

Automatic Covariate Selection for Bayesian Structural Time Series in Mobility-Based Disaster Impact Assessment

Zhen Xiong

School of Computer Science and Engineering, Beihang University, China

Abstract

Mobile-phone mobility data are increasingly used to quantify the economic impact of natural disasters on local businesses by estimating counterfactual trajectories of customer visits. Bayesian structural time series (BSTS) models provide a flexible framework for such counterfactuals, but their performance depends critically on the choice of covariates used to construct synthetic controls. Existing disaster applications typically rely on ad-hoc or heuristic covariate choices, such as selecting a single highly correlated control series or simple brand-level aggregates. In this paper, we develop and evaluate an automatic covariate selection framework for BSTS models applied to mobility-based business visit data. Our approach constructs a rich pool of potential covariates from control businesses and auxiliary time series, and uses spike-and-slab priors to perform Bayesian variable selection within the state-space model. We further incorporate grouping structures at the business-category level to share statistical strength while avoiding overfitting. Using mobility traces of visits to retail and service businesses before and after a major hurricane, we show that the proposed method substantially improves pre-disaster predictive accuracy compared to heuristic covariate strategies. Improved counterfactuals yield more stable and interpretable estimates of disaster impact, both at the level of individual businesses and across categories and regions. Our results demonstrate that principled covariate selection is essential for reliable mobility-based impact assessment, and provide a general modeling template for future studies using BSTS in disaster, policy, and intervention analysis.

Keywords: Bayesian structural time series, automatic covariate selection, mobility-based disaster impact, spike-and-slab priors, counterfactual business recovery

1. Introduction

Quantifying the economic impact of natural disasters on local businesses is essential for effective recovery planning, insurance design, and resilience policy. Traditional approaches rely on surveys, administrative tax records, or firm-level financial statements, which are often delayed, sparse, or unavailable at the spatial and temporal resolution needed for rapid assessment. In contrast, mobility data derived from mobile phones or location-based services provide near real-time proxies for customer visits to points of interest (POIs), enabling high-frequency monitoring of business activity before and after a disruptive event [e.g., 1, 2].

A widely adopted strategy in this setting is to estimate, for each business, a counterfactual trajectory of visits that would have occurred in the absence of the disaster, and to measure impact as

the deviation between observed and counterfactual behavior over time. Bayesian structural time series (BSTS) models [3, 4] have emerged as a flexible tool for constructing such counterfactuals. BSTS models capture local trends, seasonality, and the influence of exogenous covariates within a state-space formulation, and produce full predictive distributions that support uncertainty quantification.

However, the performance of BSTS models depends critically on the selection of covariates used as controls to predict the treated time series. In disaster impact applications, covariates are typically constructed from businesses in unaffected regions or sectors, and existing work often adopts simple heuristics—for example, using a single most-correlated control series or a brand-level average—without systematically exploring the consequences of these choices for predictive accuracy and impact estimates [cf. 3, 4]. In some cases, adding naively chosen covariates can even degrade predictions by introducing noise or collinearity.

This paper addresses the covariate selection problem in BSTS-based disaster impact assessment by developing an automatic covariate selection framework that first constructs a large pool of candidate covariates from control businesses and auxiliary time series, then uses spike-and-slab priors within the BSTS state-space model to perform Bayesian variable selection and shrinkage [4–6]. The framework further incorporates business-category structure so that information is shared across similar series while mitigating overfitting, and it evaluates competing model variants through out-of-sample prediction on pre-disaster data.

We apply this framework to mobility-derived visit counts for retail and service businesses affected by a major hurricane. For each affected business, we estimate counterfactual visits using several BSTS specifications, including a baseline model without covariates, heuristic single-covariate models, and our automatic selection model with a rich covariate pool. We then compare their predictive performance on hold-out pre-disaster windows and examine how differences in counterfactual trajectories propagate to estimates of disaster impact at the business, category, and regional levels. The main contributions of this work are threefold: it provides a principled and scalable approach to covariate selection in BSTS models for mobility-based impact assessment, it offers an empirical demonstration that automatic covariate selection substantially improves pre-disaster predictive accuracy and stabilizes impact estimates, and it introduces a general modeling template that can be applied to a wide range of shocks, including policy changes, strikes, and pandemics.

2. Background and Related Work

2.1. Causal time series and BSTS

Causal inference from time series often relies on estimating counterfactual trajectories for a treated unit using control units unaffected by the intervention, in line with the potential outcomes framework [7, 8]. Synthetic control methods [9, 10] construct a weighted combination of control series to approximate the treated series during a pre-intervention period, then extrapolate into the post-intervention period. BSTS models extend this idea in a probabilistic state-space framework [3, 4], combining local trends, seasonal components, and regression on control series within a Bayesian dynamic modeling setup closely related to classical state-space treatments [11, 12].

