




Bayesian time-varying meta-analysis via hierarchical mean-variance random-effects models

Kohsuke Kubota^{1,2}  · Shonosuke Sugasawa³ · Keiichi Ochiai¹ · Takahiro Hoshino³

Received: 6 March 2025 / Revised: 21 May 2025 / Accepted: 17 June 2025
© The Author(s) 2025

Abstract

Meta-analysis is widely used to integrate results from multiple experiments to obtain generalized insights. Since meta-analysis datasets are often heteroscedastic due to varying subgroups and temporal heterogeneity arising from experiments conducted at different time points, the typical meta-analysis approach, which assumes homoscedasticity, fails to adequately address this heteroscedasticity among experiments. This paper proposes a new Bayesian estimation method that simultaneously shrinks estimates of the means and variances of experiments using a hierarchical Bayesian approach while accounting for time effects through a Gaussian process. This method connects experiments via the hierarchical framework, enabling “borrowing strength” between experiments to achieve high-precision estimates of each experiment’s mean. The method can flexibly capture potential time trends in datasets by modeling time effects with the Gaussian process. We demonstrate the effectiveness of the proposed method through simulation studies and illustrate its practical utility using a real marketing promotions dataset.

Keywords Bayesian estimation · Meta-analysis · Markov Chain Monte Carlo · Gaussian process · Shrinking both means and variances

✉ Kohsuke Kubota
kohsuke.kubota@keio.jp

¹ NTT DOCOMO, INC., Tokyo, Japan

² Graduate School of Economics, Keio University, Tokyo, Japan

³ Faculty of Economics, Keio University, Tokyo, Japan

1 Introduction

Meta-analysis is used to integrate results from multiple experiments to obtain generalized insights in fields such as epidemiology (Kleinbaum et al., 1991), economics (Stanley & Doucouliagos, 2012), and engineering (Sohn, 1999). In meta-analyses, heterogeneous effects among experiments are modeled using meta-regression. To address between-study heterogeneity from subgroups with different true effects, hierarchical multilevel models (e.g. Pastor & Lazowski, 2018) particularly hierarchical Bayesian models incorporating random effects (e.g. Higgins et al., 2009; Browne and Jones, 2017), are commonly used. While these models typically allow individual study effects to vary, they generally consolidate what could be multiple, distinct between-study variances into a single summary parameter (e.g., τ^2). This approach inherently assumes that this single, representative level of dispersion is constant throughout the meta-analysis (e.g. Higgins et al., 2009; Williams et al., 2021).

A central challenge in real-world meta-analysis of promotional marketing experiments is predicting future experimental effects from accumulated past results. This challenge is complicated by datasets often involving multiple experiments with time-varying properties, leading to high heteroscedasticity and temporal heterogeneity. For example, in marketing, promotional experiments on digital platforms by retailers and media providers often exhibit such heterogeneity because they are conducted with diverse participants and media at different times. These conditions result in varying experimental means and variances, even among similar promotions. Consequently, predicting future experimental effects from these complex accumulated datasets is crucial for designing effective future promotional strategies. Since promotional effects can exhibit seasonality or other periodic patterns, capturing these periodic structures within a predictive framework can be particularly effective for accurate forecasting. Standard random-effects meta-analyses, however, typically assume homoscedasticity and neglect temporal heterogeneity, potentially inadequately addressing these complexities and thus limiting their predictive performance. Therefore, it is necessary to appropriately model high heteroscedasticity and complex time trends, including periodic components, to enhance understanding and predictive performance in such contexts.

Several approaches have been proposed to address high heteroscedasticity including between-study variances and complex time trends. Thompson and Becker (2020) introduced weakly informative priors for subgroup-specific standard deviations. Bowater and Escarela (2013) suggested making variance proportional to experiment-specific variables, such as participant numbers. Williams et al. (2021) developed a Bayesian fixed-effect model using experiment-related variables for each experiment's variance. Nevertheless, these methods can lead to unreliable estimates due to large standard errors in small-scale experiments with limited sample sizes. In meta-analysis literature addressing time effects, some methods compare time-varying functions across experiments using dose-response and time-response models (Mawdsley et al., 2016; Pedder et al., 2019). Others treat experiment timing as fixed effects (Miller & Hosanagar, 2020; Salanti et al., 2009; van Enst et al., 2015). However, the former does not account for treatment effects changing over time, limiting its applicability in scenarios with time-varying effects among experiments. The latter uses discrete

dummy variables for time effects, leading to unstable estimates when time effects vary continuously. Consequently, these approaches can be unsuitable for high heteroscedasticity and temporal heterogeneity.

To solve the aforementioned problems, we propose a new Bayesian estimation method called *Shrinkage estimation with Time-trend and Random-Effects for meta-Analysis of Means and variance* (STREAM). STREAM comprises two components: First, shrinkage estimation of means and variances through a hierarchical Bayesian approach that applies “borrowing strength”. In this component, while estimates of variances are produced, the principal motivation for robustly modeling these variances is to improve the performance and reliability of the mean effect estimation. Second, estimation of time effects using the Gaussian process to flexibly model continuous-time trends in datasets.

We review related work on the two main components of our proposed method. Modeling estimated variance is critical in small area estimation using random-effects (e.g. Sugawara & Kubokawa, 2020). Notably, Sugawara et al. (2017) and Gene Hwang et al. (2009) employed random-effects models for both means and variance to achieve stable variance estimation under small sample sizes, resulting in accurate mean estimation. Such approaches are partially adopted in meta-analysis (e.g. Wang et al., 2021). Observing parallels between small area estimation with area-level data and meta-analysis with experiment-level data, we develop the method to improve effect estimation in meta-analysis by shrinkage estimation of both means and variances. We also address temporal heterogeneity between experiments, an aspect less commonly considered in small area estimation. Roberts et al. (2013) introduces the usefulness of Gaussian processes in time series analysis. This approach is particularly beneficial when specific domain knowledge about the target time series is lacking but general knowledge exists, such as knowing that functions are smooth, continuous, or exhibit periodic variation. In meta-analysis literature, LeBlanc and Banks (2022) investigated treatment resistance in resistant bacteria using Bayesian network meta-analysis, addressing temporal heterogeneity with a Gaussian process to model time trends. However, this approach neglected between-study heterogeneity by focusing on homogeneous experiments. We deal with substantial between-study heterogeneity by introducing random effects for both mean and variance.

The remainder of this paper is organized as follows. Section 2 proposes a new Bayesian model for shrinkage estimation of experimental means and variances by incorporating hierarchical structures and accounting for time effects via a Gaussian process in meta-analysis. Section 3 demonstrates the effectiveness of the proposed method through simulation studies. Section 4 illustrates its practical advantages using a real-world dataset, and Sect. 5 concludes the discussions.

