuan Gao, Kaiyu Yang, Yuanlong Chen, Min Liu, and Noureddine El Karoui. 2022. Bidding agent design in the linkedin ad marketplace. *arXiv preprint arXiv:2202.12472* (2022).

[18] Paul Grigas, Alfonso Lobos, Zheng Wen, and Kuang-chih Lee. 2017. Profit maximization for online advertising demand-side platforms. In *Proceedings of the ADKDD'17*. 1–7.

[19] Takanori Maehara, Atsuhiro Narita, Jun Baba, and Takayuki Kawabata. 2018. Optimal Bidding Strategy for Brand Advertising.. In *IJCAI*. 424–432.

[20] Aranyak Mehta, Amin Saberi, Umesh Vazirani, and Vijay Vazirani. 2007. Adwords and generalized online matching. *Journal of the ACM (JACM)* 54, 5 (2007), 22–es.

[21] Alessandro Nuara, Francesco Trovò, Nicola Gatti, and Marcello Restelli. 2022. Online joint bid/daily budget optimization of internet advertising campaigns. *Artificial Intelligence* 305 (2022), 103663.

[22] Haozhe Wang, Chao Du, Panyan Fang, Shuo Yuan, Xuming He, Liang Wang, and Bo Zheng. 2022. ROI-Constrained Bidding via Curriculum-Guided Bayesian Reinforcement Learning. In *KDD*. doi:10.1145/3534678.3539211

[23] Jun Wang, Weinan Zhang, Shuai Yuan, et al. 2017. Display advertising with real-time bidding (RTB) and behavioural targeting. *Foundations and Trends® in Information Retrieval* 11, 4-5 (2017), 297–435.

[24] XiaoYu Wang, Bin Tan, Yonghui Guo, Tao Yang, Dongbo Huang, Lan Xu, Nikolaos M Freris, Hao Zhou, and Xiang-Yang Li. 2022. CONFLUX: A Request-level Fusion Framework for Impression Allocation via Cascade Distillation. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 4070–4078.

[25] Di Wu, Xiujun Chen, Xun Yang, Hao Wang, Qing Tan, Xiaoxun Zhang, Jian Xu, and Kun Gai. 2018. Budget Constrained Bidding by Model-free Reinforcement Learning in Display Advertising. In *CIKM*. 1443–1451. doi:10.1145/3269206.3271748

[26] Jian Xu, Kuang-chih Lee, Wentong Li, Hang Qi, and Quan Lu. 2015. Smart pacing for effective online ad campaign optimization. In *Proceedings of the 21th ACM SIGKDD international conference on knowledge discovery and data mining*. 2217–2226.

[27] Xun Yang, Yasong Li, Hao Wang, Di Wu, Qing Tan, Jian Xu, and Kun Gai. 2019. Bid optimization by multivariable control in display advertising. In *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*. 1966–1974.

[28] Weinan Zhang, Yifei Rong, Jun Wang, Tianchi Zhu, and Xiaofan Wang. 2016. Feedback control of real-time display advertising. In *Proceedings of the Ninth ACM International Conference on Web Search and Data Mining*. 407–416.