

Yabe, Takahiro; Zhang, Yunchang; Ukkusuri, Satish V.

Working Paper

Quantifying the economic impact of disasters on businesses using mobility data

ADB Working Paper Series, No. 1182

Provided in Cooperation with:

Asian Development Bank Institute (ADBI), Tokyo

Suggested Citation: Yabe, Takahiro; Zhang, Yunchang; Ukkusuri, Satish V. (2020) : Quantifying the economic impact of disasters on businesses using mobility data, ADB Working Paper Series, No. 1182, Asian Development Bank Institute (ADBI), Tokyo

This Version is available at:

<https://hdl.handle.net/10419/238539>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



<https://creativecommons.org/licenses/by-nc-nd/3.0/igo/>



ADB Working Paper Series

**QUANTIFYING THE ECONOMIC IMPACT
OF DISASTERS ON BUSINESSES USING
MOBILITY DATA**

Takahiro Yabe, Yunchang Zhang,
and Satish V. Ukkusuri

No. 1182
September 2020

Asian Development Bank Institute

Takahiro Yabe is a PhD candidate at Lyles School of Civil Engineering of Purdue University. Yunchang Zhang is a PhD student at Lyles School of Civil Engineering of Purdue University. Satish V. Ukkusuri is a professor at Lyles School of Civil Engineering of Purdue University.

The views expressed in this paper are the views of the author and do not necessarily reflect the views or policies of ADBI, ADB, its Board of Directors, or the governments they represent. ADBI does not guarantee the accuracy of the data included in this paper and accepts no responsibility for any consequences of their use. Terminology used may not necessarily be consistent with ADB official terms.

Working papers are subject to formal revision and correction before they are finalized and considered published.

The Working Paper series is a continuation of the formerly named Discussion Paper series; the numbering of the papers continued without interruption or change. ADBI's working papers reflect initial ideas on a topic and are posted online for discussion. Some working papers may develop into other forms of publication.

Suggested citation:

Yabe, T., Y. Zhang, and S. Ukkusuri. 2020. Quantifying the Economic Impact of Disasters on Businesses Using Mobility Data. ADBI Working Paper 1182. Tokyo: Asian Development Bank Institute. Available: <https://www.adbi.org/publications/quantifying-economic-impact-disasters-businesses-using-mobility-data>

Please contact the authors for information about this paper.

Emails: tyabe@purdue.edu

Asian Development Bank Institute
Kasumigaseki Building, 8th Floor
3-2-5 Kasumigaseki, Chiyoda-ku
Tokyo 100-6008, Japan

Tel: +81-3-3593-5500
Fax: +81-3-3593-5571
URL: www.adbi.org
E-mail: info@adbi.org

© 2020 Asian Development Bank Institute

Abstract

In recent years, extreme shocks (such as natural disasters) have increased in both frequency and intensity. Consequently, many cities have experienced significant economic losses. Quantifying the economic cost to local businesses after extreme shocks is important for post-disaster assessment and pre-disaster planning. Conventionally, surveys have been the primary source of data to quantify the damages that are inflicted on businesses by disasters. However, surveys often suffer from high cost and their implementation can take a long time. They also suffer from spatio-temporal sparsity in observations and limitations in scalability. Recently, large scale human mobility data (e.g., mobile phone GPS) have been used to observe and analyze human mobility patterns in an unprecedented spatio-temporal granularity and scale. In this work, we use location data that were collected from mobile phones to estimate and analyze the causal impact of hurricanes on business performance. To quantify the causal impact of the disaster, we use a Bayesian structural time series model to predict the counterfactual performance of affected businesses (what if the disaster did not occur?), which may use the performance of other businesses outside the disaster areas as covariates. We have tested our method by quantifying the resilience of 635 businesses across nine categories in Puerto Rico after Hurricane Maria. Furthermore, hierarchical Bayesian models are used to reveal the effect of business characteristics (e.g., location and category) on the long-term resilience of these businesses. This study presents a novel and more efficient method to quantify business resilience, which could assist policy makers in disaster preparation and relief processes.

Keywords: disaster resilience, mobile phones, human mobility, causal inference

JEL Classification: Q54, C54, J6

1 Introduction

Natural hazards are currently increasing in both frequency and intensity in many parts of the world. The economic losses caused by these extreme events have exceeded a total of \$2.5 trillion across the globe since 2000, and are rising each year due to rapid urbanization in many cities [1]. With the intensifying threat of significant economic damage, the question of how to improve the resilience of cities has attracted interest from a wide range of fields, including public policy, urban planning, complex systems, and economics [2]. Among the various dimensions of disaster resilience, the ability of businesses to bounce back afterwards is a critical component that significantly contributes to the economic recovery of cities after disasters.

Previous studies have analyzed the post-disaster recovery of businesses through the means of surveys and interviews. These studies have identified factors such as the pre-disaster size of the business and category of business to partly explain the reopening and demise of businesses after disasters, including Hurricanes Katrina [3, 4], Andrew [5], and (more recently) Harvey [6]. Although these studies provide a general understanding of the effect of various characteristics of businesses that affect the post-disaster recovery performance, they suffer from two critical drawbacks. First, observations are limited to discrete measurements at a few number of timings, which fails to give a quantifiable, continuous and longitudinal understanding of the recovery process of businesses. Second, the applied methods fail to model the causal effect of the disaster, which requires a statistical framework that predicts the performance of businesses if the disaster did not occur.

With the emergence of novel and often large-scale data collected from mobile sensors and online social platforms, we are now capable of observing and analyzing the dynamics of people, goods, and information at an unprecedented spatio-temporal granularity [7]. In particular, location data collected from mobile phones (e.g. call detail records, GPS trajectories) have enabled us to observe individual mobility patterns at an unprecedented high spatio-temporal granularity [8, 9]. These datasets are now utilized for a wide range of applications to solve urban challenges, including population density estimation [10, 11], traffic estimation [12, 13], predicting poverty [14], and modeling the spread of epidemics [15]. In the context of extreme events, several studies have used mobile phone data to analyze mobility patterns during and after disasters, such as earthquakes [16, 17, 18], cyclones [19], and other anomalous events [20]. Despite this progress, none of the previous studies have used large scale mobility data to analyze the recovery of businesses after disasters.

Recent advances in statistical models, in particular Bayesian structural time series (BSTS) models, allow us to make flexible predictions of time series data, which can be used to estimate the causal impact [21]. BSTS models have several advantages over conventional difference in differences models [22], including their flexibility to model the causal impact over a longitudinal time horizon rather than across two time points. A recent study using website click-through data applied BSTS models to quantify the causal impact of an online advertisement [23]. We aim to take advantage of this recently proposed methodology to quantify the causal impact of hurricanes on businesses in Puerto Rico.

This study makes several contributions to overcome the drawbacks in the previous studies on business recovery after disasters. First, this is the first work to utilize large scale mobility data collected from mobile phones to estimate the popularity of businesses before, during, and after a disaster. Second, a Bayesian structural time series model combined with an inter-city matching scheme is proposed to infer the causal impact of the disaster on businesses. Third, the proposed methodology is applied on mobile phone data collected from Puerto Rico to quantify the resilience of businesses after Hurricane Maria. Figure 1 gives an overview of this study. The causal inference procedure is composed of three steps. i) To measure the causal impact of the disaster on business i , we first identify a similar business j in another region that was not affected by the disaster. ii) We then predict the counterfactual (“*what-if the disaster did not occur?*”) visit count of i after the disaster timing using observed data from j , via a Bayesian structural time series model. iii) Finally, we can quantify the causal impact of the disaster by taking the difference between the predicted and observed visit counts in i .

