Evaluating Precision, Privacy, and Representation with Cell Phone Data: Evidence from Cape Cod Beaches

Cristina Connolly Sandro Steinbach Mike Vo Xibo Wan[1]

Abstract

While mobility data has emerged as a promising alternative for assessing the economic value of recreation, their reliability depends critically on how data are processed and protected. This paper systematically evaluates how key data-handling practices affect the accuracy of recreation demand estimation. Using a random utility model to analyze recreation visits at Cape Cod beaches from 2019-2022, we evaluate how methodological choices influence the marginal willingness to pay (MWTP) for avoiding fecal bacteria contamination. We find an average MWTP of $8.92 per visit when using the proposed practices, such as refined visit definitions, sampling weights, and long-term choice sets. Deviations from these practices can introduce significant biases: relaxing the minimum dwell time and applying differential privacy reduce MWTP by 57% and 65%, respectively, while short-term choice set definition inflates it by 10%. By demonstrating the sensitivity of welfare estimates to data-processing decisions, this study highlights the importance of transparent and judicious mobility data practices for credible environmental valuation and evidence-based policymaking.

Cell Phone Mobility Data, Data Practices, Recreational Demand, Differential Privacy

C55, C81, C83, Q26

1 Introduction

Non-market valuation studies have long relied on survey-based approaches to estimate the economic value of outdoor recreation (Lupi et al. 2020). While these methods provide direct insights into individual preferences and behaviors, they suffer from several well-documented limitations. Traditional survey approaches often generate localized estimates that fail to

1 : Postdoctoral Associate, Department of Agricultural and Resources Economics, University of Connecticut, xibo.wan@uconn.edu, Corresponding Author; Mike Vo: Postdoctoral Associate, Department of Agricultural and Resources Economics, University of Connecticut, mike.vo@uconn.edu; Cristina Connolly: Assistant Professor, Department of Agricultural and Resources Economics, University of Connecticut, cristina.connolly@uconn.edu; Sandro Steinbach: Associate Professor & Director for Agricultural Policy and Trade Studies, Department of Agribusiness and Applied Economics, North Dakota State University, sandro.steinbach@ndsu.edu.

doi:10.3368/le.102.3.012126-0010

generalize to larger regions, limiting their applicability for policy-making at broader scales (Loomis et al. 1995; Rolfe et al. 2015; Rosenberger and Loomis 2017). Furthermore, survey data are prone to selection bias and recall errors, as respondents may misreport their past visits or systematically differ from non-respondents (Dillman 2017; Tarrant et al. 1993; Rylander et al. 1995; Connelly and Brown 2011). Additionally, the high cost and time-intensive nature of survey data collection constrain the feasibility of conducting frequent and large-scale studies. Another critical challenge in survey based recreation demand estimation is the potential endogeneity of individual choice sets. Researchers typically assume a set of alternative sites, but variations in how choice sets are defined can lead to significantly different welfare estimates associated with environmental changes and policy interventions (Parsons and Hauber 1998; Parsons et al. 2000; Parsons et al. 2021).

Mobility data, which are location information collected from GPS-enabled devices or cell phone tracking, have emerged as a promising alternative for assessing the economic value of recreation sites (Merrill et al. 2020; Gonyo et al. 2024; Newbold et al. 2022; Kubo et al. 2020; Zhang et al. 2024; Lu et al. 2023). An example of such data includes aggregated foot traffic datasets from companies like Veraset, or Advan (Previously SafeGraph), which track anonymized device movements to recreational sites such as national parks, beaches, and hiking trails. These datasets offer several advantages over more traditional surveys. First, they provide large-scale, high-frequency data that enable a more comprehensive and cost-effective analysis of recreational behavior across broad geographic regions. Second, mobility data capture actual visitation patterns rather than hypothetical responses, improving the reliability of behavioral estimates. Third, individual choices are directly observed, reducing some of the biases inherent in self-reported data. As a result, mobility data have the potential to enhance the accuracy and spatial coverage of recreation demand models.

While high-resolution data on recreation visits exist for specific sites, such as through park permits, visitor surveys, or administrative counts, these sources are fragmented and inconsistent across regions. Mobility data uniquely provide nationwide coverage at high spatial and temporal resolution, enabling direct comparison of welfare outcomes across thousands of locations while retaining site-level behavioral detail. This scalability makes mobility data particularly valuable for identifying generalizable patterns in recreation demand and welfare estimation. Despite these advantages, mobility data introduce their own methodological complexities that require careful consideration before being used for welfare analysis.

We categorize these challenges into two broad categories. The first involves classic problems familiar to all recreation demand studies, including defining the relevant choice set, determining the appropriate level of trip aggregation, and ensuring the representativeness of the sample. Mobility data address these long-standing issues by providing finer spatial and temporal detail, however, they also introduce new forms of selection and coverage bias. For example, defining realistic choice sets from revealed mobility behavior requires balancing observed visitation histories with the potential for unobserved but feasible alternatives. Similarly, while large-scale data alleviate small-sample limitations, they may still underrepresent specific demographic groups, such as older adults or lower-income individuals, due to differences in smartphone usage (Cook 2025; Gonzalez et al. 2008).

The second category encompasses new problems unique to mobility data. First, researchers must infer trip purpose and recreational intent from raw pings without direct information about motivation or activity type, raising questions about what constitutes a true recreation visit (Gibbs et al. 2023). Second, privacy-preserving mechanisms, particularly differential privacy, introduce structured statistical noise and truncation that distort true visitation frequencies (Savi et al. 2023).[2] These privacy adjustments, while essential for protecting user anonymity, can bias recreation demand estimates in systematic ways that differ from traditional random measurement error. Understanding how such data-processing decisions affect model outcomes is critical for reliable welfare analysis.

This paper systematically examines how these new and classic problems affect the accuracy of recreation demand estimation. We develop and test a set of proposed practices for processing cell phone data, such as defining visits with a minimum dwell time threshold, consolidating multiple stops within a day, constructing long-term choice sets, and applying sampling weights to address sampling representativeness, and compare them against alternative approaches that relax or modify these assumptions. Using a case study of beach recreation on Cape Cod, Massachusetts, we examine how each data-handling decision affects estimated demand parameters and welfare measures, focusing on the marginal willingness to pay (MWTP) for avoiding fecal bacteria contamination advisories.[3]

Cape Cod is an ideal setting for this analysis due to its popularity as a recreational destination, with visitors frequently traveling long distances to access its diverse outdoor opportunities. Its extensive coastline, hiking and biking trails, and other attractions draw a high volume of visitors, making it a valuable case for studying demand estimation.[4] Moreover, the availability of high-quality environmental and visitation data allows for a rigorous examination of how changes in site conditions, such as water quality or beach advisories, shape recreational choices.

We focus on visitors' choices for Cape Cod beaches from 2019 to 2022 using large-scale mobility data from Veraset.[5] We estimate recreational demand using a multinomial logit model, which enables us to also identify visitors' preferences. To examine the effects of fecal bacteria contamination advisories on recreation demand, we integrate contamination advisories from the Beach Advisory and Closing Online Notification system (BEACON 2.0), spatially linking them to Cape Cod beaches and estimating the MWTP for avoiding contamination events. By systematically evaluating the effects of precision, privacy, and representation in cell phone data, through modifications such as adding measurement noise, restricting choice set definitions, and

[2] To protect user anonymity, mobility data vendors implement privacy-preserving procedures that modify raw visit counts before release. These include techniques such as differential privacy and statistical noise injection, which intentionally obscure individual-level information.
[3] During our study period, no formal beach closures were recorded in Cape Cod. All observed contamination events were issued as water quality advisories rather than closures. Accordingly, our analysis focuses exclusively on advisory effects.
[4] Cape Cod is known for its extensive ocean-facing coastline, accessible public beaches, and diverse outdoor recreation opportunities from hiking and biking trails to beachfront resorts and yachting. See more details in here.
[5] Veraset is a data provider specializing in anonymized population movement data collected through GPS signals from mobile devices across the United States. Each data entry includes a unique device identifier, latitude, longitude, and timestamp information, enabling a detailed examination of movement patterns and outdoor recreation behaviors.

applying different sampling strategies, we assess how these factors influence demand estimates and welfare measures.

We begin by laying out a set of *proposed practices* to operationalize the mobility data and estimate recreation demand. These practices include defining a visit based on a specific minimum dwell time of ten minutes, treating multiple stops within a day as a single visit, constructing individual choice sets based on long-term visitation patterns, and applying sampling weights to correct for representativeness issues.[6] These choices are designed to minimize potential biases from mobility data and improve the reliability of recreational demand estimation. We then systematically experiment with alternative practices by modifying each factor individually. This includes redefining visits with a fixed shorter threshold, considering each stop as an independent visit, assuming choice sets based on short-term behavior, omitting sampling weights, and implementing differential privacy techniques to protect user data.[7] By comparing results across these practices, we quantify the sensitivity of recreation demand models to different data-handling approaches and offer insights into enhancing estimation accuracy. Table 1 provides a complete side-by-side summary of these approaches.

Our findings reveal that using the proposed practices, the estimated average MWTP for avoiding contamination is \$8.92 per visit. However, deviations from these proposed practices significantly impact the estimates, leading to a 52% underestimation when applying alternative methods. Specifically, relaxing the minimum dwelling requirement for defining a visit results in a 58% reduction in MWTP, while implementing differential privacy procedures lowers the estimate by 65%, highlighting the sensitivity of demand models to these factors. In contrast, assuming a shorter-term choice set inflates the MWTP estimate by 10%, suggesting that restrictive definitions of available alternatives may overstate site preferences. Notably, treating each stop as a separate visit and omitting sampling weights have minimal effects on the MWTP estimate, indicating that these factors are less critical in shaping demand outcomes, at least in our samples.

Our paper contributes to the literature in four ways. First, our work complements the literature on the impact of water quality on beach visits by providing empirical evidence on how contamination advisories influence beach selection using high-frequency mobility data. Previous studies have consistently shown that deteriorating water quality reduces beach visitation, with substantial economic implications for coastal recreation (Murray et al. 2001; Parsons et al. 2009). More recent research, such as Merrill et al. (2020), Mazzotta et al. (2022), and Furey et al. (2022), highlights the importance of water quality in shaping regional beach visitation patterns, demonstrating that short-term water pollution advisories deter visitors. Our findings suggest an average MWTP of \$8.92 per visit to avoid contamination, highlighting the economic value of water quality improvements and emphasizing the need for effective environmental management policies to reduce contamination risks.

[6] These practices align with established standards in mobility data processing and recreation demand modeling. Approaches like including defining visits based on a specific minimum dwell time to capture beach activities, constructing individual choice sets using long-term visitation patterns, and applying sampling weights to adjust for representativeness issues, are widely adopted by researchers in recreation demand modeling to enhance the accuracy and reliability of site choice estimation (Hindsley 2011; Phaneuf and Smith 2015; DeShazo et al. 2003).

