Dynamic Bidding for Sponsored Products: Value-Based Bid Scaling with Per-SKU Normalization

Wanchen Shao
Wayfair LLC
Mountain View, CA, USA
wshao@wayfair.com

Aditya Gadam
Wayfair LLC
Seattle, WA, USA
sgadam@wayfair.com

Dongyue Xie
Wayfair LLC
Mountain View, CA, USA
dxie@wayfair.com

Kurt Zimmer
Wayfair LLC
Boston, MA, USA
kzimmer@wayfair.com

Matt Lindsay
Wayfair LLC
Boston, MA, USA
mlindsay@wayfair.com

Sergey Kolbin
Wayfair LLC
Seattle, WA, USA
mlindsay@wayfair.com

Jinjin Zhao
Wayfair LLC
Seattle, WA, USA
jzhao5@wayfair.com

Jean-David Ruvini
Wayfair LLC
Mountain View, CA, USA
jruvini@wayfair.com

Abstract

Sponsored-product advertising on e-commerce platforms typically uses static, per-campaign click bids that ignore the variable value of ad slots across contexts such as user intent, device, placement, and time. We present a production dynamic bidding system that adjusts CPC (Cost-Per-Click) bids in real time based on predicted post-click conversion probability.

The system is built around three design choices: (i) it models post-click conversion probability, aligning the prediction target with the CPC charging event; (ii) it normalizes predictions at the campaign-SKU level, preserving each advertiser's spending intent as the average operating point; (iii) it calibrates heterogeneous placement-specific signal to a shared CTR interface, enabling a single bidder to serve across incompatible surfaces. We provide a step-by-step derivation of the bid multiplier equation, detail the complete system architecture (model, feature pipeline, low-latency serving, and normalization design), and present results from production deployment on a large home-furnishings marketplace.

Production experiments show positive and statistically significant gains in attributed conversions per click, first on the initial placements and later after expansion to additional placement types, with neutral customer-level conversion impact, further supporting the effectiveness of the approach.

CCS Concepts

• **Information systems → Online advertising**.

Keywords

Sponsored products; dynamic bidding; enhanced CPC; value-based bidding; conversion prediction; bid normalization

1 Introduction

E-commerce platforms run multi-slot CPC auctions for sponsored products, where advertisers typically set a single fixed bid per campaign. Static bids do not adapt to context—time of day, device, user intent, placement—and therefore misallocate spend because the value of each impression varies. This creates two inefficiencies relative to an ideal auction where each advertiser bids their true expected value per insertion [24]: (1) advertisers cannot bid for value directly, only for clicks, yet they value *sales*, not clicks; (2) some clicks are far more likely to convert than others, but fixed bids treat all clicks equally.

Dynamic bidding addresses both problems by scaling the advertiser's click bid with a predicted conversion-value multiplier, so that high-value impressions receive higher bids and low-value ones receive lower bids—all within a bounded range. The idea is widely recognized and aligns with Amazon's "Dynamic Bids – Up and Down" [3], which can raise or lower bids by up to 100%, and with Google's Enhanced CPC, which likewise relies on predicted conversion likelihood. The challenge lies in 1) scaling quickly to new ad opportunities when an expansion in the ad footprint is expected—particularly in environments with heterogeneous auction surfaces, where different pages and placements use distinct scoring methods; 2) advertisers express spending intent through campaign-SKU default bids that the system must preserve on average, where SKU (stock keeping unit) is the product-level advertising unit on which suppliers configure campaign bids.

Contributions. This paper makes four contributions. First, we formulate dynamic bidding in sponsored products as a contextual repricing around the campaign-SKU default bids. This framing constrains deviation from the supplier's default bid while allowing auction-time bid adaptation. Second, we show that the normalization granularity must match the unit at which advertiser intent is expressed. In our marketplace, this implies campaign-SKU normalization; coarser levels such as campaign-level or global averages distort spend allocation. Third, we present a production architecture for cross-surface deployment in which heterogeneous placement-specific upstream scores are aligned to a shared probabilistic score and consumed by a single downstream bidder. This makes the bidder reusable with different feature richness, scoring semantics, and latency budgets. Fourth, we provide empirical evidence from production deployment, including shadow evaluation and online A/B tests, showing that our method improves ads-attributed conversion efficiency.