2 Gaussian shrinkage hierarchical meta-analysis

2.1 Setup and observed data

We formulate the problem of meta-analysis for estimating experimental effects from a collection of m individual experiments. For each experiment i ($i = 1, \dots, m$), let θ_i

denote the true underlying effect or parameter of interest. The observed data typically consist of y_i , which is an estimate of θ_i , and an associated measure of its precision. Specifically, σ_i^2 is defined as the true unknown sampling variance of the estimate y_i . The data available from each experiment i , based on a sample of size n_i , includes y_i and S_i^2 , where S_i^2 is an observed estimate of this true sampling variance σ_i^2 .

To build intuition, consider a common scenario where y_i is the sample mean of n_i independent observations drawn from a population with true mean θ_i and true variance ψ_i^2 (representing the variance of an individual observation in study i). In this case, the true sampling variance of y_i (which we denote as σ_i^2 throughout this paper) is given by $\sigma_i^2 = \psi_i^2/n_i$, and this σ_i^2 is of order n_i^{-1} . The observed S_i^2 would then be an estimate of this σ_i^2 , typically calculated as s_i^2/n_i , where s_i^2 is the sample variance of the n_i observations (an estimate of ψ_i^2).

In conventional meta-analysis, the observed estimates S_i^2 are often treated as if they were the true sampling variances σ_i^2 and are used as fixed, known weights (typically the inverse of these variances) in pooling procedures. Some simpler approaches might even disregard individual S_i^2 values if they assume a common variance across studies, though this is less common when study-specific variance estimates are available. In contrast, our proposed model explicitly acknowledges S_i^2 as an estimate of an unknown σ_i^2 and models the relationship and uncertainty, as detailed in Eq. (1).

As auxiliary information, let \mathbf{x}_i be a q -dimensional vector representing the characteristics of each experiment, and \mathbf{z}_{ai} and \mathbf{z}_{bi} be J -dimensional and K -dimensional dummy variables, respectively, indicating the specific groups to which experiment i belongs. For example, when focusing on company promotions, \mathbf{z}_{ai} can represent the company, and \mathbf{z}_{bi} can represent the type of industry. Although we consider two dummy variables as in light of the setting in data analysis in Sect. 4, extensions to handling more than two dummy variables are straightforward. Further, we assume that time information $t_i \in \{t_1, \dots, t_L\}$ indicating the time point of each experiment (e.g., the date it was conducted) is available. Accordingly, let \mathbf{z}_{ci} denote an L -dimensional dummy variable representing which specific time point from the set $\{t_1, \dots, t_L\}$ experiment i corresponds to.

For the observed data (y_i, S_i^2) from each experiment i , we assume the following distributions:

$$\begin{aligned} y_i \mid \theta_i, \sigma_i^2 &\sim N(\theta_i, \sigma_i^2), \\ S_i^2 \mid \sigma_i^2 &\sim \Gamma\left(\frac{n_i - 1}{2}, \frac{n_i - 1}{2\sigma_i^2}\right). \end{aligned} \quad (1)$$

Here, $\Gamma(\alpha, \beta)$ denotes the gamma distribution with a probability density function proportional to $x^{\alpha-1} \exp(-\beta x)$ for $x > 0$. The distributional assumption for S_i^2 in (1) is standard and implies that $(n_i - 1)S_i^2/\sigma_i^2$ follows a χ^2 -distribution with $n_i - 1$ degrees of freedom. These assumptions for y_i and S_i^2 are generally considered reasonable, particularly when the sample size n_i for each study is not excessively small. Furthermore, we assume that y_i and S_i^2 are conditionally independent given their respective underlying parameters θ_i and σ_i^2 .

2.2 Models for heterogeneous means and variances

To estimate θ_i and σ_i^2 , we employ models with random effects estimated through a hierarchical Bayesian approach (Lindley & Smith, 1972). In particular, we model θ_i and σ_i^2 with effects of the auxiliary information and unobserved time effects by using a Gaussian process (Brahim-Belhouari & Bermak, 2004; Williams & Rasmussen, 2006).

First, we express θ_i as follows.

$$\begin{aligned} \theta_i &= \alpha_\theta + \boldsymbol{\theta}_a^\top \mathbf{z}_{ai} + \boldsymbol{\theta}_b^\top \mathbf{z}_{bi} + \boldsymbol{\theta}_c^\top \mathbf{z}_{ci} + \boldsymbol{\beta}_\theta^\top \mathbf{x}_i, \\ \theta_{aj} &\sim N(0, \tau_a^2), \quad j = 1, \dots, J, \quad \theta_{bk} \sim N(0, \tau_b^2), \quad k = 1, \dots, K, \\ \boldsymbol{\theta}_c &\sim N(0_L, C(\sigma_p^2, l_p)), \end{aligned} \tag{2}$$

where α_θ represents the baseline value common to each experiment, $\boldsymbol{\theta}_a$ and $\boldsymbol{\theta}_b$ denote the random effects for the groups to which each experiment belongs, expressed as J -dimensional and K -dimensional vectors, respectively. Moreover, $\boldsymbol{\theta}_c$ represents the time effect. $\boldsymbol{\beta}_\theta$ represents the regression coefficients for the experiment characteristics \mathbf{x}_i . All elements of $\boldsymbol{\theta}_a = (\theta_{a1}, \dots, \theta_{aJ})$ and $\boldsymbol{\theta}_b = (\theta_{b1}, \dots, \theta_{bK})$ are assumed to be independent, and $\alpha, \boldsymbol{\beta}_\theta, \tau_a^2$ and τ_b^2 are unknown parameters. Here $C(\sigma_p^2, l_p)$ is the $L \times L$ variance-covariance matrix induced by the Gaussian process assumption on $\boldsymbol{\theta}_c$ as a function of t_i , and its (l, l') -element is $K(t_l, t_{l'}; \sigma_p^2, l_p)$ for some kernel function with parameters σ_p^2 and l_p . The kernel function, which determines the covariance structures of the Gaussian process, has various options. Since the data we will consider in our application contains time information as date, we use the following periodic covariance kernel (Roberts et al., 2013; Williams & Rasmussen, 2006) to take account of annual periodicity:

$$K(t, t'; \sigma_p^2, l_p) = \sigma_p^2 \exp\left(-\frac{2 \sin^2(\pi|t - t'|/p_e)}{l_p^2}\right),$$

where p_e is the length of period and l_p and σ_p are unknown parameters to characterize the kernel function.

