

Causal Inference Under Threshold Manipulation: Bayesian Mixture Modeling and Heterogeneous Treatment Effects

Kohsuke Kubota¹, Shonosuke Sugasawa²

¹NTT DOCOMO, INC., Tokyo, Japan

²Faculty of Economics, Keio University, Tokyo, Japan

kousuke.kubota.xt@nttdocomo.com, sugasawa@econ.keio.ac.jp

Abstract

Many marketing applications, including credit card incentive programs, offer rewards to customers who exceed specific spending thresholds to encourage increased consumption. Quantifying the causal effect of these thresholds on customers is crucial for effective marketing strategy design. Although regression discontinuity design is a standard method for such causal inference tasks, its assumptions can be violated when customers, aware of the thresholds, strategically manipulate their spending to qualify for the rewards. To address this issue, we propose a novel framework for estimating the causal effect under threshold manipulation. The main idea is to model the observed spending distribution as a mixture of two distributions: one representing customers strategically affected by the threshold, and the other representing those unaffected. To fit the mixture model, we adopt a two-step Bayesian approach consisting of modeling non-bunching customers and fitting a mixture model to a sample around the threshold. We show posterior contraction of the resulting posterior distribution of the causal effect under large samples. Furthermore, we extend this framework to a hierarchical Bayesian setting to estimate heterogeneous causal effects across customer subgroups, allowing for stable inference even with small subgroup sample sizes. We demonstrate the effectiveness of our proposed methods through simulation studies and illustrate their practical implications using a real-world marketing dataset.

Extended version —

<https://doi.org/10.48550/arXiv.2509.19814>

1 Introduction

Estimating causal effects of thresholds is important for the effective design of marketing strategies. Many marketing applications, including loyalty programs (e.g. Kivetz, Urminsky, and Zheng 2006; Nastasoiu et al. 2021) and credit card incentive programs, offer rewards to customers who exceed specific spending thresholds to encourage increased consumption. Regression discontinuity design (RDD) is a representative method to estimate the causal effect of thresholds on customer behavior (Hahn, Todd, and der Klaauw 2001; Thistlethwaite and Campbell 1960). This approach estimates the causal effect by using local randomization around the

threshold and is widely applied as one of the most credible research designs for causal inference, particularly because its identification assumptions are considered weak and plausible (Cattaneo, Idrobo, and Titiunik 2024; McCrary 2008).

However, the local randomization assumption of RDD can be violated when customers, aware of the thresholds, strategically manipulate their spending to qualify for the rewards. This strategic behavior of customers at the thresholds breaks the local randomization assumption supported by the *continuity* condition (Hahn, Todd, and der Klaauw 2001), which states that the potential outcomes of customers just below and above the threshold are the same. This condition cannot be satisfied when customers can manipulate their behavior to reach the thresholds because they may have different potential outcomes depending on whether they are above or below the threshold (Ishihara and Sawada 2024; McCrary 2008; Lee 2008). Although the economics literature proposes the bunching estimation (e.g. Bertanha et al. 2024) to address this issue, this framework relies on the assumption that a segment of the distribution just before the threshold, where the probability density function becomes zero, is what is transformed into the observed “bunch” at the threshold. Consequently, the practical applicability of these methods can be limited in scenarios where such a zero-density region cannot be assumed, which often arises in marketing applications.

To address causal inference under manipulation, we propose a novel Bayesian framework called Bayesian Modeling of Threshold Manipulation via Mixtures (BMTM). This framework employs a Bayesian mixture model to distinguish between customers strategically affected by a threshold and those who are not, enabling robust causal effect estimation with uncertainty quantification (see Figure 1). To fit the mixture model, we adopt a two-step Bayesian approach consisting of modeling non-bunching customers and fitting a mixture model to a sample around the threshold. In addition, we extend BMTM to a hierarchical Bayesian setting to estimate heterogeneous causal effects across subgroups of customers, which we call hierarchical BMTM (HBMTM). This hierarchical structure ensures stable estimation, even for small subgroups, by borrowing information across them to meet the increasing demand in marketing for estimating heterogeneous causal effects.

We demonstrate the effectiveness of our proposed frame-




Figure 1: Conceptual illustration of our proposed mixture model (red). Modeling the observed distribution (skyblue) as a mixture of a non-bunching distribution (green, unaffected by the threshold) and a strategic bunching distribution (blue, distorted near the threshold).

work through simulation studies and analysis of a real-world marketing dataset. The results of the simulation studies indicate that our proposed methods achieve more accurate and reliable estimates of causal effects than conventional RDD methods. In particular, HBMTM provides stable estimates of heterogeneous causal effects across customer subgroups, even in scenarios with small sample sizes. The real-world case study further illustrates the practical implications and utility of our framework.