[7] Treating each stop within a day as a single visit and applying differential privacy protections are standard practices used by Advan (Data 2025). Setting a minimum dwell time threshold, such as four minutes, is a common practice employed by SafeGraph to filter meaningful visits from incidental movement.

Second, this paper contributes to the growing literature on the application of cell phone mobility data in non-market valuation (Merrill et al. 2020; Newbold et al. 2022; Kubo et al. 2020; Zhang et al. 2024; Lee et al. 2023; Lu et al. 2023; Cheng and Wan 2025). While traditional surveys have long been the standard for estimating recreational preferences, mobility data provide large-scale, high-frequency observations of actual visitation behavior. By leveraging this data, we are able to capture real-world site choices at a finer resolution and estimate economic values associated with environmental changes more accurately. However, despite these advantages, mobility data also introduce challenges related to privacy protections, measurement errors, sampling biases, and choice set definitions, all of which can influence demand estimation outcomes. Our study systematically evaluates these trade-offs, offering insights into how data processing choices affect recreation demand modeling.

Third, our work adds to the literature on best data practices in recreation demand modeling (Lupi et al. 2020; Englin et al. 2003; Cameron et al. 1996) by offering practical guidance on the sensitivity of data practices in mobility-based recreation demand modeling. By systematically evaluating the effects of measurement errors, privacy protections, choice set definitions, and sampling biases, we identify the conditions under which mobility data yield reliable welfare estimates. Our findings emphasize the importance of defining visits appropriately, accounting for long-term choice sets, and addressing biases introduced by differential privacy. These insights help inform future studies seeking to use mobility data for non-market valuation and improve the robustness of demand estimation techniques.

Fourth, this study contributes to the literature of how privacy protections impact data accuracy and economic valuation (Acquisti et al. 2016; Totty and Watson 2024; Santos-Lozada et al. 2020; Kenny et al. 2021; Hauer and Santos-Lozada 2021; Winkler et al. 2021). As differential privacy and other anonymization techniques become standard in large-scale behavioral datasets, their unintended consequences on demand estimation remain under-explored. Our findings demonstrate that privacy-preserving procedures can substantially bias willingness-to-pay estimates, leading to significant underestimations of recreational value. By quantifying these effects, we highlight the need for careful consideration of privacy adjustments in policy-relevant economic research to balance privacy protection with data accuracy in non-market valuation studies.

The remainder of this paper is structured as follows. Section 2 introduces the study area and data, and discusses the proposed data processing practices and alternative approaches. Section 3 provides descriptive evidence on how these varying data practices influence visit counts and site selection patterns. Section 4 presents the recreation demand model and describes the estimation of MWTP for avoiding environmental hazards. Section 5 reports the main results, comparing how different data practices impact demand estimates. Section 6 discusses the implications of our findings for non-market valuation and provides recommendations for future applications of mobility data. Finally, Section 7 concludes with key takeaways and potential directions for further research.

2 Study Area, Data Collection, and Data Processing

2.1 Study Area: Cape Cod Beaches

Cape Cod, Massachusetts, serves as our setting for analyzing recreation demand due to its diverse recreational opportunities, substantial visitor base, and well-documented environmental conditions. The region is a major coastal destination known for its extensive shoreline, picturesque landscapes, and a range of outdoor activities, including beachgoing, hiking, and biking (Bartlett 2019). Its accessibility from major metropolitan areas, particularly Boston and New York, contributes to a steady influx of visitors with varied demographic and socioeconomic backgrounds, providing a reliable setting for studying differences in recreational preferences and demand patterns.

The region's coastal features include a mix of public and private beaches, estuaries, and protected natural areas, creating a range of site characteristics that influence visitor behavior. Cape Cod's reliance on seasonal tourism further highlights the importance of understanding demand fluctuations across time, as well as how visitors respond to factors such as congestion, accessibility, and site amenities. The diversity of recreational settings, from highly developed beaches with extensive facilities to remote, natural shorelines, allows for an assessment of how site-specific attributes drive visitation choices.

Cape Cod also presents a critical opportunity to examine the role of environmental conditions in shaping recreational decisions. Periodic beach closures due to bacterial contamination, algal blooms, and other water quality concerns have raised concerns about the economic impact of environmental degradation on tourism. The availability of high-resolution environmental monitoring data on water quality, beach conditions, and ecological health facilitates a rigorous evaluation of how site attributes affect visitation patterns and recreational welfare (Merrill et al. 2022; Mazzotta et al. 2022; Furey et al. 2022). These factors make Cape Cod a valuable case study for quantifying the economic benefits of maintaining coastal environmental quality.

2.2 Data Collection

Our analysis relies on multiple data sources to track recreational visits to Cape Cod beaches from 2019 to 2022 and assess how contamination advisories influence beach selection. The primary dataset comes from GPS-based location mobility data, complemented by environmental, weather, and contamination advisory datasets to enhance the accuracy of demand estimation.

The mobility data are sourced from Advan and Veraset. Using Advan, we identify points of interest (POIs) classified as nature parks (NAICS 712190) within Cape Cod and refine the selection to include only those explicitly designated as beaches, resulting in a final sample of 155 Cape Cod beaches.[8] The POI polygons are constructed to capture visits to Cape Cod beaches, which exclude adjacent non-recreational areas like roads, homes, or businesses. While polygon overlap can be a concern in dense urban settings, the distinct boundaries of outdoor sites like beaches make such misclassification less likely in our context.

[8] Advan provides detailed polygon geometries that delineate the actual physical footprint of each beach, rather than generic circular buffers.

It is important to note that the scope of POI classifications can vary across mobility data providers. For example, Furey et al. (2022) report 465 coastal access points in Cape Cod from a different provider, but that dataset includes a much broader set of water access locations such as inland freshwater ponds, estuarine inlets, river mouths, and small boat landings that are not beaches. Our sample is intentionally narrower because our research focuses specifically on beach recreation behavior that is linked to marine water quality monitoring. To ensure our sample was appropriate, we cross-referenced our 155 Advan beach polygons with the U.S. EPA's Beach Advisory and Closing Online Notification (BEACON 2.0) system, which is the source of our contamination data.

Advan provides detailed POI polygons, allowing us to link each location to visitor counts, forming the foundation of our recreational demand analysis. Veraset provides individual-level mobility data, capturing daily visits from home census block groups (CBGs) to the selected beaches. These data allow us to track travel patterns, infer visitor origins, and assess how different groups respond to contamination advisories. To ensure trips represent recreational visits rather than flight travel, we restrict the sample to individuals traveling within a 300-mile radius of Cape Cod. This threshold aligns with Dundas and Haefen (2020), who assume that a 300-mile distance (6-hour drive one way) represents the upper limit for a feasible single-day recreational trip, which is the focus of our analysis.[9] As a sensitivity check on this distance assumption, we also test alternative distance cutoffs, which we report in the Appendix. We then aggregate the visit counts to CBG-beach monthly level for our demand estimation.

We integrate contamination advisory data from the BEACON 2.0 system, which documents water quality alerts and beach closures due to contamination events.[10] BEACON provides monitoring-site point coordinates rather than polygons, so we perform a spatial join between BEACON's beach points and a 50-meter buffer around each Advan beach polygon to ensure reliable matching. This ensures that each BEACON beach is reliably linked to its corresponding Advan POI. We then use the BEACON Beach ID to connect each advisory record to the relevant beach, allowing us to track the effect of contamination warnings on visitation patterns. We generate an advisory dummy if a contamination advisory occurred in a given month for a beach. These advisories are critical for evaluating how environmental hazards influence beach selection and estimating the MWTP for avoiding contaminated sites. We acknowledge that an advisory issued late in one month may primarily affect visitation in the following month. To account for this, we also conduct a robustness check using a one-month lagged advisory dummy, which we report in the 10.

To incorporate environmental conditions, we use PRISM climate data, which provides 1km grid-level monthly precipitation and temperature measures. We match these data to the 155 selected beaches by assigning weather observations to each beach's geographic centroid. This process enables us to account for weather-related factors that influence visitation behavior, capturing both seasonal trends and short-term fluctuations in recreational demand.

---

[9] Some of the visits to Cape Cod are part of multi-day trips, which involve more complex decisions (e.g., accommodation choice) that are beyond the scope of this methodological paper.
[10] The BEACON 2.0 system, managed by the U.S. Environmental Protection Agency (EPA), provides data on beach water quality monitoring and contamination advisories. Access it here: BEACON 2.0.

In addition to mobility, environmental, and advisory data, we incorporate economic and demographic information at the CBG level to evaluate the representativeness of the mobility sample. Using data from the American Community Survey (ACS), we examine key socioeconomic indicators, including median household income, educational attainment, racial and ethnic composition, and age distribution. This allows us to assess sampling rates by comparing the share of observed mobile device users in each CBG to the total population residing in that area. With these linkages and assumptions, we construct aggregated monthly data of actual beach choices for each CBG from 2019 to 2022.

2.3 Data Processing Practices

Accurately estimating recreation demand using mobility data requires defining visits, constructing appropriate choice sets, and ensuring representative sampling. Differences in these data-handling choices can introduce substantial variations in estimated visitation patterns and welfare measures. This section details our proposed practices for data processing, outlines alternative approaches for sensitivity analysis, and discusses their potential implications for demand estimation. Again, Table 1 provides a clear summary of these approaches.

2.3.1 Proposed data practices.

To minimize biases and enhance the reliability of demand estimation, we propose several data practices to processing mobility data. Visit identification follows a minimum dwell time threshold of 10 minutes, ensuring that brief stops unrelated to recreation are excluded. This threshold follows conventions used in prior mobility-based recreation studies (Liang et al. 2022) and represents a pragmatic balance: it is long enough to remove incidental pings from individuals merely passing through or stopping momentarily, yet short enough to preserve meaningful short-duration recreational visits. By filtering out these transient observations, we capture more purposeful site engagement, thereby improving the accuracy of our demand estimates.

The plausibility of this 10-minute cutoff is further supported by both external benchmarks and our own data. The American Time Use Survey (ATUS) reports that individuals engaging in "sports, exercise, or recreation" spend on average 81 minutes on weekdays and 105 minutes on weekends and holidaysconditional on participation (U.S. Bureau of Labor Statistics 2025). Most sessions last between one and two hours, and activities shorter than ten minutes are rare, indicating that our threshold effectively distinguishes incidental from intentional recreation.[11] Consistent with these patterns, our observed dwell-time distribution is strongly right-skewed, with a median of 41 minutes, an interquartile range of 12-151 minutes, and a long tail extending up to 24 hours (Figure 7). This distribution indicates that most visits are short to moderate in duration, while a smaller share of visitors engage in multi-hour stays, further validating the behavioral plausibility of our threshold.

While our dwell-time threshold filters out transient observations, mobility data may still reflect variation in how devices record location pings. For example, visitors who frequently check their phones may appear to stay longer than those who keep their phones inactive, potentially introducing measurement noise in observed dwell durations. Our minimum dwell-time rule is

[11] These benchmarks suggest that brief activities, especially those under ten minutes, reflect incidental rather than purposeful recreation, providing an empirical basis for our 10-minute cutoff.

designed to mitigate non-recreational noise rather than to precisely measure time on site. Robustness checks using alternative thresholds confirm that our welfare estimates remain stable across reasonable variations in this rule.