Regarding the variance parameter σ_i^2 , we consider the following model:

$$\begin{aligned} \log \sigma_i^2 &\sim N(\alpha_\sigma + \boldsymbol{\delta}_a^\top \mathbf{z}_{ai} + \boldsymbol{\delta}_b^\top \mathbf{z}_{bi} + \boldsymbol{\beta}_\sigma^\top \mathbf{x}_i - \log n_i, \tau_\sigma^2), \\ \delta_{aj} &\sim N(0, \tau_c^2), \quad j = 1, \dots, J, \quad \delta_{bk} \sim N(0, \tau_d^2), \quad k = 1, \dots, K, \end{aligned} \tag{3}$$

where α_σ represents the baseline value common to each experiment, $\boldsymbol{\delta}_a$ and $\boldsymbol{\delta}_b$ denote the random effects for the groups to which each experiment belongs, expressed as J -dimensional and K -dimensional vectors, respectively, and $\boldsymbol{\beta}_\sigma$ represents the regression coefficients for the experiment characteristics \mathbf{x}_i . Further, $\alpha_\sigma, \boldsymbol{\beta}_\sigma, \tau_\sigma^2, \tau_c^2$ and τ_d^2 are unknown parameters. Note that $-\log n_i$ in the mean term of (3) corresponds to an offset term to ensure that $E[\sigma_i^2] = O(n_i^{-1})$. The proposed model consists of (1), (2) and (3), and its advantages can be summarized as follows:

- Shrinkage estimation for both the mean and variance allows us to account for the dataset’s heteroscedasticity and achieve robust mean estimates when predicting θ_i for unknown experiments.
- The time trend modeled by a Gaussian process enables us to predict future effects of experiments.

2.3 Posterior inference and predicting future effects

The proposed model (1), (2), and (3) includes several unknown parameters. By assigning prior distributions to them, we consider Bayesian inference on the model parameters and the random effects (latent variables). As the default choice, we adopt relatively standard prior distributions as follows.

$$\begin{aligned} \alpha_\theta &\sim N(m_{\alpha\theta}, s_{\alpha\theta}^2), \quad \alpha_\sigma \sim N(m_{\alpha\sigma}, s_{\alpha\sigma}^2), \quad \beta_\theta \sim N(m_{\beta\theta}, s_{\beta\theta}^2), \\ \beta_\sigma &\sim N(m_{\beta\sigma}, s_{\beta\sigma}^2), \quad \tau_a^2 \sim \text{HC}(\eta_a), \quad \tau_b^2 \sim \text{HC}(\eta_b), \quad \tau_c^2 \sim \text{HC}(\eta_c), \\ \tau_d^2 &\sim \text{HC}(\eta_d), \quad \tau_e^2 \sim \text{HC}(\eta_e), \quad \sigma_p^2 \sim \text{HC}(\eta_\sigma), \quad l_p \sim \text{HC}(\eta_l), \end{aligned} \tag{4}$$

where $\text{HC}(\eta)$ denotes the half-Cauchy prior (e.g. Gelman, 2006) with scale parameter η . The values in the above prior distributions (hyperparameters) are user-specified, and specific choices will be given in Sects. 3 and 4. We note that the log-normal model for σ_i^2 in (3) is not conjugate under the sampling model (1), which may complicate the posterior inference. To efficiently generate posterior samples, we rely on the Markov chain Monte Carlo (MCMC) algorithm, implemented in the probabilistic programming stan (Carpenter et al., 2017). Stan is a probabilistic programming language widely used for Bayesian modeling, which leverages the Hamiltonian Monte Carlo algorithm (Duane et al., 1987) to efficiently sample from complex probability distributions.

We calculate the predictive distribution of the estimates $\tilde{\mathbf{y}} = (\tilde{y}_1, \dots, \tilde{y}_{m_{te}})$ for m_{te} new future experiments with specific time point from the set $\{t'_1, \dots, t'_{L'}\}$. Specifically, we use m_{tr} experiments as observational data and the posterior samples of each parameter generated by MCMC. In constructing the predictive distribution of $\tilde{\mathbf{y}}$, we use the posterior samples generated by MCMC for various parameters in Eqs. (1), (2), and (3). We provide additional explanation since constructing the predictive distribution for θ_c differs from other parameters. Our goal is to construct the predictive distribution of the parameters $\tilde{\theta}_c = (\tilde{\theta}_{c1}, \dots, \tilde{\theta}_{cL'})$ associated with the future experiments $\tilde{\mathbf{z}}_c = (\tilde{z}_{c1}, \dots, \tilde{z}_{cL'})$, using the posterior samples. To achieve this, we aim to obtain the predictive distribution $p(\tilde{\theta}_c | D_{tr}, \tilde{\mathbf{z}}_c)$. Here, $D_{tr} = (\mathbf{z}_c, \theta_c)$, where $\mathbf{z}_c = (z_{c1}, \dots, z_{cL})$ and $\theta_c = (\theta_{c1}, \dots, \theta_{cL})$. Since the joint distribution $p(\theta_c, \theta_c | \mathbf{z}_c, \tilde{\mathbf{z}}_c)$ is given by the Gaussian process assumption in Eq. (2), it is formulated as the following multivariate Gaussian distribution.

$$p(\theta_c, \tilde{\theta}_c | \mathbf{z}_c, \tilde{\mathbf{z}}_c) \sim N(0_{L+L'}, C'(\sigma_p^2, l_p)) \tag{5}$$

where $C'(\sigma_p^2, l_p)$ is the $(L + L') \times (L + L')$ the following variance-covariance matrix:

$$C'(\sigma_p^2, l_p) = \begin{pmatrix} \Phi_{LL} & \Phi_{LL'} \\ \Phi_{LL'}^\top & \Phi_{L'L'} \end{pmatrix} \tag{6}$$

whose elements are given by

$$\begin{aligned} \Phi_{LL} &= (K(t_i, t_j; \sigma_p^2, l_p)), \quad 1 \leq i, j \leq L, \\ \Phi_{LL'} &= (K(t_i, t'_j; \sigma_p^2, l_p)), \quad 1 \leq i \leq L, \quad 1 \leq j \leq L' \\ \Phi_{L'L'} &= (K(t'_i, t'_j; \sigma_p^2, l_p)), \quad 1 \leq i, j \leq L', \end{aligned} \tag{7}$$

By using the property of conditional multivariate Gaussian (e.g. Williams & Rasmussen, 2006), we can obtain the predictive distribution $p(\tilde{\theta}_c | D_{tr}, \tilde{z}_c)$ as follows.

$$p(\tilde{\theta}_c | D_{tr}, \tilde{z}_c) = N(\Phi_{LL'}^\top \Phi_{LL}^{-1} \theta_c, \Phi_{L'L'} - \Phi_{LL'}^\top \Phi_{LL}^{-1} \Phi_{LL'}). \tag{8}$$

Then, using the posterior samples of $\alpha_\theta, \theta_a, \theta_b, \tilde{\theta}_c$ and β_θ , the posterior samples of the future effect $\tilde{\theta}_j$ can be generated as

$$\tilde{\theta}_j = \alpha_\theta + \theta_a^\top \tilde{z}_{aj} + \theta_b^\top \tilde{z}_{bj} + \tilde{\theta}_c^\top \tilde{z}_{cj} + \beta_\theta^\top \tilde{x}_j, \quad j = 1, \dots, m_{te},$$

where $\tilde{z}_{aj}, \tilde{z}_{bj}$ and \tilde{x}_j are covariate vectors for the test data. This parameter can be regarded as the true effect in the future treatment. Furthermore, given covariates $\tilde{z}_{aj}, \tilde{z}_{bj}, \tilde{x}_j$ and the sample size n_j , the posterior samples of the future variance $\tilde{\sigma}_j^2$ can be generated as

$$\log \tilde{\sigma}_j^2 \sim N(\alpha_\sigma + \delta_a^\top \tilde{z}_{aj} + \delta_b^\top \tilde{z}_{bj} + \beta_\sigma^\top \tilde{x}_j - \log n_j, \tau_\sigma^2),$$

so that the posterior predictive distribution of the observed treatment effect \tilde{y}_j can be generated from $N(\tilde{\theta}_j, \tilde{\sigma}_j^2)$.