2 Related Work

This section reviews related work on bunching estimation and methods to estimate heterogeneous causal effects using RDD. The bunching estimation examines the discontinuity in the distribution of the running variable. In the literature, there are two types of discontinuities: kinks, where the slope of an incentive schedule changes (e.g. Saez 2010), and notches, where the level of the incentive changes discontinuously (e.g. Slemrod 2013). Marketing incentive programs, which offer a discrete reward upon reaching a threshold, typically fall into the notch category. In addition to the literature in Section 1, the existing bunching estimation methods (e.g. Kleven and Waseem 2013) assume sharp bunching, where individuals can precisely target and locate themselves at the threshold. However, this assumption is often too strong in real-world marketing scenarios because customers may not be fully aware of the exact threshold value or face constraints and imperfect control over their spending behavior. In contrast to these existing approaches that assume sharp bunching, our proposed framework explicitly models the entire distribution of customer behavior around the threshold. By employing a mixture model for density estimation, our proposed method naturally captures the diffuse nature of bunching that arises from customers' imperfect control, allowing for causal effect estimation under more plausible assump-

tions for marketing applications.

In RDD (without threshold manipulation), there are some methods to estimate the heterogeneous treatment effect across subgroups. For example, Sugawara, Ishihara, and Kurisu (2023) developed a hierarchical model to estimate the subgroup-specific causal effect, and Alcantara et al. (2024) and Tao, Wang, and Ruppert (2025) employed the Bayesian additive regression tree to capture heterogeneous effects. However, these methods cannot be applied under threshold manipulation. To our knowledge, no attempts have been made to estimate the heterogeneous causal effects under threshold manipulation. Therefore, this paper would be the first to provide an effective method for estimating heterogeneous causal effects under threshold manipulation.

3 Bayesian Modeling of Threshold Manipulation via Mixtures

3.1 Causal Inference Under Threshold Manipulation

Consider the problem of estimating the causal effect of a threshold K on customers with the potential outcome framework (Imbens and Rubin 2015; Rubin 1974). Let $Y(1)$ be the potential outcome with threshold K and $Y(0)$ be the potential outcome without a threshold. Assume that there are two types of customers: those who strategically exceed a specific threshold K (i.e., bunching customers) and those who do not exceed the threshold (i.e., non-bunching customers). Bunching customers increase or adjust their spending to exceed the threshold and gain greater benefits than without it. In contrast, non-bunching customers do not exceed the threshold or gain extra benefits. Consequently, it is assumed that the bunching customers have potential outcomes $Y(t)$ (for $t \in \{0, 1\}$) which fall within the neighborhood of the threshold K , while non-bunching customers have the same potential outcomes, $Y(1) = Y(0)$ which lie outside the neighborhood. Under the framework, the average treatment effect of the threshold on customers is defined as

$$\Delta \equiv \mathbb{E}[Y(1) - Y(0) \mid Y(t) \in N_K, t \in \{0, 1\}], \quad (1)$$

where $N_K = [y, \bar{y}]$ is the interval as the neighborhood of the threshold K . This definition can be interpreted as the average treatment effect of the threshold on bunching customers, which is conceptually similar to the average treatment effect on the treated (ATT) in the potential outcome framework. In our application, we will deal with the situation with multiple thresholds, but we will describe our methodology under a single threshold for simplicity and discuss the extension in our application.

3.2 Mixture Model for Bunching and Non-bunching Distributions

Suppose we observe an independent and identically distributed observation, Y_i for $i = 1, \dots, n$, where n is the number of customers. We assume that there are no bunching customers outside the neighborhood N_K and they are not affected by the thresholds, namely $Y_i = Y_i(0)$ for $Y_i \in N_K^c$, where N_K^c is the complement of N_K . On the other hand, within the neighborhood N_K , bunching and non-bunching

customers are mixed, that is, $Y_i = Z_i Y_i(1) + (1 - Z_i) Y_i(0)$, where Z_i is an unobserved indicator of bunching ($Z_i = 1$) and non-bunching ($Z_i = 0$) customers. The key difference from typical causal inference frameworks is that the binary indicator Z_i is unobserved. Although a precise estimation of the indicator Z_i would be difficult, the ATT defined in (1) can be identified and estimated using a mixture model.

