Constrained Target-ROAS Bidding for Sponsored Products

via Value-Based Bids and Feedback Control

Wanchen Shao   Sergey Kolbin   Kurt Zimmer
Jinjin Zhao   Jean-David Ruvini

Wayfair LLC

AdKDD '26 • Jeju

Introduction and Contributions

- **Business context**:
  - Wayfair is a home-furnishings e-commerce platform.
  - Wayfair Sponsored Products run in multi-slot CPC auctions.
  - Auto-bids like tROAS bids and manual CPC bids must enter the same score-based auction and pricing system
- **Supplier goal**: simplify setup by replacing thousands per-SKU manual bid tuning with a single campaign-level target ROAS
- **Key challenge**: low-frequency purchases (e.g., furniture versus everyday grocery items) lead to sparse conversion, making realized ROAS noisy and delayed, feedback-driven optimization more difficult

Contributions

1. Production tROAS bid/rank/charge/control interface for CPC sponsored products
2. Campaign-level feedback controller with adaptive learning, cold-start handling, and stability guardrails
3. Offline AA noise floor quantifies intrinsic stochasticity and what target attainment is statistically achievable
4. Large-scale deployment online evidence: attainment, efficiency, and randomized stability improvements.

Methodology: Value-Based Bids + Campaign Feedback

Constrained optimization view

$$\mathrm{ROAS}_i(b) = \frac{\sum_{a,j} x_{i,j,a}(b) v_{i,j,a}}{\sum_{a,j} x_{i,j,a}(b) c_{i,j,a}(b)}$$

$$\max_b \sum_{a,j} x_{i,j,a}(b) v_{i,j,a} \quad \text{s.t.} \quad \mathrm{ROAS}_i(b) = T_i$$

Auction-time value and bid rule

$$b_{i,j,a} = CP_i \cdot \frac{v_{i,j,a}}{T_i} = CP_i \cdot \frac{\pi_j \hat{p}_{j|a}}{T_i}$$

Campaign-level error and multiplicative update

$$e_i^{(t)} = \frac{\overline{\mathrm{ROAS}}_i - T_i}{T_i}$$

$$\mathrm{CP}_i^{(t)} = \mathrm{CP}_i^{(t-1)} \cdot \left(1 + \alpha_i^{(t)} \cdot e_i^{(t)} + k_d \cdot (e_i^{(t)} - e_i^{(t-1)})\right)$$

