Causal Recreation Demand Estimation with Cellphone Mobility Data: Evidence from the 2021 Huntington Beach Oil Spill

Nieyan Cheng* Xibo Wan†‡

January 31, 2026

Abstract

Accurately estimating the welfare impacts of environmental changes in recreation demand modeling requires robust causal inference methods. However, a persistent challenge remains as the zero market share issue restricts the causal inference application to broader regions and longer periods. To address this, we integrate empirical Bayes posterior mean estimation into a two-step random coefficient logit model, ensuring that sites with low or zero visitation are properly incorporated without distorting demand estimates. We apply this framework to the 2021 Huntington Beach oil spill, using high-frequency cellphone data to track changes in beach visits. By combining the Synthetic Difference-in-Differences (SDID) approach with empirical Bayes-adjusted market shares, we examine the causal effects of temporary beach closures on visitor welfare. Our findings reveal that Huntington Beach experienced the largest and most prolonged welfare loss, with an estimated aggregate loss of $1.0 million and weekly losses of $83 thousand persisting beyond the initial closure. In contrast, Newport Beach and Laguna Beach exhibited faster recoveries. This study advances recreation demand modeling by refining demand estimation for low-visit sites and strengthening causal inference techniques for environmental disruptions, ultimately providing a more reliable framework for assessing the economic costs of beach closures and other environmental policy interventions.

Keywords: Cell Phone Mobility Data, Recreation Demand, Huntington Oil Spill, Causal Inference, Zero Market Shares

JEL Codes: Q53, Q26

* Nieyan Cheng: Assistant Professor, Department of Economics and Management, China University of Petroleum (Beijing), cny@cup.edu.cn.

† Xibo Wan: Postdoctoral Associate, Department of Agricultural and Resources Economics, University of Connecticut, xibo.wan@uconn.edu.

‡ Corresponding Author. The views expressed herein are those of the authors and do not necessarily reflect the views of Spectus. We gratefully acknowledge valuable feedback from Kenneth Train, Yingkang Xie, and participants of the 2025 CES conference. This research was supported by the Illinois-Indiana Sea Grant, Award Number NA22OAR4170654-T1-01.

doi:10.3368/le.102.3.013026-0020

1. Introduction

Recreation demand models are essential for valuing environmental amenities and quantifying welfare losses from pollution events or policy changes. Recent advances in revealed-preference modeling increasingly rely on large-scale mobility data, which provide unprecedented spatial and temporal details on outdoor recreation. Yet these data also introduce a pervasive problem well-known in the industrial organization literature—the occurrence of zero market shares. As shown by Berry et al. (2004), Dubé et al. (2021), Li (2019), and Gandhi et al. (2023), sparse demand can bias structural estimates and distort welfare calculations. This challenge is especially relevant in recreation settings, where many site-week combinations record no observed visits despite being available choices.

In cellphone-based recreation data, zeros often arise because the observed sample represents only a fraction of the true population, and many potential census block group (CBG)-site-week combinations exhibit zero counts due to sampling error and behavioral sparsity. Common fixes such as dropping zeros, aggregating markets, or adding small constants alter the effective choice set and can bias elasticity and welfare estimates. Building on Li (2019), we adapt the empirical Bayes (EB) correction to the context of recreation demand, transforming observed visitation shares into posterior mean estimates that remain strictly positive and theoretically consistent while preserving the full choice set. Building on this correction, we establish a unified structural framework that links demand estimation with causal inference: the EB-adjusted shares feed into a random-coefficients logit model to recover mean utilities, which then serve as the foundation for evaluating environmental shocks through a panel-data causal design. This framework bridges the gap between structural recreation demand modeling and causal identification of policy impacts.

Our empirical analysis proceeds in three stages. In the first step, we apply EB estimation to adjust zero visit shares, ensuring that sites with no observed visits are meaningfully incorporated into the analysis. Next, we use a random coefficient logit model to recover site-week mean utilities through a contraction mapping and generalized method of moments (GMM) method, incorporating the EB-adjusted market shares. In the last step, we estimate counterfactual mean utilities under alternative environmental policy conditions using the Synthetic Difference-in-Differences (SDID) approach. By comparing expected utility under observed and counterfactual scenarios, we quantify the welfare impacts of beach closures while ensuring that all sites, including those with low or zero visitation, are fully incorporated into demand estimation.

We apply this framework to the 2021 Huntington Beach crude oil spill, which resulted in temporary beach closures and public health advisories. Using anonymized cellphone mobility data from Spectus, we track weekly trips from over 200,000 census block groups to Southern California beaches before, during, and after the spill. The estimated model quantifies welfare changes by comparing observed utilities during the closure with counterfactual utilities under a no-spill scenario. We find that the oil spill led to substantial welfare losses across affected beaches, reflecting significant disruptions to recreational activity and visitor welfare.

Furthermore, our results indicate that the oil spill caused significant short-term declines in recreational demand, with varying impacts across beach clusters. Huntington Beach experienced the largest and most prolonged decline in mean utility, with sustained welfare losses even several months after the closure. In contrast, Newport Beach and Laguna Beach saw nearly no declines, suggesting that the enjoyment visitors obtained in these sites was less decreased compared to that from Huntington Beach. Using the SDID method, we estimate an aggregate welfare loss of $1.00 million for Huntington Beach, with a weekly loss of $80 thousand persisting beyond the initial closure. To ensure the robustness of our findings, we conduct several checks, including adjusting Bayesian posterior estimation with alternative prior sizes, replacing the second-stage technique with a traditional event study approach, conducting several placebo tests, varying car-trip thresholds, and examining the sensitivity of SDID weighting. Our findings highlight the substantial economic costs of temporary beach closures and the varying ability of visitors to adapt to environmental disruptions.

Our work makes three key contributions. First, we contribute to the literature on the recreational demand model and estimation methods (Tourangeau et al., 2017; Alvarez et al., 2014; Whitehead et al., 2018; English et al., 2018) by addressing zero market share issues and integrating a panel data approach within a travel cost random utility model. Leveraging cellphone mobility data, which offers rich spatial-temporal variation at high frequencies (Earle and Kim, 2024; Newbold et al., 2022; Kubo et al., 2020), we apply a contraction mapping technique (Berry, 1994; Murdock, 2006; Timmins and Murdock, 2007) and combine it with panel data causal inference methods following Earle and Kim (2024). Unlike traditional approaches that aggregate data, add small constants, or drop observations methods that can reduce precision or introduce bias (Li, 2019), we employ empirical Bayes posterior mean estimation to generate strictly positive posterior estimates while preserving the integrity of the discrete-choice framework. We further validate the performance of this correction through a Monte Carlo simulation, demonstrating that the EB estimator consistently recovers unbiased travel-cost parameters and stable mean utilities even under severe data sparsity. The integration of Bayes share correction, contraction mapping, and panel-data causal inference enhances welfare estimation by incorporating all sites and controlling for unobserved heterogeneity, providing a robust foundation for analyzing the economic impacts of environmental disruptions.

Our study enriches the literature on the negative impacts of oil spills, which have been extensively examined across various dimensions, including human health (Morris Jr et al., 2013; Meo et al., 2009; D'Andrea and Reddy, 2014), ecological damage (Jackson et al., 1989; Barron et al., 2020; Murphy and Delucchi, 2005), and housing market (Cano-Urbina et al., 2019; Winkler and Gordon, 2013; Grossman et al., 2019; Cheng et al., 2024; Herrnstadt and Sweeney, 2024). In addition to these areas, prior research has investigated the effects of oil spills on recreational activities, particularly following the 2010 BP Gulf oil spill and other major incidents. Studies by Tourangeau et al. (2017), Alvarez et al. (2014), and Whitehead et al. (2018) utilize methods such as aerial photography, interviews, and travel cost models to quantify declines in beach visits, boating, and fishing, highlighting significant welfare losses. Further refinements in these estimates have been made by Glasgow and Train (2018), English et al. (2018), and Lopes and Whitehead (2023), while Egan et al. (2024) extend the analysis by assessing potential recreational value losses under hypothetical spill scenarios in Fiordland National Park. Building on prior research, our paper contributes to the literature by providing a detailed assessment of the 2021 Huntington Beach oil spill, leveraging high-frequency datasets and advanced estimation methods to capture its impact on beach visits with greater precision. By refining welfare loss estimates specific to oil spills, our study deepens the understanding of how these environmental disasters disrupt recreation and provides policymakers with critical insights to improve response strategies and mitigation efforts for future spills.

Third, this study contributes to the use of cellphone mobility data in non-market valuation, a method increasingly applied in areas such as outdoor recreation (Zhang et al., 2024; Merrill et al., 2020; Kubo et al., 2020; Lu et al., 2023), water quality measurement (Newbold et al., 2022; Knittel et al., 2023), and flooding (Lee et al., 2023). Mobility data offer high-frequency, large-scale observations of actual visitation patterns, enabling precise estimation of short-term events like oil spill-induced beach closures by capturing temporal variations. This contrasts with traditional surveys (Carson, 2003; Loureiro, 2009), which, while useful, often rely on cross-sectional data. By leveraging the panel structure of cellphone data and applying causal inference techniques, our paper provides a more accurate estimation of recreation demand shifts and welfare losses resulting from oil spills.

The remainder of this paper is structured as follows. Section 2 introduces the modeling framework, including the integration of EB estimation, the random coefficient logit model, and the SDID approach. Section 3 describes the study area and the 2021 Huntington Beach oil spill event. Section 4 details the data sources and preparation, and key variables. Section 5 presents the demand estimation, welfare analysis results, and robustness checks. Section 6 concludes with findings and policy implications.