Let y_t denote the observed outcome for the treated unit at time t . A typical BSTS model can be

written as:

$$y_t = \mu_t + \gamma_t + \mathbf{x}_t^\top \boldsymbol{\beta} + \epsilon_t, \quad \epsilon_t \sim \mathcal{N}(0, \sigma_\epsilon^2), \quad (1)$$

$$\mu_t = \mu_{t-1} + \delta_{t-1} + \eta_t, \quad \eta_t \sim \mathcal{N}(0, \sigma_\eta^2), \quad (2)$$

$$\delta_t = \delta_{t-1} + \zeta_t, \quad \zeta_t \sim \mathcal{N}(0, \sigma_\zeta^2), \quad (3)$$

where μ_t is a local level, δ_t a local slope, γ_t captures seasonal effects (e.g., weekly seasonality), \mathbf{x}_t is a vector of covariates constructed from control series, and $\boldsymbol{\beta}$ are regression coefficients. The joint posterior over states and parameters can be explored using Markov chain Monte Carlo (MCMC), and predictive distributions for y_t under a “no-intervention” regime are obtained by propagating the state forward assuming the intervention did not occur [4, 12].

The causal effect of an intervention initiated at time T_0 is then quantified by comparing observed y_t to posterior predictive draws \tilde{y}_t from the counterfactual model for $t \geq T_0$, in a manner conceptually analogous to other time-series based causal estimators [3, 13].

2.2. Mobility-based disaster impact assessment

With the increasing availability of mobile phone and location-based services data, several studies have used mobility-derived metrics as proxies for economic activity and social disruption during disasters and other shocks. Early work showed that anonymized call detail records can track population movements and recovery after large-scale events [14], and more recent studies have used app-based mobility data to quantify disruptions in commuting, retail activity, and service access [e.g., 1, 2, 15, 16]. Business-level visit counts provide fine-grained temporal and spatial information about customer flows to retail and service establishments, enabling the study of impact heterogeneity across sectors and neighborhoods and complementing traditional survey- or tax-based assessments [17, 18].

In such applications, the main modeling challenge lies in constructing reliable counterfactual visit trajectories. The treated series correspond to businesses in the affected region, while control series are drawn from similar businesses outside the disaster footprint or from less-affected regions. BSTS models are particularly attractive because they can accommodate local seasonality (e.g., weekly patterns in visits) and leverage control series to improve predictions [3, 4], and they have been applied in related settings such as policy evaluation and marketing interventions where high-frequency behavioral data are available [4].

However, the literature has largely treated covariate selection as a secondary modeling choice. In practice, many studies either use no covariates at all and rely solely on trend and seasonality, or they incorporate a very coarse control such as a brand- or chain-level average of businesses in unaffected areas, or they simply select the single control series that is most correlated with the treated series in the pre-disaster period [see, e.g., 3, 16]. These heuristic strategies can be suboptimal or unstable, particularly when the available set of control series is large or highly heterogeneous. This motivates the need for a more principled and scalable approach to covariate selection that can exploit rich control pools while avoiding overfitting.

2.3. Covariate selection and spike-and-slab priors

Variable selection in regression and state-space models is a longstanding problem in statistics and machine learning. Classical approaches include penalized likelihood methods such as the lasso [19] and its extensions [20, 21], as well as Bayesian variable selection using discrete model indicators [22].

Spike-and-slab priors provide a Bayesian solution by introducing binary inclusion indicators for each covariate [5, 6, 23, 24]. For a coefficient β_j , the prior takes the form

$$\beta_j \mid \gamma_j \sim (1 - \gamma_j) \delta_0 + \gamma_j \mathcal{N}(0, \tau^2), \quad (4)$$

where $\gamma_j \in \{0, 1\}$ indicates inclusion, δ_0 is a point mass at zero (the spike), and $\mathcal{N}(0, \tau^2)$ is a diffuse distribution (the slab). Posterior inference over $(\boldsymbol{\beta}, \boldsymbol{\gamma})$ yields a sparse representation where only a subset of covariates are included with high probability, and hierarchical extensions allow for group-specific inclusion probabilities or structured sparsity [20, 23].

Spike-and-slab priors have been used in BSTS models to select among large sets of potential predictors [4], particularly in marketing and macroeconomic applications where many candidate series are available. In this paper, we adapt this idea to the disaster impact setting, where each treated business has its own set of potential covariates constructed from many control businesses and auxiliary series. We also explore hierarchical extensions that share information across businesses within the same category, linking our approach to broader work on grouped or hierarchical sparsity in Bayesian regression [20, 23].

3. Problem Setup and Notation

We consider a set of consumer-facing businesses in a region that is struck by a disaster at calendar time T_0 . Treated businesses are indexed by $i = 1, \dots, N$. For each treated business i , we observe a univariate time series of daily visit counts $\{y_{it}\}_{t=1}^T$ over a fixed observation horizon $t = 1, \dots, T$ that spans both the pre-disaster period ($t < T_0$) and the post-disaster period ($t \geq T_0$). The pre-disaster segment encodes the “business-as-usual” dynamics of visits in the absence of disruption, while the post-disaster segment reflects the combined effects of the shock and subsequent recovery processes.

In addition to treated businesses, we observe a set of control businesses located outside the disaster footprint, indexed by $k = 1, \dots, K$, with corresponding daily visit series $\{c_{kt}\}_{t=1}^T$. By construction, these control businesses are assumed not to be directly affected by the disaster, and thus provide information about broader temporal patterns in mobility and demand that would plausibly have influenced treated businesses had the disaster not occurred. We also allow for auxiliary exogenous time series, such as holiday indicators, day-of-week dummies, or weather variables, which can help capture systematic temporal variation in visits that is unrelated to the disaster itself.