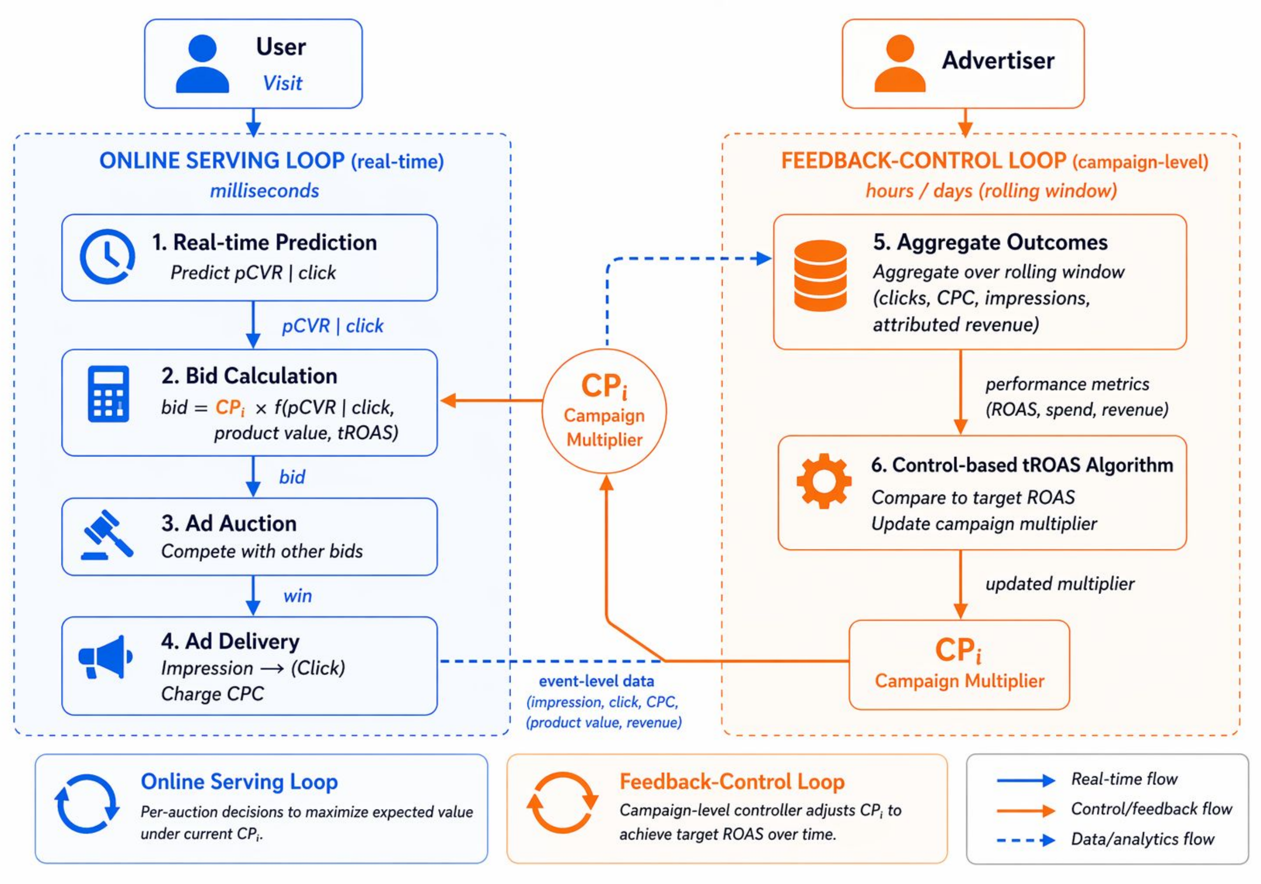

Split-sample AA benchmark

Purpose: separate controller error from irreducible finite-sample variation

$$D_i = \frac{|R_{iA} - R_{iA'}|}{2R_i}, \qquad D_i \leq 0.20$$

- N = 32,140 campaigns with positive spend.
- Take historical data and randomly split customers into A and A prime.
- A campaign is measurable in the +/-20% target band when Di <= 0.20.

Target-band achievability by order volume

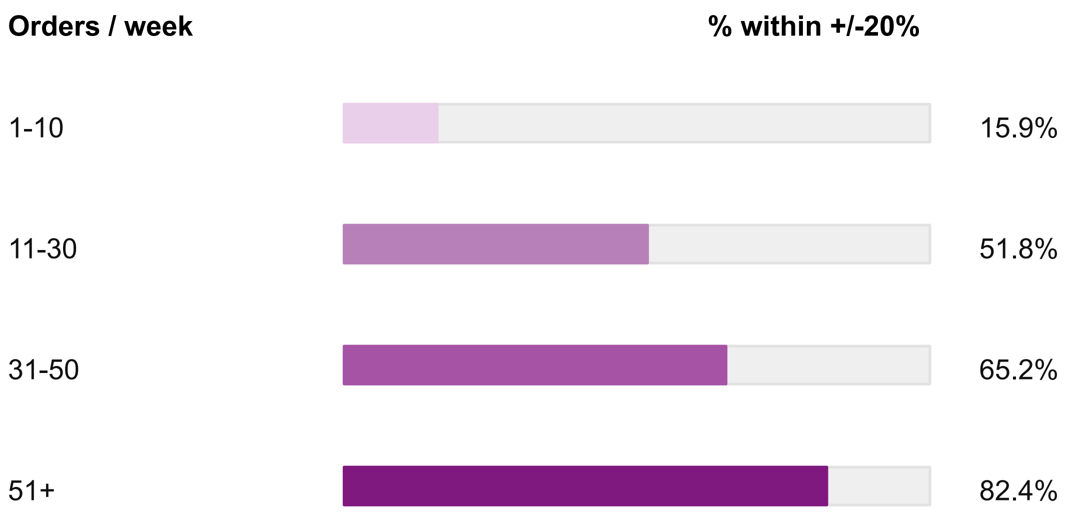

Takeaways

- Eligibility should be volume-aware; otherwise noisy ROAS makes any controller look unstable.
- Guardrails motivated by sparsity

Production Evidence and Observational Efficiency

Rollout and post-deployment readout

- In steady state, evaluable campaigns reached 90% meet-or-beat (>80% tROAS) and 74% attainment (within +/- 20% margin of tROAS)
- Efficiency analysis showed observational 11–15% efficiency gain vs. manual bidding

Supplier response

- High-tROAS adopters: +17% faster ad-spend growth and +8% faster active-SKU participation.