2 Related Work

Our work connects to three areas: autobidding, post-click conversion prediction, and calibration. Autobidding frames the platform-side problem of translating advertiser goals or constraints into auction-time bids [1]. In CPC sponsored products, the closest line is bid adjustment around an advertiser-provided base bid: multiplicative bidding languages scale a base bid by feature-dependent adjustments [4], while bid-multiplier caps control targeting intensity [14]. Related constrained autobidding and optimized CPC systems refine coarse bidding inputs using predicted response under budget, CPC, or return-on-spend constraints [2, 10, 13, 29]. WDyB (Wayfair Dynamic Bidding) differs in that it does not infer bids from scratch; suppliers already specify campaign-SKU default bids, and the bidder performs contextual repricing while preserving those defaults as the average operating point. Sponsored-search auction theory provides the pricing foundation for this setting [9, 22].

On the prediction side, WDyB is closer to post-click conversion estimation than to generic CVR prediction. Prior work formulates post-click conversion as estimating $P(\text{conversion} \mid \text{click}, \text{context})$ [21] and describes industrial CVR pipelines for advertising systems [17, 18]. Production CVR modeling must address click-space training, sparse positive labels, and delayed feedback. ESMM models the impression→click→conversion cascade to reduce sample-selection bias and sparsity [19]; causal and propensity-based methods further debias post-click CVR estimation [27, 28]; and delayed-feedback methods improve calibration when conversions arrive late [5, 12]. For our heterogeneous tabular feature space, gradient-boosted trees remain a strong industrial baseline [7].

Prior work studies ranking and calibration for click-attributed purchases in performance advertising [6]. Calibration matters because bidding consumes predictions as economic quantities, not merely ranking scores. In WDyB, calibration plays a systems role: heterogeneous placement-specific upstream scores are aligned to a shared $pCTR$ interface so that one bidder can serve multiple surfaces, while separate calibrated $pCVR|click$ is consumed by downstream bidders. WDyB further introduced campaign-SKU normalizer, so the contribution is not a new calibrator but a production design showing how calibration and normalization coexist in the bidding stack.

Large-scale recommendation models suggest future extensions but are orthogonal to the present contribution. Wukong studies scaling laws using stacked factorization-machine-style interaction blocks [26], and sequential transducer architectures reformulate recommendation as large-scale sequence modeling [25]. These models may improve future user and session representations, while WDyB focuses on the current production bidding design: shared $pCTR$ alignment, post-click prediction, and campaign-SKU normalization around supplier default bids.

3 System Objective

We study dynamic bidding for multi-slot CPC sponsored search and sponsored product auctions. The goal is to increase qualified spend and expected conversion value at auction time while preserving the advertiser's default bid for each campaign–SKU pair as the operating point. Because SKUs within the same campaign can differ substantially in price, margin, purchase frequency, and user intent, the bidder must adjust bids by context without globally shifting spend or leaking budget across heterogeneous SKUs.

3.1 Constrained Optimization Objective

We formulate dynamic bidding as a constrained optimization problem anchored to the supplier's per-SKU default CPC bid. The objective is to use contextual conversion signals to maximize expected value, while preserving the supplier's intended CPC level in expectation. Because realized CPC is determined by GSP-like (Generalized-Second-Price) pricing rather than by the submitted bid alone, this constraint preserves supplier CPC intent while allowing context-sensitive bid adjustments.

We index campaigns by i, SKUs by j, and auction contexts by a containing page, request, user, query, placement, and time information. Let $\hat{p}_{i,j,a}$ denote the predicted probability of conversion given a click, and let π_j denote the bid-time estimate of supplier value from a sale. The expected value per click is

$$v_{i,j,a} = \pi_j \cdot \hat{p}_{i,j,a}. \tag{1}$$

Let $x_{i,j,a}(b)$ denote whether campaign–SKU pair (i, j) receives a click in auction context a under submitted bids b, and let $c_{i,j,a}(b)$ denote the corresponding expected CPC. Let W denote the measurement window. The objective is then to maximize expected value subject to the supplier-preserving constraint: the expected CPC should match the advertiser's default bid b^0_{ij} for each campaign–SKU pair over W:

$$\begin{aligned} \max_{b} \quad & \sum_{j \in \mathcal{J}_i} \sum_{a \in W} x_{i,j,a}(b)\, v_{i,j,a} \\ \text{s.t.} \quad & \frac{\sum_{a \in W} x_{i,j,a}(b)\, c_{i,j,a}(b)}{\sum_{a \in W} x_{i,j,a}(b)} = b^0_{ij}, \quad \forall j \in \mathcal{J}_i, \end{aligned} \tag{2}$$

Constraint (2) is the central design decision: it enforces that the expected charged CPC equals the advertiser's default bid. This constraint functions not only as a mathematical regularizer but also as a practical trust region in operation.

In mature systems, b^0_{ij} often already encodes advertiser willingness-to-pay together with historical compensation for auction discount. If a new bidder applies another correction without respecting this anchor, a legacy campaign may be effectively corrected twice.