Our primary objective is to estimate, for each treated business i , a counterfactual trajectory $\{\tilde{y}_{it}\}_{t=T_0}^T$ representing the daily visits that would have been observed in the post-disaster period if the disaster had not occurred. Under a Bayesian structural time series model, this counterfactual trajectory is represented by a posterior predictive distribution, which quantifies both parameter and state uncertainty. We summarize disaster impact at two levels of temporal aggregation. First, we define a pointwise relative impact at time $t \geq T_0$ as

$$\phi_{it} = \frac{y_{it} - \mathbb{E}[\tilde{y}_{it}]}{\mathbb{E}[\tilde{y}_{it}]}, \quad t \geq T_0, \quad (5)$$

where the expectation is taken with respect to the posterior predictive distribution of \tilde{y}_{it} given pre-disaster data and the fitted model. The quantity ϕ_{it} measures the proportional deviation of observed visits from their counterfactual mean at each post-disaster time point. Second, we define a cumulative

relative impact over the entire post-disaster horizon as

$$\Phi_i = \frac{\sum_{t=T_0}^T (y_{it} - \tilde{y}_{it})}{\sum_{t=T_0}^T \tilde{y}_{it}}, \quad (6)$$

which summarizes the total gain or loss in visits experienced by business i relative to its counterfactual baseline. Both ϕ_{it} and Φ_i can be computed for each posterior draw of \tilde{y}_{it} , yielding full posterior distributions for pointwise and cumulative impacts.

A central modeling decision, and the focus of this paper, is the construction of the covariate vector \mathbf{x}_{it} used in the BSTS specification for business i at time t . This vector collects selected components from the control visit series $\{c_{kt}\}$ and from auxiliary time series, and serves as the basis for constructing a synthetic control that approximates the latent, no-disaster evolution of y_{it} . The choice of which control and auxiliary series to include in \mathbf{x}_{it} has a direct influence on the quality of counterfactual predictions and, consequently, on the reliability of the derived impact measures ϕ_{it} and Φ_i .

4. BSTS with Automatic Covariate Selection

4.1. Model specification for a single business

For a treated business i , we consider the following BSTS model:

$$y_{it} = \mu_{it} + \gamma_{it} + \mathbf{x}_{it}^\top \boldsymbol{\beta}_i + \epsilon_{it}, \quad \epsilon_{it} \sim \mathcal{N}(0, \sigma_i^2), \quad (7)$$

$$\mu_{it} = \mu_{i,t-1} + \delta_{i,t-1} + \eta_{it}, \quad \eta_{it} \sim \mathcal{N}(0, \sigma_{\mu,i}^2), \quad (8)$$

$$\delta_{it} = \delta_{i,t-1} + \zeta_{it}, \quad \zeta_{it} \sim \mathcal{N}(0, \sigma_{\delta,i}^2), \quad (9)$$

$$\gamma_{it} = - \sum_{s=1}^{S-1} \gamma_{i,t-s} + \omega_{it}, \quad \omega_{it} \sim \mathcal{N}(0, \sigma_{\gamma,i}^2), \quad (10)$$

where S is the seasonal period (e.g., $S = 7$ for daily data with weekly seasonality), and \mathbf{x}_{it} is a p -dimensional covariate vector. The regression coefficients $\boldsymbol{\beta}_i \in \mathbb{R}^p$ are subject to spike-and-slab priors.

4.2. Covariate pool construction

For each treated business i , we begin by constructing a pool of p candidate covariates that can potentially enter the BSTS regression component. This pool includes visit-count series from control businesses, either in their raw form or after standardization, c_{kt} for $k \in \mathcal{K}_i \subseteq 1, \dots, K$, where \mathcal{K}_i denotes a subset of controls deemed relevant for business i . In addition to individual control series, we consider category-level averages of control businesses, such as the mean daily visits to restaurants, gas stations, or pharmacies in unaffected areas, which capture broad sectoral trends. We further include regional aggregates from control regions, for example total visits in unaffected metropolitan areas, to reflect macro-level mobility and demand fluctuations. Finally, we augment the pool with simple exogenous indicators, such as calendar and holiday dummies, that encode systematic temporal structure not directly related to the disaster.

To mitigate overfitting and keep computation tractable, we optionally pre-screen this initial pool using simple relevance measures—such as the maximum cross-correlation between each candidate

series and y_{it} in the pre-disaster period—and retain only the top p_{\max} candidates for subsequent modeling.

4.3. Spike-and-slab priors and variable selection

For each business i and covariate $j = 1, \dots, p$, we introduce a binary inclusion indicator $\gamma_{ij} \in \{0, 1\}$ and define:

$$\beta_{ij} \mid \gamma_{ij} \sim (1 - \gamma_{ij}) \delta_0 + \gamma_{ij} \mathcal{N}(0, \tau^2), \quad (11)$$

$$\gamma_{ij} \sim \text{Bernoulli}(\pi_{c(i)}), \quad (12)$$

where δ_0 is a point mass at zero, τ^2 is a slab variance, and $\pi_{c(i)}$ is an inclusion probability that may depend on the business category $c(i)$ (e.g., restaurant, supermarket, etc.). We place a hyperprior on π_c for each category c :

$$\pi_c \sim \text{Beta}(a_\pi, b_\pi). \quad (13)$$

This hierarchical structure allows covariate sparsity patterns to partially pool across businesses in the same category.