2. Model

We specify a model of recreator behavior to estimate the welfare effects of beach closures. We use a discrete-choice model following the framework of Berry (1994). Each week, visitors make a decision to visit one of the beaches in their choice sets or any other activities that are not defined as beach activities. The choice set is defined using a 300-mile straight-line distance as a car trip cutoff following Dundas and

Von Haefen (2020). A market is defined as a CBG by week combination.[1] The market size is the number of residing devices times 7, assuming people make 7 recreational choices over a week. We assume the demand model is static in that visitors choose myopically, without taking into account the future evolution of travel costs and other site characteristics. More specifically, the utility function of visitor from CBG c to site j in week t is as follows:

$$U_{icjt} = \delta_{cjt} + \mu_{icjt} + \epsilon_{icjt}, \tag{1}$$

where δ_{cjt} represents the mean utility of visiting site j in week t for visitors from CBG c, and μ_{icjt} denotes the individual-specific preference deviation from the mean utility. ϵ_{icjt} follows i.i.d. Type-I extreme value distribution. The mean utility δ_{cjt} is defined as:

$$\delta_{cjt} = \alpha p_{cjt} + \gamma_c + \zeta_{jt} + \xi_{cjt}, \tag{2}$$

where p_{cjt} denotes the travel cost between the centroid of CBG c and recreation site j in week t. We also include site-by-week fixed effects ζjt to capture time-varying unobserved site attributes and CBG fixed effects γ_c to absorb unobserved demographic variables that influence the demand for outdoor recreation.

The individual deviation from the mean utility is defined as

$$\mu_{icjt} = \sigma_{tc} p_{cjt} \upsilon_i^p + \sigma_{cons} \upsilon_i^{cons}, \tag{3}$$

where υ_i^p and υ_i^{cons} are standard normal draws. Each visitor i makes a discrete choice and chooses the site j in week t with the highest utility, U_{icjt}. Specifically, each trip is treated as a separate choice decision such that visitor i may make up to seven independent site choices in week t. The utility function U_{icjt} represents the utility from one such trip, and aggregating across trips reconciles the market size assumption with the choice framework. The discrete choice model allows opting out, which introduces an "outside good" into the model. The outside good in our case is all other activities that do not qualify as beach recreational activities. The mean utility of the outside good is normalized to $U_{ic0} = \epsilon_{ic0}$.

Assuming that the random utility terms follow the extreme value distributional assumption, the probability that a utility-maximizing visitor i will choose site $j = 1 \ldots J$ in week t takes the following form:

$$share_{cjt} = \int \frac{\exp\left(\delta_{cjt} + \mu_{icjt}\right)}{\sum_k \exp(\delta_{ckt} + \mu_{ickt})} dv\,, \tag{4}$$

We classify the estimation procedure into three steps: 1) dealing with the zero market share issue before diving into the random utility model; 2) utilizing the random utility model to calculate the mean utility and

[1] In contrast to standard IO applications that define markets by geographic region and time, we define markets at the origin–time level. Each CBG in a given week constitutes a distinct demand market, reflecting the fact that residents of the same CBG face a common choice set of recreational sites but with origin-specific travel costs. Under this definition, market shares are defined over alternative sites within a CBG–week and sum to one across inside and outside options.

site-by-week fixed effects; 3) applying the causal inference technique, that is, SDID, to estimate the impact of beach closure on the mean utility via site-by-week fixed effects.

2.1 Solving Zero Market Share Issue

Before we enter the random coefficient model, a challenge needs to be handled in modeling beach visitation, which is the occurrence of zero visit shares: certain beaches record no observed visits from a portion of CBGs in particular weeks. Zero shares in our context arise at the CBG-site-week level rather than the site-week or CBG level alone. While site-by-week and CBG fixed effects can be estimated wherever any visits are observed, the structural demand model requires a consistent choice set across all CBG-site-week combinations, including those with no recorded visits in a given week. Because we work at a fine spatial-temporal resolution and cellphone data capture only a subset of the population, zeros are frequent. Many of these reflect both true non-visitation and sampling errors. Hence, dropping zeros or aggregating markets would implicitly change the choice set and bias demand and welfare estimates, making a principled correction necessary.

To address this, prior researchers often aggregated or selectively trimmed data, sometimes unintentionally, leading to potential selection bias and distorted demand estimates that impact policy simulations (Dubé et al., 2021; Li, 2019). In our context, zero market shares raise concerns about bias from two scenarios: (1) sites are available but unvisited, and (2) visits occurred but were not captured in the data. The first scenario may introduce selection bias, as unvisited sites could systematically differ from visited ones in accessibility, attractiveness, or unobserved characteristics, skewing inferences about preferences and closure impacts. The second scenario, where visits are missed due to sampling limitations, introduces measurement error, potentially underestimating demand and distorting the effects of key variables. In our study, we adjust the first issue by defining realistic choice set and address the second issue by using empirical Bayes posterior mean estimation, as proposed by Li (2019).[2] This approach refines visit probability estimates, ensuring compatibility with the estimation framework of Berry (1994) and Berry et al. (1995), which requires nonzero market shares.

Specifically, we construct a prior distribution for each CBG-site-week pair using visitation data from a set of nearby markets. For each focal CBG-site-week observation, we pool information from the 50 most comparable CBGs, where comparability is defined by having similar travel distance to the same site. The choice of 50 balances two goals: (i) borrowing enough observations to estimate a stable Beta prior, and (ii) keeping the pooled markets geographically and behaviorally close so that implicit travel costs are comparable. Travel distance here serves as a proxy for the "price" of visiting a site, so restricting to CBGs

[2] Section 4 shows how we define the choice set to reduce the selection bias from unvisited sites.

with similar distances ensures that the prior reflects similar effective costs.[3] We then model visit probabilities using a Beta-Binomial distribution, which shrinks observed zeros toward expected visitation rates in these comparable markets while avoiding oversmoothing.[4] This correction stabilizes the estimates and better reflects true recreational demand. Meanwhile, this is theoretically consistent with the random-coefficient logit framework and preserves the full choice set.

Formally, we model the number of visits to each beach site in each market, K_{jc}, as a draw from a binomial distribution with N_c trials, or total visits in each market. The trip probabilities are s_{jc}^0. For simplicity, time subscripts t have been suppressed throughout this subsection. The trip probabilities s_{jc}^0 are different for each CBG-site pair and are drawn from a Beta prior distribution with hyperparameters λ_{jc}^1 and λ_{jc}^2. This is a Beta-Binomial model of market shares (which can be generalized to a Dirichlet-Multinomial). The distributions of K and s are given by:

$$K_{jc} \sim Binomial\left(N_c, s_{jc}^0\right),$$
$$s_{jc}^0 \sim Beta\left(\lambda_{jc}^1, \lambda_{jc}^2\right), \tag{5}$$

The posterior distribution of the purchase probability is also a Beta distribution,

$$s_{jc} \sim Beta\left(\lambda_{jc}^1 + K_{jc}, \lambda_{jc}^2 + N_c - K_{jc}\right), \tag{6}$$

with posterior mean given by,

$$\hat{s}_{jc} = \frac{\lambda_{jc}^1 + K_{jc}}{\lambda_{jc}^1 + \lambda_{jc}^2 + N_c}, \tag{7}$$

For reference, the observed shares estimated by the maximum likelihood estimation (MLE) are,

$$\hat{s}_{jc}^{MLE} = \frac{K_{jc}}{N_c}, \tag{8}$$

We use strictly positive posterior means, $\hat{s}_{jc}$, in demand estimation instead of the MLE estimate. In large markets, the empirical Bayes posterior is very similar to the observed shares because the observed shares dominate the prior.

[3] Table D3 shows that alternative pooling windows yield quantitatively similar mean-utility and welfare estimates, indicating that our findings are not sensitive to the exact window size.

[4] Consider a zero observed share for trips from CBG A to site B in week t. We want a prior on the latent probability of such a trip. To construct that prior, we look at other CBGs with a similar travel distance to site B as CBG A has to B. For each CBG-site pair, we select the 50 CBGs whose distances to site B are closest to the distance from CBG A to B. We use their observed trip shares to site B (which may themselves include zeros) to form the Beta prior for that CBG-site-week observation. Importantly, even when some of those similar CBGs record a zero observed share, those zeros still carry information. In a Beta-Binomial framework, a zero observation effectively signals that the latent visitation probability is likely low, and it influences the shape of the prior by pulling the expected mean downward. Including these zeros therefore provides a more realistic picture of the distribution of possible trip probabilities and avoids the upward bias that would occur if we ignored zero observations.

For each site j in market c, the Beta prior is formed using market shares from similar markets, $l \in P_c$, where l is a market from the set of similar markets P_c. The parameters of the Beta prior, λ_{jc}^1 and λ_{jc}^2, are estimated from maximizing the log of the likelihood over the outcomes in the markets that form the priors,

$$f\left(K_{jl}, l \in P_c \middle| \lambda_{jc}^1, \lambda_{jc}^2\right) = \prod_{l \in P_c} \binom{K_{jl}}{N_l} \frac{\Gamma\left(\lambda_{jc}^1 + \lambda_{jc}^2\right)\Gamma\left(\lambda_{jc}^1 + K_{jl}\right)\Gamma\left(N_l - K_{jl} + \lambda_{jc}^2\right)}{\Gamma\left(\lambda_{jc}^1\right)\Gamma\left(\lambda_{jc}^2\right)\Gamma\left(N_l + \lambda_{jc}^1 + \lambda_{jc}^2\right)}, \tag{9}$$

We estimate a pair of hyperparameters λ_{jc}^1 and λ_{jc}^2 for each site, CBG, and week. This adjustment allows us to retain all CBG-site-week cells, regardless of whether visits were observed, thereby maintaining a fixed and consistent choice set for demand estimation and welfare simulation. The resulting posterior mean estimates, used in place of MLE, ensure strictly positive market shares and are the input for the subsequent demand estimation.

To illustrate the statistical properties of this approach, we conduct a Monte Carlo (MC) experiment that mimics the stylized environment of the model rather than the empirical setting. The simulation generates artificial markets in which consumers choose among a set of sites based on travel costs and latent quality, producing realistic variation in observed shares, including a high frequency of zeros. We compare several conventional adjustments, including dropping zero shares, aggregating markets, and adding small constants, with our EB posterior-mean correction using pools information across distance or time. Across repeated samples, the EB estimator consistently yields unbiased and stable recovery of the true travel-cost parameter and mean utilities, whereas the alternative treatments produce attenuation or instability when sparsity increases. This exercise demonstrates that the posterior-mean adjustment provides a consistent small-sample approximation to the true visitation probability and maintains compatibility with the discrete-choice framework. Detailed simulation design and results are provided in **Online Appendix B**.