Previous research on constrained autobidding has demonstrated that bids can be computed using Lagrangian or primal–dual update rules[2, 4, 10, 14]. Under assumptions of truthful auctions, these bids take on a multiplicative, value-scaled structure, where a dual variable or pair-specific multiplier translates per-click value into bid units. We adopt this structure as an approximation in production and additionally introduce a campaign–SKU normalizer, analogous to the dual variable, to ensure that the average bid matches the supplier default, as described in Section 3.2.

The auction-time bid for SKU j in campaign i at auction context a is therefore:

$$b_{i,j,a} = \lambda_{ij} \cdot v_{i,j,a} \tag{3}$$

where λ_{ij} is a campaign–SKU scaling parameter that maps per-click value into CPC bid units.

3.2 From Optimization to a Practical Bid Multiplier

Solving (2) exactly requires accurate win curves and price curves for every auction, which is expensive and brittle in a fast-moving marketplace. In production, we use a simpler approximation: choose the campaign–SKU scale λ_{ij} so that the average sub- mitted bid over the historical window remains anchored to the supplier's default max CPC bid:

$$\frac{1}{|W|}\sum_{a\in W} b_{i,j,a} = b^0_{ij}. \tag{4}$$

Substituting $b_{i,j,a} = \lambda_{ij} \cdot v_{i,j,a}$ into Eq. 4 gives

$$\lambda_{i,j}\left(\frac{1}{|W|}\sum_{a\in W} v_{i,j,a}\right) = b^0_{ij}. \tag{5}$$

Defining $\bar{v}^{(W)}_{i,j} = \frac{1}{|W|}\sum_{a\in W} v_{i,j,a}$, we obtain

$$\lambda_{ij} = \frac{b^0_{ij}}{\bar{v}^{(W)}_{i,j}}. \tag{6}$$

Substituting $\lambda_{i,j}$ gives the normalized bid rule:

$$b_{i,j,a} = b^0_{i,j}\frac{v_{i,j,a}}{\bar{v}^{(W)}_{i,j}}. \tag{7}$$

Thus, the supplier's default bid is preserved as the SKU-level anchor, while the auction-time multiplier raises or lowers the bid according to whether the current context is above or below the SKU's historical average value. Since $v_{i,j,a} = \pi_j \hat{p}_{i,j,a}$ and π_j is fixed for SKU j, the value term cancels so the value normalizer reduces to a post-click $pCVR$ normalizer:

$$\frac{v_{i,j,a}}{\bar{v}^{(W)}_{i,j}} = \frac{\pi_j \hat{p}_{i,j,a}}{\pi_j \bar{p}^{(W)}_{i,j}} = \frac{\hat{p}_{i,j,a}}{\bar{p}^{(W)}_{i,j}}, \qquad \bar{p}^{(W)}_{i,j} = \frac{1}{|W|}\sum_{a\in W}\hat{p}_{i,j,a}. \tag{8}$$

Therefore the bid multiplier is defined as Equation (9) which is the normalization step that converts the post-click $pCVR$ into a contextual bid multiplier. Intuitively, the ratio measures how strong the current auction looks relative to the usual quality of the same campaign–SKU pair. If the historical average $\bar{p}^{(W)}_{ij}$ is unavailable, for example, because the campaign–SKU pair is new or has sparse history, we fall back to a global average baseline computed over all campaign–SKU pairs in the same historical window. This preserves the same normalization logic while avoiding undefined multipliers for cold-start pairs.

$$m^{\text{raw}}_{i,j,a} = \frac{\pi_j \hat{p}_{i,j,a}}{\pi_j \bar{p}^{(W)}_{i,j}}, \qquad \tilde{p}_{i,j} = \begin{cases} \bar{p}^{(W)}_{i,j}, & \text{if } \bar{p}_{i,j} \text{ is available,} \\ \bar{p}_{\text{global}}, & \text{otherwise.} \end{cases} \tag{9}$$

In production, we apply clipping to the raw multiplier in (9) to mitigate the risk of excessive overbidding or underbidding. The final auction-time bid is then

$$b_{i,j,a} = b^0_{i,j} \cdot \text{clip}\left(\frac{\hat{p}^{i,j,a}}{\tilde{p}_{i,j}}, \alpha_{\min}, \alpha_{\max}\right), \tag{10}$$

This rule is attractive because it preserves the advertiser's bid surface in expectation while making context-sensitive adjustments at auction time. If the current context is stronger than the pair's own recent average, the multiplier exceeds one and bids up. If the context is weaker, the multiplier falls below one and bids down. Under stable calibration and stationarity, the multiplier is approximately mean-preserving:

$$\mathbb{E}\left[m^{\text{raw}}_{ija} \mid i, j\right] \approx 1. \tag{11}$$

Thus, the bidder reallocates bids across contexts without systematically distorting the long-run per-pair bid level.