3 Simulation studies

We compare the performance of the proposed method with baseline methods through simulation studies. The m experimental datasets subject to meta-analysis are represented by (y_i, S_i^2, n_i) for $i = 1, \dots, m$, where y_i is the estimates of the mean effect of interest θ_i in experiment i , S_i^2 is the variance of the y_i for experiment i , and n_i is the sample size of experiment i . In the simulation experiments, to generate the final dataset (y_i, S_i^2, n_i) , we perform data generation using the following steps.

First, we generate the observations $y_{i,j}$ for experiment i from the following:

$$y_{ij} \sim N(\theta_i, n_i \sigma_i^2), \quad i = 1, \dots, m, \quad j = 1, \dots, n_i, \quad (9)$$

Using y_{ij} and n_i , we calculate the mean y_i and variance S_i^2 for experiment i as follows.

$$y_i = \frac{1}{n_i} \sum_{j=1}^{n_i} y_{ij}, \quad S_i^2 = \frac{1}{n_i(n_i - 1)} \sum_{j=1}^{n_i} (y_{ij} - y_i)^2. \quad (10)$$

We then generate the parameter of interest θ_i using the model below, which accounts for random and time effects.

$$\begin{aligned} \theta_i &= \boldsymbol{\theta}_a^\top \mathbf{z}_{ai} + \boldsymbol{\theta}_b^\top \mathbf{z}_{bi} + a_1 \sin\left(2\pi \frac{t_i}{12}\right) \\ &\quad + b_1 \cos\left(2\pi \frac{t_i}{12}\right) + c_1 \frac{t_i}{12} + 0.5x_i, \\ \theta_{aj}, \theta_{bk} &\sim N(\mu_\theta, \sigma_\theta^2), \quad \mu_\theta \sim N(3, 2^2), \\ \sigma_\theta^2 &\sim N_{[0, \infty)}(0, d_1^2), \end{aligned} \quad (11)$$

We also generate σ_i^2 using the following model that accounts for random effects.

$$\begin{aligned} \sigma_i^2 &\sim IG(2, \exp(\boldsymbol{\delta}_a^\top \mathbf{z}_{ai} + \boldsymbol{\delta}_b^\top \mathbf{z}_{bi} + 0.1x_i)), \\ \delta_{aj}, \delta_{bk} &\sim N(\mu_\sigma, \sigma_\sigma^2), \quad \mu_\sigma \sim N(1, 0.1^2), \quad \sigma_\sigma^2 \sim N_{[0, \infty)}(0, d_2^2). \end{aligned} \quad (12)$$

We generate the covariate x_i from a uniform distribution between 1 and 10. The time information t_i takes random values from 1, ..., 24, and the sample size n_i takes random values from 100, ..., 10000. In this study, we set $J = 30$, $K = 6$, and $m = 80$. We split each simulation dataset into training and test sets in an 8:2 ratio. The test set is taken in order from the largest values of t_i , so it is the future promotion compared to the training set.

In this study, we conduct simulations based on four scenarios that combine the presence or absence of time effects and the magnitude of between-study heterogeneity. The presence of time effects is determined by the existence of a_1 , b_1 , and c_1 in Eq. (11), and the magnitude of between-study heterogeneity is determined by the values of d_1 and d_2 in σ_θ^2 and σ_σ^2 . We perform simulations based on the following four scenarios: (i) temporal heterogeneity present and high between-study heterogeneity (a_1, b_1, c_1, d_1, d_2) = (1, 1, 1, 10, 2); (ii) temporal heterogeneity present and low between-study heterogeneity (a_1, b_1, c_1, d_1, d_2) = (1, 1, 1, 5, 0.5); (iii) temporal heterogeneity absent and high between-study heterogeneity (a_1, b_1, c_1, d_1, d_2) = (0, 0, 0, 10, 2); and (iv) temporal heterogeneity absent and low between-study heterogeneity (a_1, b_1, c_1, d_1, d_2) = (0, 0, 0, 5, 0.5).

For each simulation scenario, we generate 150 simulation datasets and estimate the experimental effect θ_i using seven baseline methods and the proposed STREAM: (a) Fixed Effects (FE) model, (b) FE with monthly dummy (FE-M) model, (c) FE with monthly dummy and variance modeling (FE-MV) model, (d) Random Effects (RE)

model, (e) RE with monthly dummy (RE-M) model, (f) RE with monthly dummy and variance modeling (RE-MV) model, (g) RE-Gaussian-Process (RE-GP) model, and (h) STREAM. The FE model, corresponding to the fixed effects model in general meta-analysis (Borenstein et al., 2010), models θ_a and θ_b as fixed effects in Eq. (2) and does not account for the time effect θ_c , using an uninformative prior distribution $N(0, 1000)$ for θ_a and θ_b . The FE-M model extends the FE model by including the time effect θ_c as a fixed effect, with an uninformative prior $N(0, 1000)$ for θ_c . The FE-MV model further incorporates fixed effects for δ_a and δ_b in Eq. (3), also using an uninformative prior $N(0, 1000)$. This model corresponds to existing methods by Bowater and Escarela (2013) and Williams et al. (2021). The RE model, corresponding to hierarchical Bayesian random effects modeling in general meta-analysis (Babapulle et al., 2004; Borenstein et al., 2010), models θ_a and θ_b as random effects in Eq. (2) without considering the time effect θ_c . The RE-M model extends the RE model by modeling the time effect θ_c as a fixed effect with an uninformative prior $N(0, 1000)$. The RE-MV model further incorporates random effects for δ_a and δ_b in Eq. (3). The RE-GP model models θ_a and θ_b as random effects in Eq. (2) and represents the time effect θ_c using a Gaussian process, corresponding to the method in LeBlanc and Banks (2022).

These baseline methods adopt distinct approaches for incorporating and modeling S_i^2 . Broadly, they are classified based on whether they explicitly model variance heterogeneity. For the first group of baseline methods—FE, RE, FE-M, RE-M, and RE-GP—which do not explicitly model variance heterogeneity, the observed variance estimate for each study i , S_i^2 , is treated as a known and fixed value representing the sampling variance σ_i^2 for that study. Consequently, each study’s observed effect y_i is assumed to follow a normal distribution with mean θ_i and this fixed variance S_i^2 . Subsequently, when predicting the effect of new experiments for which individual variances are unknown, these baseline methods use a single, fixed variance for all such predictions; this fixed variance is calculated as the average of the observed S_i^2 values from the training dataset.