Because cellphone data do not include information on app-level ping frequency or device-use patterns, direct correction for this source of bias is not possible. We therefore assume broadly comparable ping dynamics across visitors, an assumption consistent with prior mobility-based recreation studies and supported by the stability of our results across alternative thresholds. Nonetheless, future research could further validate this assumption by linking mobility-derived visitation metrics to independent data sources, such as administrative counts or on-site survey observations.

To further refine visit counts, multiple stops within a day at the same recreation site are treated as a single visit. Visitors may leave and return for various reasons, such as parking constraints, meals, or short errands, which, if counted separately, would inflate visitation numbers and distort demand patterns. By consolidating these into a single visit, we ensure that the data more accurately reflect actual recreational behavior.

We construct choice sets at the census block group (CBG) level based on observed visitation patterns over the four-year window (2019-2022). Specifically, a beach is included in a given CBG's choice set if at least one mobile device from that CBG was observed visiting that beach during the sample period. The average CBG-specific choice set in our sample contains 8.95 beaches (median = 4; p25 = 2; p75 = 10; p90 = 29; max = 65). As is implicit in this method, if no device from a CBG visited a particular beach in our four-year sample, that beach is excluded from the choice set. We acknowledge this trade-off: while our approach is based on revealed behavior, it may omit sites that are feasible alternatives but were not visited by our specific sample of devices. The reason we choose this method is that recreation site selection is often influenced by familiarity and past experiences rather than being purely random (Manning et al. 2000). This approach reduces biases that could arise from assuming visitors consider only the sites they visited within a brief time window, which may not fully capture their broader recreational preferences. By extending the observation period to four years, we more accurately represent the range of beaches a visitor is likely to consider, improving the reliability of demand estimation.

We also construct and apply sampling weights to partially address potential representativeness issues in the mobility data. Because individual-level demographic attributes are not observable in the Veraset dataset, we cannot directly match devices to demographic characteristics. Instead, we adopt a first-order adjustment at the CBG level, comparing the number of unique devices identified as residents in each CBG to its total population from the American Community Survey (ACS). The resulting ratio, which is defined as the observed devices divided by total CBG population, serving as a proxy for relative coverage. Each observation is then weighted by the inverse of this coverage ratio, so that CBGs that are underrepresented in the data (i.e., with low observed device shares) receive slightly higher weight in the estimation. This adjustment is intended to correct for uneven data coverage (e.g., lower smartphone usage among older or lower-income groups), not for demographic composition per se.

We acknowledge that this approach has limitations. As the share of observed devices from any given CBG could be an outcome of mobility patterns and data coverage, this weighting variable may be potentially endogenous, and therefore cannot fully correct for underlying demographic biases. We therefore frame this weighting not as a definitive solution, but as a first-order adjustment to account for observable differences in device coverage across CBGs. More robust methods for inferring demographics and constructing weights from mobility data are an important direction for future research. Our analysis in the main text explicitly tests the impact of this weighting approach by comparing the results to an unweighted model.

2.3.2 Alternative data practices.

To assess the sensitivity of recreation demand estimates to different data-handling choices, we systematically modify individual components of the data processing approach and evaluate their impact on demand estimation and welfare measures. Visit definitions are adjusted by replacing the dwell time threshold with a four-minute threshold[12]. This change may lead to an increase in visit counts, particularly at sites with frequent short stops, but could also introduce noise by capturing non-recreational stops. A shorter threshold may overestimate demand, while a stricter threshold could exclude genuine visits, affecting the robustness of site-specific demand elasticities.

We next assess an alternative data-handling approach that treats multiple stops recorded by the same device at a recreation site within a single day as separate visits, rather than consolidating them into one. While counting each stop individually increases total visitation, this practice may overstate true visit frequency if multiple stops arise from signal fragmentation during a single continuous stay. To evaluate this, we analyze the temporal distribution of consecutive same-day stops. As shown in Figure 4, over 85 percent of devices record only one stop per day, and most consecutive stops occur within two hours of one another, suggesting that many multi-stop observations likely represent fragmented signals. Nonetheless, the hourly pattern indicates that a subset of users revisit later in the day, implying that some multi-stop cases correspond to genuine within-day returns. We therefore interpret that consolidating same-day stops modestly reduces total visit counts but yields a more behaviorally consistent measure of distinct recreation trips.

We also vary the choice set construction by basing site availability on short-term visit patterns (beaches visited by the same device in the same year) rather than long-term behavior from 2019 to 2022. This change alters the reference period used to define which beaches are considered in a visitor's choice set, potentially increasing the influence of recent site preferences while discounting long-term habitual behavior.

Next, we examine how removing sampling weights affects the accuracy of demand and welfare estimates. In our baseline specification, sampling weights are applied to prevent estimates from being dominated by origin areas with disproportionately large numbers of observed devices. When these weights are omitted, markets with higher sampling intensity contribute more heavily to the estimation, potentially biasing welfare measures if visitation responses differ systematically across origins. This analysis allows us to assess how uneven sampling across Census block groups influences estimated demand elasticities and willingness to pay for improved environmental conditions.

---

[12] This follows data practices used by SafeGraph, a cell phone data company, which applies a 4-minute minimum dwell time to define visits.

Finally, we evaluate how additional privacy protections influence the reliability of recreation demand estimation. Veraset provides individual-level mobility data that is already protected through anonymization, including the removal of personally identifiable information and only identifying home locations to the CBG level, rather than through precise latitude and longitude coordinates. In contrast, other data providers, such as Advan, use the same POIs and mobility data but further aggregate it and incorporate differential privacy techniques to obscure individual visit patterns in the final product they sell.[13] These additional measures introduce random noise, truncation, and censoring to visitation data, helping to protect user privacy but potentially reducing analytical precision.

To quantify the impact of such privacy safeguards, we apply industry-standard DP procedures to our Veraset data and compare the resulting estimates to the original anonymized dataset. Specifically, we introduce Laplace noise, apply truncation, and implement censoring to modify visit counts (Dwork et al. 2006). In Advan's implementation, Laplace noise is applied with parameters (0,10/9), meaning each visit count is perturbed by drawing from a Laplace distribution centered at zero with a scale parameter of 10/9. The adjusted visit count is then rounded up to ensure non-negative values. To further enhance privacy, visit counts that are zero or below two are replaced with a value of zero, preventing the identification of rare or unique visitors. Additionally, visit counts are censored at a minimum threshold of four, ensuring that low visitation numbers do not appear in the dataset (Sakong and Zentefis 2024). By incorporating these modifications, we evaluate how privacy-preserving mechanisms influence the accuracy of estimated visitation trends and welfare outcomes, particularly in cases where mobility data providers apply stringent privacy filters. We discuss this limitation explicitly in Section 6 and provide suggestions for researchers using anonymized mobility data.

3 Descriptive Evidence

Understanding how different data-handling approaches influence the number of identified visits to Cape Cod beaches is critical. By systematically comparing visit counts under the proposed practices and various alternative methods, we provide initial descriptive insights into how methodological choices affect observed visitation patterns. These comparisons help illustrate the extent to which different processing assumptions impact data interpretation before integrating them into the recreation demand model.

We first analyze the visit counts derived using the proposed practices to establish a baseline for comparison. Figure 1a presents the average monthly visits to Cape Cod beaches, revealing clear spatial patterns, with some beaches experiencing significantly higher visitation than others. The concentration of visits at certain locations suggests that some beaches attract a disproportionate share of visitors, which may be attributed to the proximity to major population centers, or superior site characteristics.

[13] Advan provides only aggregated visitation counts, which do not allow us to apply or compare alternative data-processing choices that depend on individual-level traces (e.g., dwell-time filtering, multi-stop trip identification, sampling variation). Therefore, Advan alone cannot be used to evaluate our core research question regarding how these decisions influence demand and welfare estimates. See more details in here.

The spatial distribution of visitor origins, shown in Figure 2a, provides further insights into the demand for Cape Cod beaches across different regions. Aggregated at CBG level, the map illustrates the areas that contribute the most visitors, highlighting the relative importance of local and distant sources of demand. The distribution reveals that a substantial share of visitors originates from within Massachusetts, particularly the Boston metropolitan area, but also includes notable contributions from other states in the Northeast, including New York, Connecticut, and New Jersey. This suggests that Cape Cod functions as both a regional and interstate recreational destination, attracting visitors from a broad geographic area.[14]

To assess how different data-handling choices affect visit identification, we compare visit counts under various alternative approaches relative to the baseline estimates. Figures 1b-1d illustrate the ratio of newly identified visits relative to the baseline counts when specific data-processing assumptions are modified. Applying a uniform four-minute minimum dwell time leads to a significant increase in visit counts across many beaches, suggesting that shorter stops, some of which may not represent true recreational visits, are being included. Counting each stop as a separate visit rather than consolidating multiple stops within the same day results in a more moderate increase, primarily impacting beaches with high repeat visitor activity. Additionally, incorporating differential privacy procedures to mobility data leads to a reduction in visit counts, particularly in less frequented locations, indicating that privacy adjustments may introduce systematic biases by disproportionately underrepresenting visits to certain sites.

We extend this analysis by examining how these alternative data-handling practices influence visits originating from home CBGs. The results, shown in Figures 2b-2d, reveal a consistent pattern with the findings at the beach level. Applying a four-minute dwell time threshold results in a substantial increase in visit counts across a wide range of CBGs, indicating that many short-duration stops, potentially including non-recreational visits or brief pass-throughs, are now classified as beach visits. Consolidating same-day stops modestly reduces total visit counts but likely merges a mix of fragmented and genuine within-day return visits (see Section 3 and Figures 4). For most CBGs, this adjustment has minimal influence on overall visit counts, confirming that the default practice of consolidating same-day stops provides a reasonable approximation of individual recreational trips.

The introduction of differential privacy adjustments alters visit counts across CBGs, reducing visits in some areas while increasing them in others. The reductions are particularly evident in CBGs with lower visitor density, where DP procedures are more likely to suppress visit records due to a limited number of observations. This effect is more pronounced in rural and suburban areas, where fewer trips make data more susceptible to privacy-related noise. In contrast, some urban and high-traffic CBGs experience an increase in recorded visits, likely due to the redistribution of noise within the dataset. These shifts indicate that privacy-preserving modifications can introduce systematic changes to visit estimates, potentially altering the observed spatial distribution of recreational demand.

[14] Although the raw heatmap in Figure [fig:figure2a] appears visually flat, this pattern mainly reflects the influence of varying CBG population sizes. When we explicitly estimate the distance-decay relationship (shown in Appendix Figure 5), the fitted function exhibits a smooth, monotonically decreasing gradient with distance, consistent with theoretical expectations from recreation demand models.

