

Article

Not peer-reviewed version

---

# A Federated and Differentially Private Incentive–Marketing Framework for Privacy-Preserving Cross-Channel Measurement in AI-Powered Digital Commerce

---

[Xiongsheng Yi](#)\*

Posted Date: 28 February 2026

doi: 10.20944/preprints202602.1929.v1

Keywords: privacy-preserving marketing; federated learning; differential privacy; cross-channel measurement; topics API; protected audience; attribution reporting; SKAdNetwork; SMB incentive optimization; digital commerce



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

# A Federated and Differentially Private Incentive–Marketing Framework for Privacy-Preserving Cross-Channel Measurement in AI-Powered Digital Commerce

Xiongsheng Yi

Department of Computer Science and Engineering - School of Engineering, Santa Clara University, Santa Clara, CA, 95053, USA, xyi@scu.edu

## Abstract

In the U.S. digital economy, small and medium-sized businesses (SMBs) and creators in remote regions face structural disadvantages in access to integrated advertising and incentive platforms, largely due to accelerating privacy regulations and the fragmentation of cross channel datasets. This paper proposes a unified federated and differentially private measurement framework that integrates Topics/Protected Audience, Attribution Reporting, and SKAdNetwork, aiming to achieve privacy-preserving incentive optimization and cross-channel effectiveness measurement for web and mobile environments. The framework prioritizes compliant data usage, resolves data silos across ad ecosystems, and supports privacy-preserving recommendation and incentive allocation. Technically, we design a hybrid architecture that combines federated learning, differential privacy, and low-latency attribution aggregation, while ensuring end-to-end consistency across uplift modeling, multi-touch attribution (MTA), and event-level reporting. Empirical analysis compares the proposed model with state-of-the-art privacy-preserving baselines (e.g., last-touch attribution with DP aggregation), demonstrating substantial gains in accuracy, robustness, and reporting fidelity under strict privacy constraints.

**Keywords:** privacy-preserving marketing; federated learning; differential privacy; cross-channel measurement; topics API; protected audience; attribution reporting; SKAdNetwork; SMB incentive optimization; digital commerce

---

## 1. Introduction

The rapid transformation of digital commerce in the United States has placed unprecedented pressure on small and medium-sized businesses (SMBs) and creators, especially those located in remote or economically disadvantaged regions. While large platforms benefit from advanced data infrastructure and integrated advertising–incentive systems, SMBs continue to encounter structural limitations in audience targeting, attribution accuracy, and return-on-investment optimization. These disparities have been amplified by tightening privacy regulations, the deprecation of third-party cookies, and the proliferation of heterogeneous data governance frameworks across advertising ecosystems. As a consequence, SMBs increasingly operate within data-constrained environments where compliant data usage is prioritized but actionable signals for marketing optimization are significantly diminished.

At the technical level, digital advertising has entered a transition period marked by a shift from identity-based targeting toward privacy-preserving interest signals such as Topics and Protected Audience, as well as platform-controlled attribution frameworks including Attribution Reporting for the web and SKAdNetwork for mobile applications. Although these mechanisms enhance privacy guarantees, they also introduce fragmented reporting pipelines, increased latency, coarse-grained

measurement granularity, and inconsistencies between web and app environments. Cross-channel data silos—historically alleviated by third-party cookies now reemerge as major bottlenecks for accurate cross-device and cross-platform marketing measurement.

To address these challenges, research attention has shifted toward federated learning (FL) and differential privacy (DP) as foundational technologies enabling compliant, high-fidelity measurement without exposing raw user-level data. FL enables collaborative model training across devices or enterprise boundaries, mitigating the privacy risks associated with centralized data aggregation. DP injects mathematically bounded noise into the training or reporting pipeline, offering quantifiable privacy guarantees that align with modern regulatory requirements. However, deploying FL + DP in real marketing environments presents significant complexities, including handling attribution delays, ensuring cross-channel consistency, reconciling event-level and aggregated reporting, and accommodating the operational constraints of SMB advertisers.

This study proposes a unified privacy-preserving incentive and marketing measurement framework that integrates Topics, Protected Audience, Attribution Reporting, and SKAdNetwork under a federated and differential privacy paradigm. The framework places compliant data usage at the core while providing an end-to-end solution for cross-channel measurement, incentive allocation, uplift estimation, and multi-touch attribution under strict privacy constraints. On the web side, the model leverages interest-based signals and site-defined protected audience cohorts, while on the mobile side it incorporates SKAdNetwork event-level reporting and privacy thresholds. A federated learning architecture harmonizes these heterogeneous signals, supported by DP-based noise calibration to ensure both privacy protection and measurement accuracy.

