Profit Aware Ad Ranking with Relevance Constraint

Manavender Malgireddy, Dongyue Xie, Sergey Kolbin,
Kurt Zimmer, Masoum Mosmer, Patrick Phelps

AdKDD – Aug 2025

Introduction to WSP

Wayfair Sponsored Products

WSP is a cost-per-click advertising program designed to enhance product discovery and drive sales on Wayfair's marketplace.

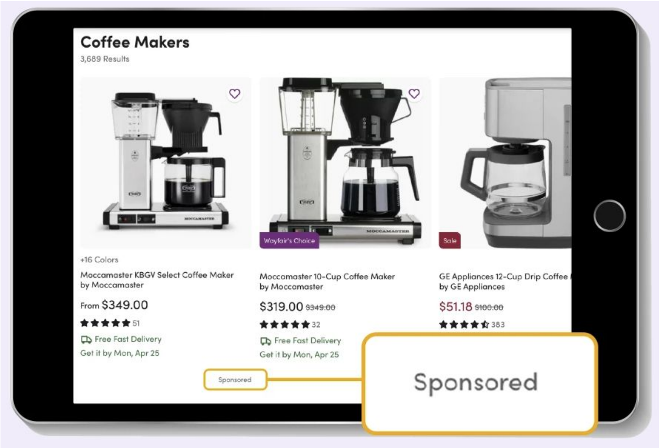

Conventional Auction Ranking

pClick * bid → Maximizes expected ad revenue

Desired Auction Ranking

- Ad Revenue + Profit from sales
- Balance profit with ad relevance

Optimization Problem Formulation

Decision Variables

Binary variables x_rs indicate whether a SKU s is selected for impression in request r (1 if shown, 0 otherwise).

$$x = \{x_{rs} \in \{0, 1\} : r \in R, s \in S_r\}$$

Expected Total Profit

Profit is computed as the sum of expected ad revenue from clicks plus weighted profit from product sales, combining advertiser bids and predicted sales profit (VCD).

$$eTotalProfit(x) = \sum_{r \in R} \sum_{s \in S_r} \Big((pClick_{rs} \cdot bid_s) + W_1 (pPurchase_{rs} \cdot VCD_{rs}) \Big) x_{rs},$$

Relevance Constraint

A minimum threshold on expected conversion rate (purchase likelihood) is enforced to maintain ad relevance and positive customer experience.

$$eCVR(x) = \frac{\sum_{r \in R} \sum_{s \in S_r} pPurchase_{rs} \cdot x_{rs}}{\sum_{r \in R} \sum_{s \in S_r} x_{rs}}$$

Optimization Objective

Maximize the expected total profit subject to the constraint that expected conversion rate exceeds the threshold, formulated as a constrained optimization problem.

$$\max_{x} eTotalProfit(x)$$
$$s.t. \quad eCVR(x) \geq b_0,$$

Dual Problem and Ranking Function

Lagrangian

The ranking optimization problem is transformed into a dual problem using the Lagrangian method, introducing a multiplier to enforce the relevance constraint effectively.

$$L(x,\lambda) = \sum_{r \in R} \sum_{s \in S_r} \Big[(pClick_{rs} \cdot bid_s) + W_1 \cdot (pPurchase_{rs} \cdot VCD_{rs}) + \lambda_0 \cdot pPurchase_{rs} - \lambda_0 \cdot b_0 \Big] x_{rs}$$

Components of the Scoring Function

The score for each ad combines expected click revenue, weighted purchase profit, and a relevance incentive term driven by the Lagrange multiplier, balancing monetization and user satisfaction.

$$Score(r,s) = (pClick_{rs} \cdot bid_s) + W_1 (pPurchase_{rs} \cdot VCD_{rs}) + \lambda_0 \cdot pPurchase_{rs}$$

How to choose b0?

- The threshold parameter $b0$ sets a minimum expected conversion rate to maintain ad relevance and user experience.
- It is defined relative to the optimal expected conversion rate $eCVR^*$ obtained by ranking SKUs purely by purchase probability.
- Formally, $b0 = \alpha \times eCVR^*$ where $\alpha \in [0,1]$ controls the strictness of the relevance constraint.
- Setting α close to 1 targets near-optimal relevance, while lower values allow more flexibility for profit-driven ranking.
- This relative formulation provides interpretability, robustness to data/model drift, and allows easy tuning to balance relevance and profitability.

