EXPLICIT BUDGET OPTIMIZATION: ACTIVATING TRAFFIC UNDER BUDGET CONSTRAINTS IN COLD-START AD BIDDING

Hongchang Wu[*], Size Wang[*], Weitong Ou, Zixin Shao, Feihong Liu, Yang Liu, Hongyan Xue, Nianhua Xie
Tencent
{stanfordwu, sizewang, wendyou, zixinshao, feynmanfliu, lioraliu, caryxue, nianhuaxie}@tencent.com

ABSTRACT

Online advertising platforms typically employ cost-constrained bidding frameworks that dynamically adjust real-time bids to optimize advertiser-defined objectives under strict cost constraints. These frameworks rely heavily on feedback signals, such as historical cost consumption and conversion outcomes, to regulate bidding behavior. While effective when sufficient feedback is available, they often struggle during the cold-start stage, where both conversion signals and reliable cost–value estimates are extremely sparse. As a result, bidding strategies tend to become overly conservative, leading to slow traffic activation and delayed signal accumulation, which further hinders subsequent optimization. To address this challenge, we propose Explicit Budget Optimization (EBO), a lightweight module that explicitly supports signal acquisition during cold-start through a small, dedicated exploration budget. Rather than optimizing immediate cost efficiency or advertiser value, EBO focuses on activating early-stage traffic and acquiring initial feedback signals under strictly bounded spending. EBO operates alongside existing cost-constrained bidding strategies and integrates with them via a dual-constraint bidding integration mechanism, in which the final bid submitted to the auction is selected as the maximum of the EBO bid and the base bid. This design enforces a lower-bound on traffic activation while preserving the original cost-control guarantees. Furthermore, EBO supports personalized exploration budgets across different advertisers, campaigns, and creatives, and employs a feedback-driven budget tracking and bid adjustment mechanism to ensure smooth and efficient consumption of the allocated exploration budget throughout the cold-start period. Extensive experiments conducted on Tencent's advertising platform demonstrate that EBO significantly accelerates early-stage signal acquisition for cold-start campaigns while maintaining strict cost control.

1 INTRODUCTION

In recent years, online e-commerce platforms have experienced rapid growth. Online advertising, as a key marketing strategy, leverages the Internet to help advertisers reach target audiences and drive conversions through auto-bidding (Maehara et al., 2018; Ren et al., 2018; Zhang et al., 2014). In practice, advertisers specify cost targets in terms of cost-per-outcome (CPX) metrics, including but not limited to cost per action (CPA) and cost per click (CPC). The platform then aims to maximize advertiser-defined value, such as clicks, conversions, or revenue, under these cost constraints (Cai et al., 2017; Borgs et al., 2007; Evans, 2009). To this end, modern bidding systems dynamically generate real-time bids to ensure that cumulative spending remains within predefined cost targets while optimizing the advertiser's value.

To achieve stable cost control, existing bidding systems commonly adopt reactive cost-control strategies, including Proportional–Integral–Derivative (PID) (Doyle et al., 2013), Model Predictive Con-

[*] indicates equal contribution.

trol (MPC) (Zhang et al., 2022) and Reinforcement Learning (RL)-based approaches (Zhao et al., 2018; Cai et al., 2017; Ye et al., 2020). These methods regulate current bids by leveraging historical feedback signals (cost-value deviations) and forecasts of future advertiser-defined outcomes or cost dynamics. However, their effectiveness critically depends on the availability of sufficient and reliable feedback signals. For new ads or ads with sparse historical data, inaccurate estimates of value or cost dynamics can lead to overly conservative bidding, delayed traffic activation, or even complete delivery stagnation, giving rise to the cold-start problem. Such scenarios correspond to the early stage of ad campaigns, where limited feedback signals are available. We refer to this early phase as the *cold-start stage*, during which the platform must initiate bidding and acquire initial feedback while historical data is sparse.

Importantly, the optimization objective during the cold-start stage differs fundamentally from that of steady-state bidding. Rather than immediately achieving cost-efficient or value-optimal bidding, the primary goal in cold-start is to activate traffic and acquire initial feedback signals while respecting budget constraints (Hillard et al., 2010; Pan et al., 2019; Zhu et al., 2025).

Existing cost-constrained bidding frameworks, which are primarily designed for reactive control based on accumulated feedback, lack explicit mechanisms to support such signal-oriented exploration. Most existing approaches to the cold-start problem instead focus on improving Click-Through Rate (CTR) prediction by generating high-quality feature embeddings for new advertisements (Pan et al., 2019; Zhu et al., 2025). While effective for mitigating cold-start issues at the prediction level, these methods do not directly address the exploration challenges inherent in cost-constrained bidding systems. In practice, the majority of prior cold-start solutions are tailored specifically to new advertisements and rely on additional platform-side interventions to stimulate early exposure. Such approaches are inherently limited in scope and do not generalize well to broader signal-sparse scenarios, where advertisers may not qualify for explicit cold-start treatment or where platform-side intervention is unavailable or undesirable.