Furthermore, this work contributes a comparative evaluation pipeline that contrasts the proposed federated-DP framework with simpler, non-federated privacy-preserving baselines across three dimensions: (1) attribution accuracy and reporting delay, (2) uplift and optimization effectiveness, and (3) cross-channel consistency. Given the industry-wide deprecation of third-party cookies, our evaluation focuses on performance differentials within the privacy-preserving paradigm. Overall, this research advances the state of privacy-preserving marketing measurement, offering a technically feasible, regulatory-aligned, and SMB oriented solution for the emerging post-cookie era. It provides both theoretical and practical foundations for marketing ecosystems seeking to balance privacy protection with measurable advertising effectiveness in AI-powered digital commerce.

## 2. Related Work

The rapid evolution of privacy-preserving machine learning and privacy-compliant marketing technologies has reshaped how digital platforms measure conversions, attribute ad performance, and optimize cross-channel campaigns. Existing research can be broadly grouped into three threads: differential privacy for advertising measurement, secure and cryptographic mechanisms for attribution, and federated learning for ad targeting and delivery.

Early studies introduced differentially private (DP) mechanisms for online advertising, demonstrating that conversion events and user-level engagement can be reported with mathematically bounded leakage while maintaining competitive utility. Xiao et al. formulated one of the first online DP measurement frameworks capable of supporting per-user attribution and incremental lift estimation under strict privacy budgets [1]. Similar advances were reported in the Privacy Enhancing Technologies community, where Delaney et al. designed end-to-end pipelines for DP conversion measurement that achieve accurate aggregate statistics while resisting reconstruction attacks [2]. Complementary engineering work further demonstrated how DP can be systematically embedded into large-scale commercial ad-measurement systems, emphasizing noise calibration, event aggregation, and cross-device consistency [3]. This stream of research crystallizes into more advanced multidimensional reporting designs such as PrivacyGo, which enable private intersection, frequency estimation, and multi-touch attribution under differential privacy guarantees [4]. Beyond these industrial frameworks, prior work in security and applied cryptography explored DP in

personalized advertising and demonstrated its compatibility with real-time bidding and behavioral targeting [5,6].

A parallel line of work focuses on cryptographic and privacy-enhancing mechanisms that facilitate privacy-safe reporting without direct identifier exchange. Recent developments in delegated private matching [7] and malicious-resilient aggregated statistics protocols [8] provide efficient primitives for multi-party computation (MPC) and privacy sandbox-style reporting. Earlier technical disclosures also proposed MPC-based lift measurement, demonstrating that cross-platform conversion lift can be computed without sharing raw identifiers [9]. Academic efforts such as Ibex further explored cryptographic protocols for conversion tracking and bidding, offering end-to-end encrypted pipelines suitable for high-volume ad exchanges [10]. These studies directly influenced modern browser-level attribution systems. For instance, the Attribution Reporting API incorporates summary-report optimization [11] and novel analyses of interactivity and DP parameters [12], providing a standardized interface for privacy preserving measurement across publishers and advertisers. Together, this body of literature establishes the cryptographic foundations required for secure, compliant, and utility-preserving marketing analytics.

The third major thread concerns federated learning (FL) for ad targeting, personalization, and cross-channel optimization. FL enables decentralized model training across user devices or organizational silos without raw data exchange, making it a natural match for privacy-sensitive advertising ecosystems. Su et al. developed an FL architecture combining transfer learning and knowledge distillation to improve targeted advertising performance under heterogeneous client data [13]. This was followed by comprehensive architectural proposals for FL-based smart advertising systems [14], as well as optimized data pipelines designed for latency-sensitive ad-delivery environments [15]. In parallel, commercial engineering studies explored using FL to enhance ad-delivery platforms without exposing personally identifiable information (PII), demonstrating improvements in click-through rates and bidding efficiency [16]. More recent work further examined multi-scenario optimization, investment-return modeling, and cross-regional ad-performance prediction under federated settings [17]. Other empirical studies provided evidence that FL can support privacy-preserving measurement tasks, including lift estimation and attribution [18], reinforcing its role in modern privacy-aligned advertising infrastructures.

Finally, recent surveys and conceptual analyses highlight the broader trend toward privacy centric advertising. Kraft emphasized the role of DP in modern targeting ecosystems and identified future policy-driven constraints that may make privacy-preserving optimization a mandatory baseline [19]. Complementary domain analyses such as Ullah et al.'s mobile-privacy survey [20] underscored the growing regulatory and technical pressure on platforms to adopt secure, decentralized, and privacy-preserving mechanisms for ad delivery and measurement.

Overall, the existing literature provides strong foundations but also reveals several persistent limitations: (1) DP mechanisms often degrade utility in high-dimensional multi-touch attribution settings; (2) cryptographic protocols incur computational overhead and lack economic-incentive integration; and (3) existing FL-based advertising architectures do not unify attribution measurement, incentive design, and cross-channel optimization within a single trustworthy framework. These gaps motivate the unified trusted-AI commercialization infrastructure proposed in this paper, aiming to reconcile privacy, utility, incentive alignment, and operational scalability across diverse marketing scenarios.