To perform a mixture model, we first assume the forms of distributions for $Y_i(t)$. Let $g(\cdot|\theta)$ and $f(\cdot|\gamma)$ be distributions for $Y_i(0)$ (non-bunching) and $Y_i(1)$ (bunching), respectively, where θ and γ are unknown parameters. The bunching distribution $f(\cdot|\gamma)$ is expected to exhibit a high density within N_K , allowing us to assign a relatively simple model. In the numerical examples, we employ a skew-normal distribution (Azzalini 1985), as it offers the flexibility to model a density peak that may be slightly offset from the threshold K . On the other hand, the support of the non-bunching distribution extends over the entire range of the observed data, so it is appropriate to adopt a flexible model. In our numerical illustration, we use the Singh-Maddala distribution (Singh and Maddala 1976), which is widely recognized for its flexibility and excellent fit to income and expenditure data (see Appendix A.3 for the density function and adjusted likelihood).

For the latent binary variable Z_i , we assume that $P(Z_i = 1) = \pi = 1 - P(Z_i = 0)$, where $\pi \in (0, 1)$ is an unknown mixing parameter. This gives the generative distribution of Y_i as

$$p(Y_i|\Psi) = \pi f(Y_i|\gamma) + (1 - \pi)g(Y_i|\theta). \quad (2)$$

for $Y_i \in N_K$, where $\Psi = \{\pi, \theta, \gamma\}$ is a set of unknown parameters. Note that the distribution of Y_i can be expressed as $p(Y_i|\Psi) = g(Y_i|\theta)$ for $Y_i \in N_K^c$.

Under the formulation (2), the ATT can be expressed as

$$\Delta(\Psi) \equiv \frac{\int_{y \in N_K} y f(y|\gamma) dy}{\int_{y \in N_K} f(y|\gamma) dy} - \frac{\int_{y \in N_K} y g(y|\theta) dy}{\int_{y \in N_K} g(y|\theta) dy}. \quad (3)$$

The above quantity is a simple difference between the expectations of a bunching and a non-bunching distribution within N_K . Since we assume that most of the mass density of the bunching distribution $f(y|\gamma)$ is included in N_K , the first term of $\Delta(\Psi)$ will be approximately equal to the expectation of the density.

3.3 Two-Step Bayesian Inference

In our framework, the supports of bunching and non-bunching distributions are expected to differ, indicating that joint estimation of two distributions would be difficult. Therefore, we first estimate the non-bunching distribution $g(Y_i|\theta)$ from the set of observations $\mathcal{D}_{K^c} \equiv \{Y_i; Y_i \in N_K^c\}$, and fit the mixture model (2) with estimated $g(Y_i|\theta)$, using the set of observations $\mathcal{D}_K \equiv \{Y_i; Y_i \in N_K\}$. In what follows, we consider a Bayesian approach by assigning a prior distribution on unknown parameters. Such a Bayesian approach enables us to obtain point estimates and uncertainty measures, and is useful for handling heterogeneous treatment effects as considered in the subsequent section.

Regarding the first step, we obtain the posterior distribution of θ by assigning a prior distribution $\pi(\theta)$ for θ . Since

$g(Y_i|\theta)$ has a considerable mass outside N_K^c , the likelihood based on the sample \mathcal{D}_K^c should be adjusted taking into account the probability falling into N_K^c . We then consider the following posterior:

$$p(\theta|\mathcal{D}_K^c) \propto p(\theta)L(\theta; \mathcal{D}_K^c), \quad (4)$$

where $L(\theta; \mathcal{D}_K^c)$ is the adjusted likelihood given by

$$L(\theta; \mathcal{D}_K^c) = \prod_{Y_i \in N_K^c} \frac{g(Y_i|\theta)}{\int_{y \in N_K^c} g(y|\theta) dy}. \quad (5)$$

As the second step, we define the posterior of (π, γ) as

$$\begin{aligned} p(\pi, \gamma|\mathcal{D}) &= \int p(\pi, \gamma|\mathcal{D}_K, \theta) p^*(\theta|\mathcal{D}_K^c) d\theta \\ &= p(\pi, \gamma) \int \prod_{Y_i \in N_K} p(Y_i|\Psi) p^*(\theta|\mathcal{D}_K^c) d\theta, \end{aligned} \quad (6)$$

where $p(\pi, \gamma)$ is a joint prior for (π, γ) and $p^*(\theta|\mathcal{D}_K^c)$ is a distribution to reflect the information from the posterior $p(\theta|\mathcal{D}_K^c)$. A possible example of p^* is a Dirac measure on $\theta = \hat{\theta}$ with $\hat{\theta}$ being the posterior mean from (4), which will be used in our numerical examples. The posterior computation of both (4) and (6) can be performed by the Markov Chain Monte Carlo (MCMC) algorithm using the probabilistic programming language `stan` (Carpenter et al. 2017), which takes advantage of the Hamiltonian Monte Carlo algorithm (Duane et al. 1987) to sample complex probability distributions efficiently. Given the posterior samples of Ψ , we can obtain posterior samples of $\Delta(\Psi)$ defined in (3), which give not only a point estimate but also a measure of uncertainty such as a 90% credible interval.