2 Related Works

2.1 Resilience of businesses after disasters

The economic impact of disasters on businesses have conventionally been studied through surveys that are performed after the disaster. Studies using surveys have identified various factors that affect the reopening and demise of businesses after disasters through econometric models (e.g. logistic regression) [3, 4]. Important factors that affected the outcomes of businesses after Hurricane Katrina include the household size of the business owner, previous disaster experience, number of employees, business age, and the legal structure of the business [3]. The qualitative details in the collected data are a significant advantage of these surveys. However, these surveys suffer from several drawbacks, including the high cost and long implementation time, spatio-temporal sparsity in observations, and limitations in scalability. Due

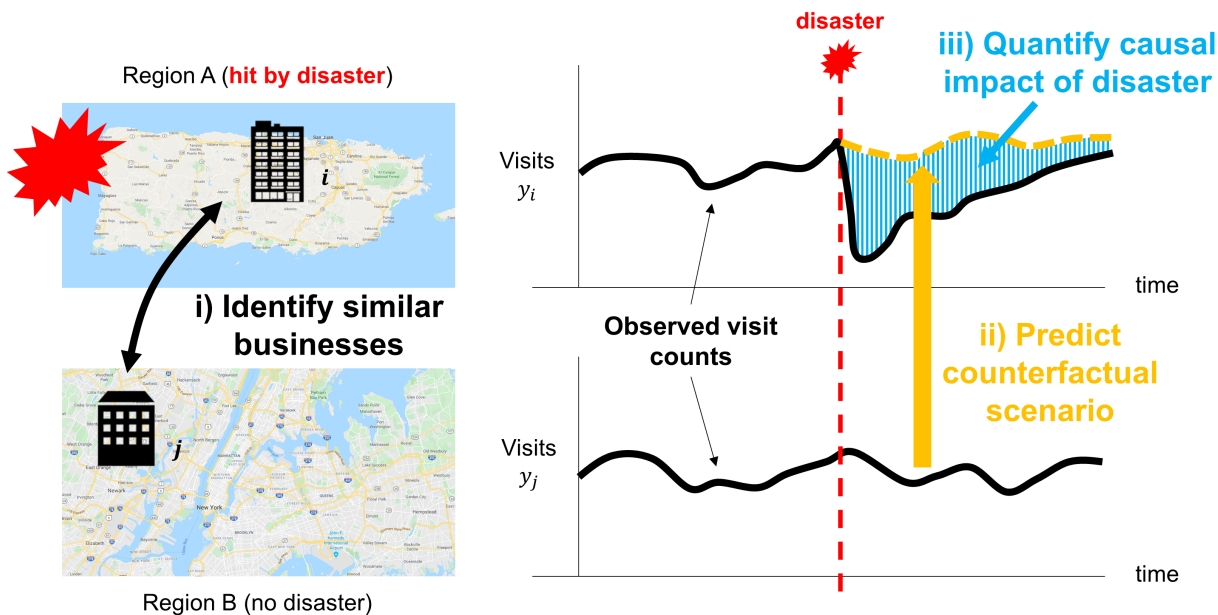


Figure 1: **Overview of study.** Our causal inference procedure is composed of three steps. i) To measure the causal impact of the disaster on business i , we first identify a similar business j in another region that was not affected by the disaster. ii) We then predict the counterfactual (“what-if the disaster did not occur?”) visit count of i after the disaster timing using observed data from j . iii) Finally, we can quantify the causal impact of the disaster by taking the difference between the predicted and observed visit counts in i .

to these limitations, it is difficult to obtain a quantifiable, continuous, and longitudinal understanding of the recovery process of these businesses. Moreover, the applied methods fail to model the causal effect of the disaster, which requires a statistical framework that predicts the performance of businesses if the disaster did not occur.

2.2 Mobility analysis using mobile phone data

With the emergence of novel and often large-scale data collected from mobile sensors and online social platforms, we are now capable of observing and analyzing the dynamics of people, goods, and information at an unprecedented spatio-temporal granularity [7]. In particular, location data collected from mobile phones (e.g. call detail records, and GPS trajectories) have enabled us to observe individual mobility patterns at an unprecedented high spatio-temporal granularity [8, 9]. These new datasets are becoming new standards for population level studies, and they are used to understand the population distribution in cities [10]. Furthermore, these datasets are now utilized for a wide range of applications to solve urban challenges, including population density estimation [11], estimation of dynamic traffic flows [12, 13], predicting poverty in developing counties [14], and modeling the impact of human mobility patterns on the spread of epidemics [15]. In the context of extreme events and disasters, several studies have used mobile phone data to analyze mobility patterns during and after disasters [16, 17, 18]. Studies using this large scale data have revealed important insights on the evacuation and migration patterns of the affected people [16, 19]. However, despite this progress, none of the previous studies have used large scale mobility data to analyze the recovery of businesses after disasters. A recent study using mobile phone GPS data (which is the same data as used in this study) has revealed the impact of the recent policy regarding the usage of bathrooms in Starbucks on the visit behavior of people to the cafe chain [24]. They validated that the spatio-temporal granularity of the mobile phone GPS data is of sufficient detail to analyze the store level visit behavior. In this study, we apply a similar approach and estimate the visit behavior of people to stores and businesses using mobile phone GPS data.

2.3 Statistical methods for causal inference

2.3.1 Difference in Differences

The difference in differences (DiD) method is a statistical method that is used to estimate the treatment effects between the "treatment" group versus the "control" group. For a before-and-after study, DiD compares the average change over time between treatment and control groups. This provides us with a classical method to estimate the causal effects of natural experiments without strict randomization [22, 25]. However, DiD has several limitations: First, it follows the parallel trends assumption, which requires that the differences between treatment and control group are invariant overtime in absence of the treatment [26, 27]. In a before-and-after study, the parallel trends assumption necessitates that the means of two groups should be balanced overtime. Consequently, issues such as time-correlated responses will contaminate the causal inference with DiD [28]. Second, only two time steps (i.e., pre-treatment time and post-treatment time) are considered in the classical DiD, which merely captures the static causal effects for a specific before-and-after study. This can be implausible and useless if the outcome of interest dynamically changes over time, such as recovery patterns after disaster, radioactive decay and so on [23].

2.3.2 Bayesian Structural Time Series Models

Compared with the classical DiD model, a structural time series model promisingly relaxes the parallel trends assumption and captures the variations of time-varying local trends and seasonality for time-correlated response variables [21, 29]. In addition, structural time series models encompass a flexible model structure, which enables us to analyze the dynamic effects of the outcome of interest during a time period [30]. Due to a large number of predictors in structural time series models, a Bayesian approach was introduced to sparse the estimation of coefficients. Scott and Varian [31, 32] proposed a spike-and-slab prior to the regression coefficients in a Google search query study, which significantly reduces the size of the problem. Nakajima and West [33] elicited a dynamic spike-and-slab prior that sparsified the estimation of time-varying parameters for a Bayesian macroeconomic time series model. The most recent Google study for causal inference of a market intervention [23] slightly revised the dynamic version of pike-and-slab prior [33] with a weakly informative prior. In addition, Bayesian structural time series models (BSTS) have been constructed to strengthen causal inference for time series data. To address the fundamental problem in causal inference [34], pre-treatment observations are trained and tested via BSTS. Consequently, the fitted BSTS can simulate the counterfactual as the synthetic post-treatment controls via posterior predictive samples. This method is extensively applied in causal inference throughout a wide variety of fields, such as socio-economics [35, 36], political science [37, 38], and environmental studies [39, 40].