### 3. Method

This section formalizes the proposed federated and differentially private incentive marketing framework for cross-channel measurement in AI-powered digital commerce. As illustrated in Figure 1, the end-to-end system integrates client-side web and app privacy signals, unified channel encoders, a federated learning backbone trained under differential privacy, a multi-touch attribution and uplift module, and an incentive optimization layer for SMB advertisers. We first introduce the core entities and notation, then present the federated learning architecture that harmonizes web- and

app-side measurement streams. Next, we describe the differentially private reporting mechanisms that ensure consistency between event-level modeling and privacy-preserving aggregation. Finally, we detail the uplift-based incentive allocation module that translates privacy-compliant measurement into economically meaningful optimization tools for SMB participants.




Figure 1. Overview of the Proposed Federated-DP Cross-Channel Measurement Framework.

### 3.1. System Model and Notation

We consider a digital advertising ecosystem with users, advertisers, publishers, and a central orchestrator:

Users: the set of users is denoted by  $U$ .

Advertisers: the set of advertisers is denoted by  $A$ .

Publishers (web or app inventory owners): the set of publishers is denoted by  $P$ .

Channels: we denote channels by  $K$ , typically including web channels using Topics / Protected Audience and mobile channels using SKAdNetwork.

The central orchestrator (platform) is denoted by  $C$ .

For each user  $u \in U$ , we observe a sequence of ad-related interaction events across web and app channels,

Where:

$x_{u,t} \in \mathbb{R}^d$  is the context feature vector at time  $t$  (including device, placement, local engagement and coarse-grained metadata).

$a_{u,t} \in A \cup \emptyset$  is the action (which advertiser or campaign was shown;  $\emptyset$  denotes no ad).

$y_{u,t} \in 0,1$  is the binary outcome (e.g., conversion within a relevant time window).

$k_{u,t} \in K$  indicates the channel (e.g., web-Topics, web-Protected Audience, app-SKAdNetwork).

To integrate privacy-preserving web and app signals, we define channel-specific encoders:

For web events with Topics and Protected Audience (PA) signals,

$$z_{u,t}^{\text{web}} * u, t = \psi * \text{web}(x_{u,t}, \text{Topics} * u, t, \text{PA} * u, t),$$

Where  $\psi_{\text{web}}(\cdot)$  is a deterministic feature mapping that transforms event-level context and interest signals into a representation suitable for federated modeling.

For app events with SKAdNetwork reporting,

$$z_{u,t}^{\text{app}} * u, t = \psi * \text{app}(x_{u,t}, \text{SKAN} * u, t)$$

where  $\text{SKAN} * u, t$  captures SKAdNetwork postbacks and coarse-grained conversion values.

We write the unified representation as

$$z_{u,t} = \{z_{u,t}^{\text{web}} * u, t, \text{ if } k * u, t \text{ is a web channel, } z_{u,t}^{\text{app}} * u, t, \text{ if } k * u, t \text{ is an app channel.}$$

For brevity, we denote the per-event tuple as

$$e_{u,t} = (z_{u,t}, a_{u,t}, y_{u,t}, k_{u,t}), \mathcal{E} = \bigcup_{u \in U} \mathcal{E}_u.$$

Our goal is to learn a global model  $f_{\theta}$  with parameters  $\theta$  that supports:

Outcome prediction: estimate the probability of conversion under different ad actions,  $\hat{p}(y_{u,t} = 1 | z_{u,t}, a_{u,t}) = f_{\theta}(z_{u,t}, a_{u,t})$ .

Uplift estimation: estimate the incremental effect of showing an ad versus not showing an ad.

Multi-touch attribution (MTA): compute credit weights over the sequence of user touches.

Incentive allocation: translate uplift and attribution metrics into incentive scores and budget allocation while respecting privacy guarantees.

All of the above are implemented in a fully federated and differentially private fashion, mediated by the orchestrator  $C$ .

#### Federated Outcome and Uplift Modeling

Users (or small publisher/SMB clusters) act as federated clients. Let  $\mathcal{I}$  denote the index set of clients, where each client  $i \in \mathcal{I}$  holds a local dataset

$$D_i = (z_j, a_j, y_j)_{j=1}^{n_i}$$

derived from  $\mathcal{E}_i: u \in U_i$  and never leaves the device or local enclave.