3.4 Theoretical Analysis

Here, we discuss the theoretical justification of the proposed two-step Bayesian method in terms of posterior contraction. Let Δ_0 be the true ATT. We assume that the two models, $f(\cdot|\gamma)$ and $g(\cdot|\theta)$, are correctly specified, that is, there exists a true parameter Ψ_0 such that $\Delta_0 = \Delta(\Psi_0)$. We also assume that $|\mathcal{D}_K| = \rho n$, namely $|\mathcal{D}_K^c| = (1 - \rho)n$ for some $\rho \in (0, 1)$, which means that the sample sizes in the neighborhood and the other region are not highly unbalanced.

Theorem 1. *Assume regularity conditions in Appendix A.1. Under $n \rightarrow \infty$, it holds that*

$$P(|\Delta(\Psi) - \Delta_0| > Mn^{-1/2} | \mathcal{D}) \rightarrow 0,$$

for a large constant M .

The proof is provided in Appendix A.2. Most of the regularity conditions are standard for the model (2) and adjusted likelihood (5), such as the identifiability of the two density functions, $f(\cdot|\gamma)$ and $g(\cdot|\theta)$, which is typically required to establish the posterior contraction rate (Ghosal, Ghosh, and Van Der Vaart 2000). The above theorem indicates that the proposed two-step Bayesian method holds posterior convergence to the true value at a rate $O(n^{-1/2})$. Hence, the information from the data can be successfully accumulated according to the sample size. This property is essential to justify the hierarchical Bayesian formulation in Section 4.

4 Heterogeneous Treatment Effect via Hierarchical Formulation

The advantage of the proposed Bayesian framework is its flexibility for model extensions, such as the estimation of heterogeneous treatment effects. Here, we discuss hierarchical Bayesian modeling to estimate heterogeneous treatment effects among customer subgroups. Let customers be divided into G subgroups according to their characteristics (e.g., age and sex), and let $Y_g(t), t \in \{0, 1\}$ be the potential outcomes in the g th subgroup. We assume that all groups share the same threshold K . Then, we can define ATT in each group, Δ_g , in the same way as (3). Let $\mathcal{D}^{(g)} = \{Y_{gi}, i = 1, \dots, n_g\}$ be a set of observations in the g th group.

We assume that the distributions of all the groups belong to the same family of distributions with different parameters, denoted by $f(\cdot|\gamma_g)$ and $g(\cdot|\theta_g)$. Then, we consider the following generative distribution

$$p_g(Y_{gi}|\Psi_g) = \pi_g f(Y_{gi}|\gamma_g) + (1 - \pi_g)g(Y_{gi}|\theta_g), \quad (7)$$

for $Y_{gi} \in N_K$, where π_g is a heterogeneous mixing proportion. A direct application of the two-step approach given in the previous section to each group can suffer from instability when the group-level sample size, n_g , is small. Instead, we consider simultaneous estimation by introducing a random-effects structure $\theta_g \sim H_\theta(\alpha_\theta)$ and $\gamma_g \sim H_\gamma(\alpha_\gamma)$ with unknown parameters α_θ and α_γ . For π_g , we introduce a logit-normal distribution, namely $\log\{\pi_g/(1-\pi_g)\} \sim N(\mu_\pi, \sigma_\pi^2)$ with unknown μ_π and σ_π^2 . According to Theorem 1, the group-specific parameter Ψ_g can be consistently estimated as n_g increases, making more precise group-level information available. Consequently, the shrinkage effect induced by the random effect becomes weaker when n_g is large, while it remains strong when n_g is small. This behavior justifies using random effects, as it allows borrowing strength across groups adaptively depending on the sample size.

As the first step, the joint posterior of $\Theta = \{\theta_1, \dots, \theta_G\}$ and α_θ based on $\mathcal{D}_K^c = \{(\mathcal{D}_K^{(g)c}, g = 1, \dots, G)\}$ is given by

$$p(\Theta, \alpha_\theta|\mathcal{D}_K^c) \propto p(\alpha_\theta) \prod_{g=1}^G h(\theta_g|\alpha_\theta)L(\theta_g; (\mathcal{D}_K^{(g)c}), \quad (8)$$

where $h(\theta_g|\alpha_\theta)$ is the density of $H_\theta(\alpha_\theta)$. In the second step, the joint posterior of $\Psi = \{\Psi_1, \dots, \Psi_G\}$ and $\alpha = \{\alpha_\gamma, \mu_\pi, \sigma_\pi^2\}$ is defined as

