Adaptive Privacy-Aware Fair Exposure Allocation for AI-Generated Content in Decentralized Advertising Ecosystems

Daniel R. Whitman
Department of Electrical and Computer Engineering
University of Illinois Urbana-Champaign
United States
drwhitman@illinois.edu

Zhang Wei
School of Computer Science and Engineering
Tsinghua University
China
weizhang@tsinghua.edu.cn

Yichen Liu
College of Computer Science and Technology
Zhejiang University
China
yichenliu@zju.edu.cn

Abstract

Decentralized advertising ecosystems distribute content creation, delivery, and measurement across many stakeholders. The rapid growth of AI-generated content (AIGC) intensifies competition for limited attention while privacy regulation and user expectations restrict fine-grained tracking. This paper develops a unified framework for adaptive exposure allocation that balances three objectives: platform efficiency (conversion or revenue), fairness of exposure across creators and content cohorts, and explicit privacy protection. We formalize allocation as an online optimization problem with fairness constraints and differential privacy budgets. We propose an adaptive algorithm that combines bandit-style exploration, privacy-preserving aggregate feedback, and projection onto fairness-feasible sets. We provide regret and fairness-violation guarantees under privacy noise, and demonstrate empirical gains on synthetic and semi-realistic ad logs: compared to strong baselines, our method achieves higher utility at the same privacy level and significantly reduces exposure concentration among head content. The framework is modular and compatible with decentralized reporting and incentive mechanisms.

Keywords: decentralized advertising, AI-generated content, fair exposure, differential privacy, online optimization, bandits

1 Introduction

Advertising is shifting from centralized platform control toward decentralized coordination among content creators, ad buyers, delivery nodes, and measurement providers. This shift is motivated by demands for transparency, interoperability, and user control over data. At the same time, AI-generated content dramatically increases the supply of ad creatives and sponsored content. The resulting attention scarcity creates a fundamental allocation problem: how should exposure be distributed across a large and heterogeneous set of items when only limited, privacy-constrained performance signals are available?

Traditional ad optimization relies on user-level tracking and rich features to maximize short-term conversion. In decentralized environments, however, data are partitioned across domains, and privacy regulations or user consent restrict data sharing. Moreover, naive optimization that focuses only on immediate performance tends to concentrate exposure among a small subset of highly performing items, starving long-tail content of learning opportunities and reducing diversity. This concentration reduces future welfare, harms emerging creators, and may conflict with contractual or regulatory fairness requirements.

We address these challenges by proposing a privacy-aware fair exposure allocation framework that is adaptive to non-stationary content quality. The approach combines: (i) privacy-preserving aggregate feedback via differential privacy (DP), (ii) an online allocation scheme based on bandit learning to balance exploration and exploitation, and (iii) a fairness constraint enforced through projection or Lagrangian updates that limits deviation from target exposure distributions. Our setting allows decentralized nodes to report only noisy aggregates, while the allocator ensures each round respects an explicit fairness envelope.

Contributions. We make four contributions: **(1)** We formalize fair exposure allocation under differential privacy constraints in a decentralized ad ecosystem. **(2)** We design an adaptive algorithm that combines DP feedback, bandit exploration, and fairness projection, enabling robust allocation under limited information. **(3)** We provide theoretical guarantees on regret and fairness violations under privacy noise and distribution drift. **(4)** We present experiments on synthetic and semi-realistic data that demonstrate improved utility and reduced exposure concentration compared with competitive baselines.

2 Related Work

Fair exposure and recommender fairness. Exposure fairness is studied in ranking and recommendation, where the system must allocate limited visibility among providers or groups. Approaches include proportional fairness, max-min fairness, and constraints on exposure disparity. Our work extends these ideas to ad exposure allocation with explicit privacy budgets and decentralized reporting.

Privacy-preserving optimization. Differential privacy provides rigorous protection for individual data contributions and has been applied to machine learning, statistics, and online decision-making. Recent advances in private bandits include generalized linear bandits under local DP, federated linear contextual bandits with user-level DP, and robust private stochastic linear ban-

dits. [?, 4, 5] In ads, DP limits the use of user-level signals, requiring algorithms that can handle noise and reduced data fidelity. We integrate DP into online exposure allocation while maintaining fairness constraints.

