

Hierarchical Bayesian Experimental Estimation of Incremental Local Advertising Effects with Secure Aggregation and Differential Privacy

Que Yang

LinseyQuilici062@outlook.com

Carnegie Mellon University, Pittsburgh, 15213, United States

Yukun Du

University of Chicago, Chicago, 60637, United States

Yuang Hu

Carnegie Mellon University, Pittsburgh, PA 15213, United States

Research Article

Keywords: local business advertising, differential privacy, regional randomized experiments, hierarchical Bayesian models, incremental conversions, advertising effectiveness measurement

Posted Date: September 23rd, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11122128/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: The authors declare no competing interests.

Hierarchical Bayesian Experimental Estimation of Incremental Local Advertising Effects with Secure Aggregation and Differential Privacy

Que Yang*

Carnegie Mellon University, Pittsburgh, 15213, United States, quey@alumni.cmu.edu

Yukun Du

University of Chicago, Chicago, 60637, United States, yukundu@uchicago.edu

Yuang Hu

Carnegie Mellon University, Pittsburgh, PA 15213, United States, yuanghai@alumni.cmu.edu

Abstract

Local commercial advertising measurement on a U.S. digital media platform is computationally constrained by sparse campaign-level conversions, the attribution bias of last-click rules, and privacy restrictions that limit the availability of user-level behavioral data. A paired-area randomized experimental framework was implemented for 1,860 local businesses across 12 U.S. states, covering 4,720 advertising campaigns, 1.26 billion impressions, 19.4 million website visits, and 2.18 million purchase or appointment records. The computational pipeline uses Z-score normalization, KD-tree candidate retrieval, Mahalanobis-distance matching, and a SHA-256-derived random seed for paired-area assignment; regional weekly records are then processed through secure aggregation and a Gaussian differential privacy mechanism with $\epsilon = 2$ and $\delta = 10^{-6}$. On the privacy-protected regional statistics, a hierarchical Bayesian negative binomial inference module was constructed with business, industry, metropolitan area, and ad placement week modeled as four random-effect levels. The likelihood explicitly incorporates the recorded privacy-noise variance, enabling the inference of incremental conversions and incremental return on ad spend without treating privacy-perturbed counts as noise-free observations. Experimental evaluation showed that the proposed computational framework reduced the mean absolute error of incremental conversions from 71.4 for the standard difference-in-differences baseline to 43.6, achieved 91.2% coverage for the 90% interval, and reduced the absolute error of incremental return on ad spend from 0.39 to 0.17. For campaigns with fewer than 500 conversions, hierarchical information sharing reduced the minimum detectable effect from 8.6 to 5.9 percentage points and the median posterior standard deviation from 0.084 to 0.061; Gaussian privacy perturbation introduced an additional error of 6.3 conversions, while 83% of enterprise-level budget decisions remained unchanged. The resulting framework integrates regional randomized computation, secure aggregation, differential privacy, hierarchical Bayesian inference, and posterior decision rules into a privacy-preserving advertising measurement pipeline for incremental-effect estimation and budget optimization.

CCS CONCEPTS:

Information systems~Information systems applications~Multimedia information systems~Multimedia content creation

Keywords

local business advertising; differential privacy; regional randomized experiments; hierarchical Bayesian models; incremental conversions; advertising effectiveness measurement

1 Introduction

Digital advertising is now a way for local businesses to reach people in the area, expand their online market, and spend marketing funds. Local restaurants, retailers, healthcare institutions and home-service companies are generally characterised by low advertising investment, few conversion indicators and short campaign cycles. The old model of last-click attribution cannot distinguish between ad-driven conversions and organic demand, and thus often overstates the incremental effect. At the same time, user-level behavioural tracking is increasingly subject to privacy restrictions. Although secure aggregation and

differential privacy can reduce the risk of individual data leakage, they add noise and weaken the statistical identification capability for small-sample campaigns; thus, this paper selects the U.S. local business advertising measurement scenario. A zip-code-level paired randomised experiment is employed as the foundation for identification, and a Bayesian negative binomial estimation framework incorporating secure aggregation, differential privacy noise correction and multi-layer parameter sharing is constructed. We focus on the joint inference of incremental conversions, incremental return on ad spend (ROAS) and budget allocation under privacy-disturbed conditions. By conducting model comparison, privacy-sensitivity tests, and small-sample identification analysis, this paper verifies the applicability of the framework in terms of estimation accuracy, interval calibration and decision stability, and provides a statistical basis for evaluating the effectiveness of local business advertising and optimising limited budgets in a privacy-preserving environment.