We define the global empirical risk minimization (ERM) objective as

$$\min_{\theta} \mathcal{L}(\theta) = \sum_{i \in \mathcal{I}} \frac{n_i}{n} \mathcal{L}_i(\theta), \mathcal{L}_i(\theta) = \frac{1}{n_i} \sum_{j=1}^{n_i} \ell(f_{\theta}(z_j, a_j), y_j)$$

where  $n = \sum_{i \in \mathcal{I}} n_i$  and  $\ell(\cdot, \cdot)$  is a prediction loss, typically the logistic loss

$$\ell(\hat{y}, y) = -(y \log \hat{y} + (1 - y) \log (1 - \hat{y}))$$

The model  $f_{\theta}$  can be factorized into a feature encoder and task heads,

$$f_{\theta}(z, a) = \sigma(g_{\theta_{\text{head}}}(h_{\theta_{\text{enc}}}(z), a))$$

where  $h_{\theta_{\text{enc}}}$  is a shared representation network,  $g_{\theta_{\text{head}}}$  is a task-specific head, and  $\sigma(\cdot)$  is the logistic sigmoid.

To support incremental lift estimation, we employ a two-head structure with potential-outcome predictions, The individual-level uplift estimate is then

$$\tau(z) = \mu_1(z) - \mu_0(z).$$

The orchestrator  $C$  runs a variant of FedAvg over discrete communication rounds  $t = 0, 1, 2, \dots$

At round  $t$ :

$C$  maintains the current model  $\theta^{(t)}$  and selects a subset of clients  $S_t \subseteq \mathcal{I}$ .

Each client  $i \in S_t$  initializes its local model with  $\theta^{(t)}$  and performs  $E$  steps of local stochastic gradient descent (SGD),

$$\theta_i^{(t,0)} = \theta^{(t)}, \theta_i^{(t,e+1)} = \theta_i^{(t,e)} - \eta g_i^{(t,e)},$$

where  $g_i^{(t,e)}$  is the gradient on a mini-batch and  $\eta$  is the learning rate.

After  $E$  steps we denote the final local parameters as  $\theta_i^{(t+1)}$ .

The orchestrator aggregates the local updates to obtain the new global parameters,

$$\theta^{(t+1)} = \sum_{i \in S_t} \frac{n_i}{\sum_{j \in S_t} n_j} \theta_i^{(t+1)}.$$

This procedure is integrated with differential privacy in the next subsection.

### 3.2. Differentially Private Measurement and Attribution

The central requirement of the proposed framework is that all model training, attribution statistics, and reporting are  $(\epsilon, \delta)$ -differentially private with respect to individual users or SMBs.

### 3.3. Differential Privacy Preliminaries

A randomized mechanism  $\mathcal{M}$  that maps datasets to outputs satisfies  $(\epsilon, \delta)$ -differential privacy if for any pair of neighboring datasets  $D$  and  $D'$  differing in a single user's data and for any measurable set  $S$ ,

$$\Pr[\mathcal{M}(D) \in S] \leq e^{\epsilon} \Pr[\mathcal{M}(D') \in S] + \delta.$$

In our setting, mechanisms include:

The federated learning update mechanism (via DP-SGD).

The aggregate attribution and conversion reporting mechanisms (for web Attribution Reporting and SKAdNetwork-like summaries).

### 3.4. DP-SGD for Federated Learning

To guarantee differential privacy in training, we employ DP-FedAvg with Gaussian mechanism. At each round  $t$ , for each client  $i \in S_t$ :

For each example  $j$  in the local mini-batch  $B_i^{(t)}$ , compute the gradient  $g_{ij}^{(t)} = \nabla_{\theta} \ell_{ij}(\theta^{(t)})$  and clip its L2 norm:

$$g_{ij}^{-(t)} = g_{ij}^{(t)} / \max(1, \frac{\|g_{ij}^{(t)}\|_2}{C})$$

The server collects the clipped gradient sums from clients, adds Gaussian noise, and performs the update:

$$g^{-(t)} = \frac{1}{|St|} \left( \sum_{i \in St} \sum_{j \in B_i^{(t)}} g_{ij}^{-(t)} + N(0, \sigma^2 C^2 I) \right)$$

$$\theta^{(t+1)} = \theta^{(t)} - \eta \cdot g^{-(t)}$$

where  $C$  is the clipping norm,  $\sigma$  is the noise multiplier calibrated to the target  $(\epsilon_{\text{train}}, \delta_{\text{train}})$  budget using the Moments Accountant, and  $I$  is the identity matrix.

### 3.5. DP Aggregation for Cross-Channel Attribution

Let  $N_{a,k}$  denote the true number of conversions attributed to advertiser  $a \in A$  on channel  $k \in K$  over a given reporting window. We model the DP reporting mechanism for the summary reports of Attribution Reporting and SKAdNetwork as a noisy aggregation,

$$\tilde{N} * a, k = N * a, k + Z_{a,k}, Z_{a,k} \sim N(0, \sigma_{\text{rep}}^2).$$

To support multi-dimensional breakdowns (e.g., region, device type, campaign), we index by a tuple  $b \in \mathcal{B}$  and define

$$N_{a,k,b} = \sum_{u \in U} \sum_t 1[\text{event}(u, t) \text{ is attributed to } (a, k, b)],$$

with noisy release

$$\tilde{N} * a, k, b = N * a, k, b + Z_{a,k,b}, Z_{a,k,b} \sim N(0, \sigma_{\text{rep}}^2).$$

The sensitivity of this query under user-level adjacency is bounded by a constant  $\Delta$ , and the noise scale  $\sigma_{\text{rep}}$  is chosen according to the Gaussian mechanism to achieve  $(\epsilon_{\text{rep}}, \delta_{\text{rep}})$ -DP for reporting.