In contrast, the FE-MV and RE-MV explicitly model the observed sample variance S_i^2 conditional on the true sampling variance σ_i^2 using a Gamma distribution, as specified in Eq. (1). This modeling strategy for $S_i^2 \mid \sigma_i^2$ is equivalent to the standard statistical assumption that $(n_i - 1)S_i^2/\sigma_i^2$ follows a Chi-squared distribution. Such an approach to modeling sample variances via a Gamma distribution has precedent in the meta-analysis literature (Chowdhry et al. 2016; Dakin et al. 2011).

Bayesian estimation is performed for these models by setting uninformative priors for each relevant parameter.

$$\begin{aligned}
 m_{\alpha\theta} &= m_{\alpha\sigma} = 0, & m_{\beta\theta} &= m_{\beta\sigma} = 0_q, \\
 s_{\alpha\theta} &= s_{\alpha\sigma} = 1000, & s_{\beta\theta} &= s_{\beta\sigma} = 1000_q, \\
 \eta_a &= \eta_b = \eta_c = \eta_d = \eta_e = \eta_l = \eta_\sigma = 2.5.
 \end{aligned}
 \tag{13}$$

We set $p_e = 12$ for the periodic covariance kernel of the Gaussian process. For each simulation and model, we executed the MCMC algorithm with four chains, each comprising 2,000 warm-ups and 8,000 sampling iterations. The point estimates of

the model parameters are calculated using the median of the posterior samples for each model.

We evaluate the models in terms of their prediction performance in point and interval estimation. The point estimation performance of the models is compared based on the mean average percentage error (MAPE) given by

$$\text{MAPE} = \frac{100}{n} \sum_{i=1}^n \left\| \frac{\theta_i - \hat{y}_i}{\theta_i} \right\|, \quad (14)$$

and the scaled mean squared error (scaled MSE) given by

$$\text{scaled MSE} = \frac{1}{n} \sum_{i=1}^n \frac{(\theta_i - \hat{y}_i)^2}{\theta_i^2}, \quad (15)$$

where \hat{y}_i is the predicted treatment effect of the experiment i , and n is the number of experiments to be predicted.

We use the interval score (IS) (Gneiting & Raftery, 2007) to evaluate the interval estimation performance of the models. Specifically, the IS for the $100(1 - \alpha)\%$ prediction interval, with lower and upper endpoints that are the predictive quantiles at level $\alpha/2$ and $(1 - \alpha)/2$, respectively, are defined below.

$$\text{IS} = \frac{1}{n} \sum_{i=1}^n \left\{ (u_i - l_i) + \frac{2}{\alpha} (l_i - \theta_i) \mathbf{1}(\theta_i < l_i) + \frac{2}{\alpha} (\theta_i - u_i) \mathbf{1}(\theta_i > u_i) \right\}, \quad (16)$$

where l_i and u_i are the lower and upper bounds of the $100(1 - \alpha)\%$ highest posterior density intervals, respectively, and $\mathbf{1}(\cdot)$ is the indicator function. The IS is a measure designed to quantify the quality of uncertainty in the model predictions. It penalizes the model if the prediction interval it outputs is vast and penalizes the model if the observed values are outside the prediction interval it outputs. In this paper, α is 0.05.

Table 1 presents the simulation results for scenarios (i) and (ii), which include temporal heterogeneity. In simulation (i), STREAM demonstrated the best performance in both point estimation (MAPE and Scaled MSE) and interval estimation (IS). In simulation (ii), STREAM achieved the best performance in interval estimation (IS), while FE-MV showed the best performance in point estimation (MAPE and Scaled MSE). Although STREAM did not secure the top spot in point estimation, it ranked third overall in MAPE and Scaled MSE.

Table 2 presents the simulation results for scenarios (iii) and (iv), which do not include temporal heterogeneity. In simulation (iii), STREAM performed best in both point estimation (MAPE and Scaled MSE) and interval estimation (IS). In simulation (iv), STREAM showed the highest performance in point estimation (MAPE, Scaled MSE) and third in IS.

Table 1 Simulation results including temporal heterogeneity. The boldface values indicate the best performance among all models, while underlined values represent the second-best results

	Method	MAPE	Scaled MSE	IS
(i)	FE	152.8	1894.84	34.8
	FE-M	65.4	263.62	23.5
	FE-MV	87.5	839.62	> 1000
	RE	78.1	158.80	35.1
	RE-M	61.3	147.14	21.9
	RE-MV	<u>36.9</u>	<u>23.75</u>	<u>13.8</u>
	RE-GP	69.2	58.27	38.4
	STREAM	29.7	8.12	10.0
(ii)	FE	48.1	46.61	33.4
	FE-M	28.5	23.57	18.4
	FE-MV	21.1	18.11	> 1000
	RE	36.1	23.57	23.4
	RE-M	22.8	<u>21.22</u>	8.2
	RE-MV	<u>21.6</u>	21.51	9.0
	RE-GP	23.9	21.74	<u>6.1</u>
	STREAM	22.0	21.50	5.4

Table 2 Simulation results without temporal heterogeneity. The boldface value indicate the best performance among all models, while underlined values represent the second-best results

	Method	MAPE	Scaled MSE	IS
(iii)	FE	49.9	67.10	19.5
	FE-M	63.9	159.56	22.8
	FE-MV	38.4	31.72	> 1000
	RE	43.7	41.12	18.3
	RE-M	52.1	74.31	21.6
	RE-MV	<u>28.0</u>	<u>5.21</u>	<u>13.6</u>
	RE-GP	429.3	32983.78	34.7
	STREAM	22.2	4.36	9.6
(iv)	FE	31.0	21.41	14.2
	FE-M	30.0	11.61	17.8
	FE-MV	19.6	5.76	> 1000
	RE	16.6	4.13	4.2
	RE-M	19.4	3.83	8.0
	RE-MV	18.3	<u>3.50</u>	9.0
	RE-GP	<u>16.4</u>	3.92	<u>4.4</u>
	STREAM	15.9	3.40	5.1

From these four simulation results, we found that STREAM consistently outperforms other baseline methods in scenarios with high between-study heterogeneity (simulations (i) and (iii)), irrespective of temporal heterogeneity. Furthermore, when between-study heterogeneity is low (simulations (ii) and (iv)), STREAM maintains stable performance in both point and interval estimation.

In scenarios with high between-study heterogeneity, STREAM outperforms other baseline methods, with RE-MV being the second-best model. This demonstrates the effectiveness of Gaussian process modeling for more flexible time effect modeling compared to RE-MV. Conversely, RE-GP, which does not account for variance modeling, performs similarly to STREAM when between-study heterogeneity is low but significantly deteriorates when heterogeneity is high. This deterioration is likely due

to RE-GP's inability to adequately model the data's variance structure. This causes the Gaussian process for time effects to overfit the data as it attempts to compensate for incorrect parameter structures, thereby reducing predictive performance. These findings show the importance of our proposed method, STREAM, which addresses high between-study heterogeneity in meta-analysis by modeling both means and variances through a hierarchical Bayesian approach while flexibly representing time effects with a Gaussian process. In simulation (ii), despite its excellent performance in some aspects, the FE-MV showed the worst performance in IS across all simulations. This is likely because the parameters for unknown categories, θ_a , θ_b , δ_a , and δ_b , were not updated and were strongly influenced by non-informative priors. This result illustrates the importance of the hierarchical Bayesian approach, which shares information to achieve stable estimation even for unknown categories.