Another key factor affecting visit identification is the representativeness of mobility data. The sampling ratio, defined as the proportion of residing devices relative to the census-reported population, varies considerably across CBGs, reflecting differences in smartphone adoption, data coverage, and mobility tracking. Figure 3a illustrates the geographic distribution of sampling ratios, which have an average value of 0.16 and exhibit a right-skewed distribution, meaning that while most areas have relatively low coverage, some CBGs have disproportionately high representation.

Further analysis of the correlation between sampling ratios and demographic characteristics, presented in Figure 3b, reveals systematic biases in data coverage. The sampling ratio is negatively correlated with income and education, indicating that CBGs with higher sampling ratios tend to have lower average income and educational attainment. This relationship reflects coverage variation at the CBG level, that is, devices from lower-income neighborhoods appear somewhat more frequently in the dataset relative to their population size, not differences in who actually visits Cape Cod beaches. The correlation should therefore be interpreted as a feature of data representation, not individual visitation behavior. Differences in smartphone operating systems, location-sharing opt-in rates, and vendor panel composition may contribute to this pattern. More broadly, prior work shows that the direction of socioeconomic bias in smartphone-based mobility data can vary by platform and geography; while many datasets tend to overrepresent higher-income and urban users, others exhibit more balanced or even reversed patterns depending on data-sharing mechanisms (Björkegren and Milusheva 2020; Schlosser et al. 2021). Recognizing that the sampling ratio captures device coverage rather than behavioral participation helps contextualize potential representativeness limitations in mobility-based analyses.

The descriptive analysis demonstrates that data-handling choices can affect visit counts, particularly when modifying visit definitions and implementing privacy protections. Reducing the minimum dwell time increases visit estimates by capturing brief stops, potentially inflating demand if non-recreational visits are included, while privacy protections that introduce noise lead to underestimation, disproportionately impacting less-frequented locations. Additionally, mobility data exhibit representativeness biases, with lower-income populations being overrepresented, highlighting the need for demographic adjustments to ensure accurate demand estimation. These findings underscore the importance of carefully considering methodological choices in visitation analysis.

4 Model and Estimation

We employ a standard random utility model (RUM) framework to estimate the effect of contamination advisories on site selection and derive the MWTP for avoiding contamination. We then outline how we test the effects of alternative data-handling practices on MWTP estimates.

4.1 Recreation Demand Model

We model recreation demand using a random utility framework, where individuals choose among alternative beach sites based on site attributes and travel costs. The utility function for an individual from CBG c visiting site j at year-month t is specified as:

$$U_{cjt}=\beta_1 TC_{cjt}+CBG_c+Site_{jt}+\epsilon_{cjt},$$

where TC_{cjt} represents the travel cost from CBG c to site j at year-month t, which captures the economic burden associated with site selection. The term CBG_c controls for demographic characteristics of the visitor's home CBG, which may influence recreation preferences. The term $Site_{jt}$ represents site-specific attributes that affect the desirability of a given beach, and ϵ_{cjt} follows an i.i.d. Type I extreme value distribution.

For each CBG, we define the total potential market as all resident devices observed in the data, multiplied by the assumed number of recreational opportunities per month.[15] The market share to each recreational site for a CBG is calculated as the number of trips from a CBG to that site divided by this potential market size. Assuming the extreme value distribution for the error term, the probability that a visitor from CBG c selects site j is given by:

$$\text{share}_{jt}=\frac{\exp(\beta_1 TC_{cjt}+CBG_c+Site_{jt})}{1+\sum_{k=1}^{J}\exp(\beta_1 TC_{ckt}+CBG_c+Site_{kt})},$$

To account for individuals who choose not to visit any beach, we introduce an outside option representing alternative leisure activities or non-recreation choices. The outside option share is defined as the remaining portion of the market: one minus the sum of all observed site shares within that CBG. This treatment provides a consistent baseline across CBGs, ensuring that welfare and demand estimates remain comparable and conceptually aligns with the random utility framework, where the outside option encompasses all alternative leisure activities. The probability that a visitor from CBG c selects outside option 0 at year-month t is given by:

$$\text{share}_{0t}=\frac{1}{1+\sum_{k=1}^{J}\exp(\beta_1 TC_{ckt}+CBG_c+Site_{kt})}.$$

Taking the log-ratio of the site share relative to the outside option, we obtain:

$$\log(\text{share}_{jt})\text{-}\log(\text{share}_{0t})=\beta_1 TC_{cjt}+CBG_c+Site_{jt}.$$

This transformation expresses the market share of each beach relative to the outside option, allowing us to estimate the time-varying alternative-specific constants ($Site_{jt}$) in a way that accounts for overall recreation participation.

4.2 First-Stage Estimation: Adjusting for Zero Visit Shares

A common challenge in demand modeling is the presence of zero market shares, where certain products or sites record no observed sales or visits in specific markets (Earle and Kim 2024; Gandhi et al. 2023; Ackerberg and Rysman 2002; Quan and Williams 2018). This issue is particularly relevant in discrete-choice models of retail and recreation demand, where infrequent purchases, niche products, or limited awareness contribute to observed zeros. Traditional models often assume positive market shares, making them less suited for datasets with frequent zero

[15] We assume that each individual makes one recreation trip decision on every holiday or weekend day, which is about 10 days per month on average.

observations (Berry et al. 2004). Additionally, small samples or survey-based approaches can introduce sampling errors, further complicating demand estimation.

Researchers have used aggregation or selective trimming, sometimes unintentionally introducing selection bias (Dubé et al. 2021; Li 2019). In our case, zero market shares may arise from (1) sites being available but unvisited, leading to selection bias, or (2) visits occurring but not captured in the data, introducing measurement error and underestimating demand. To address this, we construct an empirical Bayes prior using the 50 closest markets by travel distance. The market shares of the same recreational site from these similar markets form the prior distribution, which we model using a Beta-Binomial framework. This Bayesian adjustment shrinks observed zeros toward expected positive values, reducing downward bias from sampling errors and improving demand estimation.

Formally, the number of visits from CBG c to site j, denoted K_{jc}, is modeled as a binomial random variable with N_c trials (the total observed visits from CBG c) and trip probability s^0_{jc}. The trip probability s^0_{jc} follows a Beta prior distribution:

$$K_{jc} \sim \text{Binomial}(N_c, s^0_{jc}), \quad s^0_{jc} \sim \text{Beta}(\lambda^1_{jc}, \lambda^2_{jc}).$$

Using Bayesian updating, the posterior mean of the trip probability s_{jc} is:

$$\hat{s}_{jc} = \frac{\lambda^1_{jc} + K_{jc}}{\lambda^1_{jc} + \lambda^2_{jc} + N_c}.$$

For comparison, the maximum likelihood estimate (MLE) of the trip share is:

$$\hat{s}^{MLE}_{jc} = \frac{K_{jc}}{N_c}.$$

We use the strictly positive posterior mean $\hat{s}_{jc}$ in the demand estimation instead of the MLE, particularly in cases where $K_{jc} = 0$. In large markets, the posterior closely approximates the observed share, while in smaller markets, it stabilizes estimates to reflect expected visitation probabilities based on prior information.

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^1_{jc} and λ^2_{jc}, are estimated from maximizing the log of the likelihood over the outcomes in the markets that form the priors,

$$f(K_{jl}, l \in P_c | \lambda_{1jc}, \lambda_{2jc}) = \prod_{l \in P_c} \binom{K_{jl}}{N_l} \frac{\Gamma(\lambda^1_{jc} + \lambda^2_{jc}) \Gamma(\lambda^1_{jc} + K_{jl}) \Gamma(N_l - K_{jl} + \lambda^2_{jc})}{\Gamma(\lambda^1_{jc}) \Gamma(\lambda^2_{jc}) \Gamma(N_l + \lambda^1_{jc} + \lambda^2_{jc})},$$

We estimate a pair of hyperparameters λ_{jc}^{1} and λ_{jc}^{2} for each beach, CBG, and week. The resulting posterior mean estimates, used in place of MLE, ensure strictly positive market shares for the subsequent demand estimation.[16]

4.3 Second-Stage Estimation: The Effect of Water Contamination Advisories

Once the first-stage estimates provide strictly positive visit probabilities, we proceed with the second-stage estimation to identify the effect of contamination advisories on site choice. We decompose site attributes into:

$$Site_{jt}=\beta_2 A_{jt}+\gamma_j+\lambda_t+X_{jt}+\xi_{jt},$$

where A_{jt} is an indicator for contamination advisories at site j in period t, capturing the effect of temporary environmental hazards on visitation. The term γ_j represents site fixed effects, controlling for unobserved time-invariant characteristics such as beach amenities and long-term water quality conditions. The term λ_t captures seasonal fixed effects to account for fluctuations in demand across different periods. The vector X_{jt} includes additional controls such as weather conditions, and ξ_{jt} represents unobserved time-varying site-specific factors.

4.4 Estimating the MWTP

We estimate the model using a multinomial logit framework, where the probability of an individual choosing site j is determined by the relative utility of that site compared to all other available alternatives. The estimated coefficients provide insights into how travel costs and site conditions influence site selection. The key parameter of interest is β_2 from equation ([eq:9]), which quantifies the impact of contamination advisories on site choice. To interpret this effect in economic terms, we compute the MWTP to avoid contamination, which is given by:

$$MWTP=-\frac{\beta_2}{\beta_1}.$$

This ratio represents the monetary value individuals assign to avoiding a contamination advisory, as it converts the disutility of contamination into an equivalent travel cost. A higher MWTP indicates a stronger aversion to contamination, while a lower MWTP suggests that visitors are less responsive to advisories. To statistically evaluate how data-handling choices influence MWTP estimates, we implement a bootstrap procedure that resamples the data 1,000 times for each specification (i.e., for each set of data-handling practices). While a single full-sample estimate provides a point value, it does not capture sampling uncertainty or allow us to test whether differences across processing choices are statistically significant. Bootstrapping generates a distribution of MWTP estimates for each specification, from which we derive confidence intervals.

[16] A Monte Carlo simulation in Cheng and Wan (2025) demonstrated that the empirical Bayes (EB) posterior-mean correction yields unbiased recovery of travel-cost parameters and mean utilities under sparse cellphone mobility data, whereas conventional treatments (e.g., dropping zeros or adding constants) produce attenuation or instability.

We then compare these distributions, rather than only the point estimates, to determine whether the MWTP values obtained under alternative practices fall within or outside the confidence interval of the proposed practice. This approach allows us to assess whether observed differences across data-handling choices are statistically meaningful, rather than purely descriptive.