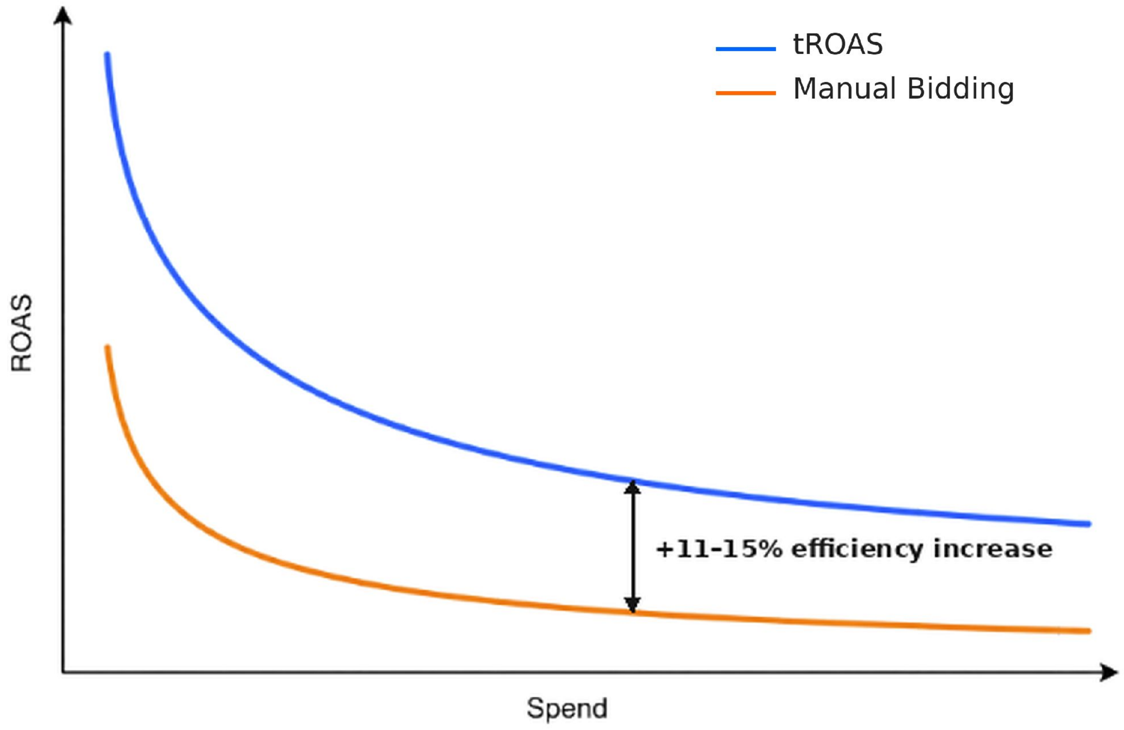

Figure: Covariate-adjusted Efficiency comparison between manual bidding and tROAS

Iterative A/B Test and Takeaways

Iterative improvement test

Metric	Lift
Lower-bound attainment	**+9.32%**
ROAS	+5.79%
CPC	**-9.28%**
CP volatility	**-49.74%**
Order revenue / campaign	+6.65%

Bold *= $p<0.05$*

Treatment B added smoothed ROAS error, decaying updates, and per-update CP caps.

Main takeaways

1. **Two-layer design:** auction-time value-based bids + campaign-level feedback control
2. Conversion sparsity creates an intrinsic ROAS noise floor. **Noise-floor analysis** define eligibility for optimal performance
3. **Near practical ceiling:** attainment approaches the offline achievability bound

Future work

1. Exploration for under-explored SKUs, e.g. New and low-volume SKUs suffer from a cold-start feed-back loop
2. Promotion / seasonality forecasting to mitigate large controller oscillations when ROAS spikes up
3. Advanced constrained optimization via online learning, RL variants and generative learning

Thank you!

Questions?