4 Application to promotion effects in marketing

We apply the proposed model to meta-analysis in promotions implemented as marketing strategies. Meta-analysis is used in promotions to predict promotion effectiveness using results from past promotions and improve corporate revenues by identifying and implementing promotions with high Return on Investment. Meta-analysis in marketing is common not only in e-commerce, where the accumulation and use of user data are actively conducted (Browne & Jones, 2017; Miller & Hosanagar, 2020; Zerbini et al., 2022), but also in retail businesses with offline stores (Blut et al., 2018; Liu-Thompkins et al., 2022). Practically, since numerous experiments are added daily and need to be frequently updated, the meta-analysis approach using summary statistics of existing experiments has a notable advantage in case of implementation and scalability.

Our proposed method enables more accurate estimates when the experiments under analysis exhibit high between-study and temporal heterogeneity. By utilizing our proposed method, marketers can make highly confident decisions about which marketing strategies are most effective for implementing more profitable campaigns.

In this study, we use a real-world dataset of promotions aimed at sales promotion conducted by participating merchants within a coalition loyalty program (a system that allows flexible acquisition and redemption of points among multiple participating stores) (e.g. Stourm & Bradlow, 2023). This promotion dataset records, for each promotion, the participating merchant, business type, promotion start/end dates, number of subjects in the experiment (with equal numbers in treatment and control groups), estimated treatment effect, variance, and promotion duration. The data analyzed in this study comprises promotions conducted from April 2022 to December 2023. From this promotion data, we extract and analyze data from companies that conducted two or more promotions during the period. The dataset includes 51 promotions carried out by 16 participating merchants in two types of businesses: retail and restaurant.

We set the promotion duration as the covariate x_i . Table 3 presents the summary statistics for the sample size n_i , the estimated mean treatment effect y_i , variance S_i^2 , and the covariate x_i of the dataset.

Table 3 Summary statistics for the sample size n_i , the estimated mean treatment effect y_i , variance S_i^2 , and the covariate x_i of the dataset

	Mean	St. Dev.
Sample size n_i	26573.0	31519.5
Mean y_i	61.7	115.0
Variance S_i^2	22.4	76.7
Covariance x_i	24.9	7.5




Fig. 1 Comparison of the distribution of the original variable y and the log-transformed variable $\log y$

This dataset consists of promotions conducted by multiple merchants at different time points, resulting in substantial between-study and temporal heterogeneity. Figure 1 shows the distribution of the original variable y and the log-transformed variable $\log y$.

As shown in Fig. 1, this dataset’s distribution of y is asymmetric due to substantial between-study heterogeneity and outliers. Also, the distribution of $\log y$ is more symmetric and reduces the effect of outliers, making it more suitable for linear modeling. Therefore, the analysis uses the log-transformed variable $\log y$ as the response variable. We apply the Delta method (Oehlert, 1992) to transform the variance associated with the log transformation of the estimator y . The variance of the log-transformed estimator $\log y$ can be approximated as $\text{Var}(\log y) \approx \text{Var}(y)/y^2$.

We split the dataset into training and test sets in an 8:2 ratio. The test set is taken in order from the largest values of t_i , so it is the future promotion compared to the training set. We executed the MCMC algorithm with four chains comprising 2,000 warm-ups and 8,000 sampling iterations. Using this real-world dataset, we evaluate seven baseline methods and the proposed method: (a) FE model, (b) FE-M model, (c) FE-MV model, (d) RE model, (e) RE-M model, (f) RE-MV model, (g) RE-GP model, and (h) STREAM.

Table 4 shows the results of the application to the promotion effect data.

From these results, the proposed method STREAM demonstrated the best performance in MAPE, Scaled MSE, and IS compared to other baseline methods. Comparisons between the FE and FE-M and the RE and RE-M revealed that accounting for time effects improves predictive performance. The RE-MV model exhibited the second-best performance, consistent with the trends observed in simulation (i) in Table 1. Furthermore, the observed deterioration in RE-GP’s performance aligns with the results of simulation (i), suggesting that the real dataset exhibits significant between-study heterogeneity and temporal heterogeneity. Simulation studies and the

Table 4 Performance metrics from the application to the promotion effect dataset. The boldface values indicate the best performance among all models, while underlined values represent the second-best results

Method	MAPE	Scaled MSE	IS
FE	77.8	0.96	531.6
FE-M	70.2	0.58	588.1
FE-MV	68.9	0.54	> 1000
RE	108.8	2.64	530.1
RE-M	67.4	0.49	603.1
RE-MV	<u>50.4</u>	<u>0.32</u>	<u>473.7</u>
RE-GP	135.5	4.04	547.8
STREAM	48.1	0.29	464.5

Table 5 Median and maximum values of Gelman-Rubin (GR) statistic \hat{R} . The boldface values indicate $\hat{R} < 1.1$, satisfying the convergence criterion for MCMC

Method	Median \hat{R}	Max \hat{R}
FE	2.2	3.8
FE-M	2.7	4.0
FE-MV	1.4	2.7
RE	2.1	3.5
RE-M	2.8	4.3
RE-MV	1.0	1.0
RE-GP	3.1	4.1
STREAM	1.0	1.0

real dataset confirmed that STREAM exhibits superior performance on datasets with high between-study and temporal heterogeneity.

To verify the convergence of each model, we calculated the Gelman-Rubin (GR) statistic \hat{R} (Gelman & Rubin, 1992) for each model. Table 5 presents each model's median and maximum values of the GR statistic \hat{R} .

According to Gelman et al. (2013), a GR statistic \hat{R} of 1.1 or below indicates convergence. Under these diagnostic criteria, the baseline methods FE, FE-M, FE-MV, RE, RE-M, and RE-GP exhibited poor MCMC convergence. In contrast, RE-MV and STREAM were confirmed to have achieved convergence. This is likely because RE-MV and the proposed STREAM effectively capture the data characteristics through hierarchical Bayesian approaches and time effect modeling, resulting in stable parameter sampling.