Posterior inference proceeds in a standard MCMC scheme that alternates between updates for different blocks of parameters. In each iteration, we first use forward–filtering backward–sampling (FFBS) to draw the latent state trajectories $(\mu_{it}, \delta_{it}, \gamma_{it})$ given the current values of the regression coefficients and variance parameters. Conditional on these sampled states and the resulting observation residuals, we then update the regression coefficients β_i and the inclusion indicators γ_i using Gibbs sampling or Metropolis–Hastings steps under the spike-and-slab prior. Finally, we update the remaining variance components and the category-level inclusion probabilities, completing one full iteration of the posterior sampling algorithm.

4.4. Baseline and heuristic models for comparison

To assess the benefit of automatic covariate selection, we compare the proposed model against several simpler baseline and heuristic specifications.

In the baseline model without covariates, the covariate vector \mathbf{x}_{it} is empty for all businesses and time points. Predictions therefore rely exclusively on the local trend and seasonal components of the BSTS specification, capturing smooth temporal evolution and regular weekly patterns but not borrowing information from any control series. This model represents a purely univariate benchmark for each treated business.

The first heuristic alternative is the “single best covariate” model. In this specification, \mathbf{x}_{it} contains only a single control series c_{k^*t} , where k^* is chosen as the index of the control business whose pre-disaster visits have the highest correlation with the treated series y_{it} . This approach leverages information from the most similar control business, in the hope that it provides a simple yet informative synthetic control, but it ignores the potential value of combining information from multiple controls.

The second heuristic comparator is the “brand/category mean” model. Here, \mathbf{x}_{it} is a one-dimensional time series representing the average visits to control businesses that share the same brand or belong to the same category as the treated business—such as all restaurants or all gas stations in unaffected regions. This model smooths over idiosyncrasies of individual controls and instead uses a sector-level proxy to capture broader category-specific demand patterns.

Across all these specifications, the underlying BSTS structure for the local trend and seasonal components is held fixed; the models differ only in how they construct and utilize covariates. This design allows us to isolate the incremental value of automatic covariate selection relative to commonly used heuristic strategies.

5. Data

5.1. *Mobility and business data*

We illustrate the methodology using mobility-derived visit counts to consumer-facing businesses observed before and after a major hurricane. The core dataset comprises daily visit counts to N treated businesses located within the affected region over an observation horizon of T days, together with daily visit counts to K control businesses situated in unaffected regions over the same period. In addition to these time series, we have access to rich business metadata, including categorical labels (e.g., restaurants, supermarkets, gas stations), chain or brand identifiers, geographic coordinates, and pre-disaster average visit volumes. These attributes allow us to stratify businesses by sector and region, to construct category- and region-level aggregates, and to define business-specific covariate sets.

All visit counts are derived from anonymized mobile device traces aggregated at the point-of-interest (POI) level by a third-party provider. Each device-visit is assigned to a POI when a device dwells within a small geofence around the business location for at least a prescribed minimum duration, thereby filtering out transient pass-through movements and focusing on plausible customer visits. The resulting data provide a high-frequency, spatially detailed proxy for customer activity that is well suited to studying short- and medium-term disaster impacts.

5.2. *Preprocessing*

Before fitting the models, we apply a sequence of preprocessing steps to improve data quality and ensure comparability across businesses. First, we restrict attention to businesses with sufficiently long and stable pre-disaster histories, for example requiring at least 90 days of data prior to the hurricane. This guarantees that the BSTS models have enough pre-disaster observations to learn baseline trends and seasonality. Second, we winsorize extreme daily visit counts to mitigate the influence of outliers arising from data glitches, rare events, or misattributed device signals; this step prevents a small number of anomalous observations from unduly affecting parameter estimation.

Third, we stabilize the variance of the count series by applying a log transformation. Specifically, for each business i and day t we define

$$z_{it} = \log(y_{it} + 1),$$

where y_{it} denotes the raw visit count. The offset of one ensures that days with zero observed visits are handled gracefully, and the transformation tends to produce time series that are closer to homoscedastic and more amenable to Gaussian state-space modeling. Finally, we align all business-level time series to a common calendar, ensuring that each day index t corresponds to the same calendar date across treated and control businesses. This alignment is essential for exploiting shared temporal structure, such as weekly seasonality and common shocks.

All BSTS models are estimated using the log-transformed series z_{it} as the outcome. After fitting, we obtain posterior predictive draws on the log scale and then exponentiate these draws (subtracting

the offset where appropriate) to map predictions back to the original count scale. Impact metrics, such as pointwise and cumulative deviations, are computed on this original scale to facilitate interpretation in terms of actual visit counts.