The causal inference methodology that was proposed in the previous section is applied on data collected from Puerto Rico before and after Hurricane Maria, which made landfall on September 20, 2019, and caused a long term devastating humanitarian and economic crisis. Although the number of fatalities as a consequence of Maria are still under investigation, recent estimates suggest that between 793 to 8,498 excess deaths occurred following the storm [41]. Heavy rainfall, flooding, storm surge, and high winds caused considerable damage to various infrastructure systems, causing power outages and water shortages for the entire island for months. The total economic losses to Puerto Rico and the US Virgin Islands are estimated to be \$90 billion, with a 90% confidence range of \pm \$25.0 billion, which makes Maria the third costliest hurricane in US history, behind Katrina (2005) and Harvey (2017) [42].

3 Data

Three main data sources are used in this study: (1) business visit data collected from mobile phones, (2) spatial distribution of housing damages due to Hurricane Maria, and (3) socio-economic factors of census blocks in Puerto Rico. In this section, we describe how these datasets were collected, processed, and used to infer the causal impact of the hurricane on businesses.

3.1 Business visit data collected from mobile phones

Establishment-level visit data are provided by Safegraph¹, which is a company that aggregates anonymized location data collected from smartphone applications to provide insights about physical places. Safegraph's location dataset covers around 10% of all smartphones in the United States, and each observation is consisted of a unique (but anonymized) user ID, longitude, latitude, and timestamp information. The longitude and latitude information are accurate to within a few meters. This allows us to analyze the visit counts to each establishment. To detect a user visiting an establishment, the location data are first cleaned by removing GPS signal drifts and jumpy observations using a spatial threshold. The

¹<https://www.safegraph.com/>

Table 1: Summary statistics of business visit time series data.

Business Category	Region		
	Puerto Rico	Downstate New York	Upstate New York
Building Material	62	36	584
Gasoline Stations	10	143	1,160
Grocery Stores	34	97	1,227
Hospitals	12	25	76
Hotels	8	81	692
Restaurants	322	585	6,352
Supermarkets	61	0	52
Telecommunication	101	9	67
Universities	25	34	199
Total	635	1,102	10,409

data are then clustered into a staypoint using a spatio-temporal DBSCAN algorithm. Then, the visited establishment is predicted from establishments nearby the clustered staypoint by using a machine learning algorithm that takes into account features such as distances from establishment to the cluster centroid, time of day, and the North American Industry Classification System (NAICS) code. Performing this procedure for all days in the dataset produces a time series data of daily visit counts for each establishment.

We use daily visit data of establishments located in Puerto Rico and the State of New York between January 2017 and March 2018 to quantify the causal impact of the hurricane on business resilience. Daily visit data of businesses in New York are used because these businesses constitute a reasonable control group that was not affected by the disruptions caused by Hurricane Maria. We will describe how we use the visit data from the control group in the causal inference model in the methods section. We limit our analysis to business categories that sell products or services directly to the customers because we will approximate business performance from the number of visits per day, as observed from mobile phone data. We also limit the analysis to medium or large sized businesses with more than 100 customers per day on average (before the disaster) because we are unable to observe visit patterns below that level using mobile phone data. As summarized in Table 1, the daily visit data of a total of 635 businesses in Puerto Rico were analyzed, along with 1,102 and 10,409 businesses in Manhattan and Up-State New York, respectively. Figure 2A plots the locations of the different types of businesses in Puerto Rico, and Figure 2B shows the three regions which were used for modeling the spatial differences in disaster impacts.

3.2 Socio-economic data

In this study, the population and income data of each county were used in the later analysis. Population data were obtained from the US National Census², and median income data were obtained from the American Community Survey³.

3.3 Spatial distribution of housing damages due to Hurricane Maria

The physical damage caused by the hurricane is measured by the housing damage rates in each county, which was provided through the ‘‘Housing Assistance Data’’ provided by the Federal Emergency Management Agency (FEMA). The raw data can be found through the following link⁴. We defined ‘‘housing damage rate’’ for each county as the total number of houses that were inspected to have had more than \$ 10,000 worth of damage due to the target hurricane, which is divided by the number of households in that county. Many of the counties in Puerto Rico experienced high housing damage rates of between 20% and 60%.

²<https://www.census.gov/>

³<https://www.census.gov/programs-surveys/acs>

⁴<https://www.fema.gov/media-library/assets/documents/34758>

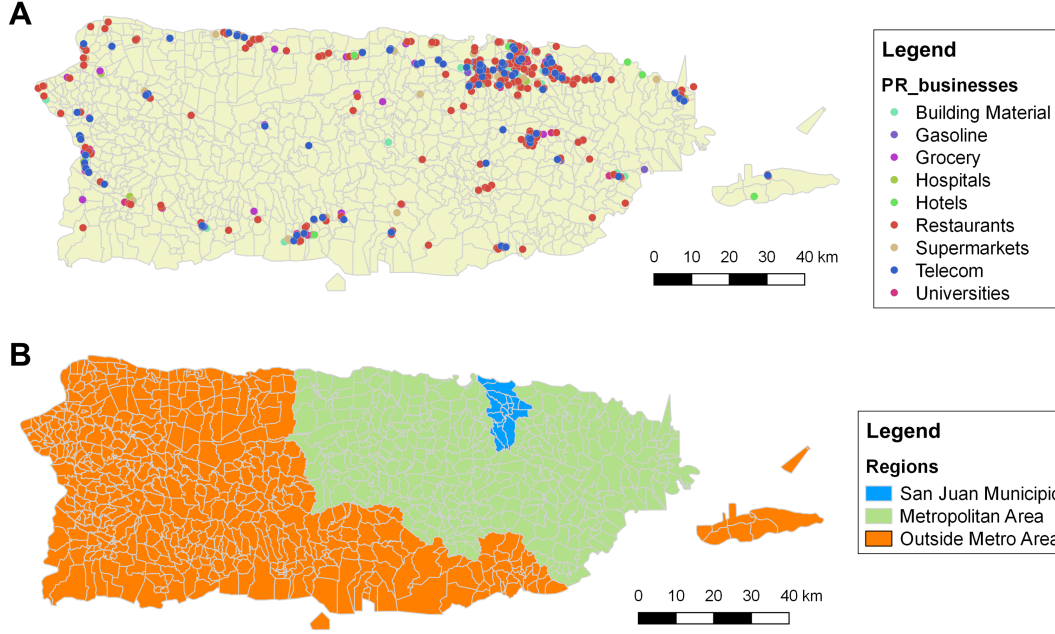


Figure 2: **Characteristics of businesses in Puerto Rico.** (A) Business locations and categories in Puerto Rico. (B) 3 regions of Puerto Rico used in this study.

4 Methods

4.1 Bayesian structural time series model

The basic structural time series model is defined as follows:

$$\begin{cases} y_{t,i} = \mu_{t,i} + \tau_{t,i} + \beta x_{t,i} + \epsilon_{t,i} \quad \forall t \\ \epsilon_{t,i} \sim \mathcal{N}(0, \sigma_y^2) \\ \sigma_y \sim \text{Cauchy}(0, 2.5) \end{cases} \quad (1)$$

where $y_{t,i}$ is the observed daily visits to business i on day t in the target region (in our case, Puerto Rico). $y_{t,i}$ is predicted by state components $\mu_{t,i}$, $\tau_{t,i}$ and $\beta x_{t,i}$ that capture critical features of the time-series data [23]. A weakly informative prior is elicited for each state component. A graphical representation of the model is shown in Figure 3.