To address this challenge, we propose an *Explicit Budget Optimization (EBO)* module that explicitly supports signal acquisition during the cold-start stage through a small, dedicated exploration budget. Rather than optimizing immediate cost efficiency or advertiser value, EBO is designed to proactively activate early-stage traffic and acquire initial feedback signals under strictly bounded spending.

EBO operates alongside existing cost-constrained bidding strategies and is designed to *complement rather than replace* them. During cold-start, EBO generates budget-feasible bids that encourage controlled traffic activation when feedback signals are scarce. As sufficient feedback is accumulated, the system naturally relies on the base cost-constrained strategy to optimize advertiser objectives under standard cost constraints, without requiring explicit switching or platform-side intervention.

At a conceptual level, EBO explicitly separates *signal acquisition* from *steady-state optimization* by introducing a dual-constraint integration mechanism. Specifically, EBO provides a lower-bound bid to ensure minimum traffic activation during cold-start, while the base bidding strategy continues to enforce cost-control guarantees. The final bid is selected as the maximum of the two bids, thereby activating early-stage traffic without violating the original cost-control constraints.

In addition, EBO supports personalized exploration budgets across different advertisers, campaigns, and creatives, allowing heterogeneous exploration intensity in diverse advertising scenarios. A feedback-driven budget tracking mechanism is further employed to regulate budget consumption over time, ensuring that the exploration budget is consumed smoothly and efficiently throughout the cold-start period.

Our contribution can be summarized as follows:

- We propose the Explicit Budget Optimization (EBO) module to address cold-start in cost-constrained advertising, ensuring early-stage signal acquisition without exceeding the allocated budget.
- We design a explicit exploration budget allocation mechanism that adapts to placement, time, and campaign characteristics, improving flexibility in heterogeneous scenarios.
- We introduce a budget-feasible bid generation mechanism, which efficiently consumes the allocated budget within the cold-start stage, and explain how the system guarantees early signal acquisition under cost constraints.

- We propose a dual-constraint bidding integration mechanism that introduces an explicit lower-bound bid for early-stage signal activation while preserving the original cost-control guarantees, enabling safe and efficient exploration under cold-start conditions.
- Extensive experiments conducted on Tencent's advertising platform demonstrate the effectiveness of EBO.

2 Related Work

Automated Bidding and Cost-Constrained Strategies Automated bidding (autobidding) has become the dominant paradigm in modern online advertising platforms (Maehara et al., 2018; Ren et al., 2018; Zhang et al., 2014). In autobidding, the platform automatically determines bids to optimize advertiser objectives under delivery and cost constraints. Early studies on real-time bidding and large-scale ad delivery systems established autobidding as the foundation of modern advertising (Grislain et al., 2019; Zhang et al., 2016). Unlike manual bidding, these systems adjust bids dynamically using historical feedback, predictive models, and control mechanisms (Doyle et al., 2013; Zhang et al., 2022; Grislain et al., 2019; Ren et al., 2018), enabling scalable campaign management across heterogeneous advertisers.

Cost-constrained bidding is central to ensuring efficient and safe delivery. A dominant line of work adopts control-theoretic methods such as Proportional–Integral–Derivative (PID) controllers (Doyle et al., 2013) and Model Predictive Control (MPC) (Zhang et al., 2022). These approaches regulate bids by tracking deviations between realized and target cost-value trajectories, often incorporating short-term predictions of candidate bids' effects. Reinforcement learning (RL) has also been explored for cost-constrained bidding, optimizing long-term objectives under budget or ROI constraints (Zhao et al., 2018; Cai et al., 2017; Ye et al., 2020). Despite methodological differences, both control-based and RL-based approaches rely on informative feedback signals. When signals are sparse or unreliable, such as during cold-start, these methods tend to behave conservatively, suppressing bids and limiting traffic activation.

Cold-Start Challenges in Online Advertising Cold-start arises for new advertisers, campaigns, or creatives with limited historical data (Hillard et al., 2010; Pan et al., 2019; Zhu et al., 2025). Prior work has focused on improving prediction under data scarcity, using representation learning, transfer learning, or auxiliary information (Pan et al., 2019; Zhu et al., 2025). Industry platforms often use heuristics to boost new ads, such as scaling eCPM or allocating extra budget. While effective in acquiring early signals, these methods require additional platform intervention and are mainly designed for new ads, limiting their general applicability. In contrast, feedback-driven bidding systems can also underperform for existing campaigns under strict cost constraints and severe signal sparsity, suppressing bids and preventing traffic activation. Our method addresses this by applying a lightweight adjustment to the advertiser's bidding strategy during cold-start. This approach is safe, requires no extra investment or ranking-side changes, and is applicable to all ads, enabling faster signal acquisition and improved overall performance.