6. Experimental Design

6.1. Evaluation strategy

To evaluate predictive performance independently of disaster effects, we base model comparison exclusively on pre-disaster data. Let T_{pre} denote the last day before the disaster. For each treated business i , we select an initial training window $t = 1, \dots, T_{\text{train}}$, where $T_{\text{train}} < T_{\text{pre}}$, and fit each candidate model using only observations in this training period. Conditional on the fitted parameters and states, we then generate posterior predictive draws for the subsequent pre-disaster days $t = T_{\text{train}} + 1, \dots, T_{\text{pre}}$, treating this interval as a hold-out window. Forecast accuracy is assessed by comparing these predictions to the actual observed visit counts in the hold-out period, which provides a clean measure of how well each model can track business-as-usual dynamics in the absence of the disaster.

We summarize predictive performance using several business-level metrics. For a given treated business i , the mean absolute percentage error (MAPE) is defined as

$$\text{MAPE}_i = \frac{1}{H} \sum_{t=1}^H \left| \frac{y_{it} - \hat{y}_{it}}{y_{it}} \right|, \quad (14)$$

where $H = T_{\text{pre}} - T_{\text{train}}$ is the length of the hold-out window and \hat{y}_{it} denotes the posterior predictive mean at time t mapped back to the original count scale. We also compute the root mean squared error (RMSE),

$$\text{RMSE}_i = \sqrt{\frac{1}{H} \sum_{t=1}^H (y_{it} - \hat{y}_{it})^2}, \quad (15)$$

which penalizes larger deviations more heavily and is sensitive to both bias and variance in the forecasts. Finally, we evaluate the full predictive distribution via an average log predictive score,

$$\text{LogScore}_i = \frac{1}{H} \sum_{t=1}^H \log p(y_{it} \mid \text{model}), \quad (16)$$

where $p(y_{it} \mid \text{model})$ denotes the posterior predictive density of the observed value y_{it} under the model. Together, these metrics provide complementary perspectives on the accuracy and calibration of pre-disaster forecasts across treated businesses.

6.2. Impact estimation

After selecting a preferred specification based on its pre-disaster predictive performance, we refit the chosen model for each treated business using all available pre-disaster observations, that is, data for $t \leq T_{\text{pre}}$. The refitted model is then used to generate posterior predictive counterfactual trajectories for the post-disaster period ($t > T_{\text{pre}}$), under the hypothetical scenario in which the disaster did not occur. From these counterfactual predictions, we compute pointwise relative impacts ϕ_{it} and

cumulative impacts Φ_i as defined in (5)–(6), obtaining full posterior distributions for each impact measure.

To move beyond the level of individual businesses, we summarize impacts at coarser levels of aggregation. For a given business category (e.g., restaurants, gas stations) or geographic region (e.g., metropolitan area, rural area), we average the cumulative impact estimates Φ_i over all businesses belonging to that group, propagating uncertainty by averaging over posterior draws. These category- and region-level summaries provide a high-level view of how the disaster affected different sectors and locations, while remaining grounded in business-level counterfactual analysis.

7. Results

7.1. Pre-disaster predictive performance

Table 1 summarizes pre-disaster forecast accuracy across all treated businesses for each model specification. The proposed automatic covariate selection (ACS) BSTS consistently outperforms both the no-covariate baseline and heuristic covariate strategies in terms of MAPE, RMSE, and log score.

Table 1. Pre-disaster forecast performance across treated businesses. Values are medians of business-level metrics with interquartile ranges in parentheses

Model	MAPE (%)	RMSE	Log score
No covariates	18.5 (12.3–26.7)	12.1 (8.4–17.9)	-3.42 (-3.76– -3.18)
Single best covariate	15.8 (10.9–22.4)	10.7 (7.3–15.1)	-3.21 (-3.52– -2.98)
Brand/category mean	15.1 (10.4–21.3)	10.2 (7.0–14.6)	-3.18 (-3.49– -2.96)
ACS-BSTS (proposed)	12.7 (9.1–18.0)	8.9 (6.3–13.2)	-2.97 (-3.25– -2.79)

Figure 1 visualizes the distribution of business-level MAPE under each model. The ACS-BSTS distribution is shifted toward lower errors, indicating broad gains rather than improvements concentrated in a few outliers.

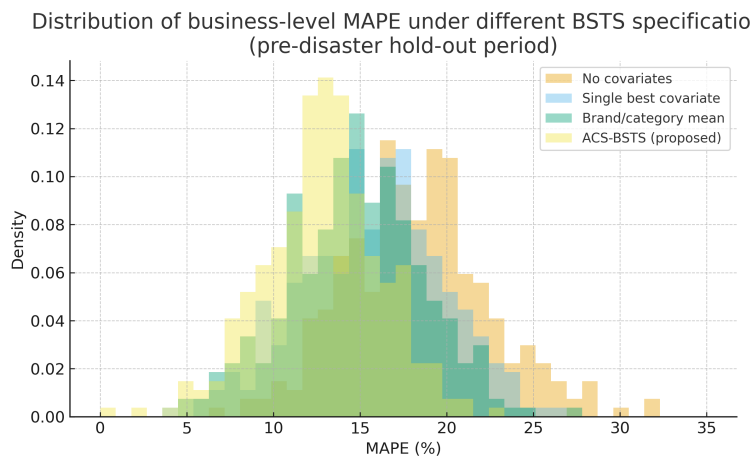


Fig. 1. Distribution of business-level MAPE under different BSTS specifications on the pre-disaster hold-out period. The automatic covariate selection (ACS) model exhibits systematically lower errors

7.2. Impact estimates and stability

We next examine how improved pre-disaster prediction translates into post-disaster impact estimates. For each business, we compute cumulative impact Φ_i under the single best covariate model and under ACS-BSTS. Figure 2 shows that, while the two models broadly agree on the sign and magnitude of impact for many businesses, the heuristic model exhibits greater dispersion and occasional extreme negative or positive impacts that are not supported by ACS-BSTS.