Figure 2 shows box plots of the mean values for merchants and business types, calculated from the posterior distributions of the estimated merchant and business type effects. As shown in Fig. 2, since the dataset comprises different experiments, there is between-study heterogeneity, and the effects vary depending on the merchant and business type. This result suggests that modeling random effects using a hierarchical Bayesian approach is appropriate. The coefficients for the business type effect are -0.108 for restaurants and 0.284 for retail, with the value being higher for retail. These results indicate that the promotions encourage consumer stockpiling, as reported in previous studies (e.g. Ailawadi et al., 2007; Mela et al., 1998), and this effect may have appeared more prominently in the retail, where stockpiling is relatively easier. Figure 3 also shows the estimated time effects arranged chronologically. As shown in Fig. 3, there is temporal heterogeneity in the dataset: the effect of promotions varies depending on the implementation period, with the highest effect in

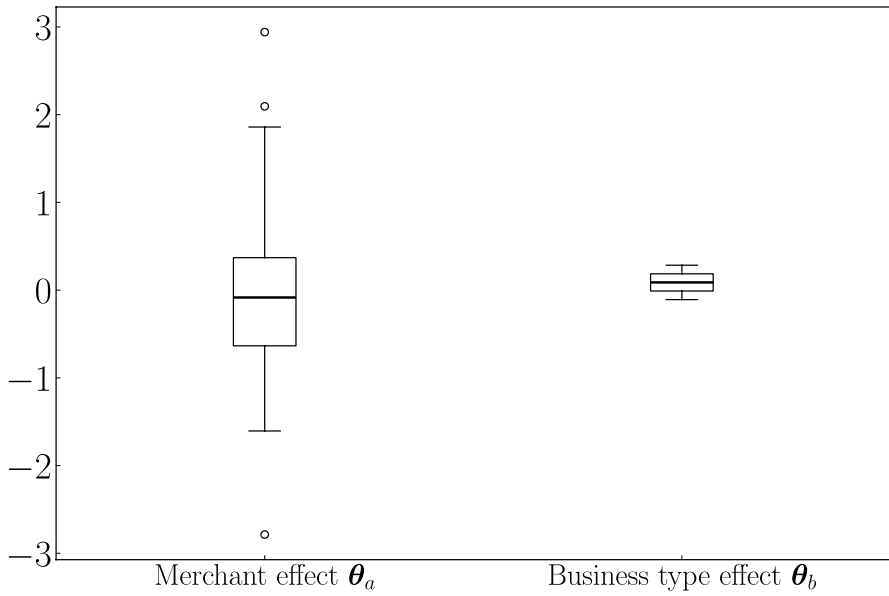


Fig. 2 Estimated merchant and business type effects




Fig. 3 Estimated time effect

July and the lowest in April. These results suggest that marketers are likely to implement more effective promotions if they carry them out during the summer.

5 Concluding remarks

This paper proposes a new Bayesian estimation method called *Shrinkage estimation with Time-trend and Random-Effects for meta-Analysis of Means and variance* (STREAM). This method performs shrinkage estimation of experiments' means and variances using a hierarchical Bayesian approach while accounting for time effects with a Gaussian process. The proposed method links related experiments by estimating the shrinkage of means and variances via the hierarchical Bayesian approach, enabling stable estimation of each experiment's parameters through "borrowing strength." Additionally, STREAM flexibly addresses complex time trends in the datasets using a Gaussian process. Through simulation experiments, we confirmed that STREAM consistently outperforms other baseline methods in both point and interval estimations when between-study heterogeneity is high, regardless of temporal heterogeneity. To demonstrate the effectiveness of our model, we applied STREAM to a real-world promotion dataset. STREAM performed the best in point and interval estimation compared to other baseline methods and showed good convergence in MCMC.

In this study, we assumed a normal distribution for the random effects to represent between-study heterogeneity in the hierarchical Bayesian model. However, assuming other distributions may enable more accurate estimation. When between-study heterogeneity is high, introducing subgroups as dummy variables relative to the number of experiments in the dataset increases dimensionality, making dimensionality reduction important. For instance, dimensionality reduction can be achieved by employing a Laplace distribution (Park & Casella, 2008) or Horseshoe prior (Carvalho et al., 2009, 2010; Qi et al., 2022) as prior distributions for random effects.

In this study, we employed the Gaussian process with a periodic kernel to represent temporal heterogeneity; however, incorporating other kernel functions could potentially achieve more accurate estimation. Williams and Rasmussen (2006) demonstrated that the sum and product of multiple kernel functions result in kernel functions that maintain positive semi-definiteness. It is possible to flexibly address more complex time trends using kernel functions composed of combining multiple kernel functions. Exploring methods to achieve higher precision in estimation through comparisons of prior distributions for random effects and comparisons of kernel functions in the Gaussian process within hierarchical Bayesian models remains a future challenge.

Declarations

Conflict of interest All authors report that NTT DOCOMO, INC provided financial support. Kohsuke Kubota and Keiichi Ochiai report a relationship with NTT DOCOMO, INC. that includes: employment. Shonosuke Sugasawa and Takahiro Hoshino report a relationship with NTT DOCOMO, INC. that includes: consulting or advisory.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Ailawadi, K. L., Gedenk, K., Lutzky, C., & Neslin, S. A. (2007). Decomposition of the sales impact of promotion-induced stockpiling. *Journal of Marketing Research*, *44*(3), 450–467.
- Babapulle, M. N., Joseph, L., Bélisle, P., Brophy, J. M., & Eisenberg, M. J. (2004). A hierarchical bayesian meta-analysis of randomised clinical trials of drug-eluting stents. *The Lancet*, *364*(9434), 583–591.
- Blut, M., Teller, C., & Floh, A. (2018). Testing retail marketing-mix effects on patronage: A meta-analysis. *Journal of Retailing*, *94*(2), 113–135.
- Borenstein, M., Hedges, L. V., Higgins, J. P., & Rothstein, H. R. (2010). A basic introduction to fixed-effect and random-effects models for meta-analysis. *Research Synthesis Methods*, *1*(2), 97–111.
- Bowater, R. J., & Escarela, G. (2013). Heterogeneity and study size in random-effects meta-analysis. *Journal of Applied Statistics*, *40*(1), 2–16.
- Brahim-Belhouari, S., & Bermak, A. (2004). Gaussian process for nonstationary time series prediction. *Computational Statistics & Data Analysis*, *47*(4), 705–712.
- Browne, W., & Jones, M. S. (2017). What works in e-commerce—a meta-analysis of 6700 online experiments. *Qubit Digital Ltd*, *21*, 8.
- Carpenter, B., Gelman, A., Hoffman, M. D., Lee, D., Goodrich, B., Betancourt, M., Brubaker, M. A., Guo, J., Li, P., & Riddell, A. (2017). Stan: A probabilistic programming language. *Journal of Statistical Software*, *76*, 8.
- Carvalho, C. M., Polson, N. G., & Scott, J. G. (2009, 16–18 Apr). Handling sparsity via the horseshoe. In: D. van Dyk and M. Welling (Eds.), *Proceedings of the Twelfth International Conference on Artificial Intelligence and Statistics*, Volume 5 of *Proceedings of Machine Learning Research*, Hilton Clearwater Beach Resort, Clearwater Beach, Florida USA, pp. 73–80. PMLR.
- Carvalho, C. M., Polson, N. G., & Scott, J. G. (2010). The horseshoe estimator for sparse signals. *Biometrika*, *97*(2), 465–480.
- Chowdhry, A. K., Dworkin, R. H., & McDermott, M. P. (2016). Meta-analysis with missing study-level sample variance data. *Statistics in Medicine*, *35*(17), 3021–3032. <https://doi.org/10.1002/sim.6908>.
- Dakin, H. A., Welton, N. J., Ades, A. E., Collins, S., Orme, M., & Kelly, S. (2011). Mixed treatment comparison of repeated measurements of a continuous endpoint: an example using topical treatments for primary open-angle glaucoma and ocular hypertension. *Statistics in Medicine*, *30*(20), 2511–2535. <https://doi.org/10.1002/sim.4284>.
- Duane, S., Kennedy, A., Pendleton, B. J., & Roweth, D. (1987). Hybrid monte carlo. *Physics Letters B*, *195*(2), 216–222.
- Gelman, A. (2006). Prior distributions for variance parameters in hierarchical models (comment on an article by browne and draper). *Bayesian Analysis*, *1*, 515–533.
- Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., & Rubin, D. B. (2013). *Bayesian Data Analysis* (3rd ed.). Chapman and Hall/CRC.
- Gelman, A., & Rubin, D. B. (1992). Inference from iterative simulation using multiple sequences. *Statistical Science*, *7*(4), 457–472.
- Gene Hwang, J., Qiu, J., & Zhao, Z. (2009). Empirical bayes confidence intervals shrinking both means and variances. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, *71*(1), 265–285.
- Gneiting, T., & Raftery, A. E. (2007). Strictly proper scoring rules, prediction, and estimation. *Journal of the American statistical Association*, *102*(477), 359–378.
- Higgins, J. P., Thompson, S. G., & Spiegelhalter, D. J. (2009). A re-evaluation of random-effects meta-analysis. *Journal of the Royal Statistical Society Series A: Statistics in Society*, *172*(1), 137–159.