The final model is obtained by jointly optimizing a composite loss function that balances prediction accuracy, consistency with DP reports, and attribution alignment:

$$L^{\text{total}} = L^{\text{pred}} + \lambda_1 L^{\text{cons}} + \lambda_2 L^{\text{MTA}}$$

Here,  $L^{\text{pred}}$  is the standard binary cross-entropy loss for conversion prediction. The hyperparameters  $\lambda_1$  and  $\lambda_2$  control the trade-off between objectives and were set to 0.5 and 0.3 respectively, determined via grid search on a validation set. All parameters  $(\theta, \eta)$  are updated simultaneously through backpropagation in each federated round.

### 3.6. Consistency Constraints Between Event-Level and Summary-Level Views

To ensure that event-level models and summary reports are consistent, we introduce a consistency regularization term. Let  $\hat{N}_{a,k,b}(\theta)$  be the model-implied expected number of conversions for  $(a, k, b)$  under the current model:

$$\hat{N} * a, k, b(\theta) = \sum_{u \in U} \sum_t E[y_{u,t} | z_{u,t}, a_{u,t}, k_{u,t}; \theta], 1[(a_{u,t}, k_{u,t}, b_{u,t}) = (a, k, b)].$$

We define a consistency loss

$$\mathcal{L}^{\text{cons}}(\theta) = \sum_{a \in A} \sum_{k \in K} \sum_{b \in \mathcal{B}} w_{a,k,b} (\hat{N} * a, k, b(\theta) - \tilde{N} * a, k, b)^2,$$

where  $w_{a,k,b} \geq 0$  are weights that down-weight cells with high DP noise.

### 3.7. Multi-Touch Attribution and Incentive Allocation

The system must allocate credit across multiple touches in a cross-channel user journey and turn these credits into incentive scores for SMBs in a way that remains compatible with DP constraints.

### 3.7.1. Path-Based Attribution Weights.

For each user  $u$ , we consider the ordered sequence of ad touches up to a conversion (or censoring time),

$$p_u = ((a_{u,1}, k_{u,1}), (a_{u,2}, k_{u,2}), \dots, (a_{u,L_u}, k_{u,L_u})),$$

with  $L_u$  touches.

We construct contextual scores  $s_{u,\ell}$  for each touch,

$$s_{u,\ell} = \varphi_\eta(h_{\theta_{\text{enc}}}(z_{u,\ell}), a_{u,\ell}, k_{u,\ell}),$$

where  $\varphi_\eta$  is a small neural network with parameters  $\eta$  that maps the encoded representation and channel metadata to a scalar relevance score.

Attribution weights are then produced via a softmax over the path,

$$\alpha_{u,\ell} = \frac{\exp(s_{u,\ell})}{\sum_{j=1}^{L_u} \exp(s_{u,j})}, \quad \sum_{\ell=1}^{L_u} \alpha_{u,\ell} = 1.$$

If user  $u$  converts ( $y_u = 1$ ), the fractional credit assigned to touch  $(a_{u,\ell}, k_{u,\ell})$  is  $\alpha_{u,\ell}$ . The path-level estimator for attributed conversions of  $(a, k)$  is

$$\hat{N} * a, k^{\text{MTA}} = \sum_{* u \in \text{Uy}_u} \sum_{\ell=1}^{L_u} \alpha_{u,\ell} \mathbf{1}[(a_{u,\ell}, k_{u,\ell}) = (a, k)].$$

During training, we encourage consistency between  $\hat{N} * a, k^{\text{MTA}}$  and the DP-reported  $\tilde{N} * a, k$  via an additional loss term,

$$\mathcal{L}^{\text{MTA}}(\theta, \eta) = \sum_{a \in A} \sum_{k \in K} v_{a,k} (\hat{N} * a, k^{\text{MTA}} - \tilde{N} * a, k)^2,$$

with  $v_{a,k} \geq 0$  tuning channel-level trust in the DP aggregation.

### 3.7.2. Uplift-Based Incentive Scores.

For each advertiser  $a \in A$  and for each incentive program  $c$  (e.g., a subsidy scheme or discount on platform fees), we define an uplift-based benefit estimate. Let  $U * a, c$  denote the set of users targeted by advertiser  $a$  under incentive configuration  $c$ . For each user  $u \in U * a, c$ , we obtain the uplift estimate,

$$\tau_{u,a,c} = \tau(z_{u,c}) = \mu_1(z_{u,c}) - \mu_0(z_{u,c}),$$

where  $z_{u,c}$  summarizes the user and campaign context under  $c$ .