Local Level Trend: The local level model represents local variations of the time series data. To simplify the model structure, we assume that the mean of the trend is a random walk with the initialization of μ_1 :

$$\begin{cases} \mu_{t+1,i} = \mu_{t,i} + \eta_{1,t,i} \quad \forall t > 1 \\ \mu_{1,i} \sim \mathcal{N}(\mu_0, \sigma_0^2) \\ \eta_{1,t,i} \sim \mathcal{N}(0, \sigma_\mu^2) \\ \sigma_0, \mu_0, \sigma_\mu \sim \text{Cauchy}(0, 2.5) \end{cases} \quad (2)$$

Seasonality: Let S denote the total number of seasons. The sum of seasonal effects over S time periods is assumed to be zero. In this study, weekly seasonality is taken into account ($S = 7$) with the initialization of $\tau_{1,i}, \tau_{2,i}, \tau_{3,i}, \tau_{4,i}, \tau_{5,i}$, and $\tau_{6,i}$:

$$\begin{cases} \tau_{t+1,i} = -\sum_{s=0}^{S-2} \tau_{t-s,i} + \eta_{2,t} \quad \forall t > 1 \\ \tau_{1,i}, \tau_{2,i}, \tau_{3,i}, \tau_{4,i}, \tau_{5,i}, \tau_{6,i} \sim \mathcal{N}(\mu_{\tau_0}, \sigma_{\tau_0}^2) \\ \eta_{2,t} \sim \mathcal{N}(0, \sigma_\tau^2) \\ \mu_{\tau_0}, \sigma_{\tau_0}, \sigma_\tau \sim \text{Cauchy}(0, 2.5) \end{cases} \quad (3)$$

4.1.1 Choice of Covariates

Apart from the local level model and seasonality, there are other unobserved effects (e.g., impacts of holidays and sport events) that may contaminate the estimation of the $y_{t,i}$. To capture the unobserved heterogeneity, $x_{t,i}$ in Equation (1) is

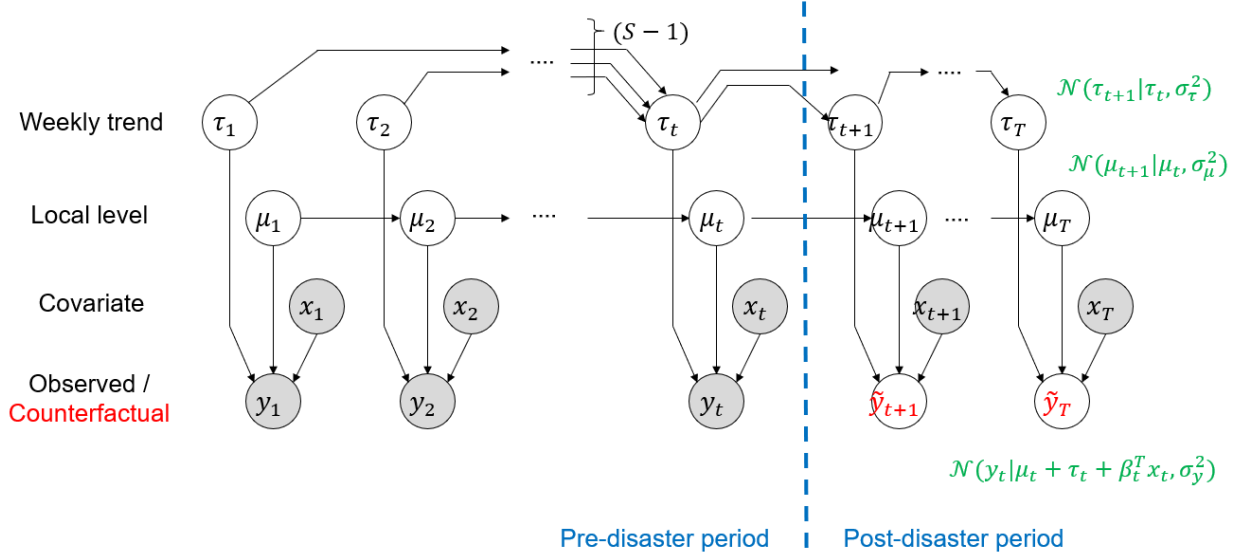


Figure 3: Graphical representation of the Bayesian structural time series model.

used as the simultaneous daily visits to a similar business type at time t in a different region that was not affected by the disaster (in our case, New York). $x_{t,i}$ accounts for the shared variance of the time series data from two different regions. The static coefficient β represents the relationship between daily visits to a specific business type from Puerto Rico and New York. In this study, we test three methods for the choice of covariates, which we will test in the experiments Section: (i) no covariate, (ii) use the average daily visit trends of the same brand businesses in the other city as covariate (e.g. if y_i was a Starbucks, we would use the average daily visit counts of all Starbucks in New York as the covariate), which we denote as $x_{category}$, and (iii) use the daily visit count of a specific business that has the highest correlation with the target business, which we denote as $x_{specific}$. For (iii), we compute the Pearson's correlation between the daily visit count data of the target business with that of all same category businesses in New York, and we use the business with the highest Pearson R.

4.2 Estimating causal impact of disasters on businesses

Let N denote the total number of days observed. We first fit the BSTS model with pre-disaster data ($n = 150$) from New York and Puerto Rico. For each business with index i , posterior predictive samples can be simulated to develop a counterfactual as the synthetic control group ($t = n + 1, \dots, N$) from Equation (4).

$$y_{t,i} \sim p(y_{t,i} | y_{t,i}) \quad t \geq n \quad (4)$$

Let $m \in [n, N]$ denote the the day when Hurricane Maria struck Puerto Rico. Point-wise comparisons estimate the impacts of hurricane on daily visits to a target business type between treatment and control groups.

$$\phi_{t,i} = \frac{y_{t,i} - \hat{y}_{t,i}}{\bar{y}_i} \quad t = m + 1, \dots, N \quad (5)$$

where, \bar{y}_i denotes the mean visit count to the visits prior to the disaster ($t < n$). The impact $\phi_{t,i}$ is a normalized measure of the disaster impact to the business. $\phi_{t,i}$ measures the number of business-as-usual days worth of impact (damage) that the disaster inflicted on the business.