5 Results

We now discuss the findings from the recreation demand estimation, beginning with the baseline results derived using the proposed practices for data processing. We then systematically assess the effects of deviations from these practices, examining how each alternative approach, from relaxing dwell time thresholds, to applying differential privacy, to redefining choice sets, and to modifying visit definitions, impacts the estimated MWTP for avoiding contamination. Finally, to explore how multiple non-optimal choices may interact, we simulate a combined case in which all alternative assumptions are applied sequentially.[17]

5.1 Baseline Results

The first-stage estimation results in Table 2 show a statistically significant negative effect of travel costs on site selection, with a coefficient of -0.039. This indicates that as travel costs increase, the probability of visiting a given beach decreases, consistent with economic theory on consumer behavior. The magnitude of this coefficient suggests a moderately elastic response to travel expenses, meaning that visitors are relatively sensitive to cost variations when deciding on recreational sites. This finding aligns with prior research on outdoor recreation demand, where negative travel cost coefficients typically range between -0.02 and -0.1 in similar random utility models (e.g., Parsons and Massey (2003; Phaneuf and Smith 2005; Murray et al. 2001)). In the second stage, the effect of contamination advisories is estimated at -0.348, indicating that environmental warnings significantly reduce the probability of a site being chosen. This suggests that visitors actively avoid beaches with contamination concerns, demonstrating a clear preference for clean water conditions.

The estimated marginal willingness to pay (MWTP) for avoiding fecal contamination is $8.92 per visit (95% CI: $1.05-$22.17). This value represents the welfare gain from accessing a beach during a month without a bacterial contamination advisory, capturing the localized welfare impact of short-term contamination events while allowing for substitution across nearby sites. Because most BEACON advisories are short-lived and confined to a few beaches, this estimate should be interpreted as the per-visit value of avoiding a single-beach advisory rather than a regional or prolonged closure.

To evaluate the robustness of this dwell-time assumption, we re-estimated both trip counts and welfare outcomes under stricter minimum dwell-time requirements of 30, 60, and 90 minutes. As shown in Appendix Table 6, increasing the threshold reduces the number of identified trips by removing a substantial share of short-duration visits, and it leads to higher willingness-to-pay (MWTP) estimates as the remaining sample becomes more concentrated among longer-stay, higher-intensity recreation. For example, raising the threshold from 10 to 30 minutes increases

[17] Again, this configuration does not represent a realistic or recommended data-processing strategy, we present it to highlight the upper bound of compounded bias when multiple non-optimal choices are made in sequence.

MWTP by approximately 40 percent, with smaller additional adjustments at 60 and 90 minutes. Crucially, however, qualitative welfare results and policy conclusions remain consistent across all thresholds. These findings confirm that while dwell-time choice influences the scale of estimated benefits, the 10-minute cutoff provides the most representative coverage of the full spectrum of beach recreation behavior.

To test the robustness of our baseline 300-mile travel distance cutoff, we also estimated the model using alternative thresholds of 100, 150, 200, and 250 miles. As shown in Appendix Table 7, the estimated coefficients on travel cost and advisory variables remain stable across specifications, indicating that the welfare estimates are not sensitive to reasonable variations in the distance cutoffs.

Finally, we examine whether advisories influence visitation with a temporal lag. This test is motivated by the possibility that some advisories, especially those issued near the end of a month, may primarily affect visitation in subsequent periods rather than immediately. To evaluate this, we include a one-month lag of the advisory variable in the baseline specification. Results from Appendix Table 8 show the contemporaneous effect remains negative and statistically significant (-0.327**), while the lagged coefficient is small and insignificant (-0.059). This pattern indicates that behavioral responses to contamination warnings occur largely within the same month of issuance, with limited delayed effects, thereby reinforcing the robustness of our main findings.

Overall, these findings confirm that beach visitors exhibit strong preferences for clean environments and that even moderate contamination advisories significantly affect site choice. The MWTP estimate provides valuable insights for policymakers assessing the economic benefits of water quality improvements, reinforcing the necessity of investments in pollution mitigation and public health monitoring programs.

5.2 Impact of Alternative Data Practices on MWTP

To assess the sensitivity of MWTP estimates to data processing choices, we re-estimate the model using several alternative data-handling practices. Each deviation from the proposed practices alters the estimation of recreation demand, affecting how visitors' preferences for avoiding contamination are quantified.

When we relax the minimum dwell time requirement for defining a visit to a fixed four-minute threshold, the MWTP drops sharply to \$3.81 (95% CI: -\$2.10 ~ \$12.46), representing a 57% underestimation compared to the baseline. This significant reduction suggests that shorter visits, many of which may not correspond to actual recreational activity, introduce measurement noise into the demand estimation process. Including brief stops as valid visits likely captures incidental or transient activity (e.g., passing through a beach without engaging in recreation), thereby weakening the estimated preference for avoiding contamination. The dilution of meaningful visits lowers the inferred economic value of clean water, biasing MWTP downward.

When we modify visit definitions by treating each stop as an independent visit, rather than consolidating multiple stops at a beach within a single day, the MWTP declines slightly to \$8.75 (95% CI: \$1.04–\$22.15). This modest change indicates that, within the Cape Cod context, consolidating same-day stops has limited influence on the estimated welfare outcomes. However, this result should not be interpreted as evidence that mobility data perfectly capture the distribution of recreational choices. Rather, it shows that this particular data-handling decision has relatively small quantitative effects on the welfare estimates in our case study. The accuracy of mobility data in representing true visitation behavior likely varies by region and activity type, and future research should test whether similar patterns hold in other recreational settings.

Applying differential privacy techniques, which introduce noise to protect user anonymity, reduces MWTP to \$3.12 (95% CI: -\$0.42 ~ \$9.73), indicating a 65% underestimation relative to the baseline. This substantial drop highlights the trade-off between privacy protection and estimation accuracy. Differential privacy mechanisms obscure precise visitation patterns, leading to misclassification of site choices and distorting inferred preferences. The suppression of true visit frequencies weakens the signal of site desirability, making it appear as though visitors are less responsive to contamination advisories. As a result, the model underestimates the economic value placed on water quality improvements.

When we redefine individual choice sets based on short-term visitation patterns instead of long-term behavior, MWTP increases to \$9.77 (95% CI: \$2.14 ~ \$20.29), reflecting a 10% overestimation. This suggests that restricting choice sets to short-term preferences may exaggerate site loyalty, thereby amplifying the estimated value of avoiding contamination. Long-term choice sets account for habitual site familiarity and broader recreational patterns, whereas short-term definitions may disproportionately emphasize recent visits. As a result, the model may overstate the degree to which visitors value site cleanliness, leading to an inflated MWTP.

Omitting sampling weights, which adjust for demographic representativeness, results in a MWTP estimate of \$9.00 (95% CI: \$2.72 ~ \$19.53), a minor deviation from the baseline. This suggests that demographic weighting does not significantly influence MWTP estimates in this context. The minimal impact indicates that mobility data, even without representativeness adjustments, reasonably capture the distribution of visitors' recreation choices. However, while this result implies that sampling biases may be relatively small for Cape Cod beaches, the necessity of applying weights should still be evaluated on a case-by-case basis, particularly in studies involving more diverse or underrepresented populations.

6 Discussion

Our findings have essential implications for non-market valuation studies, particularly in the context of using mobility data to estimate recreational demand. By structuring our analysis around the new problems unique to mobility data and the classic problems of recreation demand, we can draw clearer recommendations for researchers and policymakers. The results highlight the sensitivity of economic valuation estimates to data-processing choices, demonstrating how deviations from the proposed practices can lead to substantial biases in MWTP estimates. Given that non-market valuation plays a crucial role in informing environmental policy, resource management, and public investment decisions, ensuring methodological rigor in mobility-based valuation is essential.

6.1 Trip Definitions and Privacy

First, defining a recreational trip is a critical decision that can introduce significant measurement error. When we relaxed the minimum dwell time requirement from ten minutes to a fixed four-minute threshold, the MWTP for avoiding contamination dropped to $3.81, a 57% underestimation compared to our baseline. This significant reduction suggests that including brief stops as valid visits captures incidental or transient activity (e.g., passing through a beach without engaging in recreation), which dilutes the pool of meaningful visits and weakens the estimated preference for clean water.

Second, the privacy-accuracy trade-off presents a hard challenge. Applying industry-standard differential privacy techniques, which introduce noise to protect user anonymity, had the most substantial impact, reducing the MWTP by 65%. This result highlights a critical tradeoff: while privacy protections are important, they can obscure precise visitation patterns to such an extent that they severely distort inferred preferences and lead to the systematic undervaluation of environmental resources. Policymakers and researchers relying on mobility data must be aware that the methods used to protect privacy can directly impact the reliability of economic assessments, creating a need for strategies that balance data protection with estimation accuracy (Connolly et al. 2025; Sakong and Zentefis 2024).

6.2 Individual Choice Sets and Representativeness

When we redefined individual choice sets based on short-term visitation patterns (a single year) instead of long-term behavior (four years), the MWTP increased by 10%. Our results demonstrate the importance of constructing meaningful choice sets in non-market valuation. The overestimation of MWTP when restricting choice sets to short-term behavior suggests that omitting historical site preferences may distort demand estimates. Future applications of mobility data should incorporate long-term site selection patterns to improve the accuracy of valuation models.

Finally, while sample representativeness is a known issue for nearly all survey methods, its effect in our study was minimal. Omitting the sampling weights, which adjust for demographic representativeness, resulted in an increase of 0.89% in MWTP. While sampling weights had a limited effect on MWTP in our study, demographic biases in mobility datasets may vary across contexts. Researchers should assess the representativeness of mobility samples before applying them to valuation studies, particularly when studying diverse populations with varying access to outdoor recreation opportunities.

6.3 Recommendations for Future Applications

To enhance the reliability of mobility-based non-market valuation, we recommend several practices for future applications. First, researchers should carefully define visits based on context-specific behavioral patterns rather than applying fixed dwell time thresholds, which may introduce measurement error. Second, long-term choice sets should be used to capture habitual site preferences, ensuring that demand estimates reflect realistic decision-making processes. Third, when applying privacy-preserving techniques, efforts should be made to quantify their impact on valuation estimates, potentially by running sensitivity analyses or supplementing the anonymized data with additional sources that provide context or missing details. Finally, integrating mobility data with survey-based approaches may offer a way to cross-validate demand estimates, allowing for a more comprehensive assessment of recreational behavior.

As mobility data continue to be incorporated into environmental valuation, ensuring methodological transparency and standardization will be essential. Future research should explore how different privacy frameworks, data aggregation techniques, and sampling strategies influence non-market valuation outcomes. Additionally, extending the application of mobility data to other recreational settings, such as national parks, urban green spaces, and inland water bodies, will provide further insights into the generalizability of these methods. By addressing these methodological challenges, mobility data can serve as a powerful tool for improving non-market valuation and guiding policies that promote sustainable resource management and equitable access to outdoor recreation.

We also acknowledge that our dwell time analysis is subject to potential biases from heterogeneous ping frequencies. Different applications or user behaviors may generate location pings at different rates. A higher ping frequency more accurately captures the true duration of a visit, reducing the measurement error. Consequently, a fixed dwell time threshold could systematically bias the sample by preferentially including users of more active, high-frequency applications while excluding those whose true visits are underestimated due to infrequent pings. While the aggregation of our data across multiple applications on each device helps to mitigate this concern, the potential for such data generation processes to affect sample representativeness remains a limitation and a subject for future research.