3.3 Why Campaign–SKU Normalization

In principle, λ_{ij} could be updated with a feedback controller, as in prior bid-control systems [23]. In our setting, however, this approach is poorly suited to campaign–SKU-level dynamic bidding, due to characteristics of the low-frequency, high-order-value furniture market: conversion-related feedback is infrequent, arrives with long delays, and is extremely noisy, while the realized CPC emerges from a stochastic auction process influenced by competitor actions and discounting. Consequently, a controller operating at a highly granular campaign–SKU level would chase transient noise instead of reflecting stable contextual quality. Our problem is therefore better cast as feed-forward contextual repricing around a fixed default bid, not as closed-loop outcome control at the pair level. Feedback control remains useful at coarser levels, such as campaign tROAS bidding and budgeting, but auction-time bid adjustment at campaign–SKU granularity is more stable when driven by model predictions and pair-specific normalization.

The normalization granularity should match the economic object that the bidder is designed to preserve. In our setting, the advertiser's default bid b^0_{ij} is defined at the campaign–SKU level, and the constraint in Eq. (2) penalizes deviation between the expected charged CPC and that same pair-specific default bid.

A coarser denominator, such as a campaign-level baseline, would mix heterogeneous SKU and create cross-SKU leakage. Strong SKUs would raise the baseline for weak SKUs, while weak SKUs would inherit an anchor not implied by their own default bids. In that case, the multiplier would no longer represent a pure contextual correction to b^0_{ij}, and the economic meaning of the default-gap constraint would be weakened. Campaign–SKU normalization is therefore the natural choice because it aligns the adjustment rule with the unit at which advertiser intent is expressed.

3.4 System Design

3.4.1 Conversion Model. The post-click $pCVR$ model is trained on logged impression–click–order data under a fixed attribution definition. It is designed to capture contextual conversion variation, not only item identity. The main feature groups are auction context, customer demographics, customer engagement, product context, and product relevance.

3.4.2 Scalable Model Layer. WDyB exposes a unified calibrated $pCTR$ interface to the bidder while allowing different upstream $pCTR$ models by surface. High-volume, feature-rich surfaces can use deeper models, while sparse or latency-sensitive surfaces can use lighter models. Let $s = \sigma(a)$ denote the page or surface:

$$f_s = \begin{cases} f^{\text{DNN}}_s, & s \in \mathcal{S}_{\text{rich}}, \\ f^{\text{LW}}_s, & s \in \mathcal{S}_{\text{tail}}, \end{cases} \qquad \hat{p}^{\text{ctr}}_{ija} = g_s\big(f_s(x_{ija})\big), \tag{12}$$

where g_s maps surface-specific raw model scores to calibrated CTR probabilities on a shared scale. The calibrated $pCTR$ is then used as an input to the post-click $pCVR$ model:

$$\hat{p}^{\mathrm{cvr}}_{i,j,a} = g^{\mathrm{cvr}}\Big(f^{\mathrm{cvr}}(x_{i,j,a}, \hat{p}^{\mathrm{ctr}}_{i,j,a})\Big)\,. \tag{13}$$

This design keeps the bidding logic independent of surface-specific model choices. As long as each upstream model produces calibrated $pCTR$ under a shared label definition, the same downstream bidder can serve heterogeneous placements without modification.

3.4.3 Auction-Time Inference and Serving. At request time, the bidder retrieves the campaign-SKU default bid $b^0_{i,j}$, computes calibrated $pCTR$, obtains post-click $pCVR$, looks up the campaign-SKU baseline $\bar{p}^{(W)}_{i,j}$, applies the clipped multiplier, and emits the final bid from Eq. (10). Thus, the online bidder is a thin layer that converts contextual conversion estimates into bid adjustments around the advertiser's default bid.

The resulting bid can be consumed by the existing relevance- or value-aware ranker without changing the bidder. In our system, candidates are ranked using bid, quality, and profit terms, following profit-aware ranking frameworks [20] and unified valuation for ads and sales profit [16]:

$$\mathrm{PAE}_j = f\big(b_{i,j,a} \cdot pClick_{j,a},\ \hat{p}_{j|a} \cdot \mathrm{profit}_j\big). \tag{14}$$

Pricing remains GSP-like: the winner pays the minimum bid required to maintain position, holding auxiliary score terms fixed. This separates bidder responsibility from ranker responsibility: WDyB performs contextual repricing around advertiser controls, while the ranker manages relevance, ad load, and organic–ad tradeoffs.