3 Method

3.1 Problem Formulation

We consider an online advertising system in which the platform bids on behalf of advertisers (or campaigns), indexed by entity u. At each discrete time step t, the system selects a bid $a_{u,t}$ based on both historical observations and predicted outcomes. The bidding decision aims to balance accumulated cost–value deviations and the predicted cost and value induced by the current bid.

Specifically, $\mathrm{Cost}_{u,t}$ denotes the realized advertising cost (i.e., spending) incurred by entity u at time t, and $\mathrm{Value}_{u,t}$ denotes the advertiser-defined outcome, such as conversions or revenue. The cumulative cost and value up to time t are given by $\sum_{i=1}^{t} \mathrm{Cost}_{u,i}$ and $\sum_{i=1}^{t} \mathrm{Value}_{u,i}$, respectively.

To support forward-looking decision making, the system relies on a prediction component, denoted as `Bid2X`, which estimates the expected cost and value induced by a candidate bid. Given a bid $a_{u,t}$, `Bid2X` outputs the predicted next-step outcomes $\hat{\mathrm{Cost}}_{u,t+1}(a_{u,t})$ and $\hat{\mathrm{Value}}_{u,t+1}(a_{u,t})$. In this work, `Bid2X` is treated as a black-box predictor, and our method does not depend on its internal modeling assumptions.

In this work, we focus on the *cold-start stage* of an advertising campaign, where meaningful feedback signals are scarce or entirely absent. Specifically, feedback signals refer to advertiser-defined value outcomes (e.g., conversions or revenue) and the associated cost realizations that are used to calibrate bidding decisions. During cold-start, historical accumulations of cost and value are insufficient to reliably estimate cost–value dynamics, and predictive components such as `Bid2X` lack informative input context, resulting in highly uncertain or biased estimates. Under this condition, existing cost-constrained bidding frameworks tend to behave conservatively: to avoid potential overspending under uncertainty, the system suppresses bids, leading to low traffic exposure and delayed value acquisition. This conservative response further reduces the chance of observing new feedback signals, forming a self-reinforcing loop of low spending, sparse data, and stalled optimization. Consequently, the cold-start stage poses a fundamental challenge that cannot be resolved by improving prediction accuracy alone, but requires an explicit mechanism to ensure early-stage signal acquisition under strict cost constraints.

3.2 Explicit Budget Optimization Module

To address the lack of feedback signals during the cold-start stage, we introduce a module called *Explicit Budget Optimization* (EBO). Unlike standard cost-constrained bidding strategies that aim to optimize cost efficiency or advertiser value, EBO is designed specifically for the cold-start phase, where the primary objective is to activate traffic and acquire initial feedback signals. To this end, EBO allocates a small, dedicated exploration budget B exclusively for early-stage exploration and focuses on spending this budget in a controlled manner, without considering immediate cost efficiency. The overall structure of the EBO module is illustrated in Figure 1.

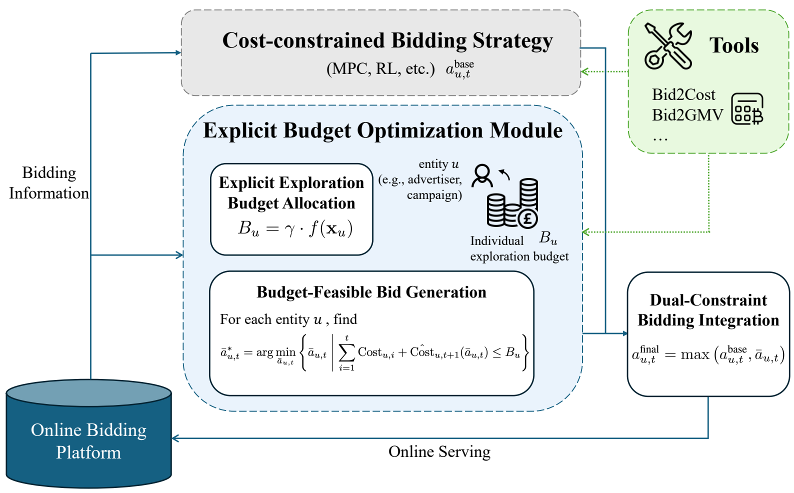

Figure 1: The Structure of Explicit Budget Optimization Module

Explicit Exploration Budget Allocation The exploration budget B is explicitly separated from the main campaign budget and is only active during the cold-start period. This budget serves as a safeguard that enables proactive signal acquisition while maintaining strict upper bounds on total spending. Once the exploration budget is exhausted, the system fully relies on standard cost-constrained bidding strategies.