2 Analysis of Local Commercial Advertising Effectiveness Evaluation Under Privacy Constraints

Local advertising uses the postal area-week as the computational unit, with the advertising group $Z=1$ and the control group $Z=0$. The 4-dimensional matching vector includes population size, historical sales, industry category, and ad spend[1]. The event-level purchase or reservation count $Y \in \mathbb{N}_0$ exhibits a concentration of zero values, variance greater than the mean, and inter-week correlation, while the secure aggregation interface can only return noisy statistics:

$$\tilde{Y}_{ipt} = Y_{ipt}(Z_{ipt}) + \eta_{ipt}, \eta_{ipt} \sim N(0, \sigma_{DP}^2) \quad (1)$$

where i 、 p 、 t represent the company, matched region, and ad placement week, respectively; $\tilde{Y}_{ipt}$ is the accessible count; $Y_{ipt}(Z_{ipt})$ is the potential conversion volume under the processing state; and η_{ipt} and σ_{DP}^2 are the privacy noise and its variance, respectively. Thus, two measures — incremental conversions and incremental returns on advertising — need to be determined over a four-level hierarchy that includes companies, industries, metropolitan areas, and ad placement weeks, thus providing constraints for subsequent paired sampling, noise correction, and stratification[2].

3 Design of a Stratified Experimental Estimation Method for Incremental Advertising Effectiveness Under Privacy Constraints

3.1 Advertising Effect Sampling Method Based on Paired Regional Randomized Experiments

The sampling module uses postal areas as randomization units and creates separate pools for each of the nine ad categories, such as food and beverage. The three continuous variables — population size, historical sales, and advertising spending — are standardized with Z-scores, and industry categories are defined as exact-match constraints; a KD-tree retrieves the top 5 candidates, and the top 1 is selected according to Mahalanobis distance, so that the matching process can be completed in 1:1 without substitution. The system extracts the 64-bit SHA-256 hash of the campaign ID as a random seed, assigning $Z_p=1$ as the ad zone and $Z_p=0$ as the control zone. Only 5 fields are left in the sample table: Enterprise ID, Campaign ID, Area ID, Week ID, and Processing Flag. To avoid duplication of locations and cross regions, a regional hash index is used [3]. In order to visually illustrate the differences between the different types of advertisements, Figure 1 illustrates nine typical advertising images. Subsequent counts are aggregated by region and week in the Security Aggregation module.



Figure 1: Examples of 9 Categories of Local Commercial Advertisements

3.2 Method for Generating Differential Privacy Statistics Based on Secure Aggregation

A three level index is set up for outgoing records according to treatment group, post area p , and placement week t . The local nodes encode impressions, visits, conversions, and ad spend into 64, 4 dimensional fixed point vectors, and perform L2 contribution truncation at the user – week level [4]. A 256-bit key is used by each node to negotiate and produce paired masks, the sum of which is expressed as:

$$s_{pt} = \sum_{k=1}^K \left(a_{kpt} + \sum_{\ell > k} m_{k\ell} - \sum_{\ell > k} m_{\ell k} \right) \quad (2)$$

where s_{pt} is the regional weekly aggregation vector, K is the number of shards participating in the calculation, a_{kpt} is the sum of truncated contributions within shard k , and $m_{k\ell}$ is the random mask shared by shards k and ℓ ; positive and negative masks cancel each other out during summation, preventing the server from reading data from individual shards. After aggregation is complete, statistics are published using a Gaussian mechanism:

$$\tilde{s}_{pt} = s_{pt} + \xi_{pt}, \xi_{pt} \sim N(0, \sigma^2 C^2 I_4), \sigma \geq \frac{\sqrt{2 \ln(1.25 / \delta)}}{\epsilon} \quad (3)$$

where $\tilde{s}_{pt}$ is the noisy output, ξ_{pt} is the Gaussian noise, C is the upper bound on user contribution sensitivity, I_4 is the 4×4 identity matrix, σ is the noise multiplier, and ϵ and δ are the privacy budget and failure probability, respectively. The system uses a Rényi differential privacy counter to accumulate weekly queries, constraining the total budget to $\epsilon = 2$ and $\delta = 10^{-6}$, while preserving C , σ , and the noise covariance. The negative noise value is not truncated so that the full error distribution is reproduced in the probability function [5].