Because bid computation runs on the auction path, online serving is limited to feature retrieval, model inference, baseline lookup, and constant-time arithmetic. Heavier optimization and artifact refresh are handled offline or asynchronously. This keeps the bidder simple enough for cross-placement deployment while maintaining p99 latency under 40 ms.

3.5 Architecture Overview

Figure 1 summarizes WDyB as a thin online bidder backed by an offline refresh loop. The online path begins after candidate generation: each campaign-SKU arrives with request, customer, product, and surface-specific signals. WDyB aligns the upstream signal to a shared $pCTR$ interface, computes $pCVR$, retrieves the campaign-SKU baseline, applies normalization and clipping, and emits the final bid.

The offline path maintains the artifacts required by online serving. Logged auctions, model outputs, clicks, orders, and feature tables are joined to refresh calibration artifacts, the post-click model, auxiliary last-touch calibration, and campaign-SKU baselines. The baselines are materialized from recent logged model outputs and refreshed daily.

This online–offline split captures the main system design: shared cross-surface signal alignment, campaign-SKU bid preservation through normalization, and a lightweight auction-time bidder that satisfies production latency constraints.

Relative to static bidding, opaque autobidding, and SKU-level feedback control, WDyB offers a practical middle ground: it reacts to auction context, preserves advertiser defaults as the economic anchor, and avoids chasing sparse realized feedback.

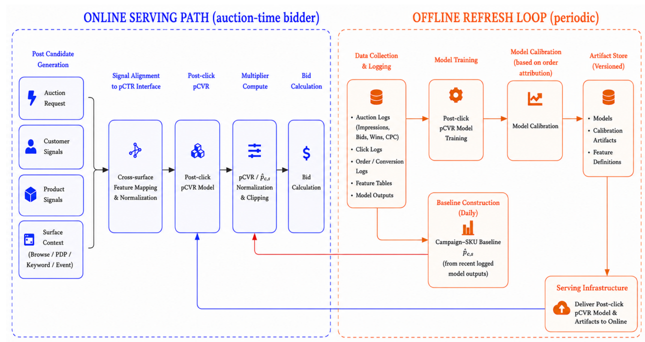

Figure 1: Production architecture of WDyB. The online serving path remains latency-critical and performs signal alignment, post-click prediction, multiplier computation, and bid calculation. The offline refresh loop maintains logs, model artifacts, calibration artifacts, and the daily campaign–SKU baseline used by the online bidder.

Table 1: Shadow test: relative lift vs. static-bid control.

Metric	Relative Lift
Click-through rate (CTR)	+2.1%
Conversion probability	+6.3%
CPC	−5.5%
Attributed revenue / auction	+4.4%
ROAS	+7.9%

4 Experiments

4.1 Offline Simulation

Before live deployment, we conducted a two-week shadow test via traffic mirroring: every live auction request was replayed through the dynamic bidder offline, and hypothetical outcomes were evaluated using customer-response counterfactual models (click and conversion probability applied consistently to both control and shadow results).

Table 1 reports relative lift over the control (static bids). The dynamic bidder increased expected conversion probability by +6.3% and attributed revenue per auction by +4.4%, while ROAS improved +7.9%.

Bid multiplier distribution. The interquartile range of observed multipliers was 20.3% (p25= −13.6%, p50= −3.6%, p75= +6.7%), indicating moderate adjustments well within the $[\alpha_{\min}, \alpha_{\max}]$ bounds. Fewer than 4% of bids hit the cap ceiling, suggesting the clipping range is not overly restrictive.

Normalization impact. The most consequential design decision was normalization granularity. Our initial version (V0) used a *global-average* baseline: $\bar{p}$ was the same for all SKUs across all campaigns. The result was systematic bid inflation for advertisers whose SKUs had above-average conversion rates, and systematic deflation for

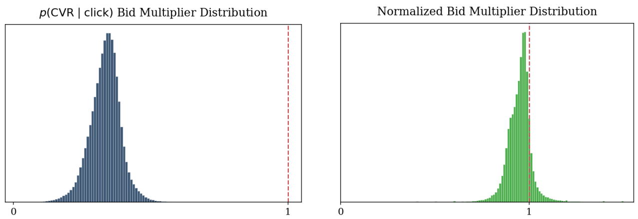

Figure 2: Why campaign-SKU normalization matters. Left: raw distribution is not centered around one. Right: normalized distribution is tightly centered around one.

Table 2: Online A/B test: initial deployment (browse grid).

Metric	Relative Lift	Significance
Ads-attributed orders/click	**+3.84%**	$p < 0.05$
Click-through rate	**+0.87%**	$p < 0.05$
Cost per click	**+2.34%**	$p < 0.05$
ROAS	+1.28%	n.s.
Customer CVR (guardrail)	−0.07%	n.s.

others. Shadow testing revealed class-level spend deviations exceeding ±50%: e.g., area rugs saw +10% spend while sectionals sofa saw −19% under the global-normalization variant.