In practice, the severity of cold-start and signal sparsity varies across advertisers, campaigns, and creatives. To accommodate such heterogeneity, we design EBO to support personalized exploration budgets. Each entity u is assigned an individualized exploration budget B_u:

$$B_u = \gamma \cdot f(\mathbf{x}_u), \tag{1}$$

where $\mathbf{x}_u$ represents entity-specific features available at deployment time, $f(\cdot)$ is a budget scaling function, and γ controls the overall exploration intensity.

Budget-Feasible Bid Generation During the cold-start stage, feedback signals such as clicks or conversions are extremely sparse, making it difficult to directly optimize value-based objectives. Instead, EBO focuses on activating sufficient traffic under a strict exploration budget. We therefore formalize the bid selection process of EBO as a budget-feasibility problem for each entity u (e.g., advertiser, campaign, or creative). Let $\bar{a}_{u,t}$ denote the bid suggested by the EBO module for entity u at time t, and let $\mathrm{Cost}_{u,i}$ denote the realized cost for entity u at time i. EBO aims to identify the **minimum feasible bid** that satisfies the time-varying budget constraint. Using overly aggressive bids may quickly exhaust the exploration budget on low-quality traffic before sufficient learning signals are obtained, while excessively conservative bids slow down feedback acquisition. Among all feasible bids, EBO therefore prefers the smallest bid that keeps the cumulative spending within the exploration budget.

This perspective can be expressed through the following constrained formulation for a given entity u:

$$\bar{a}^*_{u,t} = \arg\min_{\bar{a}_{u,t}} \left\{ \bar{a}_{u,t} \;\middle|\; \sum_{i=1}^{t} \mathrm{Cost}_{u,i} + \hat{\mathrm{Cost}}_{u,t+1}(\bar{a}_{u,t}) \le B_u \right\}. \tag{2}$$

where $\hat{\mathrm{Cost}}_{u,t+1}(\bar{a}_{u,t})$ is the predicted cost associated with bid $\bar{a}_{u,t}$ for entity u, provided by a black-box prediction component such as `Bid2X`. This formulation ensures that each EBO bid respects both the historical spending and predicted future expenditure for the specific entity u.

To efficiently and stably consume the exploration budget, EBO adjusts its bid through a feedback-driven budget tracking mechanism. Rather than reacting directly to the remaining budget, EBO compares the *expected cumulative cost* with the *realized cumulative cost* and regulates bids based on their discrepancy.

Let $\mathrm{Cost}_{u,i}$ denote the realized cost for entity u at time i, and let $\mathrm{ECost}_{u,t}$ denote the expected cumulative cost from the exploration start time to time t, estimated from global traffic consumption patterns. The realized cumulative cost is given by $\sum_{i=1}^{t} \mathrm{Cost}_{u,i}$. EBO computes a normalized budget tracking signal:

$$\Delta_{u,t} = B_u \cdot \frac{\mathrm{ECost}_{u,t}}{\mathrm{ECost}_{u,\text{total}}} - \sum_{i=1}^{t} \mathrm{Cost}_{u,i}, \tag{3}$$

which measures whether the current spending lags behind or exceeds the expected budget consumption trajectory.

Based on this tracking signal, EBO determines a bounded bid adjustment factor:

$$\mathrm{bid\_step}_{u,t} = \alpha \cdot \tanh(\Delta_{u,t}), \tag{4}$$

where $\alpha \in [0, 1]$ controls the maximum adjustment magnitude and the hyperbolic tangent enforces a smooth bound $\mathrm{bid\_step}_{u,t} \in [-\alpha, \alpha]$ to ensure stable bid updates.

The EBO bid is then updated multiplicatively as

$$\bar{a}_{u,t} = a^{\text{init}}_{u,t} \cdot \left(1 + \mathrm{bid\_step}_{u,t}\right), \tag{5}$$

where $a^{\text{init}}_{u,t}$ denotes the initial bid suggestion at time t.

When $\mathrm{ECost}_{u,t} > \sum_{i=1}^{t} \mathrm{Cost}_{u,i}$, the bid adjustment is positive, encouraging more aggressive bidding to accelerate budget consumption. Conversely, when realized spending exceeds the expected trajectory, the bid adjustment becomes negative, suppressing future bids to prevent premature budget exhaustion.