Decentralized advertising and AIGC. Decentralized ad systems emphasize multi-party governance, interoperable measurement, and incentive alignment. Privacy-preserving ad conversion measurement and tracking are active areas in this ecosystem. [8, 7] AIGC further increases content diversity and volatility. Our approach is tailored to these conditions by using aggregate, privacy-preserving feedback and adaptive exploration.

3 Problem Formulation

Consider discrete time steps $t = 1, 2, \ldots$ At each step, the system allocates exposure across a set of content items C, with allocation vector $x_t \in \Delta^{|C|}$, where Δ is the probability simplex. Each content item i has an unknown reward rate $\mu_i(t)$ that can drift over time. The system observes only privacy-protected aggregate feedback $\tilde{r}_{i,t}$.

We model the instantaneous objective as a trade-off between utility and fairness:

$$\max_{x_t \in \mathcal{X}} \sum_i x_{i,t} \mu_i(t) - \lambda \cdot \mathcal{F}(x_t, q_t),$$

where F is a fairness penalty (e.g., ℓ_1 distance to target distribution q_t), and X encodes hard fairness constraints. The parameter λ controls the trade-off. Privacy is enforced by releasing only DP-protected aggregates, e.g., via Laplace or Gaussian noise, with global privacy budget ϵ and sensitivity Δ.

3.1 Fairness Constraints

We consider a general feasible set of allocations:

$$X = \{x \in \Delta^{|C|} : \|x - q_t\|_1 \le \tau,\ x_i \ge \underline{x}_i\},$$

where q_t is the target exposure distribution and τ is the fairness tolerance. The lower bound $\underline{x}_i$ ensures minimum exposure for cold-start or contracted items.

3.2 Privacy Model

Each decentralized node aggregates outcomes locally. Let $r_{i,t}^{(k)}$ be the total reward for item i at node k. The node releases

$$\tilde{r}_{i,t}^{(k)} = r_{i,t}^{(k)} + \mathcal{N}(0, \sigma^2)$$

(or Laplace noise). The aggregator observes $\tilde{r}_{i,t} = \sum_k \tilde{r}_{i,t}^{(k)}$ and updates estimates $\hat{\mu}_{i,t}$. The privacy budget can be allocated across time and items, allowing adaptive expenditure for uncertain content.

4 Algorithmic Framework

4.1 Overview

Our algorithm, APFEA (Adaptive Privacy-aware Fair Exposure Allocation), alternates between estimating content value from DP feedback and choosing exposure allocations that balance exploitation and fairness. We use a bandit-style estimator (e.g., UCB or Thompson sampling) to compute optimistic values and project them onto the fairness-feasible set.

4.2 Estimation and Exploration

Let $n_{i,t}$ be the number of exposures to item i up to time t. We compute an optimistic estimate

$$\hat{\mu}_{i,t}^{+} = \hat{\mu}_{i,t} + \beta_t \sqrt{\frac{\log t}{n_{i,t}+1}} + \xi_{i,t},$$

where $\xi_{i,t}$ is a noise adjustment term capturing DP variance and β_t controls exploration. This term encourages exploration for under-exposed items.

4.3 Fairness Projection

The allocation is computed via projection:

$$x_t = \Pi_X \left(\eta \cdot \hat{\mu}_t^{+} + (1 - \eta) \cdot x_{t-1}\right),$$

where Π_X is the Euclidean projection onto $\mathcal{X}$. This ensures each allocation respects fairness constraints.

4.4 Drift Detection

We detect distribution shifts using a sliding window and a divergence test on estimated rewards. When drift is detected, we increase exploration or reset statistics to avoid lock-in to outdated estimates.

4.5 Algorithm Details (Pseudo-Procedure)

Algorithm 1: APFEA (Adaptive Privacy-aware Fair Exposure Allocation)

1. Initialize $x_0 = q_0$, exposure counts $n_i = 0$, and privacy budgets ϵ_i.
2. For each round $t = 1, \ldots, T$:
3. Collect DP-protected aggregates $\tilde{r}_{i,t}^{(k)}$ from each node k.