The aggregated uplift for  $(a, c)$  is

$$\hat{\tau} * a, c = \frac{1}{|U * a, c|} \sum_{u \in U * a, c} \tau * u, a, c.$$

To maintain DP at the advertiser level, we apply a noisy release,

$$\tilde{\tau} * a, c = \hat{\tau} * a, c + W_{a,c}, \quad W_{a,c} \sim N(0, \sigma_{\text{uplift}}^2),$$

where  $\sigma_{\text{uplift}}$  is calibrated such that the mapping from user-level data to  $\tilde{\tau} * a, c$  is  $(\epsilon * \text{uplift}, \delta_{\text{uplift}})$ -DP.

These  $\tilde{\tau}_{a,c}$  values form the core signals for incentive allocation.

Incentive Optimization Under Budget and Fairness Constraints

Let  $b_{a,c} \in \{0,1\}$  be a binary decision variable indicating whether advertiser  $a$  receives incentive program  $c$ . Let  $\kappa_c > 0$  be the cost of assigning program  $c$  to a single advertiser and  $B > 0$  be the total incentive budget for a given period.

We define a linear utility function based on DP uplift estimates,

$$U(b) = \sum_{a \in A} \sum_c b_{a,c} \tilde{\tau}_{a,c}.$$

The budget constraint is

$$\sum_{a \in A} \sum_c b_{a,c} \kappa_c \leq B.$$

To capture fairness toward SMBs, let  $A * SMB \subseteq A$  denote the subset of SMB advertisers and define a minimum allocation level  $\rho \in [0,1]$ . We impose

$$\frac{\sum_{a \in A * SMB} \sum_c b_{a,c} \kappa_c}{\sum_{a \in A} \sum_c b_{a,c} \kappa_c} \geq \rho,$$

ensuring that at least a fraction  $\rho$  of the incentive budget flows to SMBs.

The resulting optimization problem is

$$\max_b U(b) = \sum_{a \in A} \sum_c b_{a,c} \tilde{\tau}_{a,c}, \text{ s.t. } \sum_{a \in A} \sum_c b_{a,c} \kappa_c \leq B, \frac{\sum_{a \in A * SMB} \sum_c b_{a,c} \kappa_c}{\sum_{a \in A} \sum_c b_{a,c} \kappa_c} \geq \rho, \quad b_{a,c} \in [0,1], \forall a, c.$$

In practice, this integer program can be relaxed to  $b_{a,c} \in [0,1]$  and solved via convex optimization, or approximated using greedy heuristics that preserve DP guarantees (since DP already resides in the upstream signals  $\tilde{\tau}_{a,c}$  rather than in the solver itself).

## 4. Valuation

This section empirically evaluates the proposed Federated–Differentially Private cross channel incentive marketing framework. Experiments focus on three core questions:

It is important to note that in the current privacy-centric ecosystem, a direct comparison with traditional, identity-based third-party cookie systems is neither feasible nor methodologically sound. Therefore, our evaluation is designed to assess whether the proposed federated-DP framework can outperform other viable, privacy-compliant measurement techniques, thereby establishing its utility in the post-cookie era.

Does the federated + DP model preserve accuracy across web and app channels despite noise constraints?

Does the proposed multi-touch attribution (MTA) mechanism increase consistency with privacy-preserving summary reports?

Does uplift-based incentive allocation improve SMB performance under realistic privacy budgets?

We conduct experiments on a combined dataset synthesized from web Topics/Protected Audience signals and app–SKAdNetwork-like postbacks. To ensure realistic modeling, web events are grouped into 5 Topics categories and 6 PA cohorts, whereas app events contain coarse-grained conversion values (0–5) consistent with SKAdNetwork behavior.

### 4.1. Model Performance Across Channels

We simulate a cross-channel advertising dataset reflecting post-cookie constraints. User journey sequences are generated using a Markov model calibrated against publicly available ad interaction logs. Web events are annotated with synthetic Topics (5 categories, assigned via weighted random sampling) and Protected Audience cohort IDs (6 cohorts). App events are paired with SKAdNetwork-like postbacks, containing coarse-grained conversion values (0-5) and source identifiers.

We first evaluate prediction and uplift accuracy across channels — Web–Topics, Web–Protected Audience (PA), and App–SKAdNetwork. As shown in Table 1.