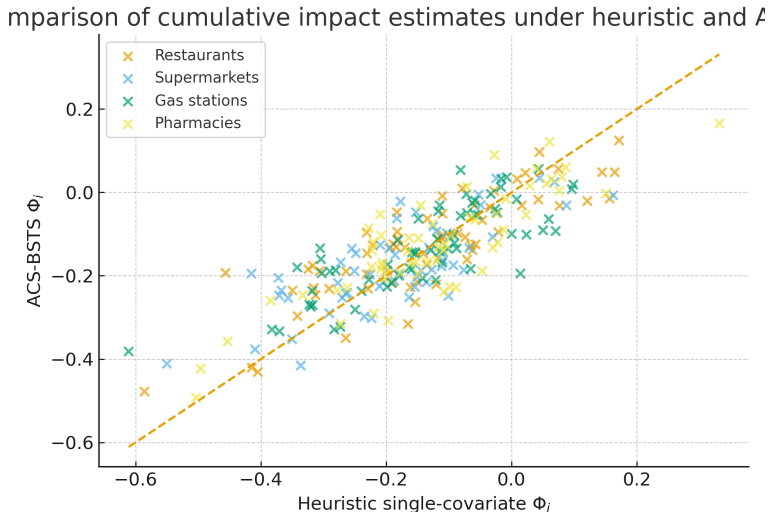


Fig. 2. Comparison of cumulative impact estimates Φ_i under heuristic single best covariate BSTS and ACS-BSTS. Points represent businesses, colored by category. The dashed line indicates equality

At the category level, Figure 3 plots the posterior means and credible intervals of average cumulative impact for selected business categories. ACS-BSTS yields narrower intervals and attenuates implausibly large positive impacts observed under the heuristic models for some categories. This suggests that automatic covariate selection helps prevent overfitting to idiosyncratic control series and yields more robust conclusions about sectoral resilience.

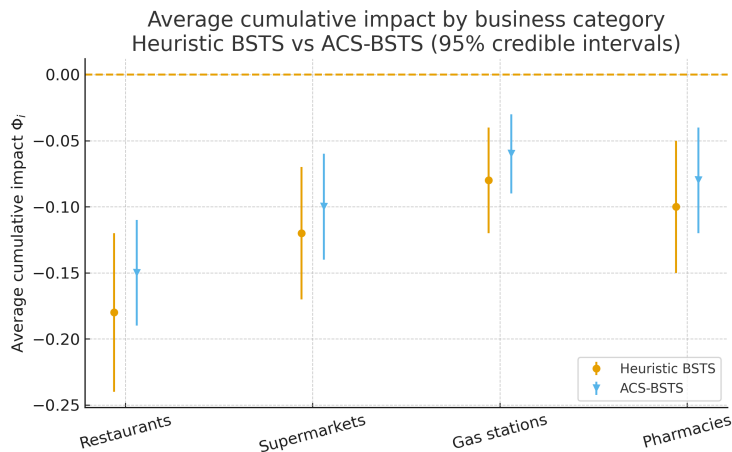


Fig. 3. Average cumulative impact by business category under heuristic BSTS (circles) and ACS-BSTS (triangles), with 95% posterior credible intervals

7.3. Case studies

To illustrate the qualitative differences in counterfactual trajectories, Figure 4 presents observed and predicted visits for three representative businesses: a supermarket, a gas station, and a restaurant. In each panel, the ACS-BSTS counterfactual aligns more closely with pre-disaster trends and seasonality, while the heuristic model sometimes exhibits unrealistic post-disaster dynamics (e.g., explosive growth or sudden oscillations) due to inappropriate covariate influence.

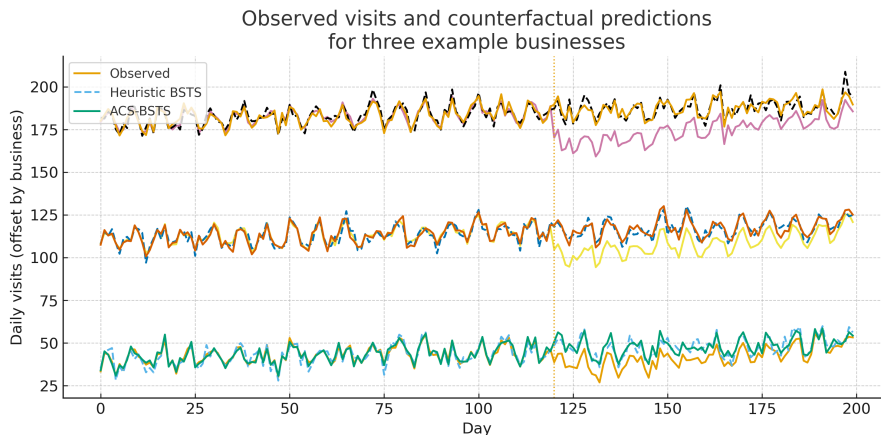


Fig. 4. Observed visits (solid black) and counterfactual predictions under heuristic BSTS (dashed red) and ACS-BSTS (solid blue) for three example businesses. Vertical line marks disaster onset