- Kleinbaum, D. G., Kupper, L. L., & Morgenstern, H. (1991). *Epidemiologic research: Principles and quantitative methods*. John Wiley & Sons.
- LeBlanc, P.M., & Banks, D. (2022). Time-varying bayesian network meta-analysis. arXiv preprint [arXiv:2211.08312](https://arxiv.org/abs/2211.08312).
- Lindley, D. V., & Smith, A. F. M. (1972). Bayes estimates for the linear model. *Journal of the Royal Statistical Society. Series B (Methodological)*, 34(1), 1–41.
- Liu-Thompkins, Y., Khoshghadam, L., Attar Shoushtari, A., & Zal, S. (2022). What drives retailer loyalty? a meta-analysis of the role of cognitive, affective, and social factors across five decades. *Journal of Retailing*, 98(1), 92–110.
- Mawdsley, D., Bennetts, M., Dias, S., Boucher, M., & Welton, N. J. (2016). Model-based network meta-analysis: A framework for evidence synthesis of clinical trial data. *CPT: Pharmacometrics & Systems Pharmacology*, 5(8), 393–401.
- Mela, C. F., Jedidi, K., & Bowman, D. (1998). The long-term impact of promotions on consumer stockpiling behavior. *Journal of Marketing Research*, 35(2), 250–262.
- Miller, A. P., & Hosanagar, K. (2020). An empirical meta-analysis of e-commerce A/B testing strategies. Technical report, Working paper.
- Oehlert, G. W. (1992). A note on the delta method. *The American Statistician*, 46(1), 27–29.
- Park, T., & Casella, G. (2008). The bayesian lasso. *Journal of the American Statistical Association*, 103(482), 681–686.
- Pastor, D. A., & Lazowski, R. A. (2018). On the multilevel nature of meta-analysis: A tutorial, comparison of software programs, and discussion of analytic choices. *Multivariate Behavioral Research*, 53(1), 74–89.
- Pedder, H., Dias, S., Bennetts, M., Boucher, M., & Welton, N. J. (2019). Modelling time-course relationships with multiple treatments: Model-based network meta-analysis for continuous summary outcomes. *Research Synthesis Methods*, 10(2), 267–286.
- Qi, X., Zhou, S., Wang, Y., & Peterson, C. (2022). Bayesian sparse modeling to identify high-risk subgroups in meta-analysis of safety data. *Research Synthesis Methods*, 13(6), 807–820.
- Roberts, S., Osborne, M., Ebden, M., Reece, S., Gibson, N., & Aigrain, S. (2013). Gaussian processes for time-series modelling. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 371(1984), 20110550.
- Salanti, G., Marinho, V., & Higgins, J. P. (2009). A case study of multiple-treatments meta-analysis demonstrates that covariates should be considered. *Journal of Clinical Epidemiology*, 62(8), 857–864.
- Sohn, S. Y. (1999). Meta analysis of classification algorithms for pattern recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21(11), 1137–1144.
- Stanley, T. D., & Doucouliagos, H. (2012). *Meta-regression analysis in economics and business*. Routledge.
- Stourm, V., & Bradlow, E. T. (2023). Cross-reward effects in a coalition loyalty program: The impact of a point currency devaluation. *International Journal of Research in Marketing*, 40(2), 276–293.
- Sugasawa, S., & Kubokawa, T. (2020). Small area estimation with mixed models: A review. *Japanese Journal of Statistics and Data Science*, 3, 693–720.
- Sugasawa, S., Tamae, H., & Kubokawa, T. (2017). Bayesian estimators for small area models shrinking both means and variances. *Scandinavian Journal of Statistics*, 44(1), 150–167.
- Thompson, C. G., & Becker, B. J. (2020). A group-specific prior distribution for effect-size heterogeneity in meta-analysis. *Behavior Research Methods*, 52, 2020–2030.
- van Enst, W. A., Naaktgeboren, C. A., Ochodo, E. A., de Groot, J. A. H., Leeftang, M. M. G., Reitsma, J. B., Scholten, R. J. P. M., Moons, K. G. M., Zwinderman, A. H., Bossuyt, P. M. M., & Hooft, L. (2015). Small-study effects and time trends in diagnostic test accuracy meta-analyses: A meta-epidemiological study. *Systematic Reviews*, 4, 66.
- Wang, Z., Lin, L., Hodges, J. S., MacLehose, R., & Chu, H. (2021). A variance shrinkage method improves arm-based Bayesian network meta-analysis. *Statistical Methods in Medical Research*, 30(1), 151–165.
- Williams, C. K., & Rasmussen, C. E. (2006). *Gaussian processes for machine learning*. MIT Press.
- Williams, D. R., Rodriguez, J. E., & Bürkner, P. C. (2021). Putting variation into variance: Modeling between-study heterogeneity in meta-analysis. PsyArXiv.
- Zerbin, C., Bijmolt, T. H., Maestriperri, S., & Luceri, B. (2022). Drivers of consumer adoption of e-commerce: A meta-analysis. *International Journal of Research in Marketing*, 39(4), 1186–1208.