Moreover, we hope to estimate the cumulative causal effects of the hurricane on a target business type over time, which represents the resilience of a business after a hurricane. The cumulative sum of causal increments is a practical quantity when the response variable $y_{t,i}$ is measured over time. We calculate the total impact of the disaster to business i using the following equation:

$$\phi_i = \sum_{t=m}^N \phi_{t,i} \quad (6)$$

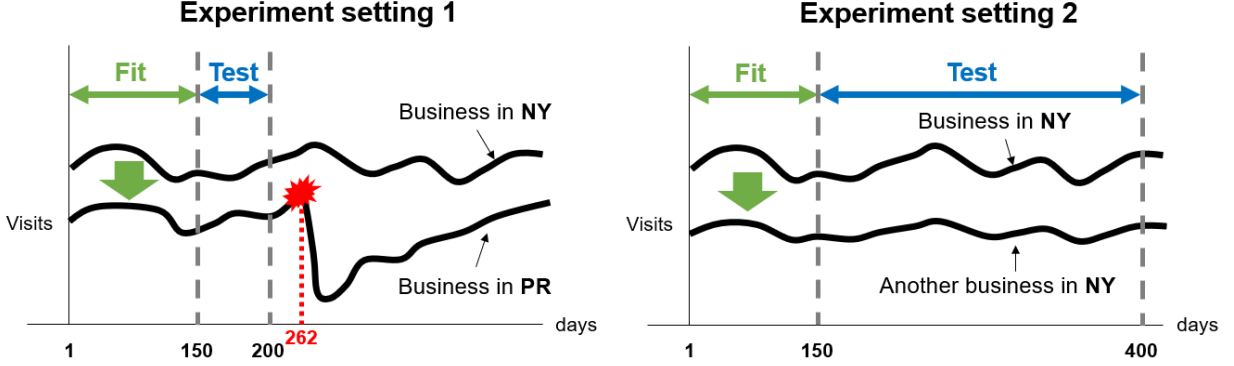


Figure 4: Experiment settings for model validation.

The cumulative sum of causal increments can be further transformed into the estimated total economic loss by multiplying average spending in dollar(s) per customer.

4.3 Regression via hierarchical Bayesian model

To further interpret the estimated cumulative disaster impacts, we apply a hierarchical Bayesian model to regress the cumulative disaster impacts ϕ_i with regional (spatial), business categorical, and disaster damage features. Hierarchical Bayesian models (HBMs) allow us to flexibly model the group-level effects on the estimand by introducing hyper-prior distributions on the model parameters. This is a significant difference from regular linear regression models, which can only either i) assign one global parameter for all groups, or ii) estimate parameters separately for each group. For further details on the advantages of HBMs, readers should refer to [43].

To estimate the cumulative disaster impact of all businesses, we construct the HBM as follows:

$$\begin{cases} \phi_i \sim N(\beta X_i + \delta_{r(i)} + \gamma_{c(i)}, \sigma^2) \\ \delta_{r(i)} \sim N(0, \tau_\delta^2), \quad \forall r \in \{0, 1, 2\} \quad \#region \\ \gamma_{c(i)} \sim N(0, \tau_\gamma^2), \quad \forall c \in \{0, 1, 2, 3, 4, 5, 6, 7, 8\} \quad \#category \\ \beta, \sigma, \tau_\delta, \tau_\gamma \sim Cauchy(0, 2.5) \end{cases} \quad (7)$$

where, $r(i) \in \{0, 1, 2\}$ and $c(i) \in \{0, 1, 2, 3, 4, 5, 6, 7, 8\}$ denote the region index and category index for business i . We assume that the cumulative disaster impact on business i , denoted by ϕ_i , can be modeled as a linear sum of the effects of exogenous features X_i (which include pre-disaster business mean visits, housing damages caused by the disaster), regional effects δ_r , and business categorical effects γ_c . The model is Bayesian in the sense that the model parameters (β, δ, γ) all have priors. The model is also hierarchical because the hyper-parameters in the prior distributions $(\tau_\delta, \tau_\gamma)$ come from another higher level distribution. We assume that the hyper-parameters are drawn from weakly informative priors (Cauchy distribution). The hierarchical prior distributions allow us to model the dependencies across different groups (i.e., regional groups and categorical groups).

5 Model Validation

5.1 Experiment Setup

We analyze daily visits to businesses in Puerto Rico and New York from January 2017 to March 2018 (400 days). As explained in the methods section, we will test three methods of selecting the covariate: no covariate, $x_{category}$, and $x_{specific}$. To verify which type of covariate improves the prediction accuracy the most, two different model settings (as shown in Figure 4) will be explored:

- Setting 1 (Inter-State prediction): Pre-disaster data will be used from Puerto Rico and New York. The model will be fitted using data until day 150, and tested using data between days 151 and 200.
- Setting 2 (Intra-State prediction): To test the accuracy of long-term predictions, data from businesses in Manhattan will be used to predict the visit counts of businesses in Up-State New York, using the whole observation period (train: 0-150, test: 151-400).

Table 2: Model validation results of two experimental settings.

		Evaluation Metric	Use of Covariates		
			No Covariates	$x_{category}$	$x_{specific}$
Setting 1	Train	MAPE	12.35 (± 16.67)	10.50 (± 14.03)	10.66 (± 14.46)
		Pearson R	0.539 (± 0.222)	0.696 (± 0.169)	0.626 (± 0.136)
	Test	MAPE	8.568 (± 14.37)	8.518 (± 15.85)	8.888 (± 15.30)
		Pearson R	0.351 (± 0.238)	0.354 (± 0.239)	0.295 (± 0.257)
	Selected (%)		34.9	40.4	24.7
Setting 2	Train	MAPE	0.229 (± 0.257)	0.249 (± 0.251)	0.257 (± 0.252)
		Pearson R	0.855 (± 0.144)	0.742 (± 0.145)	0.744 (± 0.115)
	Test	MAPE	0.704 (± 0.811)	0.475 (± 0.612)	0.477 (± 0.538)
		Pearson R	0.420 (± 0.189)	0.512 (± 0.181)	0.466 (± 0.183)
	Selected (%)		40.3	25.1	34.6

5.2 Evaluation Metrics

The prediction tasks will be evaluated using two different metrics: i) Pearson’s R, which captures the correlation between the predicted and true time series values; and ii) mean absolute percentage error (MAPE), which captures the relative magnitude of the absolute error between the predicted and true time series values. MAPE is calculated by the following equation:

$$MAPE_i = \frac{1}{n} \sum_{t=1}^n \left| \frac{y_{t,i} - \hat{y}_{t,i}}{y_{t,i}} \right| \quad (8)$$

We measure the performance of the methods using these two distinct metrics, where Pearson’s R measures the relative correlation between the two sequences, while MAPE measures the absolute magnitude of discrepancy between the two vectors.

5.3 Validation Results

The performance of the BSTS models with three types of covariates, were tested on the two experimental settings, using Pearson’s correlation and MAPE as evaluation metrics. Table 2 shows the performance of the three BSTS models on both settings. Surprisingly, although the model with business-category covariates performs the best on average in both experimental settings, the predictive performance of the three methods are quite similar. Using extra covariates do not always improve the prediction model. We also see that over 34% of the businesses in experiment setting 1 had the best performance when not using extra covariates (similarly, over 40% of businesses in experimental setting 2). Extra covariates, which are aimed to capture the long term trends and anomalies (e.g. New Years, Christmas), are not effective when making predictions of businesses that have less long-term variation and a relatively stable periodicity in visit counts. From experiment 1, we determine the best performing model out of the three for each business. We then use that business to predict the counterfactual daily visit counts after the disaster period.

Figure 5 shows an example of how the disaster impact is quantified. As shown in panel (A), we first predict the counterfactual daily visit counts after the disaster (blue plot) using the best performing model identified in the model validation experiment. Then, as shown in (B), we calculate the point-wise disaster impact $\phi_{t,i}$ by subtracting the observed daily visit count sequence from the predicted sequence and then normalizing it by the pre-disaster mean daily visits. The cumulative disaster impact ϕ_i can be calculated by aggregating the point-wise disaster impacts over time. Panel (C) shows the cumulative disaster impact over time, starting from the time of the landfall of the hurricane. In this particular business, we observe that there was a significant negative impact until around day 300 with around $\phi_i = -25$, meaning that by day 300 this business had lost 25 business-as-usual days worth of customers due to the hurricane. We actually see positive impacts of the hurricane before the two hurricanes; however, the positive impacts are significantly negated by the negative impacts. Gradually, we see an increase in visits compared to pre-disaster levels after 1 month from the hurricane’s landfall, which decreases the negative disaster impact. As a result of the BSTS modeling, we are able to obtain the quantified disaster impact for each of the businesses in Puerto Rico over time. In the next section, we will analyze the obtained results to further understand which business categories and which locations suffered disaster impact in Puerto Rico after Hurricane Maria.