Dual-Constraint Bidding Integration For safe online deployment, EBO operates in parallel with an existing cost-constrained bidding strategy. At each bidding opportunity, the final bid submitted to the auction is computed as:

$$a^{\text{final}}_{u,t} = \max\left(a^{\text{base}}_{u,t}, \bar{a}_{u,t}\right), \tag{6}$$

where a^{base}_{t} denotes the bid suggested by the standard cost-constrained strategy. This *maximum* operator induces a dual-constraint mechanism:

- The EBO bid $\bar{a}_{u,t}$ enforces a *minimum activation constraint*, ensuring that sufficient traffic is obtained to generate feedback signals during cold-start.
- The base bid $a_{u,t}^{\text{base}}$ enforces a *maximum cost constraint*, preserving the original cost-control guarantees of the system.

Thus, the dual-constraint mechanism ensures that the exploration budget is spent efficiently and safely, while maintaining the original cost-control guarantees. Conceptually, EBO provides a lower-bound bid to activate early-stage signals, and the final bid selection ensures proper integration with existing bidding strategies without violating budget constraints.

Algorithmic Summary The workflow of EBO is summarized in Algorithm 1, which computes a feedback-adjusted bid, enforces the maximum budget constraint, combines it with the base bid, and submits the final bid to the auction.

Algorithm 1: Explicit Budget Optimization (EBO) for Entity u

Input: Initial bid $a_{u,t}^{\text{init}}$; Exploration budget B_u; Traffic consumption map M; Hyperparameter $\alpha \in [0, 1]$

Output: Adjusted EBO bid $\bar{a}_{u,t}$

Step 1: Traffic Trend Modeling Construct a global traffic consumption map M, where each minute m of the day is mapped to an expected consumption rate $\mathrm{M}[m]$. This map is shared across entities and captures the temporal traffic pattern of the platform.

Step 2: Budget Consumption Estimation for Entity u

(a) Realized cost. Compute the realized cost of entity u from the campaign start to time t:

$$\text{real\_cost}_{u,t} \leftarrow \sum_{i=1}^{t} \text{Cost}_{u,i}.$$

(b) Expected cost. Let ECost_t denote the cumulative expected consumption rate from campaign start to time t, and $\text{ECost}_{\text{total}}$ denote the expected total consumption rate over the entire exploration horizon. The expected cost of entity u up to time t is computed as:

$$\text{expect\_cost}_{u,t} \leftarrow B_u \cdot \frac{\text{ECost}_{u,t}}{\text{ECost}_{u,\text{total}}}.$$

Step 3: Feedback-Driven Bid Adjustment

(a) Compute bid adjustment step. Calculate the normalized budget deviation for entity u:

$$\Delta_{u,t} = \text{expect\_cost}_{u,t} - \text{real\_cost}_{u,t}.$$

Apply a bounded nonlinearity to obtain the bid adjustment step:

$$\text{bid\_step}_{u,t} = \alpha \cdot \tanh(\Delta_{u,t}),$$

(b) Update bid. Adjust the EBO bid for entity u:

$$\bar{a}_{u,t} = a_{u,t}^{\text{init}} \cdot \left(1 + \text{bid\_step}_{u,t}\right).$$

return $\bar{a}_{u,t}$

4 Intuition for Budget-Feasible Bid Generation

To better understand the behavior of the feedback-driven bid adjustment used in EBO, we provide an intuitive interpretation of the exploration process during the cold-start stage. As described in Section 3.2, EBO aims to identify a bid that activates sufficient traffic while ensuring that the cumulative spending does not exceed the exploration budget. This objective can be interpreted through the lens of constrained optimization, where the bid must satisfy a budget-feasibility constraint.

Introducing a Lagrange multiplier $\lambda_{u,t} \geq 0$ for the budget constraint, the corresponding Lagrangian can be written as

$$\mathrm{L}(\bar{a}_{u,t}, \lambda_{u,t}) = \bar{a}_{u,t} + \lambda_{u,t} \left(\sum_{i=1}^{t} \mathrm{Cost}_{u,i} + \hat{\mathrm{Cost}}_{u,t+1}(\bar{a}_{u,t}) - B_u \right). \tag{7}$$

The Karush–Kuhn–Tucker (KKT) conditions provide useful intuition for understanding the behavior of the budget-constrained bid selection process. In particular:

- **Primal feasibility:** ensures that the selected bid $\bar{a}_{u,t}$ never causes the cumulative cost for entity u to exceed the allocated exploration budget, i.e.,

$$\sum_{i=1}^{t} \mathrm{Cost}_{u,i} + \hat{\mathrm{Cost}}_{u,t+1}(\bar{a}_{u,t}) \leq B_u.$$

- **Dual feasibility:** the Lagrange multiplier $\lambda_{u,t} \geq 0$ can be interpreted as representing the *budget pressure* for entity u, which increases as the remaining exploration budget decreases.
- **Complementary slackness:**

$$\lambda_{u,t} \left(\sum_{i=1}^{t} \mathrm{Cost}_{u,i} + \hat{\mathrm{Cost}}_{u,t+1}(\bar{a}_{u,t}) - B_u \right) = 0.$$

  When the budget constraint is not binding, $\lambda_{u,t} = 0$ and the bid can decrease freely. When the budget becomes tight, $\lambda_{u,t} > 0$ prevents further bid reduction, ensuring budget safety.
- **Stationarity:** the optimal solution balances the objective of selecting a conservative bid with the requirement that the exploration budget constraint remains satisfied.