Solving for the Lagrange Multiplier

Initialize Parameters

Set initial lambda (λ) value, iteration count, step size (μ), and best objective values for dual and primal problems. Prepare for iterative updates.

Iterative Dual Optimization

At each iteration, solve the ranking optimization problem using current λ. Update λ by increasing or decreasing it based on whether the relevance constraint is violated, steering toward feasibility.

Convergence Monitoring

Track dual and primal objective values to monitor progress. Adjust step size if improvement stalls, ensuring stable convergence towards optimal λ.

Finalize Optimal Lambda

Terminate iterations when the difference between primal and dual objectives is within tolerance. Output optimal $\lambda 0$ used in the scoring function for ad ranking.

Initial lambda value
Iteration counters
Step size parameter
Best dual and primal objective values

Updated lambda values
Feasibility checks on relevance constraint
Iterative solutions for ranking decisions

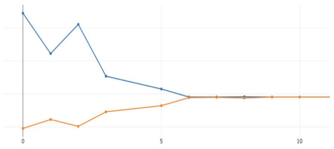

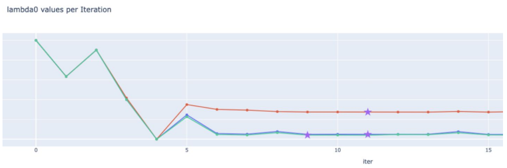

Offline Simulation Methodology

01

Data Collection

Gather logged auction data from the production environment, capturing all candidate SKUs and their associated features such as bids, click probabilities, and purchase likelihoods.

02

Re-ranking Process

Apply the proposed ranking function with different parameter configurations to re-rank SKUs for each auction, simulating alternative ranking outcomes.

03

Metric Computation

Calculate key performance metrics including ad revenue, total profit, conversion rates, and bid sensitivity using counterfactual estimates of clicks and purchases.

04

Comparative Analysis

Compare the simulation results across various parameter settings to understand trade-offs between profitability, relevance, and customer experience.

05

Parameter Selection Guidance

Use insights from simulations to select optimal parameters that balance total profit and user satisfaction ahead of online testing and deployment.

Key Metrics for Evaluation

System-Driven Metrics

- Display Price: Average listed price of products shown in ranked results.
- Ratings and Review Count: Indicators of product quality and popularity.
- Bid Sensitivity: Measures ranking position change when bids increase by 50%.
- Higher sensitivity means ranking is more bid-driven; lower means more relevance-driven.

Customer Interaction Metrics

- Ad Revenue: Profit from clicks on ads.
- WSP Product VCD: Profit from purchases attributed to sponsored ads.
- Total VCD: Sum of Ad Revenue and WSP Product VCD, representing overall profit.
- ROAS: Return on ad spend, ratio of ad revenue to ad cost.
- CPC: Cost per click, reflecting advertiser cost efficiency.

Simulation Results Showing Relevance Constraint Trade-offs

Higher α → Stronger relevance enforcement (↑λ_0).

VCD & ROAS improve sharply when α > 0.7.

Ad profit declines beyond α ≈ 0.7, indicating a trade-off with product profit.

CTR & CVR rise (better engagement), while **CPC** and **bid sensitivity** drop.

Overall: Adjusting α balances **relevance vs. profitability** in ranking strategy.