2.2 Random Coefficient Logit Model

After replacing the zero market shares with the empirical Bayes mean, we then estimate a static demand model where visitors make myopic decisions based on current conditions. Following the BLP framework, we partition the parameter vector into θ_1 and θ_2, which govern the linear and non-linear components of the model, respectively. The vector θ_1 contains parameters that enter the mean utility δ_{jt} additively and are estimated in the linear GMM step. In our specification, θ_1 includes the mean coefficients on travel cost (α). By contrast, θ_2 governs the random-coefficient component μ_{icjt} that captures heterogeneous preferences across consumers. Specifically, θ_2 includes the standard deviations of the random coefficients associated with travel cost, which determine how individual tastes vary around the population means. These parameters are identified through the non-linear GMM contraction mapping that equates simulated market shares with observed ones. To account for spatial variations and unobserved factors, we use panel data and a set of fixed effects.

As highlighted by recent work on recreation demand, one potential concern is that the observed travel-cost regressor may be endogenous for measurement error in per-mile fuel costs (i.e., arising from local fuel price heterogeneity or imprecise price assignment). To address both channels, we implement an instrumental-variable specification following Bradt (2025). Specifically, we use travel distance and crude-oil prices interacted with state dummies as instruments for the travel costs. These interactions capture broad fuel-price shocks that shift travel costs across time and space but are plausibly orthogonal to localized recreation demand shocks. Because we include CBG, site, and week fixed effects separately, distance remains identified and enters the utility function directly as the spatial component of travel cost, while the IV addresses potential correlation between the fuel-cost component and unobserved, time-varying utility. Formally, for instruments Z_{tc}:

$$E[Z_{tc}\xi(\theta_2)] = 0, \tag{10}$$

The following equation sets the basis for the estimation of the demand model. We estimate the market share system with a general method of moments estimator. For every parameter guess, we invert the market system using a contraction mapping to obtain $\xi(\theta_2)$. Define Z_{tc} as the matrix of instruments and A_{tc} as a weighting matrix. We estimate θ_2 by:

$$\min \xi(\theta_2)' \, Z_{tc} A_{tc} Z'_{tc} \xi(\theta_2) \tag{11}$$

Once we estimate the CBG-site-specific utility from the demand model, we use these utility measures to analyze how external factors, such as beach closures, influence recreational choices. Specifically, the estimated site-specific utility values ζ_{jt} become the dependent variable in our causal analysis. By leveraging an SDID approach, we can isolate the effect of beach closures from broader seasonal and spatial variations.

2.3 Synthetic Difference-in-differences Application

To measure how beach closures impact recreational utility, we apply a Synthetic Difference-in-Differences approach. This method allows us to model the site-specific utility ζ_{jt} as a function of various fixed effects, two kinds of weights, and an indicator variable for closures:

$$\{\hat{\beta}, \hat{\phi}_j, \hat{\eta}_t\} = argmin \left\{ \sum_{j=1}^{N} \sum_{t=1}^{T} \left(\zeta_{jt} - \alpha - \phi_j - \eta_t - \beta D_{jt}\right)^2 \hat{\omega}_j \hat{\lambda}_t \right\}, \tag{12}$$

In this specification, β measures the effect of beach closures D_{jt} on site-specific utility ζ_{jt}, representing the reduction in expected beach visitation utility due to closure. The term ϕ_j represents site fixed effects, capturing time-invariant characteristics of each beach, such as amenities, geographic attributes, and historical visitation patterns. The week fixed effects η_t control for temporal fluctuations in visitation due to seasonality, weather patterns, or broader trends affecting recreational demand. The weights $\hat{\omega}_j$ and $\hat{\lambda}_t$ are site-specific and week-specific weights, which are applied to sites and weeks to ensure robustness against potential selection bias and balance systematic differences between closed and unclosed beaches.

By constructing counterfactual visitation trends for closed beaches using weighted comparisons with similar open beaches, this approach strengthens causal identification. Conditioning on beach and week fixed effects allows us to isolate the direct effect of closures from other confounding influences. This framework is particularly useful for evaluating policies that affect public access to coastal areas, such as temporary closures, environmental restoration efforts, or disaster response measures (Earle and Kim, 2024). The estimated β coefficient provides empirical evidence on how closures impact the recreational utility, offering valuable insights for policymakers balancing conservation goals with recreational access.

3. Application

The 2021 Huntington Beach oil spill serves as a case study to illustrate how our model estimates the causal effects of beach closures on recreational demand. By tracking visitation patterns before, during, and after the spill, we evaluate the economic impact of temporary closures on affected and nearby beaches. This section begins with an overview of the environmental and economic impacts of oil spills, followed by a detailed account of the Huntington Beach spill, including its timeline, government responses, and effects on recreational areas. Finally, we explain why Huntington Beach provides an ideal setting for studying the recreational consequences of environmental disruptions.

3.1 Oil Spill Facts

Oil spills pose significant threats to ecosystems, human health, and local economies, making them a critical environmental and socio-economic concern. The U.S. Department of Energy estimates that approximately 1.3 million gallons of oil are spilled into the environment annually, endangering marine and coastal ecosystems (Live Science, 2010). These spills contaminate water bodies, destroy marine habitats, and harm wildlife, with recovery often taking years or even decades.

The health effects of oil spills on nearby communities are well-documented and severe. The 2010 Deepwater Horizon spill, one of the largest environmental disasters in history, caused widespread health issues among cleanup workers and Gulf Coast residents, particularly affecting vulnerable populations such as children, the elderly, and those with pre-existing conditions (D'Andrea and Reddy, 2018; Cano-Urbina et al., 2019). Economically, oil spills disrupt industries reliant on fisheries, tourism, and coastal resources. Tourism-dependent communities, in particular, suffer from declines in visitors as beaches and recreational areas become polluted and unsafe.

Given the persistent risks from oil spills—whether from collisions, equipment failures, natural disasters like hurricanes, or operational discharges during oil transport—understanding their wide-ranging impacts is essential for developing effective prevention and response strategies.

3.2 2021 Huntington Beach Oil Spill

On October 1, 2021, a pipeline connected to the offshore platform Elly, operated by Beta Offshore, a subsidiary of Amplify Energy, was damaged—likely dragged over 100 feet by a cargo ship's anchor—resulting in the release of approximately 588 barrels (27,996 gallons) of crude oil into the ocean (NOAA, 2021b). The spill occurred about 3.5 miles off the coast of Huntington Beach, California, eventually contaminating shorelines as far south as Dana Point and affecting key recreational areas in Huntington Beach and Newport Beach. Thirteen protected areas and one federally protected bird species were also impacted (NOAA, 2021a). Amplify Energy and its subsidiaries were fined $13 million to cover federal penalties and cleanup costs (CBS News, 2023a).

In response, the cities of Huntington Beach, Newport Beach, and Laguna Beach swiftly enacted beach closures, while the California Department of Fish and Wildlife imposed a ban on fishing and water activities from Huntington Beach to Dana Point. Although beaches in Huntington and Newport reopened on October 11 and Laguna Beach on October 14, advisories remained in place, warning visitors to avoid contact with oiled areas and tar balls. The fishing ban persisted until November 11, 2021.[5] Additionally, the final day of the popular Pacific Airshow was canceled, further disrupting local tourism.

The spill had substantial economic repercussions, particularly in Huntington Beach and Newport Beach, where tourism and fishing are vital to the local economy. To gauge public attention toward the event, we analyzed Google Trends data related to the 2021 Huntington Beach oil spill. Figure A1 illustrates a sharp spike in search activity on October 3, coinciding with widespread media coverage and the official confirmation of beach closures and the fishing ban. Elevated search interest persisted for roughly two weeks. As shown in Figure A2, the highest search activity was concentrated in the Los Angeles metro area, followed by surrounding regions.

3.3 Study Area

Huntington Beach presents a unique opportunity for studying the effects of environmental disruptions on recreational demand for several reasons. First, it is one of the most popular beach destinations in Southern California, attracting millions of visitors annually due to its extensive coastline, favorable weather, and vibrant tourism industry. This high baseline of recreational activity provides a rich dataset for detecting shifts in visitation patterns following environmental shocks, such as the 2021 oil spill.

Second, the Huntington Beach oil spill serves as a natural experiment to study the causal impact of sudden, unanticipated environmental events on recreation. The abrupt nature of the spill, coupled with

[5] Official beach closure and reopen dates can be found at here. The spill's location relative to Huntington Beach, Newport Beach, and Laguna Beach can be find in here: https://www.eastbaytimes.com/2021/10/05/heres-what-beaches-are-off-limits-as-officials-work-to-contain-massive-oil-spill/amp/.

immediate beach closures and fishing bans, created an exogenous shock to recreational access. This allows for more precise causal inference, as the spill was not influenced by visitor behavior or pre-existing trends. Third, the proximity of alternative beach sites, such as Newport Beach and Laguna Beach, enables an examination of substitution effects. By comparing visitation patterns across these neighboring beaches, we can assess how visitors adjust their behavior in response to localized environmental disruptions, offering broader insights into recreational demand dynamics during environmental crises.

Finally, the availability of high-frequency, cellphone-based mobility data for the region allows for detailed tracking of visitor movements before, during, and after the spill. This granular data enhances the accuracy of demand estimation and welfare analysis, making Huntington Beach an ideal setting for applying advanced econometric methods such as empirical Bayes estimation and SDID. Together, these factors position Huntington Beach as an exemplary case for advancing our understanding of how environmental disasters affect recreational behavior and economic welfare.

4. Data

This section discusses the data sources of Huntington Beach Oil Spill, beach locations, and their number of visits. Additionally, it outlines and details our step-by-step data-processing procedure. An introduction to key variable construction is provided. Lastly, we summarize the descriptive statistics of our key variables.

4.1 Data Sources and Preparation

To analyze the impact of the 2021 Huntington Beach Oil Spill on recreational demand, we integrate multiple data sources, including oil spill incident records, high-frequency mobility data, and geographic information on coastal zones and points of interest (POIs). We first obtain detailed information about the oil spill from the Pipeline and Hazardous Materials Safety Administration (PHMSA) incident database, which tracks incidents related to natural gas and crude oil pipeline transportation in the United States. PHMSA classifies an event as an incident if it results in fatalities, serious injuries requiring hospitalization, property damage exceeding $50,000, or the release of more than three million cubic feet of gas. The dataset provides key attributes, including the geographic coordinates, occurrence timing, cause of the event, number of fatalities, and estimated volume of oil released, making it a critical resource for assessing the scale and impact of the spill.