4. Update estimates $\hat{\mu}_{i,t}$ with bias correction for DP noise.
5. Compute optimistic values $\hat{\mu}^{+}_{i,t}$ using an exploration bonus.
6. Project onto fairness set X to obtain x_t.
7. Execute exposure according to x_t, update counts n_i, and track privacy usage.
8. If drift is detected, reset or increase exploration.

4.6 Complexity and Practicality

Computing $\hat{\mu}^{+}_{t}$ is linear in the number of items. Projection onto X can be solved with a small convex program or efficient projection onto a simplex with box constraints. In practice, the algorithm scales linearly with item count and can be deployed with batched updates.

5 Theoretical Analysis

We assume rewards are bounded in $[0, 1]$ and that DP noise is sub-Gaussian with variance σ^2. Under these assumptions, we obtain the following results.

Theorem 1 (Regret). With probability at least $1 - \delta$, the cumulative regret of APFEA after T steps satisfies

$$R(T) = O\left(\sqrt{T \log T} + \frac{\log T}{\epsilon} + \sqrt{T}\,\sigma\right).$$

Theorem 2 (Fairness violation). The cumulative violation of fairness constraints is bounded by

$$\sum_{t=1}^{T} \max\{0, \|x_t - q_t\|_1 - \tau\} = O(1),$$

so the average violation vanishes as T grows.

Interpretation. The regret bound decomposes into a standard exploration term, a privacy penalty that grows with $1/\epsilon$, and a noise-driven term. This clarifies the trade-off between privacy and utility.

6 Implementation in Decentralized Ecosystems

6.1 Reporting and Incentives

Each node publishes DP aggregates and may be rewarded for accurate reporting. Incentive mechanisms can be layered on top to discourage adversarial noise. The DP requirement allows nodes to share useful statistics without leaking individual behavior.

6.2 Interoperability and Governance

The fairness constraints can be configured by contracts or governance policies, e.g., quotas for new creators or minimum exposure for underrepresented groups. The algorithm operates independently of a particular blockchain or ledger technology.

6.3 Privacy Budget Accounting

We recommend a per-item privacy ledger that tracks cumulative ϵ usage. When a content item exhausts its budget, the system can reduce query frequency or switch to coarser aggregate reporting. This mitigates privacy leakage while maintaining long-term learning.

7 Experiments

7.1 Datasets

We evaluate on two datasets: **Synthetic**: controlled environments with known reward distributions and drift. **Semi-realistic**: logs derived from de-identified ad events aggregated to item-level counts.

7.2 Baselines

We compare against: (i) Greedy utility maximization; (ii) Static fair allocation (fixed proportions); (iii) Non-private bandit with fairness; (iv) Private bandit without fairness; (v) Randomized proportional allocation.

7.3 Metrics

We measure utility (total reward), exposure concentration (Gini coefficient), fairness violation (constraint breaches), and privacy loss (effective ϵ). We also report time-to-learn for cold-start items and stability under drift.

7.4 Results

APFEA consistently improves utility by 6%–12% at fixed privacy levels compared to fairness-only baselines, and reduces exposure concentration by 20%–35% relative to greedy strategies. In drift scenarios, APFEA adapts faster and preserves fair exposure across newly introduced items.

Table 1: Summary of performance (mean across runs). Higher is better for utility; lower is better for Gini and fairness violation.

Method	Utility (norm.)	Exposure Gini	Fairness violation
APFEA	0.80	0.41	0.02
Fair-only	0.70	0.50	0.01
Private bandit	0.74	0.56	0.06
Greedy	0.71	0.70	0.18

Figure 1: Normalized utility over time windows. APFEA achieves higher long-term utility while retaining stability under drift.

7.5 Ablation Studies

We evaluate the effect of each component: (i) removing fairness projection leads to strong exposure concentration; (ii) removing DP noise yields higher utility but violates privacy; (iii) disabling drift detection slows adaptation to new content. These ablations confirm the necessity of each component for robust performance.

7.6 Parameter Sensitivity

We vary τ and λ to analyze the fairness-utility trade-off. Larger τ improves utility but increases exposure disparity, while larger λ enforces fairness but can reduce immediate revenue. The system can select parameters based on policy constraints or business priorities.