Impact of Relevance Constraint (α) on Key Business Metrics

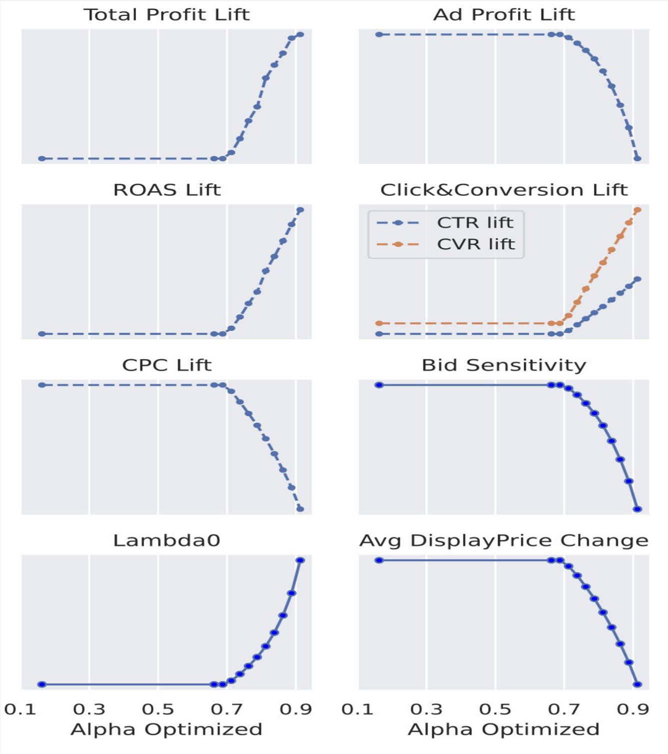

Simulation Results for Chosen Variants

Performance Comparison: Conservative vs Aggressive Variants

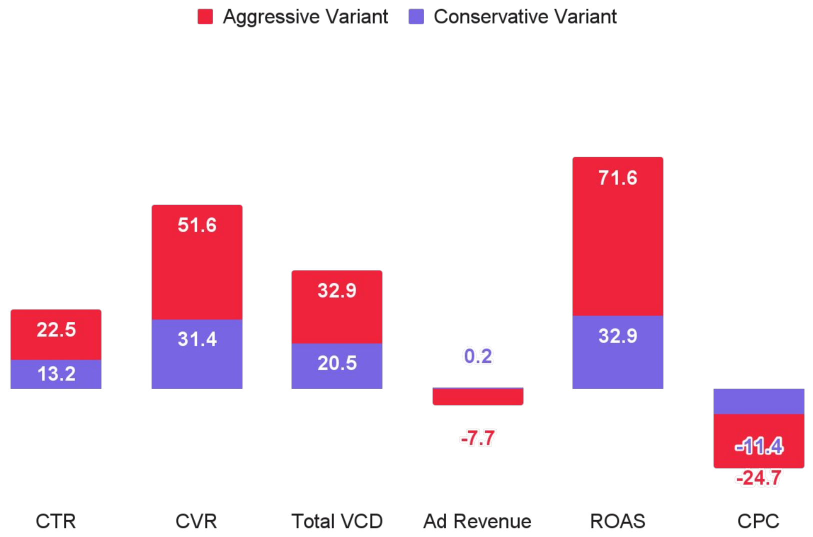

Simulation Results Summary

- Conservative variant improves CTR (+13.2%), CVR (+31.4%), and total profit (+20.5%) with minimal impact on ad revenue (+0.2%) and slight rise in display price (+4.8%).
- Aggressive variant boosts CTR (+22.5%), CVR (+51.6%), and total profit (+32.9%) but lowers ad revenue (-7.7%)
- Both variants show a trade-off between profit, relevance, ad revenue, and CPC, with the aggressive variant focusing more on relevance.

Online Test Results

Key Metrics from Online A/B Tests (% Lifts vs Baseline)