A related limitation is that mobility data may also exhibit site-specific behavioral biases beyond demographic underrepresentation. Some visitors may be more likely to generate pings at particular types of locations, for example, individuals who check their phones more frequently while at beaches than while hiking or shopping. These behavioral patterns can lead to uneven data coverage across sites, effectively weighting the sample toward individuals who interact with their devices more actively in certain settings.

Because vendor data do not include information on app usage or phone interaction behavior, we cannot directly correct for this bias. Our analysis therefore assumes relatively uniform data generation across sites, an assumption that should be tested in future work. Comparing mobility-derived visitation patterns with intercept surveys or administrative visitor counts would help quantify the magnitude of these site-specific biases and improve the interpretability of mobility-based recreation demand estimates.

7 Conclusion

Mobility data have emerged as a promising solution to longstanding challenges in recreation demand estimation. Traditional survey-based approaches, while valuable, suffer from issues such as recall bias, limited geographic scope, high costs, and difficulties in capturing real-time behavioral responses. The availability of large-scale, passively collected mobility data offers the potential to overcome these limitations, providing detailed insights into recreational visitation patterns across wide spatial and temporal scales. However, leveraging this data effectively requires careful methodological considerations, particularly in defining visits, constructing choice sets, and ensuring data representativeness. Our research examines the sensitivity of data processing choices on non-market valuation, highlighting their implications for recreation demand modeling and environmental policy.

Specifically, this study systematically evaluates how different data-handling choices influence the estimation of recreation demand and the valuation of environmental quality. We apply a random utility model to cell phone visits to Cape Cod beaches, estimating the MWTP for avoiding contamination under various data-processing scenarios. Our findings indicate that adhering to proposed practices, including longer minimum dwell time, multi-stop visits, long-term choice set definition, and demographic weighting, yields an MWTP estimate of $8.92 per visit. Additionally, while past studies primarily relied on survey data, our findings demonstrate that mobility data can yield comparable estimates when processed correctly, reinforcing the credibility of this emerging data source.

However, deviations from these practices introduce significant biases. Our results show that welfare estimates are highly sensitive to even small data-handling choices. Relaxing minimum dwell time from 10 minutes to 4 minutes results in a 57 percent reduction in MWTP, suggesting that shorter visits can introduce noise by including incidental or transient stops that do not represent meaningful recreation. Applying differential privacy leads to a 65 percent reduction in MWTP, highlighting the trade-off between protecting user privacy and maintaining estimation accuracy. Privacy-preserving noise obscures true visit patterns, biasing site preference estimation. Restricting choice sets to short-term behavior inflates MWTP by 10 percent, indicating that failing to account for habitual site familiarity overstates visitors' sensitivity to environmental conditions. Omitting demographic sampling weights has minimal impact on MWTP, suggesting that mobility data may already provide a reasonable representation of the broader population's recreational behaviors. This finding reinforces the need for standardized, transparent methodologies in mobility-based non-market valuation to ensure consistency and comparability across studies.

Accurate recreation demand estimation is essential for environmental policy, resource management, and public investment decisions, yet our findings reveal significant risks associated with misprocessing mobility data. When mobility data are not handled carefully, underestimating MWTP, such as when additional privacy protection measures are applied or visit definitions are relaxed, can lead to undervaluation of environmental resources, resulting in insufficient funding for conservation efforts. Conversely, overestimating MWTP due to restrictive choice set definitions may misallocate public funds toward environmental improvements that exceed the public's true valuation. The particularly strong impact of differential privacy on MWTP estimates raises concerns about the unintended consequences of privacy regulations, as excessive anonymization can obscure true visitation patterns and weaken the reliability of non-market valuation (Dekel et al. 2022). This trade-off between privacy protection and data accuracy underscores the need for policymakers and data providers to develop strategies that balance individual data security with the integrity of economic research.

Beyond privacy, our findings clarify how other methodological choices translate into practical consequences for applied recreation demand analysis. Shorter dwell-time thresholds and relaxed visit definitions risk including incidental stops, weakening the behavioral signal of meaningful recreation. In contrast, overly restrictive choice sets, based only on recent visits, inflate estimated welfare measures by overstating loyalty to familiar sites. Meanwhile, adjustments for demographic representativeness ensure that inferred preferences reflect the broader recreating population rather than only smartphone-active users. Together, these results provide a roadmap for researchers and practitioners: privacy adjustments primarily affect estimation precision, trip definition and choice-set construction shape behavioral realism, and weighting schemes influence representativeness and policy relevance.

While our study provides valuable insights, several limitations and avenues for future research remain. Our analysis focuses on Cape Cod beaches, a coastal destination with a distinct visitor base. Future studies should assess whether these findings generalize to other recreational settings, such as inland lakes, national parks, urban green spaces, or large coastal systems with a greater number of substitute sites, where the welfare impacts of advisories and the sensitivity to data-handling choices may differ. We assume a homogeneous response to contamination advisories, but visitor preferences may vary by demographic group, trip purpose, or prior experience with advisories. Future research could explore preference heterogeneity using random coefficients models or latent class approaches. Our analysis spans 2019-2022, but visitor behavior may evolve over longer time horizons due to climate change, policy interventions, or shifting recreation trends. Expanding the dataset to include pre- and post-policy periods could enhance our understanding of dynamic site choice responses.

Taken together, these insights underscore that rigorous and transparent data processing is essential for credible mobility-based recreation demand modeling. Rather than identifying a single “best practice," our results demonstrate that the choice of data-handling procedure is not neutral. Different choices can lead to substantially different welfare estimates. As mobility data continue to be integrated into environmental economics, ensuring transparency and standardization in data handling will be essential for generating reliable insights that inform resource management and public investment decisions. Moving forward, interdisciplinary collaboration between economists, data scientists, and policymakers will be critical in refining methods that balance privacy protection with data accuracy, ensuring that mobility data fulfill their potential as a valuable tool for understanding and managing natural resource use.

References

Ackerberg, D. A., & Rysman, M. (2005). Unobserved product differentiation in discrete-choice models: Estimating price elasticities and welfare effects. *RAND Journal of Economics*, 36(4), 771–788.

Acquisti, A., Taylor, C., and Wagman, L. (2016). The economics of privacy. *Journal of Economic Literature*, 54(2):442–492.

Bartlett, R. (2019). *Moon Cape Cod, Martha’s Vineyard & Nantucket*. Moon Travel.

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.

Björkegren, D. and Milusheva, S. (2020). The smartphone data gap. Policy Research Working Paper 9371, World Bank.

Cameron, T. A., Shaw, W. D., Ragland, S. E., Mac Callaway, J., and Keefe, S. (1996). Using actual and contingent behavior data with differing levels of time aggregation to model recreation demand. *Journal of Agricultural and Resource Economics*, pages 130–149.

Cheng, N. and Wan, X. (2025). Causal recreation demand estimation with cellphone mobility data: Evidence from the 2021 huntington beach oil spill. Working Paper.

Connelly, N. and Brown, T. (2011). Effect of recall period on annual freshwater fishing effort estimates in new york. *Fisheries Management and Ecology*, 18(1):83–87.

Connolly, C., Steinbach, S., Vo, M., and Wan, X. (2025). Estimating structural models with privacy-protected data: A correction for bias in large-scale mobility data. Working Paper.

Cook, C. (2025). Heterogeneous preferences for neighborhood amenities: Evidence from gps data. *Review of Economics and Statistics*, pages 1–43.

Data, D. (2025). Dewey academic research data platform. https://www.deweydata.io/. Platform provides datasets with DOIs.

Dekel, I., Cummings, R., Heffetz, O., and Ligett, K. (2022). The privacy elasticity of behavior: Conceptualization and application. Technical report, National Bureau of Economic Research.

DeShazo, J. R., Pendleton, L., and Cutter, W. B. (2003). Activities in models of recreational demand. Technical report, University of California, Los Angeles. Working paper.

Dillman, D. A. (2017). *The promise and challenge of pushing respondents to the web in mixed-mode surveys*. Statistics Canada.

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.

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.

Dwork, C., McSherry, F., Nissim, K., and Smith, A. (2006). Calibrating noise to sensitivity in private data analysis. *In Theory of Cryptography: Third Theory of Cryptography Conference*, TCC 2006, New York, NY, USA, March 4-7, 2006. Proceedings 3, pages 265–284. Springer.

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.

Englin, J. E., Holmes, T. P., and Sills, E. O. (2003). Estimating forest recreation demand using count data models. *Forests In A Market Economy*, pages 341–359.

Furey, R. P., Merrill, N. H., Sawyer, J. P., Mulvaney, K. K., and Mazzotta, M. J. (2022). Evaluating water quality impacts on visitation to coastal recreation areas using data derived from cell phone locations. *PLOS One*, 17(4):e0263649.

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.

Gibbs, H., Eggo, R. M., and Cheshire, J. (2023). Detecting behavioural bias in gps location data collected by mobile applications. *medRxiv*, pages 2023–11.

Gonyo, S. B., Burkart, H., and Regan, S. (2024). Leveraging big data for outdoor recreation management: A case study from the york river in virginia. *Journal of Environmental Management*, 354:120482.

Gonzalez, M. C., Hidalgo, C. A., and Barabasi, A.-L. (2008). Understanding individual human mobility patterns. *Nature*, 453(7196):779–782.

Hauer, M. E. and Santos-Lozada, A. R. (2021). Differential privacy in the 2020 census will distort covid-19 rates. *Socius*, 7:2378023121994014.

Hindsley, P. (2011). Addressing on-site sampling in recreation site choice models. Technical report, East Carolina University, Dept. of Economics. Working paper.

Kenny, C. T., Kuriwaki, S., McCartan, C., Rosenman, E. T., Simko, T., and Imai, K. (2021). The use of differential privacy for census data and its impact on redistricting: The case of the 2020 us census. *Science Advances*, 7(41):eabk3283.

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.

Lee, S., Wan, X., and Zheng, S. (2023). Estimating the indirect cost of floods: Evidence from high-tide flooding. MIT Center for Real Estate Research Paper No. 23/10, SSRN Electronic Journal.

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

Liang, Y., Yin, J., Pan, B., Lin, M. S., Miller, L., Taff, B. D., and Chi, G. (2022). Assessing the validity of mobile device data for estimating visitor demographics and visitation patterns in yellowstone national park. *Journal of Environmental Management*, 317:115410.

Loomis, J., Roach, B., Ward, F., and Ready, R. (1995). Testing transferability of recreation demand models across regions: a study of corps of engineer reservoirs. *Water Resources Research*, 31(3):721–730.

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.

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.

Manning, R., Valliere, W., Minteer, B., Wang, B., and Jacobi, C. (2000). Crowding in parks and outdoor recreation: A theoretical, empirical, and managerial analysis. *Journal of Park & Recreation Administration*, 18(4).

Mazzotta, M. J., Merrill, N. H., and Mulvaney, K. K. (2022). Coastal recreation in southern new england: results from a regional survey. *Journal of Ocean and Coastal Economics*, 9(1):1.