3.3 Hierarchical Bayesian Estimation Model Incorporating Privacy-Preserving Noise Correction

The model reads the noisy conversion volume $\tilde{Y}_{ipt}$ and the noise variance w_{pt}^2 from the output, sets

the unperturbed count Y_{ipt} as a discrete latent variable, and constructs a negative binomial count layer and a Gaussian privacy observation layer:

$$Y_{ipt} \sim \text{NB}(\mu_{ipt}, \phi), \tilde{Y}_{ipt} | Y_{ipt} \sim N(Y_{ipt}, w_{pt}^2) \quad (4)$$

where i 、 p 、 t represent the company, matched region, and campaign week, respectively; μ_{ipt} is the latent conversion mean; ϕ is a discrete parameter; and $\text{Var}(Y_{ipt}) = \mu_{ipt} + \mu_{ipt}^2 / \phi$; w_{pt}^2 are read from the differential privacy metadata. The count mean adopts a 4-level hierarchical structure:

$$\log \mu_{ipt} = \log E_{ipt} + \alpha + (\beta + b_i)Z_p + \mathbf{x}_{ipt}^T \gamma + u_i + v_{h(i)} + w_{m(i)} + s_t \quad (5)$$

In this model, E_{ipt} represents the exposure bias term, α represents the population intercept, $Z_p \in \{0,1\}$ represents the treatment label, β represents the population advertising effect, b_i represents the firm-level random slope, and $\mathbf{x}_{ipt}$ and γ represent the standardized covariates and their coefficients, respectively; u_i 、 $v_{h(i)}$ 、 $w_{m(i)}$ and s_t denote the random intercepts for firm, industry, metropolitan area, and advertising week, respectively. Let $(u_i, b_i)^T \sim N_2(0, \Sigma_b)$; the remaining random effects follow a zero-mean normal prior, with standard deviations following a half-t prior. The model uses decentralized parameters and NUTS sampling, runs 4 chains, with 4,000 iterations per chain and a warm-up of the first 2,000 iterations; the target acceptance rate is set to 0.90. From the same posterior sample, counterfactual conversion volumes are generated by setting $Z_p = 1$ and $Z_p = 0$, respectively, to compute incremental conversions and the probability of budget decisions[6].

3.4 A Method for Evaluating Advertising Decisions Based on Posterior Inference

Using 8,000 sets of posterior samples retained from four NUTS chains, and holding the firm, region, week, and covariates constant, we set the treatment variable to $Z_p = 1$ and $Z_p = 0$, respectively. We then use two-stage vectorization inference to generate the latent mean conversions for the advertising and natural states. The firm-level evaluation metric is defined as:

$$\delta_{ipt}^{(s)} = \mu_{ipt}^{(s)}(1) - \mu_{ipt}^{(s)}(0), \Delta_i^{(s)} = \sum_{p,t} \delta_{ipt}^{(s)}, R_i^{(s)} = \frac{\sum_{p,t} v_{ipt} \delta_{ipt}^{(s)}}{A_i} \quad (6)$$

where s is the a posteriori sample index; $\mu_{ipt}^{(s)}(1)$ and $\mu_{ipt}^{(s)}(0)$ represent the two latent means for firm i in region p and week t ; $\delta_{ipt}^{(s)}$ is the local increment; $\Delta_i^{(s)}$ is the cumulative incremental conversions; v_{ipt} is the conversion value; A_i is the ad spend; and $R_i^{(s)}$ is the incremental return on ad spend. Budget actions are determined by minimizing a posteriori risk:

$$a_i^* = \arg \max_{a \in A} \frac{1}{S} \sum_{s=1}^S L(a, \Delta_i^{(s)}, R_i^{(s)}) \quad (7)$$

where $S = 8000$, A include four actions: increase, maintain, reduce, and stop, and $L(\cdot)$ is the budget misallocation loss function. The decision engine outputs the 5th, 50th, and 95th percentile at the same time; the budget is increased when the posterior probability of the return is not less than 0.90, and the 90 per cent confidence limit is greater than 0; when the probability is between 0.60 and 0.90, it decreases when the probability is less than 0.60, and stops when the probability is not lower than 0.90 [7]. Figure 2 illustrates the execution process of posterior reasoning, counterfactual generation, risk assessment, and action output.