Figure 5: Example of how the disaster impact is quantified: (A) predicted and actual observed daily visit patterns for a randomly selected business; (B) point-wise impact $\phi_{t,i}$; and (C) cumulative impact ϕ_i of the disaster.

6 Analysis of Estimated Business Resilience

Using the BSTS method for predicting the counterfactual business performance, we will now quantitatively analyze the resilience of businesses after Hurricane Maria and answer the following questions:

1. How does the disaster impact evolve over time? And, do the temporal patterns vary across business categories and locations?
2. Can we explain why we observe this heterogeneity in disaster impacts across businesses in Puerto Rico?

Because it was revealed that the optimal prediction models varied across different businesses in the model validation section, we use the best performing model out of the three (either no covariate, average NY trend as covariate, or specific NY business trend as covariate) to predict the counterfactual visit time series for each of the businesses in Puerto Rico.

6.1 Quantifying Disaster Impact Patterns to Businesses

To answer the first research question, we aggregate the disaster impacts over the time horizon by business category and business location (i.e., San Juan Municipio, Metropolitan Area, Rural Area; as shown in Figure 2B).

The cumulative disaster impacts are shown in Figure 6 for three different aggregation time thresholds, as follows: (A) landfall to 1 month from landfall, (B) until 2 months from landfall, and (C) until 4 months from landfall. The numbers of ϕ_i should be interpreted as “the number of business-as-usual days worth of impact.” For example, building material businesses in San Juan experienced a median disaster impact of $\phi = -10$ during the first month. This indicates that the building material businesses in San Juan lost 10 days worth of customers who were supposed to visit if the disaster did not occur. Most of the regions and business categories experience a negative impact in the first month, except for hotels in San Juan. We also clearly observe the urban-rural disparity in disaster impacts across many of the business categories across all three temporal thresholds. However, the urban-rural gap gradually closes as time passes, and in many of the

industries we observe little difference by 4 months from landfall (e.g. building material, grocery stores, restaurants, and telecommunications).

Although the general patterns show consistent insights (e.g., the urban-rural disparity, larger impact right after the landfall, and differences in disaster impacts across business categories), we are unable to delineate the effects of each characteristic on the disaster’s impacts. In the next section, we will attempt to reveal the impacts of the business characteristics on the observed disaster impacts by applying a hierarchical Bayesian modeling technique [43].

7 Discussion

In this study, we used business visit data collected from mobile phone trajectories in Puerto Rico and New York to quantify the causal impact of Hurricane Maria on businesses in Puerto Rico. Using the Bayesian Structural Time Series (BSTS) model, we predicted the counterfactual (what if the disaster did not happen) daily visit counts to businesses in Puerto Rico. We also computed the point-wise disaster impact, as well as the cumulative disaster impact of Hurricane Maria. The performance of the BSTS models was evaluated, and we whether the covariates (information of daily visit counts of businesses not affected by the disaster) positively contributes to the prediction accuracy varied across businesses. Furthermore, the estimated disaster impacts were analyzed using hierarchical Bayesian models to understand the effects of various business characteristics on disaster impacts.

The findings in this study should be considered in the light of some limitations. First, the assumptions that we make on the data to estimate business performance have several limitations. We used daily visit data as a proxy to estimate the performance of businesses in this study. However, we note that this approximation only holds when the business is a business-to-customer (B2C) type. If the main flow of transactions of the business is with other companies, then mobile phone data would not be an appropriate data source for analysis. This is the reason why we limited the study to nine business categories, who usually have customers visit their stores to make profit. Moreover, although a previous study showed that the number of visits estimated from mobile phone data correlates well with actual business performance through the case study of a coffee chain [24], we could not fully validate this for other businesses in this study due to lack of corporate finance data. Further investigation on the relationships between business performance and customer visits after disasters would be worthy of investigation in future research. Second, although our study was able to produce more spatio-temporally granular and scalable analysis and estimations of disaster impacts on businesses compared to past studies using surveys, our study did miss some of the advantages of survey studies. Previous studies have revealed that more detailed characteristics of the businesses (e.g., years of operation, number of employees, and age of the owner) all affect the recovery performance after disasters [4]. Due to the limitations in data collection, we were unable to include such covariates in the hierarchical Bayesian model in the latter section of the analysis. Further efforts in combining different data sources (e.g. mobile phone data and survey data) to complement each other would be a very interesting research direction. Third, from a methodological point of view, one could apply a more complex method to select and generate covariates for predicting future daily visit counts. In this study, we applied a heuristic approach in choosing the covariates for prediction, where the business that was most highly correlated with the target business daily counts was selected. Empirical validation showed that, in some cases, using the covariate decreased the prediction accuracy. Efficient algorithms to detect and select appropriate covariates for future time series prediction would be of future research interest. Finally, our analysis was limited to a 160 day period after the hurricane, thus we were unable to test whether the method could perform well for longer time periods. Expanding the time horizon using additional data could yield more robust insights and implications of the model.

We finally discuss how the methods, analysis, and findings presented in this study may be applied in disaster management and policy making. As mentioned in the introduction, surveys have been the primary data source to estimate economic losses after disasters for policy makers. However, large scale mobility datasets have been more common in the decision-making processes in various domains, including epidemic control, traffic management and disaster relief. This study lays out an example of how such large scale mobility data can be used for i) post-disaster assessment and monitoring, ii) economic cost estimation, and iii) developing relief supply allocation strategies. As shown in Figures 5 and 6, we are able to quantitatively monitor the negative impact and recovery of each business using the models in this paper. It is not technically difficult to detect businesses who are struggling to recover after a disaster, and then carry out assistance programs for those businesses. Moreover, the estimated point-wise or cumulative disaster impacts can be multiplied with the average money spent per customer to easily calculate the daily or total economic loss for each business. We can also identify the business categories that have not recovered in each region to develop strategies for allocating relief supplies (e.g., for distributing gasoline across the island).