Merrill, N., Mazzotta, M., Mulvaney, K., Sawyer, J., Twichell, J., Atkinson, S., Keith, D., and Erban, L. (2022). The value of water quality for coastal recreation in new england, usa.

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.

Murray, C., Sohngen, B., and Pendleton, L. (2001). Valuing water quality advisories and beach amenities in the great lakes. *Water Resources Research*, 37(10):2583–2590.

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.

Parsons, G., Leggett, C. G., Herriges, J., Boyle, K., Bockstael, N., and Chen, Z. (2021). A siteportfolio model for multiple-destination recreation trips: Valuing trips to national parks in the southwestern united states. *Journal of the Association of Environmental and Resource Economists*, 8(1):1–25.

Parsons, G. R. and Hauber, A. B. (1998). Spatial boundaries and choice set definition in a random utility model of recreation demand. *Land Economics*, pages 32–48.

Parsons, G. R., Kang, A. K., Leggett, C. G., and Boyle, K. J. (2009). Valuing beach closures on the padre island national seashore. *Marine Resource Economics*, 24(3):213–235.

Parsons, G. R. and Massey, D. M. (2003). 12. a random utility model of beach recreation. *The New Economics of Outdoor Recreation*, 241.

Parsons, G. R., Plantinga, A. J., and Boyle, K. J. (2000). Narrow choice sets in a random utility model of recreation demand. *Land Economics*, pages 86–99.

Phaneuf, D. J. and Smith, G. C. (2015). Recreation demand models. *Review of Environmental Economics and Policy*, 9(2):136–158.

Phaneuf, D. J. and Smith, V. K. (2005). Recreation demand models. *Handbook of Environmental Economics*, 2:671–761.

Quan, T. W. and Williams, K. R. (2018). Product variety, across-market demand heterogeneity, and the value of online retail. *The RAND Journal of Economics*, 49(4):877–913.

Rolfe, J., Johnston, R. J., Rosenberger, R. S., and Brouwer, R. (2015). Introduction: Benefit transfer of environmental and resource values. In Benefit transfer of environmental and resource values, pages 3–17. Springer.

Rosenberger, R. S. and Loomis, J. B. (2017). Benefit transfer. In A primer on nonmarket valuation, pages 431–462. Springer.

Rylander, R. G., Propst, D. B., and McMurtry, T. R. (1995). Nonresponse and recall biases in a survey of traveler spending. *Journal of Travel Research*, 33(4):39–45.

Sakong, J. and Zentefis, A. K. (2024). A simulation-based method to estimating economic models with privacy-protected data. *NBER Chapters*.

Santos-Lozada, A. R., Howard, J. T., and Verdery, A. M. (2020). How differential privacy will affect our understanding of health disparities in the united states. *Proceedings of the National Academy of Sciences*, 117(24):13405–13412.

Savi, M. K., Yadav, A., Zhang, W., Vembar, N., Schroeder, A., Balsari, S., Buckee, C. O., Vadhan, S., and Kishore, N. (2023). A standardised differential privacy framework for epidemiological modeling with mobile phone data. *PLOS Digital Health*, 2(10):e0000233.

Schlosser, A., Maier, B. F., Jack, O., Hinrichs, D., Zachariae, A., and Brockmann, D. (2021). Covid-19 lockdown induces disease-mitigating structural changes in mobility networks. *Proceedings of the National Academy of Sciences*, 118(52):e2026646118.

Tarrant, M. A., Manfredo, M. J., Bayley, P. B., and Hess, R. (1993). Effects of recall bias and nonresponse bias on self-report estimates of angling participation. *North American Journal of Fisheries Management*, 13(2):217–222.

Totty, E. S. and Watson, T. (2024). Privacy protection and accuracy: What do we know? do we know things?? let's find out! Technical report, National Bureau of Economic Research.

U.S. Bureau of Labor Statistics (2025). American time use survey (atus): 2024 results. https://www.bls.gov/news.release/atus.t02.htm. Table 2. Time spent in primary activities and percent of the civilian population engaging in each activity, averages per day on weekdays and weekends, 2024.

Winkler, R. L., Butler, J. L., Curtis, K. J., and Egan-Robertson, D. (2021). Differential privacy and the accuracy of county-level net migration estimates. *Population Research and Policy Review*, pages 1–19.

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.

Figure 1: Spatial Distribution of Beach Visits and Effects of Alternative Visit Definitions on Observed Mobility Patterns

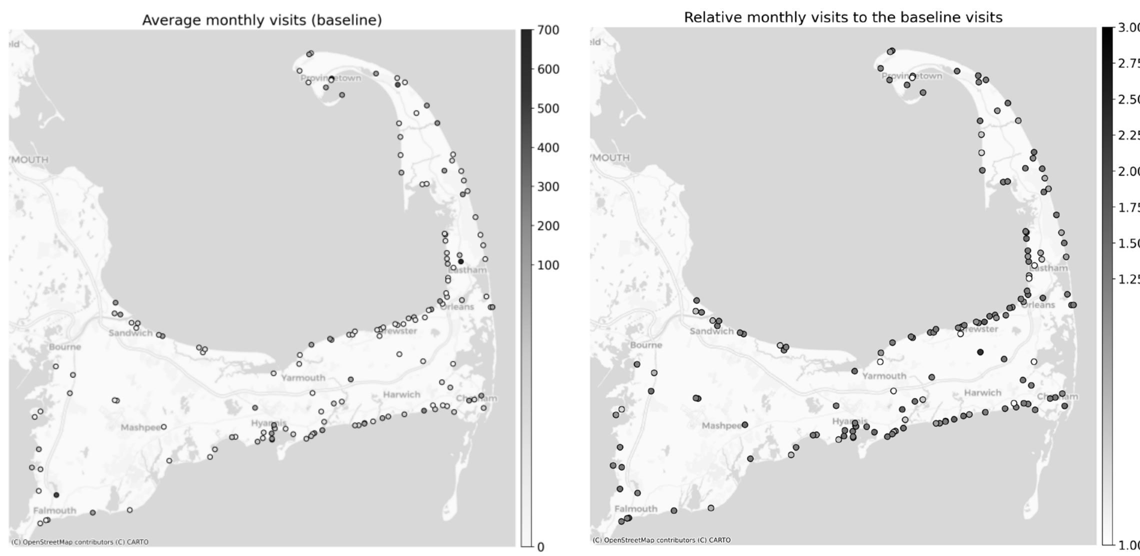

(a) average monthly visits to beaches

(b) dwelling > 4 mins

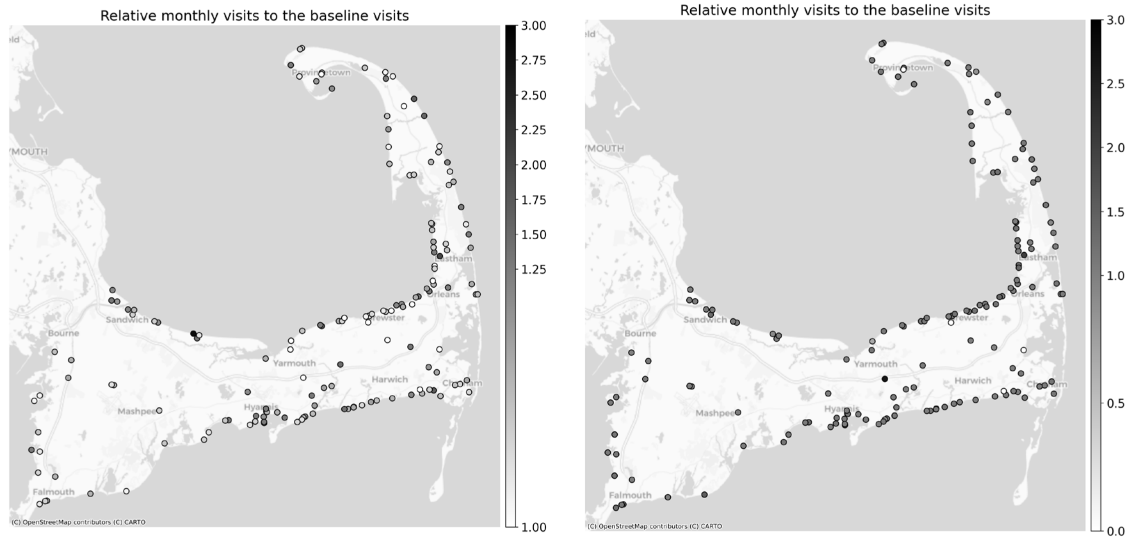

(c) use single-stop visits

(d) differential privacy

Notes: Each panel visualizes average monthly visits to 155 Cape Cod beaches based on Advan mobility data (2019-2022). Panel (a) shows baseline average monthly visits under the preferred specification. Panels (b)-(d) depict relative visit counts compared with the baseline, illustrating how alternative data-handling assumptions affect measured visitation: (b) imposing a minimum dwell time of 4 minutes, (c) restricting to single-stop visits per device, and (d) applying differential-privacy adjustments to obscure low-frequency visits. Values above 1.0 indicate higher relative visitation than in the baseline, while values below 1.0 indicate lower observed visitation.

Figure 2: Spatial Distribution of Visits from Home CBGs and Effects of Alternative Visit Definitions on Observed Outflow