Figure 2: Process for Evaluating Incremental Advertising Effectiveness and Budget Decision-Making Based on Posterior Inference

4 Experimental Results and Analysis

4.1 Construction of the Experimental Dataset and Setting of Evaluation Metrics

The experimental data come from 1,860 local businesses in the 12 U.S. states, and these four industries are: food service, retail, healthcare and home repair. A total of 4,720 advertising campaigns have been run, with 1.26 billion impressions, 19.4 million website visits and 2.18 million purchase or appointment records. On average, each business has about 2.54 campaigns, and each campaign has around 266,900 impressions, 4,110 visits and 462 conversions. The conversion rate of impressions to visits was 1.54%, and the conversion rate of visits to purchases or appointments was 11.24%. The data had a small sample size at the campaign level and cross-industry imbalances. All records were grouped by company ID to prevent leakage of activities from the same company across sets, and independent regional experiments were used as the basis for effect calibration; weekly regional statistics were generated using a Gaussian differential privacy mechanism with $\epsilon=2$ and $\delta=10^{-6}$. Evaluation metrics included the mean absolute error (MAE) of incremental conversions, 90% posterior interval coverage, absolute error in incremental return on ad spend, minimum detectable effect (MDE), median posterior standard deviation, additional error from privacy noise, and budget decision consistency rate. For the small-sample analysis, MDE was evaluated for campaigns with fewer than 500 observed conversions using the 8,000 retained posterior draws. The detectable-effect threshold was determined from the posterior distribution under the same positive-effect criterion used for posterior decision inference, so that MDE and posterior uncertainty were evaluated on a common probabilistic scale [8].

4.2 Comparison of Performance Between the Stratified Bayesian Model and Baseline Methods

Using the same weekly statistics for paired regions and enterprise-level data segmentation, we tested the standard double-difference model and the stratified Bayesian negative binomial model proposed in this paper, with incremental conversions measured in independent regional experiments as the reference. The results show that the mean absolute error (MAE) of incremental conversions decreased from 71.4 for the standard difference-in-differences model to 43.6 for the hierarchical Bayesian model, an absolute reduction of 27.8 and a relative reduction of 38.9%. The absolute error in incremental return on ad spend decreased from 0.39 to 0.17, corresponding to a 56.4% reduction. These point-error measures were considered jointly with posterior dispersion in the small-sample analysis rather than being interpreted independently, allowing the reduction in estimation error to be assessed together with the change in posterior uncertainty; and the 90% confidence interval coverage rate of the model in this paper is 91.2%, which is 1.2 percentage points lower than the theoretical uncovered level of 10.0%. The above results show that four-level parameter sharing across enterprises, industries, metropolitan areas and ad placement weeks can reduce the estimation variance of sparse events, and the negative binomial likelihood effectively

handles the highly discrete nature of conversion counts [9-10]. Specific results are as follows: Table 1

Table 1: Comparison of Estimation Performance Between the Hierarchical Bayesian Model and the Baseline Method

Evaluation Metrics	Reference Method or Nominal Value	Model in This Paper	Absolute Change	Relative Change
Conversion Increment MAE/Iteration	71.4 (standard double difference)	43.6	-27.8	-38.9%
Absolute error in incremental return on ad spend	0.39 (ordinary double difference)	0.17	-0.22	-56.4%
90% Confidence Interval Coverage / %	90.0 (nominal level)	91.2	+1.2 percentage points	—

4.3 Analysis of the Impact of Differential Privacy Mechanisms on Estimation Performance

We compared the results of the regional weekly statistics for 4,720 advertising campaigns under conditions without privacy perturbation and those with Gaussian differential privacy ($\epsilon=2$, $\delta=10^{-6}$). Given the privacy constraints, the model's incremental conversion MAE was 43.6 times, and an additional error of 6.3 times was added due to privacy noise; without noise, the MAE was 37.3 times. After adding the privacy mechanism, the error increased by 16.9% compared with the no-noise case, but it was still lower than the 71.4 times obtained by a standard double-difference model. In terms of budget output, 83% of the companies kept their original budget retention or adjustment decisions, and only 17% changed their actions; thus, the noise variance correction in the likelihood function did not directly alter the advertising effect parameters. Both the increase in error and the stability of the decision can be shown in Figure 3: the left side shows the change in error from 37.3 to 43.6, and the right side shows that 83% of the companies maintained their conclusion while only 17% changed it. Therefore, we will now investigate the problem of privacy noise amplification in low-conversion campaigns.




Figure 3: Impact of the differential privacy mechanism on incremental estimation error and budget decision stability