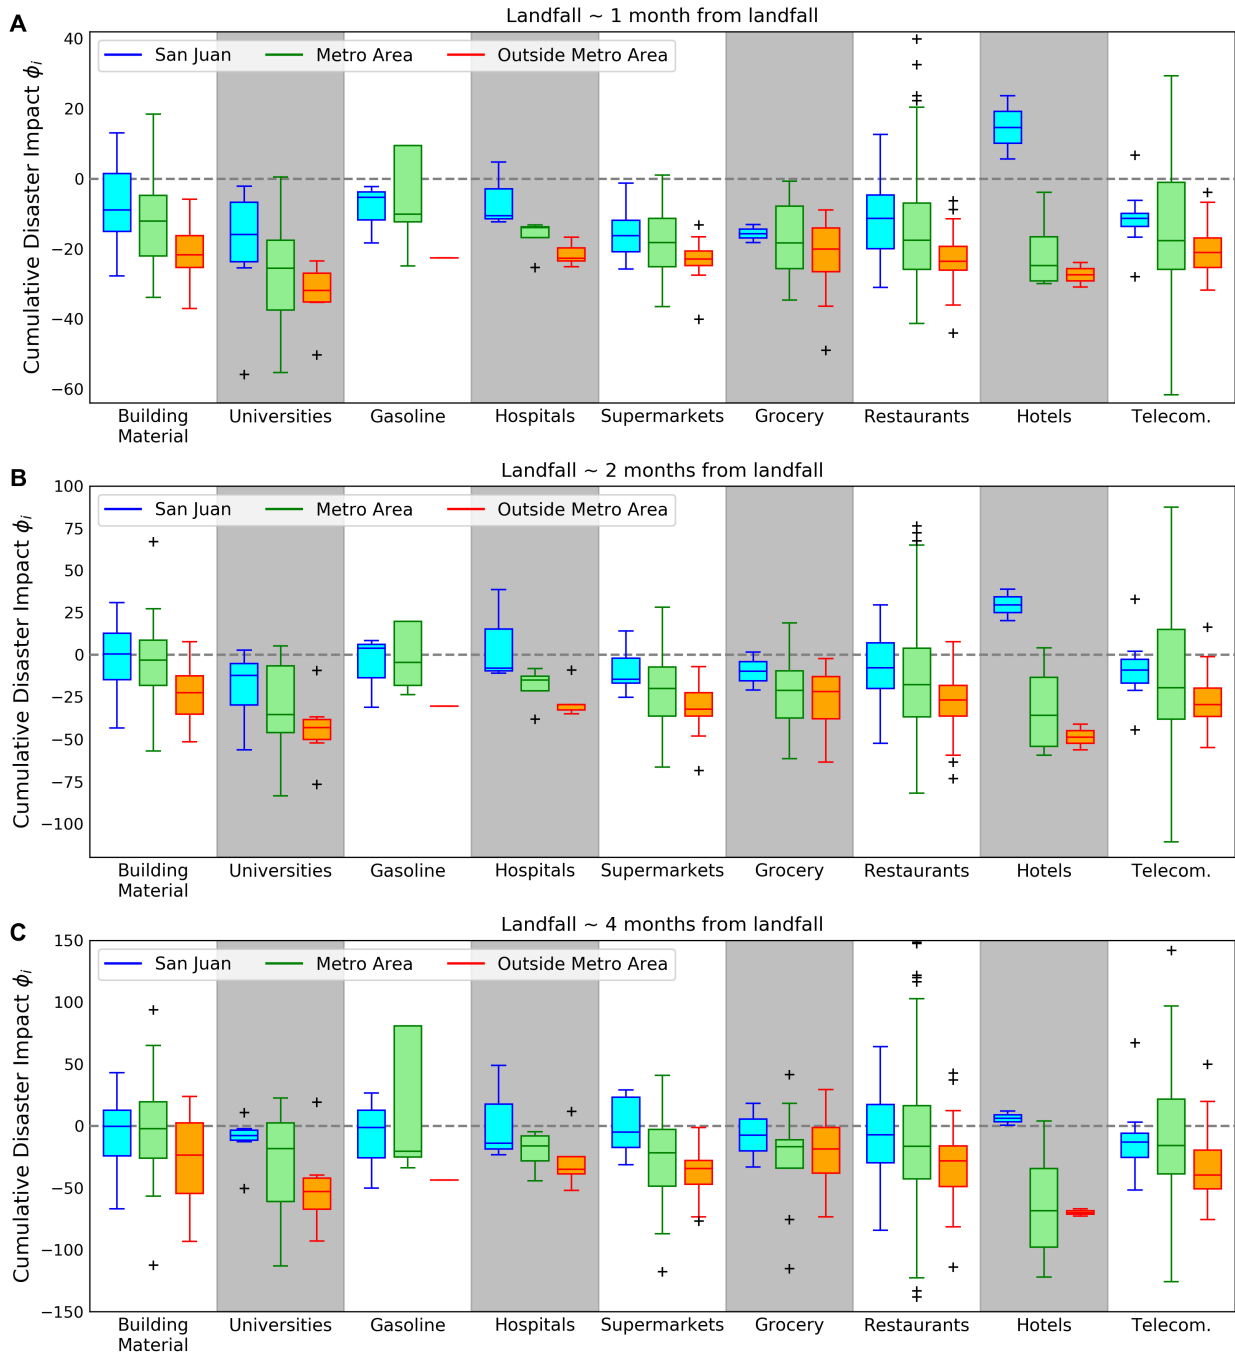


Figure 6: Cumulative disaster impacts across different business categories and regions. Results are shown for different aggregation time thresholds, as follows: (A) landfall to 1 month from landfall, (B) until 2 months from landfall, and (C) until 4 months from landfall.

8 Conclusion

Quantifying the economic impact of disasters to businesses is crucial for disaster relief and preparation. The availability of large scale human mobility data enables us to observe daily visit counts to businesses in an unprecedented spatio-temporal granularity. In this work, we presented a methodology to estimate the causal impact of disasters to businesses from mobile phone location data, using a Bayesian modeling framework. We used this methodology to quantify the causal impact of Hurricane Maria on businesses in Puerto Rico. The estimation results provide insights into what types of businesses (and their locations) are able to recovery quickly after the hurricane. These insights could assist policy makers during disaster preparation and relief processes.

References

- [1] UNDRR. Global assessment report on disaster risk reduction 2013: From shared risk to shared value: the business case for disaster risk reduction, 2013.
- [2] Hallie Eakin, Luis A Bojórquez-Tapia, Marco A Janssen, Matei Georgescu, David Manuel-Navarrete, Enrique R Vivoni, Ana E Escalante, Andres Baeza-Castro, M Mazari-Hiriart, and Amy M Lerner. Opinion: urban resilience efforts must consider social and political forces. *Proceedings of the National Academy of Sciences*, 114(2):186–189, 2017.
- [3] Maria I Marshall, Linda S Niehm, Sandra B Sydnor, and Holly L Schrank. Predicting small business demise after a natural disaster: an analysis of pre-existing conditions. *Natural Hazards*, 79(1):331–354, 2015.
- [4] Sandra Sydnor, Linda Niehm, Yoon Lee, Maria Marshall, and Holly Schrank. Analysis of post-disaster damage and disruptive impacts on the operating status of small businesses after hurricane katrina. *Natural Hazards*, 85(3):1637–1663, 2017.
- [5] Gary R Webb, Kathleen J Tierney, and James M Dahlhamer. Predicting long-term business recovery from disaster: A comparison of the loma prieta earthquake and hurricane andrew. *Global Environmental Change Part B: Environmental Hazards*, 4(2):45–58, 2002.
- [6] Jim Lee. Business recovery from hurricane harvey. *International Journal of Disaster Risk Reduction*, 34:305–315, 2019.
- [7] Michael Batty. Big data, smart cities and city planning. *Dialogues in Human Geography*, 3(3):274–279, 2013.
- [8] Marta C Gonzalez, Cesar A Hidalgo, and Albert-Laszlo Barabasi. Understanding individual human mobility patterns. *nature*, 453(7196):779, 2008.
- [9] Vincent D Blondel, Adeline Decuyper, and Gautier Krings. A survey of results on mobile phone datasets analysis. *EPJ data science*, 4(1):10, 2015.
- [10] Pierre Deville, Catherine Linard, Samuel Martin, Marius Gilbert, Forrest R Stevens, Andrea E Gaughan, Vincent D Blondel, and Andrew J Tatem. Dynamic population mapping using mobile phone data. *Proceedings of the National Academy of Sciences*, 111(45):15888–15893, 2014.
- [11] NA Wardrop, WC Jochem, TJ Bird, HR Chamberlain, D Clarke, D Kerr, L Bengtsson, S Juran, V Seaman, and AJ Tatem. Spatially disaggregated population estimates in the absence of national population and housing census data. *Proceedings of the National Academy of Sciences*, 115(14):3529–3537, 2018.
- [12] Md Shahadat Iqbal, Charisma F Choudhury, Pu Wang, and Marta C González. Development of origin–destination matrices using mobile phone call data. *Transportation Research Part C: Emerging Technologies*, 40:63–74, 2014.
- [13] Francesco Calabrese, Giusy Di Lorenzo, Liang Liu, and Carlo Ratti. Estimating origin-destination flows using opportunistically collected mobile phone location data from one million users in boston metropolitan area. *IEEE Pervasive Computing*, 10(4):36–44, 2011.
- [14] Joshua Blumenstock, Gabriel Cadamuro, and Robert On. Predicting poverty and wealth from mobile phone metadata. *Science*, 350(6264):1073–1076, 2015.
- [15] Linus Bengtsson, Jean Gaudart, Xin Lu, Sandra Moore, Erik Wetter, Kankoe Sallah, Stanislas Rebaudet, and Renaud Piarroux. Using mobile phone data to predict the spatial spread of cholera. *Scientific Reports*, 5:8923, 2015.
- [16] Xin Lu, Linus Bengtsson, and Petter Holme. Predictability of population displacement after the 2010 haiti earthquake. *Proceedings of the National Academy of Sciences*, 109(29):11576–11581, 2012.