$$p(\Psi, \alpha|\mathcal{D}) \propto p(\alpha) \prod_{g=1}^G h_\gamma(\gamma_g|\alpha_\gamma)h_\pi(\pi_g|\mu_\pi, \sigma_\pi^2) \times \int \prod_{Y_{gi} \in N_K^{(g)}} p(Y_{gi}|\Psi_g)p^*(\Psi_g|\mathcal{D}_K^c)d\theta_g, \quad (9)$$

where $h_\gamma(\gamma_g|\alpha_\gamma)$ and $h_\pi(\pi_g|\mu_\pi, \sigma_\pi^2)$ are the density functions of $H_\gamma(\alpha_\gamma)$ and $\text{logit}(\pi_g) \sim N(\mu_\pi, \sigma_\pi^2)$, respectively, and $p(\alpha)$ is a prior distribution.

In our simulation study and application, we use the following Singh-Maddala distribution (Singh and Maddala

1976) for the non-bunching distribution:

$$g(Y_{gi}|\theta_g) = \frac{a_g q_g y^{a_g-1}}{b_g^{a_g} (1 + (Y_{gi}/b_g)^{a_g})^{q_g+1}},$$

where $\theta_g = (a_g, b_g, q_g)$, a_g and q_g are shape parameters and b_g is a scale parameter. For these parameters, we assume log-normal random effects, $\log(a_g) \sim N(\mu_a, \sigma_a^2)$, $\log(b_g) \sim N(\mu_b, \sigma_b^2)$ and $\log(q_g) \sim N(\mu_q, \sigma_q^2)$, with unknown parameters, $\alpha_\theta = (\mu_a, \sigma_a^2, \mu_b, \sigma_b^2, \mu_q, \sigma_q^2)$. For the bunching distribution, we employ the skew-normal distribution

$$f(Y_{gi}|\gamma_g) = \frac{2}{\omega_g} \phi\left(\frac{Y_{gi} - \beta_g}{\omega_g}\right) \Phi\left(\delta_g \frac{Y_{gi} - \beta_g}{\omega_g}\right),$$

where $\gamma_g = (\beta_g, \omega_g, \delta_g)$, which are location, scale, and skewness parameters, respectively, $\phi(\cdot)$ and $\Phi(\cdot)$ are the density and cumulative distribution functions of $N(0, 1)$. For the group-level parameters, we employ random effects models, $\beta_g \sim N(\mu_\beta, \sigma_\beta^2)$, $\log(\omega_g) \sim N(\mu_\omega, \sigma_\omega^2)$ and $\delta_g \sim N(\mu_\delta, \sigma_\delta^2)$ with unknown parameters $\alpha_\gamma = (\mu_\beta, \sigma_\beta^2, \mu_\omega, \sigma_\omega^2, \mu_\delta, \sigma_\delta^2)$. The detailed settings of the prior distributions are provided in Appendix B. We can again use `stan` to generate random samples from the posterior distributions (8) and (9). Based on the group-level parameters, Ψ_g , we can define the group-wise heterogeneous treatment effect $\Delta(\Psi_g)$ in the same way as (3), and the posterior distributions of $\Delta(\Psi_g)$ can be approximated by the posterior samples of Ψ_g . This gives point estimates and measures of uncertainty quantification of $\Delta(\Psi_g)$.

5 Simulation Studies

We demonstrate the effectiveness of our proposed methods by comparing their performance with the standard RDD method. We assume that the customer threshold K is known and set to 50, and the observed data originate from a population structured into $G = 100$ distinct subgroups. The observed data for each subgroup g ($g = 1, \dots, 100$) are generated by a mixture model consisting of a bunching distribution and a non-bunching distribution, weighted by subgroup-specific proportions π_g and $1 - \pi_g$, respectively. The parameters defining the distributions of components, including the proportions, can vary from subgroup to subgroup. We equally divided G groups into four clusters and set the same group-specific sample sizes n_g ($g = 1, \dots, G$) to the same values within the same clusters. The sample sizes of the clusters are set to (50, 100, 200, 300). We consider two scenarios to evaluate the model performance under different conditions. Scenario (A) assumes moderate bunching with low heterogeneity across subgroups. In contrast, Scenario (B) presents a more sparse bunching and higher heterogeneity than Scenario (A). The detailed simulation settings are deferred to Appendix C.