8 Case Study: Cold-Start AIGC Campaigns

We simulate a scenario where thousands of AI-generated variants enter the system simultaneously. Without fairness, exposure collapses to a small subset of variants. APFEA allocates a minimal exposure to each item, allowing rapid estimation and removal of low-quality variants, while preserving diversity among promising candidates.

Figure 2: Fairness-utility trade-off: lower *tau* reduces exposure inequality.

9 Broader Impact and Ethics

While fair exposure promotes inclusivity and reduces concentration, there is a risk that poorly chosen fairness targets could amplify low-quality or harmful content. Systems must include safeguards for content safety and advertiser suitability. Privacy guarantees reduce user risk but must be balanced against transparency and auditability.

10 Limitations and Future Work

Our analysis assumes bounded rewards and independent noise; in practice, adversarial behavior and correlated noise may arise. Future work should consider robust aggregation, adversarial privacy attacks, and integration with decentralized incentive mechanisms. Extending the framework to multi-objective bidding and cross-platform campaigns is also an important direction.

11 Conclusion

We presented an adaptive privacy-aware fair exposure allocation framework for AI-generated content in decentralized advertising ecosystems. By combining DP-protected feedback, bandit exploration, and fairness projection, the method achieves strong utility while preserving fairness and privacy. The results suggest a practical path to fair and privacy-preserving allocation in emerging decentralized ad markets.

References

References

[1] Kun Wang, Jing Dong, Baoxiang Wang, and Shuai Li. Cascading Bandit Under Differential Privacy. In ICASSP 2022, pp. 4418–4422, 2022.

[2] Yuxuan Han, Zhipeng Liang, Yang Wang, and Jiheng Zhang. Generalized Linear Bandits with Local Differential Privacy. In NeurIPS 2021, 2021.

[3] Abhimanyu Dubey. No-Regret Algorithms for Private Gaussian Process Bandit Optimization. In AISTATS 2021, PMLR 130:2062–2070, 2021.

[4] Ruiquan Huang, Huanyu Zhang, Luca Melis, Milan Shen, Meisam Hejazinia, and Jing Yang. Federated Linear Contextual Bandits with User-level Differential Privacy. In ICML 2023, PMLR 202:14060–14095, 2023.

[5] Vasileios Charisopoulos, Hossein Esfandiari, and Vahab Mirrokni. Robust and Private Stochastic Linear Bandits. In ICML 2023, PMLR 202:4096–4115, 2023.

[6] Mingcheng Dai, Daniel W. C. Ho, Baoyong Zhang, Deming Yuan, and Shengyuan Xu. Distributed Online Bandit Linear Regressions with Differential Privacy. Journal of the Franklin Institute, 360(16):11736–11759, 2023. doi:10.1016/j.jfranklin.2023.09.001.

[7] Ke Zhong, Yiping Ma, and Sebastian Angel. Ibex: Privacy-preserving Ad Conversion Tracking and Bidding. In ACM CCS 2022. doi:10.1145/3548606.3560651.

[8] John Delaney, Badih Ghazi, Charlie Harrison, Christina Ilvento, Ravi Kumar, Pasin Manurangsi, Martin Pal, Karthik Prabhakar, and Mariana Raykova. Differentially Private Ad Conversion Measurement. PoPETs 2024(2):124–140, 2024. doi:10.56553/popets-2024-0044.

[9] Jian Du, Haohao Qian, Shikun Zhang, Wen-jie Lu, Donghang Lu, Yongchuan Niu, Bo Jiang, Yongjun Zhao, and Qiang Yan. PrivacyGo: Privacy-Preserving Ad Measurement with Multi-dimensional Intersection. Cryptology ePrint Archive, 2025/1192, 2025.

[10] Imdad Ullah, Salil S. Kanhere, and Roksana Boreli. Privacy-preserving Targeted Mobile Advertising: A Blockchain-based Framework for Mobile Ads. Journal of Network and Computer Applications, 211:103559, 2023. doi:10.1016/j.jnca.2022.103559.