Switching to per-campaign SKU baselines (V1) resolved the problem, limiting class-level deviations to ±12% and ensuring that each SKU's multiplier averages to ≈1 over the 7-day window, which validated the per-SKU design as essential for respecting advertiser allocation intent across product categories; as illustrated in Figure 2, the bid multiplier in V0 using a *global-average* baseline is not centered around one, while switching to per-campaign SKU baselines (V1) is much more tightly centered around one. This is operationally desirable because it preserves the supplier's bid on average while still allowing meaningful contextual variation.

4.2 Online Experiment

We deployed the system on the browse-grid placement and ran a four-week randomized online A/B experiment against the static-bid control. Our primary metric was ads-attributed orders per click, which is chosen to match the post-click nature of the bidding signal. Secondary metrics included click-through rate, CPC, ROAS, and conversion guardrails at the customer-level. As shown in Table 2, the dynamic bidder improved ads-attributed orders per click by +3.84% ($p < 0.01$), with a modest CPC increase of +2.34% and no detectable change in customer conversion rate (−0.07%, $p = 0.92$). These results indicate that the system improves post-click traffic quality significantly, rather than merely buying more clicks.

4.3 Placement Expansion A/B Test

To validate generalization, we expanded dynamic bidding to two additional placement types: keyword search grid and product-detail-page carousels, and ran a second A/B test over four weeks with a 50/50 customer split [15]. Ratio-metric significance was assessed with delta-method approximations [8].

Table 3: Placement expansion A/B test: dynamic bidding across three placement types.

Placement	Lift in ads-attr. orders/click	Significance
Overall (all placements)	**+1.16%**	$p < 0.05$
Browse	**+1.32%**	$p < 0.05$
Keyword search	+1.48%	n.s.
Product carousel	**+3.77%**	$p < 0.05$
Customer CVR (guardrail)	+0.02%	n.s.

Table 3 reports results. Overall attributed conversions per click improved +1.16% ($p = 0.013$). Lift varied by placement: the product-detail-page carousels showed the largest gain, while keyword search also contributed a non-statistically significant but directional lift. The customer-level conversion rate guardrail remained neutral.

5 Discussion and Future Work

Our results suggest that the main practical benefit of dynamic bidding in sponsored products is the stable contextual repricing around advertiser-provided default bids. In our setting, the key enabler is campaign-SKU normalization that keeps the supplier's default bid as the local operating point while allowing the system to bid up in stronger-than-usual contexts and bid down in weaker ones. This makes the bidder more adaptive than static bidding, while remaining more interpretable and operationally safer than a fully autobidder.

This same design choice also clarifies the system's limitations. The bidder does not estimate full auction win-price response curves online, and it relies on a fixed historical baseline window together with a global fallback for low-history campaign-SKU pairs. As a result, the system prioritizes robustness and deployability over full optimality.

This production-focused scope also limits the breadth of our empirical comparisons. We do not provide a direct empirical comparison against external auto-bidding frameworks or state-of-the-art conversion models. The closest published analogue is Taobao OCPC [29], but its objectives, attribution setup, and reported metrics are not directly comparable. At the same time, the current system leaves several important questions open.

Cold start and sparse-history robustness. The present bidder falls back to a global anchor when campaign-SKU history is insufficient. A natural next step is hierarchical normalization, in which the baseline backs off from campaign-SKU to campaign-category, supplier-category, or other structured priors.

Exploration under uncertainty. Recent autobidding work has emphasized exploration as a way to discover profitable actions while maintaining stable updates [11]. Extending our bidder with uncertainty-aware caps or confidence-conditioned multiplier widening is a natural next step.

Cross-surface and cross-dimension consistency. As deployment broadens, an important question is whether the bidder induces systematic differences in CPC efficiency across placements, time-of-day segments, devices etc. A more general future design would combine dimension-specific normalization with dimension-aware corrections so that the system preserves local supplier intent while also preventing persistent efficiency skews across traffic slices.

Deep learning conversion models A promising next step is to use neural architectures that can jointly encode richer customer, SKU, and auction-context features, while also modeling user behavior sequences to capture evolving intent over time. This would allow the bidder to move beyond mostly static feature interactions toward deeper personalization, adapting bid adjustments to each customer's unique journey across browsing, search, and product-detail interactions. [30].