For the simulated dataset, we fitted the proposed BMTM and HBMTM models based on 12,000 posterior samples consisting of four chains each with 3,000 samples after discarding the first 3,000 samples as burn-in. We use `CmdStanPy` (version 1.2.5) for model fitting. The settings

Scenario	Method	MAE	CP	AL	IS
(A)	RDD	3.31	–	–	–
	BMTM	0.78	0.91	4.03	4.65
	HBMTM	0.33	0.84	1.20	1.85
(B)	RDD	3.50	–	–	–
	BMTM	1.79	0.94	9.05	9.88
	HBMTM	0.37	0.88	1.62	2.15

Table 1. Simulation results of the proposed and baseline methods in the subgroup setting. The minimum values of MAE and IS are shown in bold.

for the prior distributions of BMTM and HBMTM are provided in Appendix B. Based on the posterior samples, we computed the posterior means and the 90% highest density interval (HDI) of the group-wise ATT. For comparison, we also applied an RDD method, which estimates the effect of treatment by measuring the discontinuity in the probability density of Y at the threshold K . This can be achieved by applying the kernel density estimation separately on either side of K and computing the difference between the density values. In our implementation, a boundary correction technique for density estimation (Jones 1993) was applied, and a bandwidth of 10, centered at the threshold $K = 50$ (covering the range $[40, 60]$), was used. Note that this RDD does not provide interval estimates.

The performance of the point estimation is evaluated using Mean Absolute Error (MAE), defined as

$$\text{MAE} = \frac{1}{G} \sum_{g=1}^G \left| \Delta_g - \hat{\Delta}_g \right|,$$

Regarding the interval estimation, we use the Coverage Probability (CP), Average Length (AL), and Interval Score (IS) (Gneiting and Raftery 2007), defined as

$$\text{CP} = \frac{1}{G} \sum_{g=1}^G \mathbf{1}(l_g \leq \Delta_g \leq u_g), \quad \text{AL} = \frac{1}{G} \sum_{g=1}^G (u_g - l_g),$$

$$\text{IS} = \frac{1}{G} \sum_{g=1}^G \left\{ (u_g - l_g) + \frac{2}{\alpha} (l_g - \Delta_g) \mathbf{1}(\Delta_g < l_g) + \frac{2}{\alpha} (\Delta_g - u_g) \mathbf{1}(\Delta_g > u_g) \right\},$$

where l_g and u_g are the lower and upper bounds of $100(1 - \alpha)\%$ HDI for each subgroup g , respectively, and $\mathbf{1}(\cdot)$ is the indicator function. Note that $\alpha = 0.1$ in our settings. The CP measures the proportion of true values covered by the interval, with an ideal value of $1 - \alpha$. The IS is a composite metric that jointly evaluates the AL and CP of the interval. The IS increases when the prediction interval generated by the model is too broad and increases when the observed value falls outside of this interval. Table 1 summarizes the above performance measures averaged over 100 Monte Carlo replications.

In both scenarios, our proposed methods, BMTM and HBMTM, substantially outperform the baseline RDD. In




Figure 2: Point estimates and 90% credible intervals of the subgroup-specific causal effects Δ_g by BMTM and HBMTM.

particular, HBMTM consistently demonstrates superior performance with the lowest MAE and IS, and the superiority is evident in Scenario (B). In Scenario (B) with a weaker bunching signal, the accuracy of BMTM degrades (MAE increases from 0.78 to 1.79), while that of HBMTM remains remarkably stable and accurate. In Scenario (A), although CP of BMTM is slightly closer to the nominal 90% than that of HBMTM, HBMTM maintains reliable CP across both scenarios. These findings demonstrate that the hierarchical structure of HBMTM successfully borrows strength across subgroups to deliver highly accurate and precise causal effect estimates, even when the identifying bunching signal is weak.

Figure 2 visually contrasts the subgroup-specific causal effects estimates Δ_g from BMTM and HBMTM in Scenario (B). It confirms that the hierarchical model, HBMTM, consistently provides more accurate point estimates and narrower 90% credible intervals. These visual comparisons further emphasize the effectiveness of our proposed HBMTM method in estimating the threshold’s causal effect on customers, even in the presence of small sample sizes within subgroups.

6 Application to Promotion with Thresholds in Marketing

In this section, we apply the HBMTM to a real-world dataset from a marketing promotion in which consumers could earn incentives by exceeding spending thresholds. A month-long promotion offered consumers entry into a prize draw for spending more than 30,000 yen. The winning odds were tripled for spending more than 50,000 yen, multiplied by five for spending more than 70,000 yen, and the winners were announced after the promotion.

Figure 3 shows the histogram of the total payment amount, Y , per participant during the promotion, with the thresholds indicated by red lines. As shown in Figure 3, the distribution of Y is asymmetric, with its peak skewed




Figure 3: Observed distribution of the total payment amount during the promotion (skyblue) and the three thresholds (red).

to the left and a long and heavy tail extending to the right. Such a right-skewed, heavy-tailed distribution is a typical characteristic of expenditure data. We then adopt the Singh-Maddala distribution for non-bunching distributions, as explained in Section 4, a flexible three-parameter distribution for positive-valued data.