To accurately identify beach locations in California, we use Advan's polygon-based POI dataset, which is updated monthly via web crawling, public APIs, and third-party licensing. Unlike point datasets, Advan defines POIs as polygons that capture spatial boundaries. A physical beach can comprise multiple distinct POIs (e.g., access points, parking lots, visitor facilities). We classify a POI as a beach if (i) it carries the "Beach" category in Advan's taxonomy based on industry classification (e.g., NAICS code) or (ii) its name contains "Beach". We also verified labels via manual checks. We refer to the collection of POIs that belong to the same-named shoreline site as a beach cluster (e.g., the Laguna Beach cluster consists of 15 POIs that together represent the beach).

We measure visits using Spectus mobility data from opt-in, anonymized smartphone signals (Bluetooth, GPS, WiFi, IoT), with roughly 15 million daily active users in the U.S. We track daily visits to selected beaches from June 1, 2021 through December 31, 2021 (weeks 32-51). For demand estimation, we aggregate to a CBG-POI-week panel that captures weekly visit counts. By integrating these diverse data sources, we build a robust dataset that enables a detailed examination of the impact of the spill on recreational demand, allowing us to quantify visitation changes, substitution effects, and economic losses resulting from beach closures.

Before linking beach POIs to the 2021 Huntington Beach oil spill site, we first verify their accuracy by cross-referencing POIs with the California Coastal Commission's coastal zone map (digitized from 1:24,000 boundary maps). Since Advan's POI classifications are sometimes imprecise, for instance, bike trails may be mislabeled as beaches, we remove any POIs that fall outside designated coastal zones. Additionally, we drop the overlapped and smaller polygons to prevent double counting. To ensure representativeness, we further refine our sample by removing beach POIs smaller than 500 square meters. Smaller POIs are more susceptible to recording zero visits due to sampling errors, which could distort our demand estimates. Excluding these POIs helps reduce the incidence of zero market shares, improving the robustness of our analysis.[6] To further explore potential measurement errors, we assess network coverage using the National Broadband Map provided by the Federal Communications Commission. This map confirms that the beaches in our study area have nearly 100% coverage from at least one of the three major service providers (AT&T, Verizon, and T-Mobile) through 4G LTE and 5G networks, minimizing concerns about device mislocation (FCC, 2023).

4.2 Basic Validity Checks for Spectus Visits Data

After data cleaning, we conduct a set of diagnostic checks to validate the Spectus CBG-POI-week visit measures shown. These analyses are based on the cleaned sample after geofencing, coverage screening, and overlap resolution. Figure A3a normalizes visits to the first week of our study window and plots average weekly visitation for treated and control beaches. We observe a sharp, immediate decline at Huntington Beach following the spill, consistent with event timing and expected behavioral responses; the temporary spike in week 39 coincides with the Pacific Airshow. Figure A3b documents a strong positive relationship between the number of sampled devices and CBG population, indicating that Spectus sampling density scales proportionally with population and supporting demographic representativeness.

[6] Very small POIs are overwhelmingly micro-amenities (access points, kiosks, restrooms, small overlooks) or misclassified polygons rather than full beach extents. Including them would expand the choice set with options that are not meaningful beach destinations and would invite EB to "rescue" zeros that arise from misclassification, not from sampling error.

Figure A3c examines the distance-decay relationship between visitation probability and travel distance under alternative functional forms (log-log, log-linear, polynomial, cubic B-spline), alongside categorical distance bins. To isolate the shape of the decay, each specification includes site, CBG, and week fixed effects, netting out baseline site popularity, home-market size, and common time shocks. Across all specifications, visitation declines smoothly and monotonically with distance, matching standard recreation-demand predictions. Also, we note that all POIs are continuously observed over the full window.

Finally, a reduced-form SDID specification yields a sharp post-spill decline at treated sites with flat pre-trends shown in Table E1, corroborating the direction and timing used in our structural analysis. More details are provided in **Online Appendix E**. Together, these diagnostics demonstrate that the Spectus data are consistent with population and behavioral expectations, providing confidence in their use for causal and welfare analyses.

4.3 Key Variable Construction

We determine affected beaches by measuring the distance from each beach POI boundary to the spill site. Since the crude oil spread southeastward, reaching Dana Point, we classify beaches within 43 km southeast of the spill site as affected. Although Long Beach is geographically close to the spill, it is excluded from the affected group because it lies north of the site and was not exposed to the oil spread. Based on these criteria, we identify 21 affected beach POIs, including 1 in Huntington Beach cluster, 5 in Newport Beach cluster, and 15 in Laguna Beach cluster. The remaining 458 unaffected beaches serve as the control group. Figure 1 shows our treatment and control beach POIs in our study area.

We then calculate travel costs by incorporating both the monetary costs of travel and the opportunity cost of time. Following standard practice in the recreation demand literature (Whitehead et al., 2018; English et al., 2018), we assume that a visitor's opportunity cost of time is one-third of the hourly wage rate implied by the median annual income of their CBGs. The round-trip travel cost is formulated as:

$$TC_{cjt} = 2 \times (g_{sct} + f_t) \times Dist_{cjt} + 2 \times \gamma \frac{Medinc_c}{2080} Time_{cjt}, \tag{13}$$

where TC_{cjt} denotes the travel costs, calculated using the one-way travel distance $Dist_{cjt}$ and time $Time_{cjt}$ between the centroid of CBG c and POI j in week t. The formula incorporates the state-level average gas cost g_{sct} and the average vehicle maintenance, repair, and depreciation costs f_t.[7] Additionally, the median annual income $Medinc_c$ in the visitor's CBG is converted to an hourly wage rate by dividing by the number of working hours per year (i.e., 40 hours/week × 52 weeks = 2080 hours). The parameter γ is set to 1/3 for the central estimate. All travel costs are expressed in 2021 dollars.

[7] We parameterize non-fuel operating costs using AAA's Your Driving Costs methodology. In the baseline, the depreciation component in f_t is set to \$0.26 per mile for a midsize vehicle—converted from AAA's annual ownership costs to a per-mile figure—while maintenance and tires follow AAA's per-mile estimates; fuel costs g_{sct} come from state-level prices (AAA, Your Driving Costs, recent editions).

Since the mobility data from Spectus captures all recreational trips taken by users from a CBG, some trips may involve air travel rather than car travel. As air travel introduces significantly different cost structures and preferences compared to car travel, our analysis focuses on trips that are plausibly taken by car. To do this, we restrict the sample to individuals traveling within a 300-mile radius of each beach in our sample. This threshold aligns with Dundas and Von Haefen (2020), who assume that a 300-mile distance represents the upper limit for a feasible single-day recreational trip. Beach sites within a 300-mile radius of a home CBG are defined as the choice set for residents in that CBG.[8]

Once the choice sets for each CBG are determined, we calculate beach trip shares in the following steps. First, we define the market for each CBG as all devices residing in that CBG, regardless of whether they made any recreational trip during the period. This definition allows us to capture both the extensive margin of recreation demand (the decision to take any beach trip) and the intensive margin (the choice among alternative beaches). We estimate the total market size for each CBG in a given year using the number of unique devices recorded by Spectus and approximate the weekly market size by multiplying the number of residing devices by seven, assuming that individuals face at most one recreational choice opportunity per day. The observed trip share for each beach POI is then calculated as the number of trips from a CBG to that POI divided by the total potential market size. The outside share represents all other non-beach activities and is derived as the difference between the total potential trips and observed recreational trips from the CBG, divided by the total market size.

Moreover, we construct a dummy variable to capture the temporal change due to the oil spill. Specifically, weeks 32 to 38 in 2021 belong to the pre-spill period and weeks 39 to week 51 are in the post-spill period. In addition, we generate a dummy variable assigning 1 to Huntington Beach in week 39 and 0 otherwise, to capture the impact of the Pacific Airshow occurred in week 39 around Huntington Beach on beach visits.

4.4 Summary Statistics

With these linkages and assumptions, we construct aggregate data of actual site choices for each CBG from June 2021 to December 2021. Our final estimation data set consists of 3,739,601 observations corresponding to CBG-POI-week combinations. Table 1 reports the summary statistics regarding the key variables in both 1st and 2nd stages in our models. Panel A presents that the average number of visits is 0.32 with a wide range of visitation frequencies, from 0 to a maximum of 2,579 visits. This suggests the existence of the zero market share issue. The travel distance is 56.14 miles in the mean, with a standard deviation of 56.92 miles, suggesting significant variability in travel distances. Travel time averages 1.32 hours, with a standard deviation of 1.34 hours. Gas costs are constant at 18.65 cents per mile, while

[8] We believe this distance-based restriction captures the most realistic choice set for residents within a CBG and helps mitigate the selection bias discussed in Section 2.

maintenance and depreciation costs are fixed at 9.55 cents and 26.00 cents, respectively. The travel cost ranges from \$0.03 to \$399.79. The observed share of visits averages 0.00060, while the Bayesian share averages 0.00059, both indicating low average visitation probabilities with some high outliers.

We further plot the empirical Bayes posterior means against the observed market shares in Appendix Figure D1. Each dot represents a CBG-POI-week observation. In Appendix Figure D1a, most dots are on the 45-degree line, implying that the majority of posterior mean estimates are the same as the original observed market shares. When zooming into a smaller scale as shown in Appendix Figure D1b, the spike at 0 on the x-axis (observed market shares) shows that observed zero shares are mapped to different empirical Bayes posterior means, suggesting that transforming zero shares into the same non-zero share is inconsistent with the demand model.

Panel B in Table 1 presents the summary statistics for the second stage, with a sample size of 9,580 observations at the POI-week level. The mean utility, derived from the first stage, has an average value of -0.2546. The treatment variable, which assigns a value of 1 to the 21 affected beaches and 0 otherwise, has a mean of 0.043. The post dummy variable, indicating the pre-spill period with 1 and the post-spill period with 0, is averaged at 0.65. Additionally, the mean of the Huntington-week 39 dummy variable is 0.0001.

5. Results

In this section, we first display the baseline results of disclosure impacts for all affected beaches, then explore the heterogeneity by beach clusters, and further calculate the welfare loss per visitor and in aggregation. A series of robustness checks are also conducted.