References

[1] Gagan Aggarwal, Ashwinkumar Badanidiyuru, Santiago R. Balseiro, Kshipra Bhawalkar, Yuan Deng, Zhe Feng, Gagan Goel, Christopher Liaw, Haihao Lu, Mohammad Mahdian, Jieming Mao, Aranyak Mehta, Vahab Mirrokni, Renato Paes Leme, Andres Perlroth, Georgios Piliouras, Jon Schneider, Ariel Schvartzman, Balasubramanian Sivan, Kelly Spendlove, Yifeng Teng, Di Wang, Hanrui Zhang, Mingfei Zhao, Wennan Zhu, and Song Zuo. 2024. Auto-bidding and Auctions in Online Advertising: A Survey. *ACM SIGecom Exchanges* 22, 1 (2024), 159–183.

[2] Gagan Aggarwal, Ashwinkumar Badanidiyuru, and Aranyak Mehta. 2019. Autobidding with Constraints. In *Web and Internet Economics (Lecture Notes in Computer Science, Vol. 11920)*. Springer, 17–30. doi:10.1007/978-3-030-35389-6_2

[3] Amazon Ads. 2024. Guide to Dynamic Bidding – Up and Down with Sponsored Products. https://advertising.amazon.com/library/guides/dynamic-bidding-sponsored-products.

[4] MohammadHossein Bateni, Jon Feldman, Vahab S. Mirrokni, and Sam Chiu-Wai Wong. 2014. Multiplicative Bidding in Online Advertising. In *Proceedings of the Fifteenth ACM Conference on Economics and Computation*. 715–732. doi:10.1145/2600057.2602874

[5] Olivier Chapelle. 2014. Modeling Delayed Feedback in Display Advertising. In *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 1097–1105. doi:10.1145/2623330.2623634

[6] Sougata Chaudhuri, Abraham Bagherjeiran, and James Liu. 2017. Ranking and Calibrating Click-Attributed Purchases in Performance Display Advertising. In *Proceedings of the 2017 ADKDD and TargetAd Workshop*. ACM, 7:1–7:6. doi:10.1145/3124749.3124755

[7] Tianqi Chen and Carlos Guestrin. 2016. XGBoost: A Scalable Tree Boosting System. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*. 785–794.

[8] Alex Deng, Ulf Knoblich, and Jiannan Lu. 2018. Applying the Delta Method in Metric Analytics: A Practical Guide with Novel Ideas. *arXiv preprint arXiv:1803.06336* (2018).

[9] Benjamin Edelman, Michael Ostrovsky, and Michael Schwarz. 2007. Internet Advertising and the Generalized Second-Price Auction: Selling Billions of Dollars Worth of Keywords. *American Economic Review* 97, 1 (2007), 242–259. doi:10.1257/aer.97.1.242

[10] Zhe Feng, Swati Padmanabhan, and Di Wang. 2023. Online Bidding Algorithms for Return-on-Spend Constrained Advertisers. In *Proceedings of the ACM Web Conference 2023*. 3550–3560. doi:10.1145/3543507.3583491

[11] Jingtong Gao, Yewen Li, Shuai Mao, Peng Jiang, Nan Jiang, Yejing Wang, Qingpeng Cai, Fei Pan, Peng Jiang, Kun Gai, et al. 2025. Generative auto-bidding with value-guided explorations. In *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 244–254.

[12] Yuyao Guo, Xiang Ao, Qiming Liu, and Qing He. 2023. Leveraging Post-Click User Behaviors for Calibrated Conversion Rate Prediction Under Delayed Feedback in Online Advertising. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*. 3918–3922. doi:10.1145/3583780.3615161

[13] Yue He, Xiujun Chen, Di Wu, Junwei Pan, Qing Tan, Chuan Yu, Jian Xu, and Xiaoqiang Zhu. 2021. A Unified Solution to Constrained Bidding in Online Display Advertising. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 2993–3001. doi:10.1145/3447548.3467199

[14] Patrick Hummel and Uri Nadav. 2018. Bid-Limited Targeting. In *Proceedings of the 2018 World Wide Web Conference*. doi:10.1145/3178876.3186038

[15] Ron Kohavi, Diane Tang, and Ya Xu. 2020. *Trustworthy Online Controlled Experiments: A Practical Guide to A/B Testing*. Cambridge University Press.