In this study, we defined subgroups based on total spending in the previous month to investigate how the impact of promotional thresholds varies across customers with different baseline spending habits. Defining subgroups based on prior spending is motivated by the premise that a customer’s proximity to the promotional thresholds is a key determinant of their behavioral response. Specifically, we constructed $G = 21$ subgroups by grouping participants into 10,000 yen intervals based on their previous month’s spending, with the tier exceeding 200,000 yen treated as a single, top-coded group.

We applied our proposed method, HBMTM, to this promotion dataset. The detailed settings for the prior distributions of HBMTM are provided in Appendix B.2. Since the dataset has three threshold values, denoted by $K_1 = 30,000$, $K_2 = 50,000$ and $K_3 = 70,000$, we first define \mathcal{D}_{K^c} by excluding observations that fall within the interval $[20,000, 80,000]$, which represents the combined range of $\pm 10,000$ yen around each of the three thresholds. Then we separately apply the mixture model (2) to each neighborhood, $N_{K_m} = [K_m - 10,000, K_m + 10,000]$ for $m = 1, 2, 3$. Here, we assume that the three thresholds are mutually independent, meaning that the presence of the others does not influence customers around one threshold.

Figure 4 shows the Singh-Maddala distribution for each subgroup estimated in the first step of HBMTM. As shown in Figure 4, the distributional shapes differ significantly among subgroups. This difference in shape suggests the existence of significant spending heterogeneity, which in turn validates our adoption of the highly flexible Singh-Maddala distribution.

Figure 5 provides the posterior median and 90% HDI for





Figure 4: The non-bunching distribution for each subgroup, estimated using the Singh-Maddala distribution.

each subgroup’s causal effect Δ_g across the three thresholds. The results presented in Figure 5 demonstrate a significant heterogeneity in causal effects across subgroups, with a consistent trend observed across all thresholds. Specifically, a clear positive causal effect was observed at every threshold in subgroups where the previous month’s spending was below or slightly above the threshold (e.g., Group index 1-4 for K1). For some groups, 90% HDI was entirely within the positive range, suggesting statistical confidence in the effect. This positive effect aligns with the intuitive hypothesis that consumers increase spending to achieve the incentive, a finding consistent with the hypotheses suggested in existing marketing research (e.g. Kivetz, Urminsky, and Zheng 2006).


In contrast, a consistent trend of diminishing and eventually negative causal effects was observed as spending in the previous month significantly exceeded each threshold (e.g., Group index ≥ 6 for K1). This spending-dampening effect on high-spending subgroups may be explained by the anchoring effect (e.g. Holst, Hermann, and Musshoff 2015; Zong and Guo 2022). It is plausible that the relatively low promotional thresholds acted as a new psychological anchor for consumers, leading those with higher baseline spending levels to adjust their expenditures downward, influenced by this reference point. Although more research is needed to elucidate the detailed mechanisms underlying these effects, these results strongly suggest that our proposed method effectively captures heterogeneous behavioral changes under threshold manipulation.

7 Concluding Remarks


This study addresses the challenge of estimating the causal effects of marketing thresholds under strategic customer manipulation, a scenario where standard RDD and bunching estimation methods face significant limitations. To overcome this issue, we propose BMTM, a novel Bayesian mixture model framework. This framework robustly estimates causal effects by modeling the observed spending distribution as a mixture of two latent distributions: one representing cus-



(a) Threshold at 30,000



(b) Threshold at 50,000



(c) Threshold at 70,000

Figure 5: Point estimates and 90% credible intervals of each subgroup’s causal effect Δ_g obtained by the proposed method HBMTM at each threshold.

tomers unaffected by the threshold and another representing those strategically affected. Furthermore, we introduce its hierarchical extension, HBMTM, which, to our knowledge, is the first method to provide stable and accurate estimates of heterogeneous causal effects across customer subgroups, even with small sample sizes. Our empirical results from both simulations and a real-world case study demonstrate the effectiveness of our methods, yielding more accurate and reliable estimates than conventional approaches. This study provides a practical tool for data-driven marketing and contributes a principled approach for causal inference under strategic manipulation. Future research could explore more flexible non-parametric models within this framework and extend its application to other domains where agents respond to algorithmic or policy-based thresholds.

References

Alcantara, R.; Wang, M.; Hahn, P. R.; and Lopes, H. 2024. Modified bart for learning heterogeneous ef-

fects in regression discontinuity designs. *arXiv preprint arXiv:2407.14365*.