5.1 Baseline Results

We present our demand estimation results in Panel A of Table 2, comparing three model specifications: base logit (Column 1), logit model with instrumental variables (IV) (Column 2), and random coefficient logit (R.C. logit) (Column 3). The coefficient on Travel Costs is negative and statistically significant across all specifications, indicating that higher travel costs reduce the probability of visiting a beach in California. In Column (2), incorporating IV estimation yields a consistent estimate, addressing potential endogeneity concerns. Our preferred specification in Column (3), which allows for individual-specific heterogeneous preferences, produces similar estimates for travel costs. The Travel Costs (sd) term in Column (3) is small and statistically insignificant, suggesting that travel cost sensitivity is relatively homogeneous across individuals.

In Panel B of Table 2, we extend the random coefficient logit model by incorporating the SDID method to evaluate the causal impact of beach closures on mean utility. Column (1) shows that the beach closure led to a decrease in mean utility of 0.009, although this is insignificant when considering all affected beaches. To capture the dynamic effects of beach closures, we employ the SDID event study method proposed by Ciccia (2024). Figure 2a shows that the differences in mean utility dropped sharply by approximately 0.05

in the week of closure. Over the following four weeks, mean utility gradually recovered to pre-spill levels, suggesting that the overall impact was short-term, lasting about one month. This pattern aligns with Google search trends in Figure A1, where elevated search intensity for the spill is observed only during the same period. These results highlight the temporary but relatively substantial disruptions in recreation demand following the Huntington Beach oil spill, particularly for directly impacted sites.

We also explore the possible heterogeneous effects of beach closures on diverse beach clusters. Huntington Beach cluster experienced the most substantial and sustained decline in mean utility following the beach closure shown in Table 2. Column (2) in Panel B presents a mean utility drop of 0.2584.[9] Also, as shown in Figure 2b, the mean utility dropped sharply in closure week, reaching its lowest point. Unlike the overall trend observed across all affected beaches, where mean utility rebounded within a month, Huntington Beach cluster continued to exhibit a persistently depressed mean utility after reopening. Even 12 weeks later, it did not return to the pre-closure level, indicating a sustained reduction in the intrinsic enjoyment of visiting Huntington. One plausible mechanism is a stigma or persistent risk-perception effect associated with Huntington's proximity and direct exposure to the spill (Taylor et al., 2016). Even though beaches were reopened, visitors might still avoid water-based activities such as swimming and fishing. Some research shows that ecological impacts may linger, which can dampen marine-related activities such as shell collection (CBS News, 2023b). In addition, Huntington is especially oriented toward direct beach-based recreation (surfing, beach volleyball, etc.), making its site-specific value more sensitive to closures and contamination.

Newport Beach cluster exhibited a more muted change, as shown in Column (3) of Panel B and Figure 2c. The decline in mean utility was substantially smaller and remained relatively close to pre-closure levels throughout the observation period. This pattern is consistent with Newport's location being less directly exposed to offshore contamination and a relatively rapid reopening as water-quality tests indicated non-toxic oil levels (City of Newport Beach, 2021). Importantly, Newport's harbor amenities provide within-site leisure opportunities beyond shoreline recreation, which may have buffered the decline in the site's intrinsic utility.

Laguna Beach cluster showed the least impact, as observed in Column (4) of Panel B and Figure 2d. Mean utility remained relatively stable, with only minor fluctuations during and after the closure period. Consistent with the fact that Laguna experienced minimal direct oil contamination compared to Huntington and Newport, and that it offers a broader amenity mix (art galleries, festivals, hiking trails, tide pooling, etc.), the site-specific enjoyment of visiting Laguna appears largely unaffected when shoreline access is restricted.

[9] We use the Jackknife method to calculate the standard errors except for Huntington Beach cluster due to only one treated unit. Thus, we apply the placebo method with seed 100. To ensure the stability of our results, we replace the seed with 200, 300, and 400. Table D1 indicates robust estimates.

Thus, the heterogeneous mean-utility decreases across these three beach clusters possibly stem from the amenity composition and the extent of oil contamination exposure. This variation underscores the importance of considering site-specific factors when assessing the economic impact of beach closures. In the next subsection, we quantify these differences by estimating the welfare loss per visitor and in aggregate.

5.2 Welfare Effects of Beach Closures

To quantify the welfare impact of beach closures, we simulate a counterfactual scenario where no beach closures were implemented. Using the estimated visitation responses, we calculate the changes in consumer surplus as the difference in compensating variation (CV) between the baseline scenario (with closures) and the counterfactual scenario (without closures). The compensating variation associated with the closure of selected beaches is as follows:

$$CV_C = -\frac{1}{\alpha}\left[log\sum_{j=0}^{J}\exp\left(V_{cj}\middle|D_1\right) - log\sum_{j=0}^{J}\exp\left(V_{cj}\middle|D_0\right)\right], \tag{14}$$

where $V_{cj}|D_1$ and $V_{cj}|D_0$ denote the deterministic components of utility under observed and counterfactual conditions, respectively. CV_C is computed at the CBG-week level (week t omitted for brevity). Aggregate welfare losses are obtained by multiplying CBG-level CV_C by CBG population and summing across all CBGs and post-spill weeks.[10]

To ensure that the welfare estimates accurately reflect heterogeneous site-level impacts, we estimate welfare using the heterogeneous SDID coefficients rather than imposing a homogeneous treatment effect across all affected beaches. Because Newport and Laguna exhibited no statistically meaningful visitation decline after the spill, a specification that constrains all affected beaches to share a common treatment effect would mechanically dilute the true magnitude of the Huntington Beach response, leading to an understated welfare loss. To address this concern, we estimate welfare losses under two complementary specifications designed to capture heterogeneous treatment effects. The first allows beach-specific SDID responses while aggregating welfare effects across all affected beaches, ensuring that differences in visitation sensitivity are preserved in the welfare calculation. The second isolates the Huntington Beach effect by assuming no spill-related impacts at Newport and Laguna, allowing us to quantify Huntington's specific contribution to the total welfare loss.

Panel A of Table 3 allows all three affected clusters (Huntington, Newport, and Laguna) to respond independently to the spill. The estimated aggregate welfare loss is \$0.99 million (95% CI: –\$0.11M to \$2.32M) across all affected beaches combined, corresponding to a weekly loss of about \$0.10 million during

[10] That is, Total Welfare Loss $=\sum_{c.t} CV_{ct} \times Pop_c$.

the closure period and $83 thousand per week thereafter.[11] The persistence of welfare losses beyond the closure window suggests lasting behavioral responses—visitors' lingering safety concerns, delayed cleanup, or habit reallocation.

Panel B isolates the welfare impact attributable specifically to Huntington Beach by restricting the treatment response to that site and assuming no spill-related impacts at Newport or Laguna. The estimated loss is $1.00 million (95% CI: $0.03M to $2.19M), or $0.10 million per week during closure and $84 thousand per week afterward. The near-identical results across Panels A and B confirm that Huntington Beach alone explains nearly the entire welfare decline.

Taken together, these findings underscore the dominant role of Huntington Beach in driving the overall welfare loss. Its higher baseline visitation, longer closure duration, and limited nearby substitutes magnified the impact. The persistence of reduced visitation even after reopening highlights how environmental disasters generate not only short-term access disruptions but also long-term economic scars through sustained behavioral shifts.

5.3 Robustness Checks

To ensure the reliability of our findings, we conduct a series of robustness checks. First, we perform a sensitivity analysis by varying the number of priors used to generate Bayesian posterior shares, testing the stability of our demand estimates. Second, we replace the SDID method with a traditional event study approach in the second stage, evaluating whether our welfare estimates hold under a different causal inference framework. Third, we apply several placebo tests to check the cross-site substitution. Fourth, we conduct a set of robustness checks using alternative travel distance cutoffs. Last, sensitivity analysis of SDID weighting and control group design is performed. These checks help confirm the robustness of our results and the validity of our methodological choices.

5.3.1 Empirical Bayesian Estimator using Alternative Priors

We then examine how demand estimates vary depending on the choice of EB priors. Appendix Table D3 reports results from a logit demand model, comparing estimates using observed market shares and empirical Bayes posterior mean market shares constructed with different prior sample sizes. Column (1) presents estimates using observed market shares, where observations with zero market share are excluded. Columns (2) through (6) show results based on empirical Bayes posterior means, each constructed with a different number of similar markets in the prior. Column (2) incorporates priors from 20 similar markets, Column (3) from 30 markets, and so on, with Column (6) using 90 markets. Across specifications, estimates using empirical Bayes priors differ from those in Column (1) but remain stable across Columns (4) through (6),

[11] The weekly welfare loss $83,000 corresponds to roughly $12,000 per beach per day. This value is higher than the $4,600 per beach-day reported by Earle and Kim (2024), which is consistent with Huntington Beach being a substantially more popular recreational destination with higher baseline visitation.

with precision improving as the number of prior markets increases. This consistency suggests that our main findings are robust to different choices of prior sample size. In the main analysis, counterfactual and welfare estimates are based on priors constructed from 50 similar markets.

Appendix Table D3 Column (1) differs from the empirical Bayes columns in two key ways: (i) it contains fewer observations because zero shares are dropped, and (ii) it relies on MLE rather than empirical Bayes posterior means. To investigate how these differences influence demand estimates, we compare the distribution of market shares under different estimation approaches. Appendix Table D4 presents summary statistics of market share estimates, with Panel A summarizing the full sample, while Panels B and C present statistics separately for nonzero and zero observed shares, respectively. In Panels A and B, empirical Bayes market shares exhibit similar means but lower variances compared to observed shares, and the distributions remain consistent across different priors. In Panel C, where observed shares are strictly zero, empirical Bayes estimates shift these values away from zero based on the prior, ensuring strictly positive shares.

5.3.2 Event-study Causal Technique

Lastly, the SDID method in the 2nd stage is replaced with the traditional event study approach to assess the sensitivity of our results to different causal inference techniques. Appendix Figure D2 illustrates that the event study method yields consistent results across different beach clusters. However, the event study results show a broader 95% confidence interval, indicating less precise estimates compared to the SDID approach. Notably, the interpretation of the results for Huntington Beach under the event study framework changes, as the 95% confidence interval includes zero after the closure week. This suggests a more short-term impact, contrasting with the findings from the SDID method. Such inaccuracies could lead policymakers to underestimate the severity of the oil spill, potentially affecting the effectiveness of their response strategies.