While the EBO update rule is not derived directly from the KKT conditions, it can be interpreted as a practical approximation of a budget-aware adjustment mechanism. In particular, the feedback-driven update

$$\bar{a}_{u,t} = \bar{a}_{u,t}^{\text{init}} \cdot \big(1 + \alpha \cdot \tanh(B_{u,\text{residual}})\big) \tag{8}$$

adjusts the bid based on the deviation between the expected and realized budget consumption. Here,

$$B_{u,\text{residual}} = B_u \cdot \frac{\mathrm{ECost}_{u,t}}{\mathrm{ECost}_{u,\text{total}}} - \sum_{i=1}^{t} \mathrm{Cost}_{u,i}$$

measures whether the current spending is ahead of or behind the expected budget trajectory.

Intuitively, this residual signal plays a role analogous to a budget pressure term: when the realized spending lags behind the expected trajectory, the adjustment increases the bid to accelerate traffic acquisition; when spending is ahead of schedule, the adjustment suppresses the bid to avoid premature budget exhaustion. The $\tanh(\cdot)$ function further stabilizes the update by bounding the adjustment magnitude.

Therefore, the proposed feedback-driven update can be viewed as a practical and stable mechanism for performing budget-aware bid adjustment during cold-start exploration.

5 Experiments

5.1 Experimental Setup

We evaluate the proposed Explicit Budget Optimization (EBO) module through large-scale online A/B experiments conducted on Tencent commercial advertising platform. The experiments cover multiple business scenarios with different advertiser objectives, including short-drama monetization, mobile game monetization, and small-shop promotion.

In all experiments, EBO is deployed as an auxiliary module alongside the existing cost-constrained bidding strategy. The control group uses the base bidding strategy only, while the treatment group enables EBO during the cold-start stage with a small dedicated exploration budget. All other components of the bidding and delivery system remain identical across groups.

5.2 Evaluation Metrics

We evaluate the effectiveness of the proposed Explicit Budget Optimization (EBO) module using metrics that jointly capture traffic activation, budget utilization, delivery quality, and platform-side revenue. Let AD denote the set of ads during the evaluation window, and $\mathrm{AD}_{\text{new}} \subset \mathrm{AD}$ denote newly launched ads in the cold-start stage. For each ad $ad \in \mathrm{AD}$, let $Cost(ad)$ and $Value(ad)$ denote its cumulative advertising cost and advertiser-defined value (e.g., conversions or revenue), and let $Budget(ad)$ denote its allocated budget.

- **Average Cost per Ad.** Measures the average spending per ad:

$$\mathrm{AvgCost} = \frac{\sum_{ad \in \mathrm{AD}} Cost(ad)}{|\mathrm{AD}|}. \tag{9}$$

- **Pricing Rate.** Evaluates delivery quality under cost constraints:

$$\mathrm{PR}(ad) = \frac{Cost(ad)}{Value(ad)}. \tag{10}$$

  A pricing rate close to 1 indicates good alignment between cost and advertiser-defined value. Following industrial practice, ads with average pricing rate in the range $[0.8, 1.2]$ are considered successfully achieved, denotes as AD_{ach}:

$$\mathrm{AD}_{\text{ach}} = \{ad \in \mathrm{AD} \mid 0.8 \le Cost(ad)/Value(ad) \le 1.2\}. \tag{11}$$

- **Cold-Start Activation Rate.** Captures early-stage traffic activation. For a time window τ, it is defined as

$$\mathrm{ActivationRate}(\tau) = \frac{1}{|\mathrm{AD}|} \sum_{ad \in \mathrm{AD}} \mathrm{I}\left[Cost_{[0,\tau]}(ad) \ge \kappa\right], \tag{12}$$

  where $Cost_{[0,\tau]}(ad)$ is the cumulative cost incurred by ad ad within the first τ hours after launch, and κ is a predefined minimum spending threshold (e.g., 100 RMB). This metric can be reported for $\tau = 24$h and $\tau = 72$h.