Azzalini, A. 1985. A class of distributions which includes the normal ones. *Scandinavian journal of statistics*, 171–178.

Bertanha, M.; Caetano, C.; Jales, H.; and Seegert, N. 2024. Bunching estimation methods. *Handbook of Labor, Human Resources and Population Economics*, 1–44.

Carpenter, B.; Gelman, A.; Hoffman, M. D.; Lee, D.; Goodrich, B.; Betancourt, M.; Brubaker, M. A.; Guo, J.; Li, P.; and Riddell, A. 2017. Stan: A probabilistic programming language. *Journal of statistical software*, 76.

Cattaneo, M. D.; Idrobo, N.; and Titiunik, R. 2024. *A practical introduction to regression discontinuity designs: Extensions*. Cambridge University Press.

Duane, S.; Kennedy, A.; Pendleton, B. J.; and Roweth, D. 1987. Hybrid Monte Carlo. *Physics Letters B*, 195(2): 216–222.

- Ghosal, S.; Ghosh, J. K.; and Van Der Vaart, A. W. 2000. Convergence rates of posterior distributions. *Annals of Statistics*, 500–531.
- Gneiting, T.; and Raftery, A. E. 2007. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477): 359–378.
- Hahn, J.; Todd, P.; and der Klaauw, W. V. 2001. Identification and Estimation of Treatment Effects with a Regression-Discontinuity Design. *Econometrica*, 69(1): 201–209.
- Holst, G. S.; Hermann, D.; and Musshoff, O. 2015. Anchoring effects in an experimental auction – Are farmers anchored? *Journal of Economic Psychology*, 48: 106–117.
- Imbens, G. W.; and Rubin, D. B. 2015. *Causal inference in statistics, social, and biomedical sciences*. Cambridge university press.
- Ishihara, T.; and Sawada, M. 2024. Manipulation-Robust Regression Discontinuity Designs. arXiv:2009.07551.
- Jones, M. C. 1993. Simple boundary correction for kernel density estimation. *Statistics and computing*, 3: 135–146.
- Kivetz, R.; Urminsky, O.; and Zheng, Y. 2006. The Goal-Gradient Hypothesis Resurrected: Purchase Acceleration, Illusory Goal Progress, and Customer Retention. *Journal of Marketing Research*, 43(1): 39–58.
- Kleven, H. J.; and Waseem, M. 2013. Using Notches to Uncover Optimization Frictions and Structural Elasticities: Theory and Evidence from Pakistan *. *The Quarterly Journal of Economics*, 128(2): 669–723.
- Lee, D. S. 2008. Randomized experiments from non-random selection in U.S. House elections. *Journal of Econometrics*, 142(2): 675–697. The regression discontinuity design: Theory and applications.
- McCrary, J. 2008. Manipulation of the running variable in the regression discontinuity design: A density test. *Journal of Econometrics*, 142(2): 698–714. The regression discontinuity design: Theory and applications.
- Nastasoiu, A.; Bendle, N. T.; Bagga, C. K.; Vandenbosch, M.; and Navarro, S. 2021. Separating customer heterogeneity, points pressure and rewarded behavior to assess a retail loyalty program. *Journal of the Academy of Marketing Science*, 49: 1132–1150.
- Rubin, D. B. 1974. Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of educational Psychology*, 66(5): 688.
- Saez, E. 2010. Do Taxpayers Bunch at Kink Points? *American Economic Journal: Economic Policy*, 2(3): 180–212.
- Singh, S. K.; and Maddala, G. S. 1976. A Function for Size Distribution of Incomes. *Econometrica*, 44(5): 963–970.
- Slemrod, J. 2013. Buenas notches: lines and notches in tax system design. *Ejtr*, 11: 259.
- Sugasawa, S.; Ishihara, T.; and Kurisu, D. 2023. Hierarchical Regression Discontinuity Design: Pursuing Subgroup Treatment Effects. *arXiv preprint arXiv:2309.01404*.
- Tao, K.; Wang, Y. S.; and Ruppert, D. 2025. Bayesian analysis of regression discontinuity designs with heterogeneous treatment effects. *arXiv preprint arXiv:2504.10652*.
- Thistlethwaite, D. L.; and Campbell, D. T. 1960. Regression-discontinuity analysis: An alternative to the ex post facto experiment. *Journal of Educational psychology*, 51(6): 309.
- Zong, Y.; and Guo, X. 2022. An Experimental Study on Anchoring Effect of Consumers' Price Judgment Based on Consumers' Experiencing Scenes. *Frontiers in Psychology*, Volume 13 - 2022.