We further compare welfare loss to those obtained using an event study method, which yields similar results. Results from Panel A of Table D5 suggest that the event study approach produces an aggregate welfare loss estimate of \$0.88 million (95% CI: –\$0.88M ~ \$3.34M), with a weekly loss of \$86 thousand (95% CI: –\$65 thousand ~ \$298 thousand) during closure and \$74 thousand (95% CI: –\$73 thousand~ \$279 thousand) per week after the initial closure. The similarity between the SDID and event study results reinforces the robustness of the estimated economic impact of beach closures. When only calculating the welfare loss for Huntington Beach, an aggregate loss of \$0.88 million (95% CI: –\$0.76M ~ \$3.21M) is estimated using the event study method, which is similar to the number produced through the SDID approach. The per week during closure week and after the initial closure are estimated as \$82 thousand (95% CI: –\$59 thousand ~ \$282 thousand) and \$73 thousand (95% CI: –\$63 thousand ~ \$267 thousand), respectively, slightly lower than that in the SDID method.

Overall, the SDID approach estimates a slightly lower welfare loss than the event study method, particularly in the post-closure period. This discrepancy suggests that the event study method may slightly underestimate the long-term persistence of economic losses, whereas SDID captures a more sustained decline in utility and visitation behavior. These findings indicate that recreational demand does not rebound as quickly as the event study method implies, underscoring the importance of choosing robust causal inference techniques to assess the economic impact of environmental disruptions.

5.3.3 Check on Cross-site Substitution

A common concern is the cross-site substitution, where visitors may redirect their trips to nearby untreated similar sites due to the closure of the affected beaches. We employ three methods to alleviate this concern. First, we follow the standard practice in the synthetic control paper (Abadie, 2021) by drawing the spaghetti plot. Specifically, we artificially assign one of control group beaches as the treatment beach and estimate the impact of these "placebo closures" on their mean utilities. We go over all 458 control beaches and compare the difference between their synthetic and observed mean utilities. Figure D3a illustrates the SDID event study results. On average, the placebo estimates (grey lines) nearly uniformly fluctuate around zero. While some exhibit temporary fluctuations similar to Huntington Beach (blue dotted line), none experience a sustained and substantial decline in mean utility. Meanwhile, we perform a placebo test by randomly assigning 50 control sites to a 'treatment' group and repeating this 100 times. If substitution were pervasive, placebo 'treated' sites would show artificial mean utility increases. However, Figure D3b reveals a distribution centered near zero, supporting our assumption of limited interference.

Additionally, we conduct robustness checks by progressively excluding beaches in the control group near the spill site (50-, 100-, 200-, and 300-mile radii). This approach directly tests whether proximity-driven substitution inflates our estimates. Table D6 shows that our results remain stable across these thresholds, suggesting minimal bias from spillovers.

This reinforces the validity of our causal inference, confirming that the observed welfare loss is driven by the oil spill-induced beach closures rather than unobserved temporal shocks affecting all beaches.

5.3.4 Sensitivity of Travel Distance Cutoff

To test the sensitivity of our results to the definition of the feasible travel market, we vary the maximum travel distance allowed in the sample. Our baseline specification follows Dundas and Von Haefen (2020) in restricting observations to visitors traveling within a 300-mile radius of each beach, representing an upper bound on plausible single-day driving trips. To examine whether this choice affects inference, we re-estimate both stages of the model using alternative cutoffs of 100, 150, 200, and 250 miles. As reported in Table D2, the estimated travel-cost coefficients and closure effects remain highly stable across these alternative definitions, indicating that behavioral responses are not sensitive to the precise market boundary. While welfare magnitudes vary somewhat, as shorter travel ranges imply lower average costs and smaller feasible choice sets, the qualitative conclusions remain unchanged. These results confirm that our findings are robust to alternative definitions of the spatial market for recreational trips.

5.3.5 SDID Weighting and Design Sensitivity

We examine the sensitivity of our SDID results to weighting and donor composition. We report both POI-specific (unit) weights and week-specific (time) weights. Following the SDID package defaults and Arkhangelsky et al. (2021) (eq. 5), we use a ridge penalty of for unit weights; time weights are lasso-penalized with the default 10−6. We do not otherwise tune or adjust the smoothing parameters.

Figure C1a plots control units, where dots are sized by weight. In the pooled analysis (all 21 treated beaches), 456 control units receive strictly positive weights (nearly 100% non-zero). Figure C1b shows week-specific pre-period weights, with the largest mass on week 38. Although near-uniform unit weights in the pooled run may appear unexpected, restricting to Huntington Beach only yields the selective, sparse donor pattern emphasized in the SDID paper: 170/458 (37.1%) controls with zero weights (Figure C2a). By contrast, Newport and Laguna exhibit near-uniform unit weights (only 6-9 zero-weight units out of 458; Figures C2b-C2c), consistent with their high pre-treatment similarity to the donor pool. These patterns illustrate a useful property of SDID: even when unit weights are diffuse, informative time weights can still recover reliable effects; in the limiting case where both unit and time weights add little structure, the estimator effectively collapses to DID, and inference reduces to DID under its standard assumptions.

To assess robustness to donor composition, we repeat 100 times by drawing 200 controls at random and re-estimate SDID. Estimates shown in Figure C3 remain stable. This confirms that our findings are not driven by a particular donor pool or weight concentration.

Thus, the weighting diagnostics and random-subsample exercise indicate that our SDID results are robust to reasonable variations in weighting and donor composition. More details are reported in **Online Appendix C**.

6. Discussion and Conclusions

This study provides robust empirical evidence on the economic impact of beach closures, with a particular focus on addressing the zero market share issue, which is a common challenge in discrete-choice models of recreation demand. By integrating a random coefficient logit model with SDID approach, we quantify the welfare losses resulting from restricted access to coastal beaches. Our approach ensures that observations with zero market share are properly accounted for, improving the reliability and stability of demand estimates.

One of the key methodological advancements in this study is our treatment of zero market shares, a pervasive issue in site-choice models where certain recreation sites receive no recorded visits in some periods. If unaddressed, these zero shares could distort welfare estimates and bias demand elasticity measures. To overcome this, we employ empirical Bayes posterior mean estimation, which transforms zero shares into small but positive values while preserving the underlying distribution of visitation probabilities.

This correction allows us to integrate all beaches into the demand estimation process, ensuring that sites with low visitation are not arbitrarily excluded. Importantly, our empirical Bayes correction not only ensures positive market shares but also preserves the full choice set at the CBG-site-week level, addressing the pervasive sparsity (88% zeros) typical of cellphone mobility data. To validate the performance of this correction, we conduct Monte Carlo simulations showing that the empirical Bayes estimator reliably recovers true parameters and stabilizes mean utility estimates under varying levels of data sparsity.

Our findings demonstrate that properly addressing the zero market share issue has meaningful implications for welfare analysis. By incorporating adjusted visitation probabilities into the random coefficient logit model, we improve the precision of mean utility estimates and generate more accurate welfare loss calculations. This is particularly crucial in assessing the effects of beach closures, where some sites may exhibit temporary zero shares due to access restrictions rather than an intrinsic lack of demand.

The results highlight heterogeneous impacts across different beach clusters. Huntington Beach experienced the most substantial and persistent decline in utility, with an estimated total welfare loss of \$1.0 million using SDID, compared to \$0.88 million to the event study method. The prolonged decline in mean utility suggests that visitors to Huntington Beach were less able to find suitable substitutes, leading to sustained welfare losses. In contrast, Newport Beach and Laguna Beach showed more moderate and short-lived impacts. While there was an initial decline in mean utility, both beach clusters exhibited quicker recoveries, with significantly lower post-closure welfare losses. These patterns indicate that visitors to these beaches could adapt more effectively, either by substituting to nearby locations or adjusting their recreation behavior. Furthermore, comparing the SDID and event study estimates provides insights into the persistence of welfare losses. The event study method tends to slightly underestimate post-closure effects, while SDID captures a more sustained decline in visitation utility, suggesting that closures had lasting economic consequences beyond the initial restriction period.

Despite the strengths of this study, some limitations remain. First, while our empirical Bayes approach effectively addresses the zero market share issue, further research could explore alternative data sources, such as survey-based stated preference data or geotagged social media check-ins, to further validate visitation patterns. Second, our analysis primarily focuses on beach site substitution, but some visitors may completely shift away from beach recreation to other forms of outdoor activity. Future studies could examine broader substitution effects, including whether visitors switch to urban parks, hiking trails, or indoor entertainment options in response to closures. Third, while cellphone mobility data provide detailed visitation patterns, they do not reveal travelers' trip purposes. Consequently, some observed beach visits may represent incidental or secondary stops rather than primary recreation trips. This limitation, noted by Lupi et al. (2020), implies that our welfare estimates could slightly overstate true recreation benefits. Lastly, our welfare estimates focus on visitor utility, but closures may have wider economic effects, including business losses, reduced tourism revenue, and changes in local employment. Future research could integrate local economic data to provide a more comprehensive view of the broader financial impact of beach access restrictions.

Overall, by integrating empirical Bayes correction with structural demand estimation and causal inference, this study advances the methodological frontier of recreation demand analysis and offers a replicable framework for quantifying the welfare consequences of environmental disruptions using large-scale mobility data.

References

[1] Abadie, A. (2021). Using synthetic controls: Feasibility, data requirements, and methodological aspects. *Journal of Economic Literature*, 59(2):391–425.

[2] Alvarez, S., Larkin, S. L., Whitehead, J. C., and Haab, T. (2014). A revealed preference approach to valuing non-market recreational fishing losses from the deepwater horizon oil spill. *Journal of Environmental Management*, 145:199–209.

[3] Arkhangelsky, D., Athey, S., Hirshberg, D. A., Imbens, G. W., and Wager, S. (2021). Synthetic difference-in-differences. *American Economic Review*, 111(12):4088–4118.

[4] Barron, M. G., Vivian, D. N., Heintz, R. A., and Yim, U. H. (2020). Long-term ecological impacts from oil spills: comparison of exxon valdez, hebei spirit, and deepwater horizon. *Environmental Science & Technology*, 54(11):6456–6467.

[5] Berry, S., Linton, O. B., and Pakes, A. (2004). Limit theorems for estimating the parameters of differentiated product demand systems. *The Review of Economic Studies*, 71(3):613–654.