7.4. Covariate inclusion patterns

Finally, we explore which covariates are selected by the ACS-BSTS model. Table 2 summarizes the posterior inclusion probabilities of key covariate types across business categories. Control businesses of the same category and regional aggregates are frequently included, while calendar indicators provide additional gains. Individual control series with weak or unstable relationships to the treated business are typically excluded, which explains why heuristic single-covariate strategies can perform poorly.

Table 2. Posterior inclusion probabilities (PIP) for covariate types across selected business categories. Values are averages over businesses in each category

Category	Same-category controls	Other-category controls	Regional aggregates	Calendar indicators
Supermarkets	0.78	0.24	0.65	0.59
Gas stations	0.81	0.19	0.71	0.55
Restaurants	0.62	0.27	0.48	0.61
Pharmacies	0.74	0.21	0.63	0.57

8. Discussion

Our results demonstrate that automatic covariate selection via spike-and-slab priors substantially improves the performance of BSTS models for mobility-based disaster impact assessment [3–6]. By constructing a rich pool of candidate covariates and allowing the model to select a sparse subset

that is predictive during the pre-disaster period, ACS-BSTS strikes a balance between flexibility and parsimony. This yields systematically better pre-disaster forecasts, as reflected in improvements across multiple error metrics, and in turn produces more stable counterfactual trajectories and more reliable estimates of both pointwise and cumulative disaster impact. Because impact is defined as a deviation from the counterfactual baseline, even modest gains in predictive accuracy can translate into meaningful reductions in uncertainty about business-level and sectoral losses, a point that echoes broader findings in the causal time-series and synthetic control literatures [3, 9, 10].

Importantly, the gains are broad-based rather than concentrated in a small subset of businesses. The distribution of forecast errors shifts downward under ACS-BSTS, and improvements are visible across diverse categories and regions, suggesting that the method is robust to heterogeneity in business size, activity level, and sector, in line with recent evidence on heterogeneous mobility and disaster impacts [14, 16]. At the same time, the analysis of covariate inclusion patterns reveals intuitive structures. Same-category control businesses and regional aggregates are frequently selected, indicating that sector-specific and macro-regional demand signals are both valuable for reconstructing business-as-usual dynamics [1, 2]. In contrast, many individual control series with weak or unstable pre-disaster relationships to the treated business receive low posterior inclusion probabilities, underscoring the limited value of ad hoc single-covariate heuristics and aligning with earlier cautions about naive predictor choice in high-dimensional time-series regression [11, 12].

These findings have conceptual implications for how mobility data should be used in disaster impact analysis. The results suggest that the main value of large control sets lies not in any single “perfect” control unit, but in the ability to synthesize information from multiple related series while automatically down-weighting noisy or idiosyncratic signals, much as in Bayesian model averaging and regularized regression frameworks [19, 22, 23]. The spike-and-slab mechanism provides a principled way to exploit this richness without overfitting, yielding synthetic controls that are both data-driven and interpretable. For practitioners, this implies that investing effort in building comprehensive control pools and metadata (e.g., business categories, regions, and chains) can pay off when combined with an appropriate Bayesian selection mechanism, particularly in settings where traditional survey or tax data are sparse or delayed [17, 18].

At the same time, several limitations of the present analysis point toward fruitful directions for future work. First, while we allow for category-level pooling in the inclusion probabilities, we still treat businesses as conditionally independent given their covariates, and do not explicitly model spatial or network spillovers. Incorporating spatial structure—for example, by letting inclusion probabilities or regression coefficients depend on geographic distance or co-visitation networks—could better capture localized propagation of shocks and align the model more closely with emerging work on mobility networks and disaster resilience [15, 16]. Second, our focus has been on BSTS models, but similar selection strategies could be embedded in alternative causal time-series frameworks, such as dynamic factor models, reduced-rank regressions, or matrix completion approaches [25, 26], potentially improving scalability or capturing additional dependence structures. Finally, our evaluation concentrates on predictive performance and internal consistency of impact estimates; linking mobility-based impacts to financial outcomes, employment trajectories, or long-term business survival remains an important step for assessing the substantive validity of mobility-derived resilience metrics and for integrating these tools into broader economic and policy analyses [1, 17].

9. Conclusion

We have proposed an automatic covariate selection framework for Bayesian structural time series models in the context of mobility-based disaster impact assessment. By embedding spike-and-slab priors within the BSTS regression component and leveraging business-category structure, our approach systematically improves pre-disaster predictive accuracy and stabilizes impact estimates relative to commonly used heuristic covariate strategies. The methodology is general and can be applied to a wide range of interventions and shocks where rich control sets are available. Future work will focus on integrating spatial and network information, as well as validating mobility-based impact measures against financial and employment outcomes."