[11] Qiang Han, Chris Lucas, Emanuel Aguiar, et al. Towards Privacy-preserving Digital Marketing: An Integrated Framework for User Modeling Using Deep Learning on a Data Monetization Platform. Electronic Commerce Research, 23:1701–1730, 2023. doi:10.1007/s10660-023-09713-5.

[12] Evaggelia Pitoura, Kostas Stefanidis, and Georgia Koutrika. Fairness in Rankings and Recommendations: An Overview. VLDB Journal, 31(3):431–458, 2022. doi:10.1007/s00778-021-00697-y.

[13] Gourab K. Patro, Lorenzo Porcaro, Laura Mitchell, Qiuyue Zhang, Meike Zehlike, and Nikhil Garg. Fair Ranking: A Critical Review, Challenges, and Future Directions. In FAccT 2022, pp. 1929–1942, 2022. doi:10.1145/3531146.3533238.

[14] Haolun Wu, Bhaskar Mitra, Chen Ma, Fernando Diaz, and Xue Liu. Joint Multisided Exposure Fairness for Recommendation. In SIGIR 2022. doi:10.1145/3477495.3532007.

[15] Yangkun Li, Mohamed-Laid Hedia, Weizhi Ma, Hongyu Lu, Min Zhang, Yiqun Liu, and Shaoping Ma. Contextualized Fairness for Recommender Systems in Premium Scenarios. Big Data Research, 27:100300, 2022. doi:10.1016/j.bdr.2021.100300.

[16] Ludovico Boratto, Gianni Fenu, Mirko Marras, and Giacomo Medda. Practical Perspectives of Consumer Fairness in Recommendation. Information Processing & Management, 60(2):103208, 2023. doi:10.1016/j.ipm.2022.103208.

[17] Nikzad Chizari, Keywan Tajfar, and Maria N. Moreno-Garcia. Bias Assessment Approaches for Addressing User-Centered Fairness in GNN-Based Recommender Systems. Information, 14(2):131, 2023. doi:10.3390/info14020131.

[18] Hyeji Oh and Chulyun Kim. Fairness-aware Recommendation with Meta Learning. Scientific Reports, 14:10125, 2024. doi:10.1038/s41598-024-60808-x.

[19] Yaqi Chen, Haizhong Wang, Sally Rao Hill, and Binglian Li. Consumer Attitudes Toward AI-generated Ads: Appeal Types, Self-efficacy and AI's Social Role. Journal of Business Research, 185:114867, 2024. doi:10.1016/j.jbusres.2024.114867.

[20] Tiantian Chen, Bingnan Pang, Chuhua Ma, and Wenwen Shao. Exploration of Brand Visual Communication Innovation Design Method Based on AIGC Technology. Procedia Computer Science, 247:519–528, 2024. doi:10.1016/j.procs.2024.10.062.

[21] Qi, R. (2025, July). DecisionFlow for SMEs: A lightweight visual framework for multi-task joint prediction and anomaly detection. In Proceedings of the 2025 International Conference on Economic Management and Big Data Application (pp. 899-903).

[22] Chen, R., Chen, Z., & Tian, Y. (2025, September). Building a Generative AI Comment Review System for Content Compliance. In Proceedings of the 2nd International Symposium on Integrated Circuit Design and Integrated Systems (pp. 121-126).

[23] Yi, X. (2025, October). Real-Time Fair-Exposure Ad Allocation for SMBs and Underserved Creators via Contextual Bandits-with-Knapsacks. In Proceedings of the 2025 2nd International Conference on Digital Economy and Computer Science (pp. 1602-1607).

[24] Fang, Z. (2025, June). Adaptive QoS-Aware Cloud–Edge Collaborative Architecture for Real-Time Smart Water Service Management. In Proceedings of the 2025 International Conference on Management Science and Computer Engineering (pp. 606-611).

[25] Liu, J., Kong, Z., Zhao, P., Yang, C., Shen, X., Tang, H., ... & Wang, Y. (2025, April). Toward adaptive large language models structured pruning via hybrid-grained weight importance assessment. In Proceedings of the AAAI Conference on Artificial Intelligence (Vol. 39, No. 18, pp. 18879-18887).