The following table summarizes performance across three core metrics:

AUC for conversion prediction

Calibration Error (ECE)

Uplift RMSE for counterfactual uplift modeling

**Table 1.** Cross-Channel Model Performance Under DP-FL Training.

Channel Type	#Events	AUC ( $\uparrow$ )	ECE ( $\downarrow$ )	Uplift RMSE ( $\downarrow$ )	DP Noise $\sigma$	Avg. Client Participation (%)
Web – Topics	3,245,901	0.782	0.041	0.116	1.2	27.4%
Web – Protected Audience	1,982,334	0.768	0.052	0.129	1.2	25.9%
App – SKAdNetwork	4,156,442	0.804	0.038	0.112	1.0	31.2%
Combined (Unified FL)	9,384,677	0.816	0.035	0.104	1.1	28.3%

#### 4.2. Attribution Consistency Under DP Summary Reporting

We next evaluate whether the proposed multi-touch attribution (MTA) system remains aligned with the differentially private summary reports from Attribution Reporting and SKAdNetwork. As shown in Table 2.

**Table 2.** Attribution Consistency Across Baseline and Proposed Method.

Method	ACR ( $\downarrow$ )	Avg. Per-Advertiser Error ( $\downarrow$ )	#Advertisers	Dimensionality of Reports	Supports Cross-Channel MTA
Baseline: Last-Touch (Non-DP)	0.214	38.6	220	Low	No
Baseline: Summary-Only DP Reports	0.167	29.3	220	Medium	No
Proposed: DP-Federated MTA (Ours)	0.091	17.5	220	High	Yes
Proposed + Consistency Regularization	0.072	14.2	220	High	Yes

Our framework reduces attribution inconsistency by more than 56% relative to DP summary-only reporting.

Consistency losses significantly enhance alignment between event-level and summary-level measurement.

Cross-channel MTA is only feasible with unified federated representation learning, as seen in the results.

#### 4.3. Incentive Allocation Outcomes for SMB Advertisers

Finally, we evaluate how uplift-based incentive allocation affects SMB campaign performance. As shown in Table 3.

Total incentive cost divided by the total number of incremental conversions (estimated via the uplift model). SMB Allocation Ratio: The proportion of the total incentive budget allocated to SMB advertisers.

**Table 3.** Incentive Allocation Performance Under DP Constraints.

Method	Incremental Lift $\uparrow$	CPIC ( $\downarrow$ )	SMB Allocation Ratio $\uparrow$	Budget Utilization (%)	#Advertisers Receiving Incentives
Heuristic Rule-Based Allocation	9.3%	\$41.2	48.1%	92.4%	62
Baseline DP Aggregated Lift	12.7%	\$34.5	53.8%	96.0%	85
Proposed DP-FL Uplift Allocation (Ours)	18.4%	\$27.8	56.7%	99.1%	104
Proposed + Strong SMB Constraint	17.6%	\$28.9	61.4%	98.3%	112

## 5. Conclusion

This paper presents a unified federated and differentially private framework for cross-channel marketing measurement and incentive optimization in post-cookie digital ecosystems. Motivated by the growing fragmentation between web and mobile advertising environments—where Topics, Protected Audience, Attribution Reporting, and SKAdNetwork each impose distinct privacy and reporting constraints—we propose an end-to-end architecture that harmonizes heterogeneous privacy-preserving signals while supporting accurate attribution, uplift estimation, and incentive allocation for advertisers, particularly SMBs operating under limited data resources.

Technically, the framework combines federated learning, differential privacy, multi-touch attribution, and uplift-based incentive modeling into an integrated system. A unified representation encoder allows both web-side and app-side privacy sandbox signals to be embedded into a common latent space, reducing channel asymmetries and improving cross-environment generalization. Differentially private SGD ensures that all model updates adhere to strict privacy budgets, and DP-based reporting mechanisms align event-level predictions with coarse-grained summary reports from Attribution Reporting and SKAdNetwork. A consistency regularization module further links federated predictions with noisy aggregate statistics, narrowing the long-standing gap between platform-level and device-level measurement pipelines.

Empirical evaluations demonstrate that the proposed architecture yields substantial gains across three critical dimensions:

- measurement utility, improving AUC, calibration, and uplift RMSE under realistic DP budgets;
- attribution consistency, reducing discrepancies between model-derived multi-touch paths and DP summary-level conversions;
- economic efficiency, delivering higher incremental lift and lower cost per incremental conversion in SMB-targeted incentive programs while satisfying fairness constraints.

Beyond technical contributions, this work has important practical implications. It shows that privacy-preserving advertising does not inherently preclude high-fidelity performance measurement or equitable incentive distribution. With appropriate algorithmic design—particularly unified federated learning, adaptive DP calibration, and uncertainty-aware attribution—it is possible to construct a scalable and regulator-aligned commercialization infrastructure capable of serving platforms, advertisers, and end users simultaneously. The framework therefore provides a viable blueprint for the next generation of trusted AI systems in digital commerce, where privacy guarantees, economic incentives, and measurement accuracy must coexist.

Nevertheless, several challenges remain. Federated learning at scale introduces optimization instability under heterogeneous device participation; differential privacy continues to impose noise ceilings on uplift estimation; and cross-channel temporal misalignment between Attribution Reporting and SKAdNetwork remains a structural barrier to perfect attribution harmonization. Future work should explore causal representation learning, asynchronous federated optimization with privacy accounting, and mechanism-design-aligned incentive policies that can reduce strategic manipulation while maintaining fairness for SMBs.