[6] Berry, S. T. (1994). Estimating discrete-choice models of product differentiation. *The RAND Journal of Economics*, pages 242–262.

[7] Berry, Steven T., James Levinsohn, and Ariel Pakes. (1995). Automobile Prices in Market Equilibrium. *Econometrica* 63 (4): 841–890.

[8] Bradt, J. T. (2025). Hotelling meets wright: Spatial sorting and measurement error in recreation demand models. *Journal of the Association of Environmental and Resource Economists*, 12(6):1563– 1600.

[9] Cano-Urbina, J., Clapp, C. M., and Willardsen, K. (2019). The effects of the bp deepwater horizon oil spill on housing markets. *Journal of Housing Economics*, 43:131–156.

[10] Carson, R. T. (2003). *Contingent valuation: A comprehensive bibliography and history*. Edward Elgar Publishing.

[11] CBS News (2023a). Amplify energy and its subsidiaries to pay $13 million for huntington beach oil spill. Available at: https://www.cbsnews.com/losangeles/news/amplify-energy-subsidiaries-13-million-oil-spill-huntington-beach/.

[12] CBS News (2023b). Huntington beach oil spill could have major long-term impacts on local wildlife, experts say. Avaliable at: https://www.cbsnews.com/losangeles/news/ huntington-beach-oil-spill-could-have-major-long-term-impacts-on-local-wildlife-experts-say.

[13] Cheng, N., Li, M., Liu, P., Luo, Q., Tang, C., and Zhang, W. (2024). Pipeline incidents and property values. *Journal of Environmental Economics and Management*, 127:103041.

[14] Ciccia, D. (2024). A short note on event-study synthetic difference-in-differences estimators. City of Newport Beach (2021). Oil spill response - ocean safety. Available at: https://www.newportbeachca.gov/trending/ocean-safety/oil-spill-reponse.

[15] D'Andrea, M. A. and Reddy, G. K. (2014). Crude oil spill exposure and human health risks. *Journal of Occupational and Environmental Medicine*, 56(10):1029–1041.

[16] Dubé, J.-P., Hortaçsu, A., and Joo, J. (2021). Random-coefficients logit demand estimation with zero-valued market shares. *Marketing Science*, 40(4):637–660.

[17] Dundas, S. J. and Von Haefen, R. H. (2020). The effects of weather on recreational fishing demand and adaptation: implications for a changing climate. *Journal of the Association of Environmental and Resource Economists*, 7(2):209–242.

[18] D'Andrea, M. A. and Reddy, G. K. (2018). The development of long-term adverse health effects in oil spill cleanup workers of the deepwater horizon offshore drilling rig disaster. *Frontiers in Public Health*, 6:117.

[19] Earle, A. and Kim, H. (2024). *Causal inference, high-frequency data, and the recreational value of water quality*. Department of Economics, East Carolina University and Department of Food and Resource Economics, Korea University.

[20] Egan, A. L., Chilvers, B. L., and Cassells, S. (2024). Expected direct costs of an oil spill in a unesco world heritage area in new zealand. *Australian Journal of Agricultural and Resource Economics*, 68(3):608–627.

[21] English, E., von Haefen, R. H., Herriges, J., Leggett, C., Lupi, F., McConnell, K., Welsh, M., Domanski, A., and Meade, N. (2018). Estimating the value of lost recreation days from the deepwater horizon oil spill. *Journal of Environmental Economics and Management*, 91:26–45.

[22] FCC (2023). Fcc broadband map – mobile coverage information. Available at: https://broadbandmap.fcc.gov/location-summary/mobile?version=dec2023&lon=-117.807662&lat=33.550761&addr_full=33.550761%2C+-117.807662&zoom=15.00&env=0&tech=tech4g.

[23] Gandhi, A., Lu, Z., and Shi, X. (2023). Estimating demand for differentiated products with zeroes in market share data. *Quantitative Economics*, 14(2):381–418.

[24] Glasgow, G. and Train, K. (2018). Lost use-value from environmental injury when visitation drops at undamaged sites. *Land Economics*, 94(1):87–96.

[25] Grossman, V., Martínez-García, E., Torres, L. B., and Sunc, Y. (2019). Drilling down: The impact of oil price shocks on housing prices. *The Energy Journal*, 40(2_suppl):59–84.

[26] Herrnstadt, E. and Sweeney, R. L. (2024). Housing market capitalization of pipeline risk: Evidence from a shock to salience and awareness. *Land Economics*, 100(4):652–667.

[27] Jackson, J. B., Cubit, J. D., Keller, B. D., Batista, V., Burns, K., Caffey, H. M., Caldwell, R. L., Garrity, S. D., Getter, C. D., Gonzalez, C., et al. (1989). Ecological effects of a major oil spill on panamanian coastal marine communities. *Science*, 243(4887):37–44.

[28] Knittel, C. R., Li, J., and Wan, X. (2023). *I love that dirty water? value of water quality in recreation sites*. Working Paper.

[29] Kubo, T., Uryu, S., Yamano, H., Tsuge, T., Yamakita, T., and Shirayama, Y. (2020). Mobile phone network data reveal nationwide economic value of coastal tourism under climate change. *Tourism Management*, 77:104010.

[30] Lee, S., Wan, X., and Zheng, S. (2023). *Estimating the indirect cost of floods: Evidence from high-tide flooding*. SSRN Working Paper.

[31] Li, J. (2019). Compatibility and investment in the us electric vehicle market. *Unpublished manuscript*, MIT.

[32] Thompson, Andrea. (2010). FAQ: The science and history of oil spills. Live Science. Accessed: https://www.livescience.com/9885-faq-science-history-oil-spills.html.

[33] Lopes, A. F. and Whitehead, J. C. (2023). Estimating the ex ante recreational loss of an oil spill using revealed and stated preference data. *Land Economics*, 99(4):490–508.

[34] Loureiro, M. L. (2009). Estimating the economic impacts of oil spills in the coast: The prestige oil spill case study. *Environmental Resource Economics*, 44:537–553.

[35] Lu, J., Huang, X., Kupfer, J. A., Xiao, X., Li, Z., Wei, H., Wang, S., and Zhu, L. (2023). Spatial, temporal, and social dynamics in visitation to us national parks: A big data approach. *Tourism Management Perspectives*, 48:101143.

[36] Lupi, F., Phaneuf, D. J., and von Haefen, R. H. (2020). Best practices for implementing recreation demand models. *Review of Environmental Economics and Policy* 14 (2): 302–323. https://doi.org/10.1093/reep/reaa007.

[37] Lyon, S. F., Merrill, N. H., Mulvaney, K. K., and Mazzotta, M. J. (2018). Valuing coastal beaches and closures using benefit transfer: an application to barnstable, massachusetts. *Journal of Ocean and Coastal Economics*, 5(1):1.

[38] Meo, S. A., Al-Drees, A. M., Rasheed, S., Meo, I. M., Khan, M. M., Al-Saadi, M. M., and Alkandari, J. R. (2009). Effect of duration of exposure to polluted air environment on lung function in subjects exposed to crude oil spill into sea water. *International Journal of Occupational Medicine and Environmental Health*, 22(1):35–41.

[39] Merrill, N. H., Atkinson, S. F., Mulvaney, K. K., Mazzotta, M. J., and Bousquin, J. (2020). Using data derived from cellular phone locations to estimate visitation to natural areas: An application to water recreation in new england, usa. *PLOS One*, 15(4):e0231863.

[40] Morris Jr, J. G., Grattan, L. M., Mayer, B. M., and Blackburn, J. K. (2013). Psychological responses and resilience of people and communities impacted by the deepwater horizon oil spill. *Transactions of the American Clinical and Climatological Association*, 124:191.

[41] Murdock, J. (2006). Handling unobserved site characteristics in random utility models of recreation demand. *Journal of Environmental Economics and Management*, 51(1):1–25.

[42] Murphy, J. J. and Delucchi, M. A. (2005). A review of environmental damage estimation studies from oil spills. *Resource and Energy Economics*, 27(3):253–273.

[43] Newbold, S. C., Lindley, S., Albeke, S., Viers, J., Parsons, G., and Johnston, R. (2022). *Valuing satellite data for harmful algal bloom early warning systems*. Technical report, RFF working paper WP 22.

[44] NOAA (2021a). Huntington beach oil spill: Response and restoration. Available at: https://response.restoration.noaa.gov/oil-and-chemical-spills/significant-incidents/ huntington-beach-oil-spill.

[45] NOAA (2021b). Pipeline P00547 huntington beach oil spill. Available at: https://darrp.noaa. gov/oil-spills/pipeline-p00547-huntington-beach-oil-spill.

[46] Taylor, L. O., Phaneuf, D. J., and Liu, X. (2016). Disentangling property value impacts of environmental contamination from locally undesirable land uses: Implications for measuring post-cleanup stigma. *Journal of Urban Economics*, 93:85–98.

[47] Timmins, C. and Murdock, J. (2007). A revealed preference approach to the measurement of congestion in travel cost models. *Journal of Environmental Economics and Management*, 53(2):230– 249.

[48] Tourangeau, R., English, E., McConnell, K. E., Chapman, D., Cervantes, I. F., Horsch, E., Meade, N., Domanski, A., and Welsh, M. (2017). The gulf recreation study: Assessing lost recreational trips from the 2010 gulf oil spill. *Journal of Survey Statistics and Methodology*, 5(3):281–309.

[49] Whitehead, J. C., Haab, T., Larkin, S. L., Loomis, J. B., Alvarez, S., and Ropicki, A. (2018). Estimating lost recreational use values of visitors to northwest florida due to the deepwater horizon oil spill using cancelled trip data. *Marine Resource Economics*, 33(2):119–132.

[50] Winkler, D. T. and Gordon, B. L. (2013). The effect of the bp oil spill on volume and selling prices of oceanfront condominiums. *Land Economics*, 89(4):614–631.

[51] Zhang, W., Wan, X., Fan, W., and Ji, Y. (2024). Uncovering disparities in water-based outdoor recreation using cell phone mobility data. *Environmental Research Letters*, 19(11):114057.