- [17] Xuan Song, Quanshi Zhang, Yoshihide Sekimoto, and Ryosuke Shibasaki. Prediction of human emergency behavior and their mobility following large-scale disaster. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 5–14. ACM, 2014.
- [18] Robin Wilson, Elisabeth zu Erbach-Schoenberg, Maximilian Albert, Daniel Power, Simon Tudge, Miguel Gonzalez, Sam Guthrie, Heather Chamberlain, Christopher Brooks, Christopher Hughes, et al. Rapid and near real-time assessments of population displacement using mobile phone data following disasters: the 2015 nepal earthquake. *PLoS currents*, 8, 2016.
- [19] Xin Lu, David J Wrathall, Pål Roe Sundsøy, Md Nadiruzzaman, Erik Wetter, Asif Iqbal, Taimur Qureshi, Andrew Tatem, Geoffrey Canright, Kenth Engø-Monsen, et al. Unveiling hidden migration and mobility patterns in climate stressed regions: a longitudinal study of six million anonymous mobile phone users in bangladesh. *Global Environmental Change*, 38:1–7, 2016.
- [20] James P Bagrow, Dashun Wang, and Albert-Laszlo Barabasi. Collective response of human populations to large-scale emergencies. *PLoS one*, 6(3):e17680, 2011.
- [21] Andrew Harvey and SJ Koopman. Structural time series models. *Wiley StatsRef: Statistics Reference Online*, 2014.
- [22] Richard Ernest Bellman and Kenneth L Cooke. Differential-difference equations. 1963.
- [23] Kay H Brodersen, Fabian Gallusser, Jim Koehler, Nicolas Remy, Steven L Scott, et al. Inferring causal impact using bayesian structural time-series models. *The Annals of Applied Statistics*, 9(1):247–274, 2015.
- [24] Umit Gurun and David Solomon. The perils of private provision of public goods, 2019.
- [25] Andrew Goodman-Bacon. Difference-in-differences with variation in treatment timing. Technical report, National Bureau of Economic Research, 2018.
- [26] Alberto Abadie. Semiparametric difference-in-differences estimators. *The Review of Economic Studies*, 72(1):1–19, 2005.
- [27] John Antonakis, Samuel Bendahan, Philippe Jacquart, and Rafael Lalive. On making causal claims: a review and recommendations. *The leadership quarterly*, 21(6):1086–1120, 2010.
- [28] Christian B Hansen. Asymptotic properties of a robust variance matrix estimator for panel data when t is large. *Journal of Econometrics*, 141(2):597–620, 2007.
- [29] Andrew C Harvey. *Forecasting, structural time series models and the Kalman filter*. Cambridge university press, 1990.
- [30] Peter SH Leeflang, Tammo HA Bijmolt, Jenny Van Doorn, Dominique M Hanssens, Harald J Van Heerde, Peter C Verhoef, and Jaap E Wieringa. Creating lift versus building the base: current trends in marketing dynamics. *International Journal of Research in Marketing*, 26(1):13–20, 2009.
- [31] Steven L Scott and Hal R Varian. Predicting the present with bayesian structural time series. *Available at SSRN 2304426*, 2013.
- [32] Steven L Scott and Hal R Varian. Bayesian variable selection for nowcasting economic time series. Technical report, National Bureau of Economic Research, 2013.
- [33] Jouchi Nakajima and Mike West. Bayesian analysis of latent threshold dynamic models. *Journal of Business & Economic Statistics*, 31(2):151–164, 2013.
- [34] Guido W Imbens and Donald B Rubin. Rubin causal model. *The New Palgrave Dictionary of Economics*, pages 1–10, 2017.
- [35] Andrew C Harvey, Thomas M Trimbur, and Herman K Van Dijk. Trends and cycles in economic time series: a bayesian approach. *Journal of Econometrics*, 140(2):618–649, 2007.
- [36] Obryan Poyser. Exploring the dynamics of bitcoin’s price: a bayesian structural time series approach. *Eurasian Economic Review*, 9(1):29–60, 2019.
- [37] Patrick T Brandt and John R Freeman. Advances in bayesian time series modeling and the study of politics: theory testing, forecasting, and policy analysis. *Political Analysis*, 14(1):1–36, 2006.
- [38] Patrick T Brandt, John R Freeman, and Philip A Schrodt. Real time, time series forecasting of inter-and intra-state political conflict. *Conflict Management and Peace Science*, 28(1):41–64, 2011.
- [39] Yu Jiang, Zhe Song, and Andrew Kusiak. Very short-term wind speed forecasting with bayesian structural break model. *Renewable Energy*, 50:637–647, 2013.

- [40] Frank de Vocht. Inferring the 1985–2014 impact of mobile phone use on selected brain cancer subtypes using bayesian structural time series and synthetic controls. *Environment international*, 97:100–107, 2016.
- [41] Nishant Kishore, Domingo Marqués, Ayesha Mahmud, Mathew V Kiang, Irmay Rodriguez, Arlan Fuller, Peggy Ebner, Cecilia Sorensen, Fabio Racy, Jay Lemery, et al. Mortality in puerto rico after hurricane maria. *New England journal of medicine*, 2018.
- [42] Richard J Pasch, Andrew B Penny, and Robbie Berg. National hurricane center tropical cyclone report: Hurricane maria. *TROPICAL CYCLONE REPORT AL152017*, National Oceanic And Atmospheric Administration and the National Weather Service, pages 1–48, 2018.
- [43] Andrew Gelman, John B Carlin, Hal S Stern, David B Dunson, Aki Vehtari, and Donald B Rubin. *Bayesian data analysis*. Chapman and Hall/CRC, 2013.