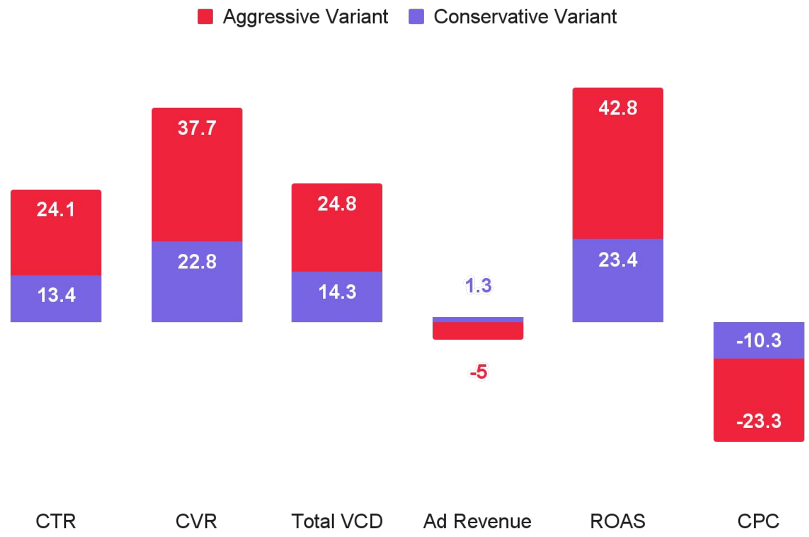

Test Outcomes and Trade-offs

- BAU is based on predicted purchase probability multiplied by the maximum bid
- Experiment ran for four weeks
- Both conservative and aggressive variants improved total profit (Total VCD) and conversion rate (CVR) compared to baseline.
- The aggressive variant achieved higher total profit lift but showed a decline in ad revenue due to reduced CPC.
- CTR and CVR increased for both variants, indicating better user engagement with the new ranking function.
- Online results are directionally consistent with offline simulations

VCD Ablation Study

- Conducted a follow-up online experiment isolating the VCD term effect by comparing a variant with W1 and λ0 terms against a variant with W1 set to zero (no VCD).
- Observed a 2.9% decrease in total variable contribution dollar (VCD) when the VCD term was removed, confirming its positive impact on total profit.
- Other advertising metrics such as CTR, CVR, ad revenue, and ROAS remained largely neutral, indicating that the VCD term primarily drives profit improvements without degrading user engagement or ad performance.
- Findings validate the importance of incorporating sales profit (VCD) into the ranking function to maximize overall profitability while maintaining balanced ad metrics.

Expansion Test Learnings

Unexpected Metric Drops

During the UK expansion, key metrics like CTR and CVR unexpectedly declined despite higher relevance constraints. This pointed to deeper calibration issues in the purchase probability model.

Calibration Issue Identified

Analysis revealed poor calibration of purchase probability with respect to variable contribution dollar (VCD). High VCD SKUs were ranked disproportionately high despite lower actual conversion rates, causing misaligned rankings.

Calibration Fix Applied

A transformation was applied to adjust purchase probabilities relative to VCD using tunable parameters. This improved calibration accuracy and simulation outcomes, facilitating plans to re-run online tests.

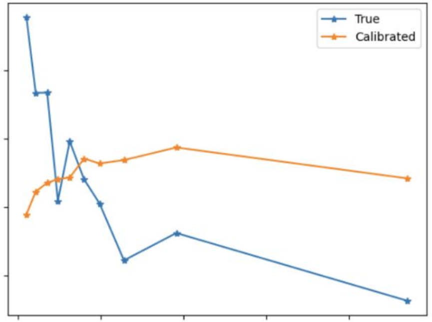

$$pPurchase' = c \cdot \frac{pPurchase}{(VCD + a)^b}$$

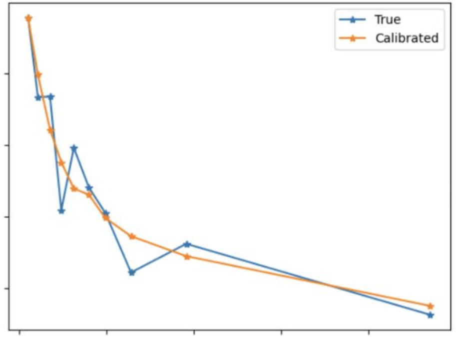

Summary

Conclusions and Future Work

01

Key Contributions

Introduced a novel ranking function balancing total profit and relevance using a constrained optimization framework, supported by a robust parameter tuning approach.

02

Empirical Results

Achieved a 14% increase in total profit and a 22% improvement in conversion rate in online tests, validating the ranking function's effectiveness and relevance trade-offs.

03

Calibration Insights

Identified the critical need for well calibrated purchase probability relative to VCD (sale profit)

04

Future Directions

Explore applying relevance constraints at finer granularity such as individual queries or product categories and enhance calibration methods to sustain performance across markets.