- **Achieved Spend Ratio.** Measures the fraction of total spend allocated to successfully achieved ads (pricing rate within $[0.8, 1.2]$):

$$\mathrm{ASR} = \frac{\sum_{ad \in \mathrm{AD}_{\text{ach}}} Cost(ad)}{\sum_{ad \in \mathrm{AD}} Cost(ad)}, \tag{13}$$

- **Budget Utilization Rate.** Captures how efficiently allocated budgets are consumed:

$$\mathrm{BUR}(ad) = \frac{Cost(ad)}{Budget(ad)}. \tag{14}$$

5.3 Results

We evaluate the effectiveness of the Explicit Budget Optimization (EBO) module across three business scenarios: short-drama monetization, mobile game monetization, and small-shop promotion. In each scenario, A/B experiments are conducted by enabling EBO for a fraction of ads while keeping the remainder as control: 50% of ads in short-drama monetization, 50% in mobile game monetization, and 30% in small-shop promotion are assigned to the EBO treatment group. The results are reported in Tables 1 and 2, focusing on key performance metrics for both overall ads and newly launched ads.

Table 1 presents the incremental impact of EBO on cost and delivery metrics. Across all scenarios, the module increases average cost per ad, particularly for new ads, indicating that EBO successfully activates traffic during the cold-start stage and brings additional revenue to the platform (+21.43% for new ads in short-drama monetization, +17.00% for new ads in mobile game monetization, and +6.50% for new ads in small-shop promotion). Achieved Spend Ratio (ASR) is also improved for new ads in short-drama and small-shop campaigns (+3.87% and +2.18%, respectively), showing that more ads are successfully delivered and contributing to increased total spending. Pricing rate changes are reported as the movement from pre-EBO to post-EBO values (1.11 $\rightarrow$ 1.09 for new ads in short-drama, 1.20 $\rightarrow$ 1.34 for new ads in mobile games), illustrating that bid adjustments remain within acceptable ranges and better align cost with value, leading to more efficient bidding.

Table 1: Incremental impact of EBO on cost and delivery performance

Metric	Short-Drama Monetization		Mobile Game Monetization		Small-Shop Promotion	
	All Ads	New Ads	All Ads	New Ads	All Ads	New Ads
Average Cost per Ad (% Change)	+15.87	+21.43	+12.14	+17.00	+15.37	+6.50
Achieved Spend Ratio (% Change)	+1.26	+3.87	-	-	+2.18	-
Average Pricing Rate (from $\rightarrow$ to)	-	1.11 $\rightarrow$ 1.09	1.05 $\rightarrow$ 1.03	1.20 $\rightarrow$ 1.34	1.06 $\rightarrow$ 1.05	-

Table 2 summarizes cold-start activation and budget utilization metrics. The 24-hour activation rate for new ads exceeds 20% in all three business scenarios (+20.60% for short-drama monetization, +48.65% for mobile game monetization, and +33.82% for small-shop promotion), showing substantial improvement and demonstrating rapid acquisition of early-stage feedback signals. The 72-hour activation rate further confirms that EBO accelerates campaign ramp-up in the longer term (+17.64% for mobile game monetization and small-shop promotion). Budget Utilization Rate also increases notably in short-drama (+22.80%) and mobile game (+35.87%) scenarios, indicating more efficient consumption of allocated budgets while maintaining strict cost compliance.

Table 2: Incremental impact of EBO on cold-start activation and budget utilization

Metric	Short-Drama Monetization	Mobile Game Monetization	Small-Shop Promotion
Cold-Start Activation Rate 24h (% Change)	+20.60	+48.65	+33.82
Cold-Start Activation Rate 72h (% Change)	-	+17.64	+17.64
Budget Utilization Rate (% Change)	+22.80	+35.87	-

These results validate the effectiveness of EBO in achieving early-stage traffic activation, maintaining budget compliance, and supporting subsequent cost-constrained bidding strategies. The consistent gains for new ads highlight its particular benefit in cold-start conditions, while the accelerated activation of ad traffic also contributes positively to overall metrics, indicating that EBO improves campaign performance beyond just the newly launched ads.

Across all scenarios, EBO consistently improves early-stage traffic activation and budget utilization, while maintaining stable pricing and risk-control behavior. The results confirm that explicitly allocating and managing a small exploration budget is an effective way to address cold-start challenges in cost-constrained bidding systems.

Notably, EBO does not aim to optimize short-term ROI during the cold-start stage. Instead, it prioritizes signal acquisition under strict cost constraints, enabling downstream bidding strategies to function more effectively once sufficient feedback becomes available.

6 Conclusion

This work addresses a critical limitation of cost-constrained bidding systems during the cold-start stage. By explicitly separating early-stage signal acquisition from steady-state optimization, the proposed Explicit Budget Optimization (EBO) module enables safe and controlled exploration under strict budget constraints. EBO complements existing bidding strategies, adaptively allocates exploration budget, and generates budget-aware lower-bound bids to activate initial traffic. Extensive experiments, including online deployment on Tencent's advertising platform, demonstrate that EBO effectively accelerates feedback signal collection and improves overall campaign performance, highlighting its practical value in real-world advertising systems.