4.4 Analysis of Small-Sample Ad Campaign Performance Identification and Budget Optimization

For advertising campaigns with fewer than 500 conversions in the observation period, the non-hierarchical and hierarchical estimators were evaluated on the same campaign subset, paired regional weekly statistics, covariates, differential-privacy setting of $\epsilon=2$ and $\delta=10^{-6}$, and posterior evaluation

procedure. The comparison therefore changed the hierarchical information-sharing component while holding the underlying observations and privacy conditions fixed, avoiding differences in sample composition from affecting the MDE comparison. Under this matched evaluation setting, the MDE decreased from 8.6 percentage points without hierarchical information sharing to 5.9 percentage points with hierarchical sharing, corresponding to an absolute reduction of 2.7 percentage points and a relative reduction of 31.4%. Because both estimators were evaluated on the same low-conversion campaigns and privacy-protected regional statistics, the reduction reflects the change in the information-sharing structure rather than a change in campaign composition or privacy budget. The median posterior standard deviation decreased from 0.084 to 0.061, an absolute reduction of 0.023 and a relative reduction of 27.4%. In variance terms, the corresponding squared posterior dispersion decreased from 0.00706 to 0.00372, a reduction of approximately 47.3%. Together with the reduction in MDE from 8.6 to 5.9 percentage points, this result indicates that hierarchical information sharing reduced posterior dispersion without obtaining the lower uncertainty by imposing a higher detectable-effect threshold. The result therefore characterizes the error–variance trade-off as a simultaneous improvement in small-sample detection sensitivity and posterior concentration. Budget optimisation no longer relies on a single-point estimate; instead, the posterior median, 5th and 95th percentiles, and the probability of achieving return targets are fed into four action classifiers—*increase*, *maintain*, *reduce*, and *stop*—to reduce budget misallocation in small-sample campaigns caused by short-term conversion spikes. Figure 4 jointly presents the change in MDE and posterior uncertainty: hierarchical information sharing reduced the MDE from 8.6 to 5.9 percentage points while reducing the median posterior standard deviation from 0.084 to 0.061, providing a joint view of detection sensitivity and posterior dispersion under the same small-sample setting.




Figure 4: Comparison of Performance Identification Capabilities and Posterior Uncertainty in Small-Sample Advertising Campaigns

5. Conclusions

Evaluating the effectiveness of local commercial advertising under privacy constraints requires striking a balance among data availability, statistical reliability, and individual privacy protection. Based on a hierarchical Bayesian estimation framework that incorporates regional randomized experiments, secure aggregation mechanisms, and fusion noise correction, this method achieves stable identification of incremental effects in scenarios with small sample sizes and highly discrete conversions, and further supports budget adjustment decisions through posterior inference. This method overcomes the reliance of traditional attribution models on user-level behavioral data, providing an interpretable statistical approach to ad measurement in privacy-protected environments. However, the regional randomized design remains limited by external market changes, cross-regional spillover effects, and scalability across industry

heterogeneity. Future research could incorporate dynamic experimental data over longer time periods, introduce time-varying causal structures, and employ multi-source privacy-preserving computation mechanisms to further enhance the adaptability and decision accuracy of incremental ad effect identification in complex market environments.

References

- [1] Barrientos A F, Williams A R, Snoke J, et al. A feasibility study of differentially private summary statistics and regression analyses with evaluations on administrative and survey data[J]. *Journal of the American Statistical Association*, 2024, 119(545): 52-65.
- [2] Zhang Y, Xu Q, Tang N, et al. Differentially private data release for mixed-type data via latent factor models[J]. *Journal of Machine Learning Research*, 2024, 25(116): 1-37.
- [3] Samartsidis P, Seaman S R, Harrison A, et al. A Bayesian multivariate factor analysis model for causal inference using time-series observational data on mixed outcomes[J]. *Biostatistics*, 2024, 25(3): 867-884.
- [4] Ye T, Westling T, Page L, et al. Nonparametric identification of causal effects in clustered observational studies with differential selection[J]. *Journal of the Royal Statistical Society Series A: Statistics in Society*, 2024, 187(4): 1013-1036.
- [5] Guo Q, Barrientos A F, Peña V. Differentially private methods for compositional data[J]. *Journal of Computational and Graphical Statistics*, 2025, 34(3): 848-858.
- [6] Imbens G, Kallus N, Mao X, et al. Long-term causal inference under persistent confounding via data combination[J]. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 2025, 87(2): 362-388.
- [7] Kallus N, Mao X. On the role of surrogates in the efficient estimation of treatment effects with limited outcome data[J]. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 2025, 87(2): 480-509.
- [8] Yu P, Jiang Y, Su Z, et al. Differentially Private Bayesian Envelope Regression via Sufficient Statistic Perturbation[J]. *Journal of Data Science*, 2026, 24(1): 187.
- [9] Runge J, Pauwels K. Open-source media and marketing mix modeling: Practice-oriented overview, challenges, opportunities[J]. *Customer Needs and Solutions*, 2026, 13(1): 3.
- [10] Luo J, Qin H, Xu Z. Differential privacy statistical inference for a directed network model with covariates[J]. *Journal of Systems Science and Complexity*, 2026, 39(3): 1285-1305.