[16] Sergey Kolbin, Arushi Jain, Manavender Malgireddy, Kurt Zimmer, Masoum Mosmer, and Patrick Phelps. 2026. More Ads, Happier Shoppers: Unified-Valuation Ad Allocation at Scale. In *Computational Advertising: 19th International Workshop, AdKDD 2025, Toronto, ON, Canada, August 4, 2025, Proceedings (Communications in Computer and Information Science, Vol. 2796)*. Springer, 85–98. doi:10.1007/978-3-032-16358-5_7

[17] Kuang-Chih Lee, Burkay Orten, Ali Dasdan, and Wentong Li. 2012. Estimating Conversion Rate in Display Advertising from Past Performance Data. In *Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 768–776. doi:10.1145/2339530.2339651

[18] Quan Lu, Shengjun Pan, Liang Wang, Junwei Pan, Fengdan Wan, and Hongxia Yang. 2017. A Practical Framework of Conversion Rate Prediction for Online Display Advertising. In *Proceedings of the 2017 ADKDD and TargetAd Workshop*. ACM, 9:1–9:9. doi:10.1145/3124749.3124750

[19] Xiao Ma, Liqin Zhao, Guan Huang, Zhi Wang, Zelin Hu, Xiaoqiang Zhu, and Kun Gai. 2018. Entire Space Multi-Task Model: An Effective Approach for Estimating Post-Click Conversion Rate. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*. 1137–1140. doi:10.1145/3209978.3210104

[20] Manavender Malgireddy, Dongyue Xie, Sergey Kolbin, Kurt Zimmer, Masoum Mosmer, and Patrick Phelps. 2026. Profit Aware Ad Ranking with Relevance Constraint. In *Computational Advertising: 19th International Workshop, AdKDD 2025, Toronto, ON, Canada, August 4, 2025, Proceedings (Communications in Computer and Information Science, Vol. 2796)*. Springer, 71–84. doi:10.1007/978-3-032-16358-5_6

[21] Rómer Rosales, Haibin Cheng, and Eren Manavoglu. 2012. Post-click Conversion Modeling and Analysis for Non-Guaranteed Delivery Display Advertising. In *Proceedings of the Fifth ACM International Conference on Web Search and Data Mining*. 293–302. doi:10.1145/2124295.2124333

[22] Hal R. Varian. 2007. Position Auctions. *International Journal of Industrial Organization* 25, 6 (2007), 1163–1178.

[23] Xun Yang, Yasong Li, Hao Wang, Di Wu, Qing Tan, Jian Xu, and Kun Gai. 2019. Bid Optimization by Multivariable Control in Display Advertising. *arXiv preprint arXiv:1905.10928* (2019). arXiv:1905.10928 [cs.GT] doi:10.48550/arXiv.1905.10928

[24] Andrew Yates. 2024. How to Build Ads Dynamic Bidding. Promoted.ai Blog, https://medium.com/promoted/how-to-build-ads-dynamic-bidding-b29cb706a9d0.

[25] Jiaqi Zhai, Lucy Liao, Xing Liu, Yueming Wang, Rui Li, Xuan Cao, Leon Gao, Zhaojie Gong, Fangda Gu, Jiayuan He, Yinghai Lu, and Yu Shi. 2024. Actions Speak Louder than Words: Trillion-Parameter Sequential Transducers for Generative Recommendations. In *Proceedings of the 41st International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 235)*. PMLR, 58484–58509. https://proceedings.mlr.press/v235/zhai24a.html

[26] Buyun Zhang, Liang Luo, Yuxin Chen, Jade Nie, Xi Liu, Shen Li, Yanli Zhao, Yuchen Hao, Yantao Yao, Ellie Dingqiao Wen, Jongsoo Park, Maxim Naumov, and Wenlin Chen. 2024. Wukong: Towards a Scaling Law for Large-Scale Recommendation. In *Proceedings of the 41st International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 235)*. PMLR, 59421–59434. https://proceedings.mlr.press/v235/zhang24ao.html

[27] Wenhao Zhang, Wentian Bao, Xiao-Yang Liu, Keping Yang, Quan Lin, Hong Wen, and Ramin Ramezani. 2020. Large-scale Causal Approaches to Debiasing Post-click Conversion Rate Estimation with Multi-task Learning. In *Proceedings of The Web Conference 2020*. 2775–2781. doi:10.1145/3366423.3380037

[28] Yuqing Zhou, Tianshu Feng, Mingrui Liu, and Ziwei Zhu. 2023. A Generalized Propensity Learning Framework for Unbiased Post-Click Conversion Rate Estimation. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management (CIKM '23)*. ACM, 3554–3563. doi:10.1145/3583780.3614760

[29] Han Zhu, Junqi Jin, Chang Tan, Fei Pan, Yifan Zeng, Han Li, and Kun Gai. 2017. Optimized Cost per Click in Taobao Display Advertising. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 2191–2200. doi:10.1145/3097983.3098134

[30] Jinfeng Zhuang, Yinrui Li, Runze Su, Ke Xu, Zhixuan Shao, Kungang Li, Ling Leng, Han Sun, Meng Qi, Yixiong Meng, et al. 2025. On the practice of deep hierarchical ensemble network for ad conversion rate prediction. In *Companion Proceedings of the ACM on Web Conference 2025*. 671–680.