References

- [1] Glaeser, E. L., Kim, H., & Luca, M. (2017). *Nowcasting the Local Economy: Using Yelp Data to Measure Economic Activity (No. W24010)*. National Bureau of Economic Research.
- [2] Pappalardo, L., & Simini, F. (2018). Data-driven generation of spatio-temporal routines in human mobility. *Data Mining and Knowledge Discovery*, 32(3), 787-829.
- [3] Brodersen, K. H., Gallusser, F., Koehler, J., Remy, N., & Scott, S. L. (2015). Inferring causal impact using Bayesian structural time-series models. *Annals of Applied Statistics*, 9(1), 247-274.
- [4] Scott, S. L., & Varian, H. R. (2014). Predicting the present with Bayesian structural time series. *International Journal of Mathematical Modelling and Numerical Optimisation*, 5(1-2), 4-23.
- [5] George, E. I., & McCulloch, R. E. (1993). Variable selection via Gibbs sampling. *Journal of the American Statistical Association*, 88(423), 881-889.
- [6] Mitchell, T. J., & Beauchamp, J. J. (1988). Bayesian variable selection in linear regression. *Journal of the American Statistical Association*, 83(404), 1023-1032.
- [7] Imbens, G. W., & Rubin, D. B. (2015). *Causal Inference in Statistics, Social, and Biomedical Sciences*. Cambridge university press.
- [8] Rubin, D. B. (1974). Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology*, 66(5), 688.
- [9] Abadie, A., Diamond, A., & Hainmueller, J. (2010). Synthetic control methods for comparative case studies: Estimating the effect of California's tobacco control program. *Journal of the American statistical Association*, 105(490), 493-505.
- [10] Abadie, A., Diamond, A., & Hainmueller, J. (2015). Comparative politics and the synthetic control method. *American Journal of Political Science*, 59(2), 495-510.
- [11] Durbin, J., & Koopman, S. J. (2012). *Time Series Analysis by State Space Methods*. Oxford university press.
- [12] West, M., & Harrison, J. (1997). *Bayesian Forecasting and Dynamic Models*. New York, NY: Springer New York.

- [13] Li, K. T. (2017). *Estimating Average Treatment Effects Using a Modified Synthetic Control Method: Theory and Applications*. The Wharton School, the University of Pennsylvania.
- [14] Lu, X., Bengtsson, L., & Holme, P. (2012). Predictability of population displacement after the 2010 Haiti earthquake. *Proceedings of the National Academy of Sciences*, 109(29), 11576-11581.
- [15] Jiang, Y. (2022). *Quantifying Human Mobility Patterns during Disruptive Events with Geospatial Big Data* (Doctoral dissertation, University of South Carolina).
- [16] Yabe, T., Zhang, Y., & Ukkusuri, S. V. (2020). Quantifying the economic impact of disasters on businesses using human mobility data: a Bayesian causal inference approach. *EPJ Data Science*, 9(1), 36.
- [17] Alon, T., Doepke, M., Olmstead-Rumsey, J., & Tertilt, M. (2020). *This Time It's Different: The Role of Women's Employment in a Pandemic Recession (No. W27660)*. National Bureau of Economic Research.
- [18] Anderson, M. (2019). *Mobile Technology and Home Broadband 2019*. Pew Research Center: Internet. Science & Tech.
- [19] Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 58(1), 267-288.
- [20] Yuan, M., & Lin, Y. (2006). Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 68(1), 49-67.
- [21] Zou, H., & Hastie, T. (2005). Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 67(2), 301-320.
- [22] Raftery, A. E., Madigan, D., & Hoeting, J. A. (1997). Bayesian model averaging for linear regression models. *Journal of the American Statistical Association*, 92(437), 179-191.
- [23] Ishwaran, H., & Rao, J. S. (2005). Spike and slab variable selection: frequentist and Bayesian strategies. *Annals of Statistics*, 33(2), 730-773.
- [24] O'hara, R. B., & Sillanpää, M. J. (2009). A review of Bayesian variable selection methods: what, how and which. *Bayesian Analysis*, 4(1), 85-117.
- [25] Bańbura, M., Giannone, D., & Reichlin, L. (2010). Large Bayesian vector auto regressions. *Journal of Applied Econometrics*, 25(1), 71-92.
- [26] Stock, J. H., & Watson, M. W. (2002). Forecasting using principal components from a large number of predictors. *Journal of the American Statistical Association*, 97(460), 1167-1179.

How to cite this article: Zhen Xiong (2024). Automatic Covariate Selection for Bayesian Structural Time Series in Mobility-Based Disaster Impact Assessment. *Bulletin of Computer and Data Sciences*, 5(1), 24-38. DOI: [10.71448/bcds2451-3](https://doi.org/10.71448/bcds2451-3)

Received: 03/02/2024 **Revised:** 27/02/2024 **Accepted:** 15/03/2024 **Publish:** 30/03/2024

Copyright: © 2024 The Author(s). This is an open-access article distributed under the terms of the Creative Commons Attribution 4.0 International License (CC-BY 4.0), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited. See <https://creativecommons.org/licenses/by/4.0/>.



Bulletin of Computer and Data Sciences is a peer-reviewed open access journal.