REFERENCES

Christian Borgs, Jennifer Chayes, Nicole Immorlica, Kamal Jain, Omid Etesami, and Mohammad Mahdian. Dynamics of bid optimization in online advertisement auctions. In *Proceedings of the 16th International Conference on World Wide Web*, WWW '07, pp. 531–540, New York, NY, USA, 2007. Association for Computing Machinery. ISBN 9781595936547. doi: 10.1145/1242572.1242644. URL `https://doi.org/10.1145/1242572.1242644`.

Han Cai, Kan Ren, Weinan Zhang, Kleanthis Malialis, Jun Wang, Yong Yu, and Defeng Guo. Real-time bidding by reinforcement learning in display advertising. In *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining*, WSDM 2017, pp. 661–670. ACM, February 2017. doi: 10.1145/3018661.3018702. URL `http://dx.doi.org/10.1145/3018661.3018702`.

John C Doyle, Bruce A Francis, and Allen R Tannenbaum. *Feedback control theory*. Courier Corporation, 2013.

David S. Evans. The online advertising industry: Economics, evolution, and privacy. *Consumer Law eJournal*, 2009. URL `https://api.semanticscholar.org/CorpusID:154745950`.

Nicolas Grislain, Nicolas Perrin, and Antoine Thabault. Recurrent neural networks for stochastic control in real-time bidding. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '19, pp. 2801–2809. ACM, July 2019. doi: 10.1145/3292500.3330749. URL `http://dx.doi.org/10.1145/3292500.3330749`.

Dustin Hillard, Stefan Schroedl, Eren Manavoglu, Hema Raghavan, and Chirs Leggetter. Improving ad relevance in sponsored search. In *Proceedings of the Third ACM International Conference on Web Search and Data Mining*, WSDM '10, pp. 361–370, New York, NY, USA, 2010. Association for Computing Machinery. ISBN 9781605588896. doi: 10.1145/1718487.1718532. URL `https://doi.org/10.1145/1718487.1718532`.

Takanori Maehara, Atsuhiro Narita, Jun Baba, and Takayuki Kawabata. Optimal bidding strategy for brand advertising. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence*, IJCAI'18, pp. 424–432. AAAI Press, 2018. ISBN 9780999241127.

Feiyang Pan, Shuokai Li, Xiang Ao, Pingzhong Tang, and Qing He. Warm up cold-start advertisements: Improving ctr predictions via learning to learn id embeddings, 2019. URL `https://arxiv.org/abs/1904.11547`.

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, April 2018. ISSN 1041-4347. doi: 10.1109/tkde.2017.2775228. URL `http://dx.doi.org/10.1109/TKDE.2017.2775228`.

Yujian Ye, Dawei Qiu, Mingyang Sun, Dimitrios Papadaskalopoulos, and Goran Strbac. Deep reinforcement learning for strategic bidding in electricity markets. *IEEE Transactions on Smart Grid*, 11(2):1343–1355, 2020. doi: 10.1109/TSG.2019.2936142.

Haoqi Zhang, Junqi Jin, Zhenzhe Zheng, Fan Wu, Haiyang Xu, and Jian Xu. Control-based bidding for mobile livestreaming ads with exposure guarantee. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*, CIKM '22, pp. 2539–2548, New York, NY, USA, 2022. Association for Computing Machinery. ISBN 9781450392365. doi: 10.1145/3511808.3557269. URL `https://doi.org/10.1145/3511808.3557269`.

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, pp. 1077–1086, New York, NY, USA, 2014. Association for Computing Machinery. ISBN 9781450329569. doi: 10.1145/2623330.2623633. URL `https://doi.org/10.1145/2623330.2623633`.

Weinan Zhang, Yifei Rong, Jun Wang, Tianchi Zhu, and Xiaofan Wang. Feedback control of real-time display advertising. In *Proceedings of the Ninth ACM International Conference on Web Search and Data Mining*, WSDM 2016, pp. 407–416. ACM, February 2016. doi: 10.1145/2835776.2835843. URL http://dx.doi.org/10.1145/2835776.2835843.

Jun Zhao, Guang Qiu, Ziyu Guan, Wei Zhao, and Xiaofei He. Deep reinforcement learning for sponsored search real-time bidding, 2018. URL https://arxiv.org/abs/1803.00259.

Wenqiao Zhu, Lulu Wang, and Jun Wu. Addressing cold-start problem in click-through rate prediction via supervised diffusion modeling, 2025. URL https://arxiv.org/abs/2504.06270.