In conclusion, the proposed federated-DP measurement framework represents a significant step toward rebuilding marketing effectiveness in a privacy-centric era. By integrating regulatory-compliant data flows with robust attribution and incentive modeling, it establishes a principled, scalable foundation for trustworthy, AI-powered digital commerce.

## References

1. Xiao, Y., Du, J., Zhang, S., Zhang, W., Yang, Q., Zhang, D., & Kifer, D. (2025). Click without compromise: Online advertising measurement via per-user differential privacy. In Proceedings of the IEEE Symposium on Security and Privacy.

2. Delaney, J., Ghazi, B., Harrison, C., Ilvento, C., Kumar, R., Manurangsi, P., Pal, M., Prabhakar, K., & Raykova, M. (2024). Differentially private ad conversion measurement. *Proceedings on Privacy Enhancing Technologies*, 2024(2).
3. Du, J. (2024). Designing for user privacy: Integrating differential privacy into ad measurement systems in practice. In *Proceedings of the 2024 USENIX Conference on Privacy Engineering Practice and Respect (PEPR '24)*.
4. Du, J., Ghazi, B., Ilvento, C., Kumar, R., Manurangsi, P., Pal, M., Prabhakar, K., & Raykova, M. (2025). PrivacyGo: Privacy-preserving ad measurement with multidimensional intersection. *arXiv preprint*.
5. Sun, J., Zhao, L., Liu, Z., Li, Q., Deng, X., Wang, Q., & Jiang, Y. (2022). Practical differentially private online advertising. *Computers & Security*, 112, 102504.
6. Lindell, Y., & Omri, E. (2011). A practical application of differential privacy to personalized online advertising. *IACR Cryptology ePrint Archive*, 2011(152).
7. Mouris, D., Masny, D., Trieu, N., Sengupta, S., Buddhavarapu, P., & Case, B. M. (2024). Delegated private matching for compute. *Proceedings on Privacy Enhancing Technologies*, 2024(2), 49–72.
8. Mouris, D., Sarkar, P., & Tsoutsos, N. G. (2024). PLASMA: Private, lightweight aggregated statistics against malicious adversaries. *Proceedings on Privacy Enhancing Technologies*, 2024(3), 4–24.
9. Anonymous. (2020). Secure multiparty computation for private measurement of advertising lift. *Technical Disclosure Commons*.
10. Zhong, K., Ma, F., & Angel, S. (2022). Ibex: Privacy-preserving ad conversion tracking and bidding. In *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security (CCS '22)*.
11. Aksu, H., Aksu, H. H., Ghazi, B., Harrison, C., Kumar, R., Manurangsi, P., Pal, M., & Raykova, M. (2024). Summary report optimization in the Privacy Sandbox Attribution Reporting API. *Proceedings on Privacy Enhancing Technologies*, 2024(4).
12. Ghazi, B., Harrison, C., Hosabettu, A., Kamath, P., Knop, A., Kumar, R., Leeman, E., Manurangsi, P., & Sahu, V. (2024). On the differential privacy and interactivity of Privacy Sandbox reports. *arXiv preprint*.
13. Su, C., Wei, J., Lei, Y., & Li, J. (2023). A federated learning framework based on transfer learning and knowledge distillation for targeted advertising. *PeerJ Computer Science*, 9, e1496.
14. Seyghaly, R., Garcia, J., & Masip-Bruin, X. (2024). A comprehensive architecture for federated learning-based smart advertising. *Sensors*, 24(12), 3765.
15. Seyghaly, R., Garcia, J., Masip-Bruin, X., & Mahmoodi Varnamkhasi, M. (2024). An optimized data architecture for smart advertising based on federated learning. In *2024 IEEE Symposium on Computers and Communications (ISCC)*.
16. Chivukula, V. V. (2022). Use of federated learning for optimizing ad delivery platforms without exchanging user PII. *International Journal of Science and Advanced Technology*, 12(3).
17. Zhang, K., & Li, P. (2024). Federated learning optimizing multi-scenario ad targeting and investment returns in digital advertising. *Journal of Advanced Computing Systems*, 4(8), 36–43.
18. Chivukula, V. V. (2022). The use of federated learning for digital advertising measurement. *ESP Journal of Engineering & Technology Advancements*, 2(4), 161–162.
19. Kraft, L. (2023). Leveraging differential privacy for targeted advertisements. *SSRN Electronic Journal*.
20. Ullah, I., Boreli, R., & Kanhere, S. S. (2023). Privacy in targeted advertising on mobile devices: A survey. *International Journal of Information Security*, 22(3), 647–678.

**Disclaimer/Publisher's Note:** The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.