Main Figure

Figure 1: Identified Affected and Unaffected Beaches

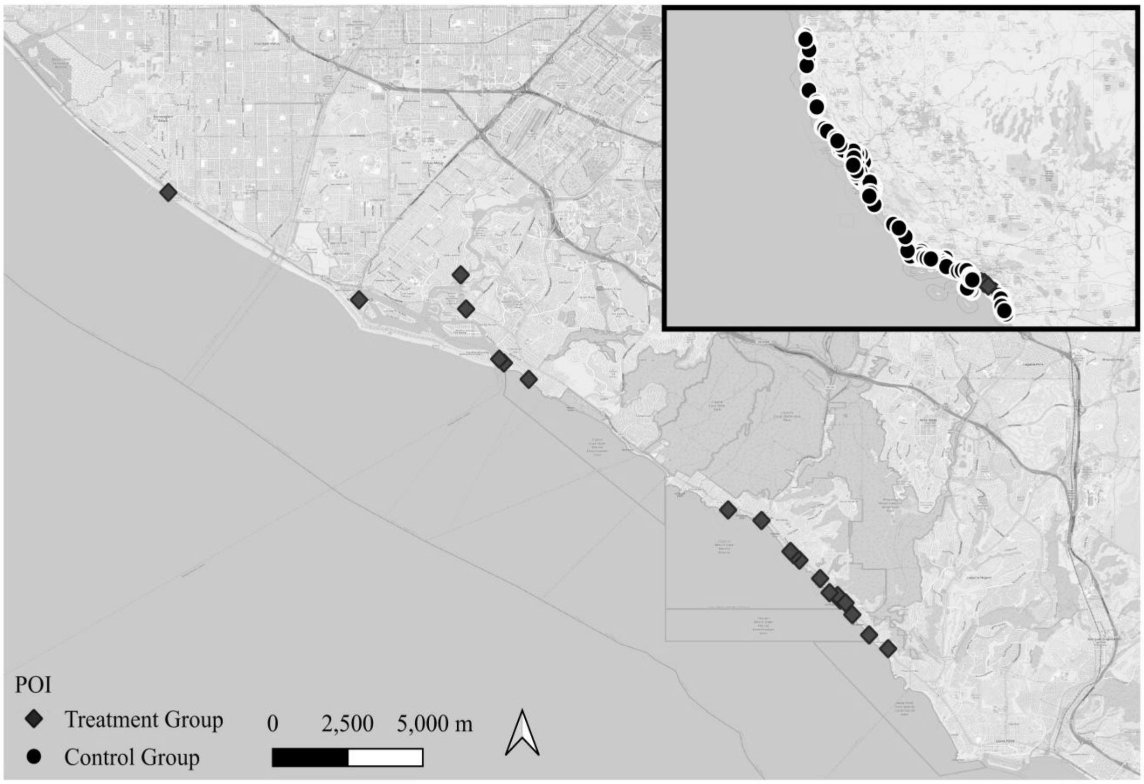

Notes: Affected beaches (diamond) were identified based on proximity to the spill site and the southeastward spread of oil, reaching Dana Point. Beaches within 43 km southeast of the spill were classified as affected, while unaffected beaches (circle) serve as the control group. Long Beach was excluded despite its proximity, as it lies north of the spill and was not exposed. A total of 21 beaches were affected (1 in Huntington Beach, 5 in Newport Beach, and 15 in Laguna Beach), while 458 unaffected beaches remain in the control group. Beach POIs smaller than 500 square meters were excluded.

Figure 2: The Difference between Observed and SDID Mean Utilities

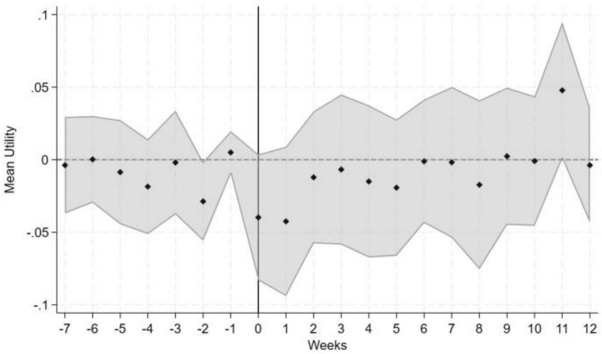

(a) All Beaches

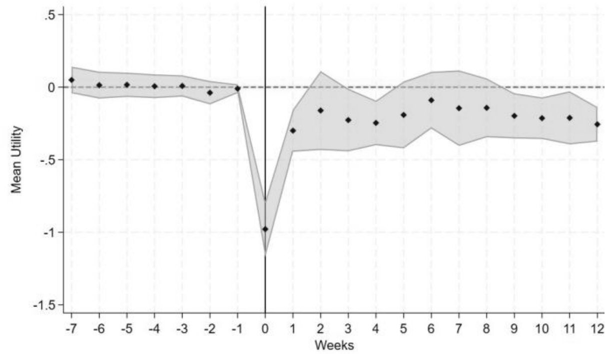

(b) Huntington Beach Cluster

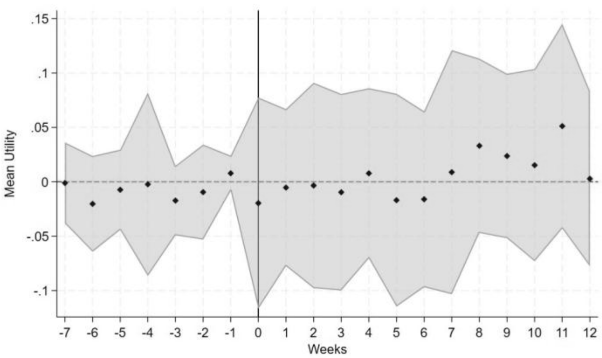

(c) Newport Beach Cluster

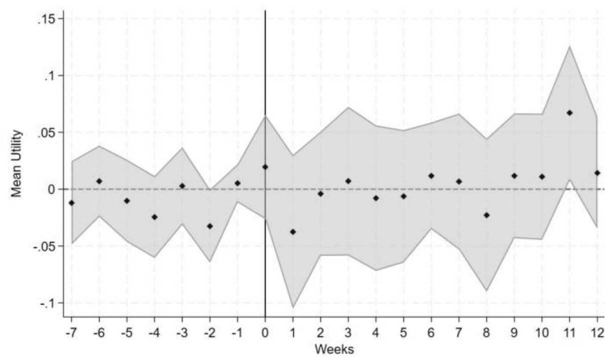

(d) Laguna Beach Cluster

Notes: This figure presents the difference between observed and SDID-estimated mean utilities across all beaches and for Huntington, Newport, and Laguna Beach separately. The vertical line at week 0 represents the week of beach closure. Blue dots indicate estimated differences, and the shaded area represents 95% confidence intervals.

Main Table

Table 1: Summary Statistics

Panel A: 1st Stage	Obs.	Mean	Std. Dev.	Min.	Max.
Visits	3,739,601	0.32	7.62	0.00	2,579.00
Travel Dist. (mile)	3,739,601	56.14	56.92	0.02	299.95
Travel Time (h)	3,739,601	1.32	1.34	0.00	7.06
Gas Costs (cent/mile)	3,739,601	18.65	0.00	18.65	18.65
Maintenance (cent)	3,739,601	9.55	0.00	9.55	9.55
Depreciation (cent)	3,739,601	26.00	0.00	26.00	26.00
Travel Costs ($)	3,739,601	67.90	68.77	0.03	399.79
Observed Share	3,739,601	0.00060	0.00529	0.00000	0.97818
Bayes Share	3,739,601	0.00059	0.00495	4.03e-16	0.97769
Panel B: 2nd Stage					
Mean Utility	9,580	-0.2546	0.4418	-1.9661	3.3783
Treatment Dummy	9,580	0.0438	0.2047	0	1
Post Dummy	9,580	0.6500	0.4769	0	1
Huntington-Week 39 Dummy	9,580	0.0001	0.0102	0	1

Notes: This table presents summary statistics for key variables used in the analysis at CBG-POI weekly level in Panel A and beach weekly level in Panel B. "Visits" represents the number of observed visits to beaches. "Observed share" refers to the direct market share of each beach, while "Bayes share" represents the adjusted market share incorporating empirical Bayes priors. "Mean Utility" refers to the mean utility derived from stage 1. "Treatment Dummy" assigns 1 to affected beaches and 0 to those unaffected. "Post Dummy" presents the periods before (assigned 0) and after (assigned 1) 2021 Huntington oil spill. "Huntington-Week 39 Dummy" captures the impacts of Airforce show happened in Week 39 at Huntington Beach.

Table 2: Demand Estimation Results

Panel A: 1st Stage	OLS logit	IV logit	R.C. logit	
	(1)	(2)	(3)	
Travel Costs	-0.0053***	-0.0053***	-0.0052***	
	(0.0001)	(0.0001)	(0.0001)	
Travel Costs (sd)			-0.0002 (0.0010)	
Observation	3,739,601	3,739,601	3,739,601	
Kleibergen-Paap		3,530		
GMM			23,162	
Panel B: 2nd Stage	**All Beaches**	**Huntington Beach**	**Newport Beach**	**Laguna Beach**
	(1)	(2)	(3)	(4)
Closure	-0.0085	-0.2584***	0.0055	0.0054
	(0.0209)	(0.0657)	(0.0251)	(0.0206)
Observation	9,580	9,180	9,260	9,460

Notes: This table presents first-stage and second-stage demand estimation results. Panel A compares OLS logit, IV logit, and random coefficient logit models in the first-stage estimation. Panel B employ the SDID approach to estimate the beach closure impacts on mean utility. Column (1) use all affected beaches as the treatment group; Column (2)-(3) includes Huntington, Newport, and Laguna beach as the treatment group. Standard errors are in parentheses. *** p < 0.01, ** p < 0.05, * p < 0.1.

Table 3: Welfare Effects of Beach Closures

Panel A: All Beaches	Value	Lower bound	Upper bound
Aggregate Loss ($M)	0.993	-0.113	2.322
Loss per closure week ($M)	0.103	-0.000	0.231
Loss per week after initial closure ($M)	0.083	-0.009	0.194
Panel B: Huntington Beach			
Aggregate Loss ($M)	1.002	0.031	2.187
Loss per closure week ($M)	0.100	0.008	0.215
Loss per week after initial closure ($M)	0.084	0.003	0.182

Notes: This table reports estimated welfare losses from beach closures using SDID method in stage 2. Lower and upper bounds are calculated using 95% confidence interval. Panel A presents results using all beaches in the second stage while Panel B reports results using only Huntington Beach.