(a) average monthly visits from home CBGs

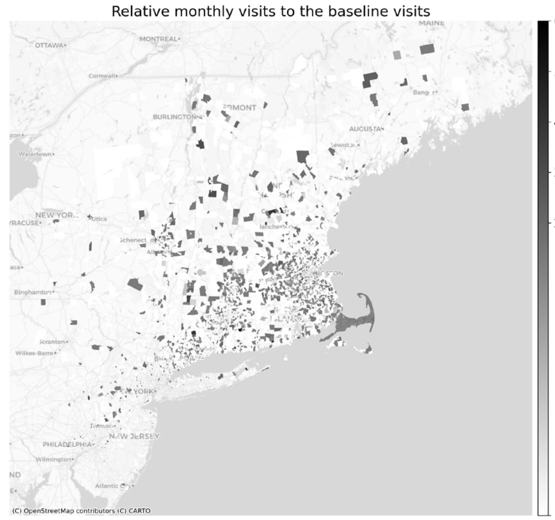

(b) dwelling > 4 mins

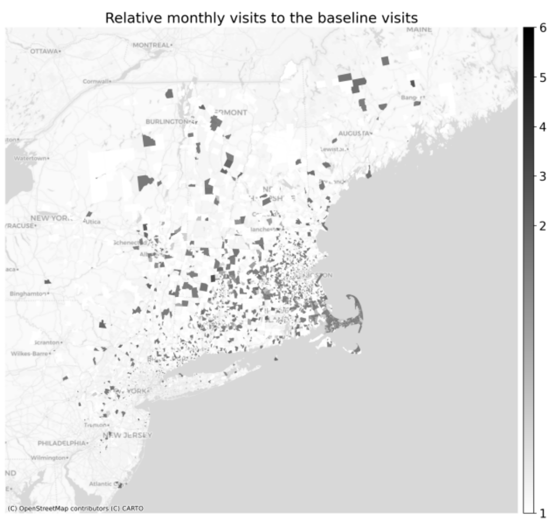

(c) use single-stop visits

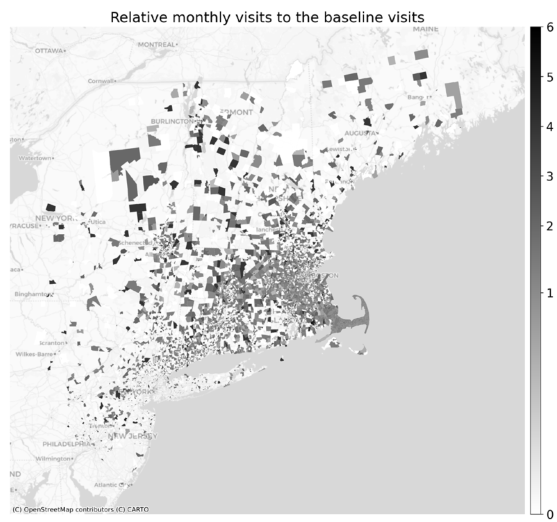

(d) differential privacy

Notes: Each panel illustrates the spatial distribution of average monthly trips from CBGs to Cape Cod beaches based on Advan mobility data from 2019-2022. Panel (a) presents baseline average monthly visits per CBG under the preferred specification. Panels (b)-(d) show relative changes in visit intensity compared with the baseline, reflecting how alternative processing choices influence estimated visitation flows: (b) imposing a minimum dwell time of 4 minutes, (c) limiting to single-stop visits per device, and (d) applying differential-privacy adjustments that obscure low-frequency trip counts. Values above 1.0 indicate greater relative visitation from a CBG compared with the baseline, whereas values below 1.0 indicate reduced observed outflows.

Figure 3: Spatial Distribution and Demographic Correlates of Sampling Ratios Across Home CBGs

(a) sampling ratio in home CBGs

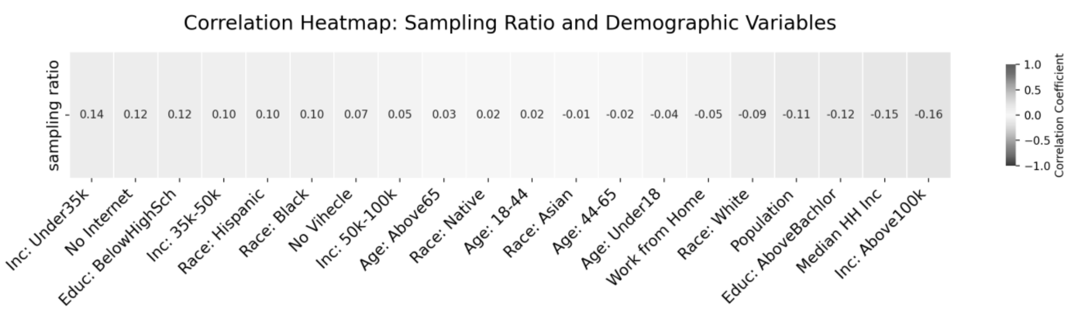

(b) correlation with demographic variables

Notes: Panel (a) maps the sampling ratio for each home CBG, calculated as the number of unique mobile devices observed in the Advan data divided by the total residential population from the American Community Survey (ACS, 2019-2022). Higher values indicate greater representation of residents in the mobility data relative to the true population. Panel (b) displays correlations between the sampling ratio and key demographic and socioeconomic variables from the ACS, including income, education, race, age, and commuting behavior. Positive correlations indicate groups that are more likely to be represented in the mobile sample, while negative correlations denote underrepresented groups.

Figure 4: Temporal Patterns and Sequencing of Same-day Beach Visits

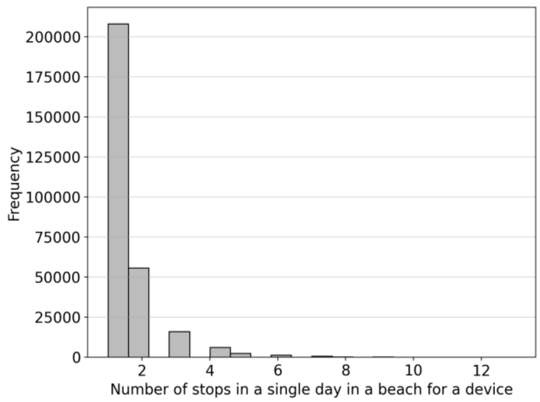

(a) distribution of stops in same-day visits

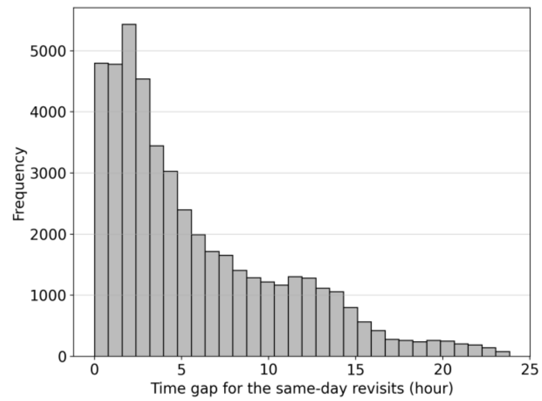

(b) time gap between consecutive stops (hrs)

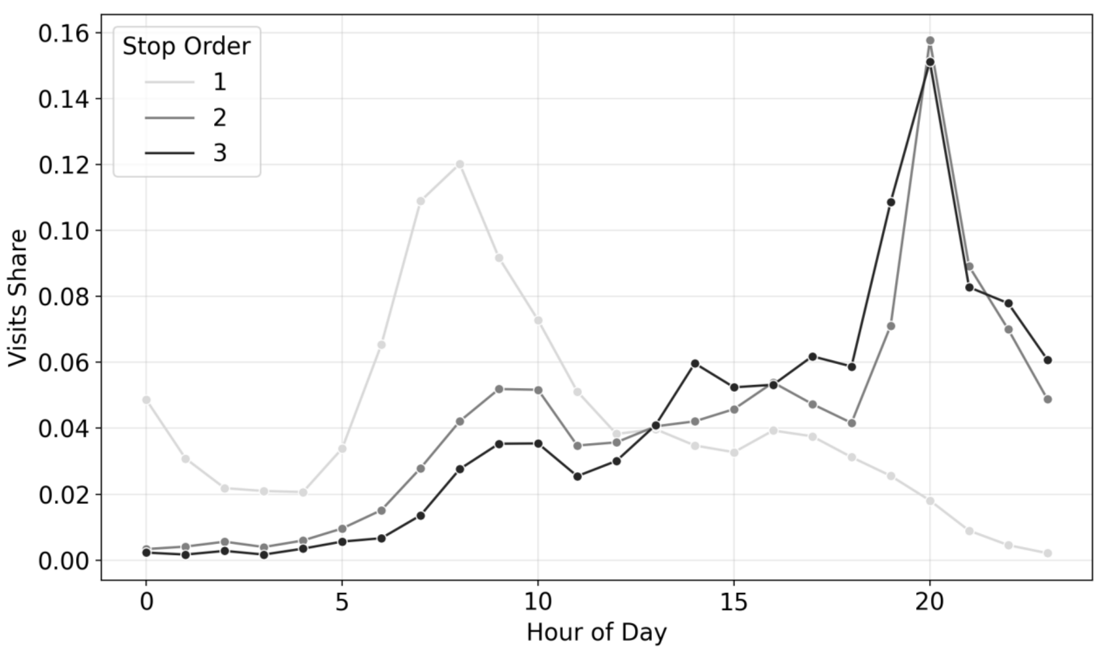

(c) hourly distribution of visits by stop order

Notes: Panel (a) shows that most devices record only one or two stops during a single-day beach visit, indicating limited within-day site switching. Panel (b) presents the time gap distribution between consecutive same-day stops, which is typically short (under 6 hours), consistent with continuous recreational activities rather than separate trips. Panel (c) plots the hourly visit shares by stop order, revealing that first stops tend to occur in the late morning and early afternoon, while second and third stops peak sharply around 7-8 PM, reflecting afternoon and early-evening beach revisits before departure.

Table 1. Comparison of Proposed and Alternative Practices

Category	**Proposed Practices**	**Alternative Practices**
Visit definition (Minimum Dwell Timing)	Apply a 10-minute dwell time threshold to exclude brief stops and consolidate multiple same-day visits.	Use a 4-minute threshold, potentially increasing visit counts but capturing non-recreational stops.
Visit definition (Multi-stop)	Consolidate multiple short-duration stops at the same site into a single visit.	Count each short-duration stop separately.
Choice Set	Define choice sets based on four years of visitation data (2019-2022) to capture long-term behavior.	Use only single-year visit data, potentially missing habitual behavior.
Weights	Apply sampling weights to adjust for underrepresented groups (e.g., older adults, lower-income individuals).	Use raw data without weighting.
Privacy	Use anonymized mobility data with home locations only at the CBG level.	Apply additional differential privacy techniques (Laplace noise, truncation, censoring) to obscure visit patterns.

Notes: This table presents a side-by-side comparison of proposed practices and alternative practices in processing mobility data for recreation demand estimation. Key differences include visit identification thresholds, choice set construction, demographic weighting, sensitivity analysis, and privacy considerations. The proposed practices are designed to minimize biases, enhance representativeness, and ensure robust demand estimation, while alternative practices highlight common variations in data processing that may affect results.

Table 2. Summary Statistics

Variable	Obs	Mean	Std. Dev.	Min	Max
Panel A: Proposed Practices					
Advisory Dummy	97,050	0.0241	0.1534	0	1
Visits	236,464	0.7441	7.1169	0	704
Total Visits	236,464	2223.813	2575.815	60	42,990
Travel Costs ($)	236,464	34.4089	49.9952	0.0570	416.4536
Observed Share	236,464	0.0006	0.0041	0	0.3138
Bayes Share	236,464	0.0005	0.0039	4.83e-18	0.3118
Panel B: Alternative Practices					
Advisory Dummy	28,553	0.0288	0.1673	0	1
Visits	67,879	3.8677	17.6057	0	904
Total Visits	67,879	2290.714	2532.237	60	42,990
Travel Costs ($)	67,879	38.0072	59.1619	0.0570	416.4536
Observed Share	67,879	0.0030	0.0122	0	0.9448
Bayes Share	67,879	0.0029	0.0119	1.77e-06	0.9436

Notes: This table presents summary statistics for key variables using proposed practices and alternative practices. The advisory variable is a binary indicator, where 1 denotes a site under a contamination advisory. Observed share refers to the proportion of total visits to a given site from a CBG in a year-quarter, while Bayesian-adjusted share corrects for zero-visit sites using an empirical Bayes adjustment.

Table 3. The Effects of Data Practices on MWTP for Avoiding Contamination

	Proposed Practices	**Alt: 4-Min Dwell**	**Alt: Diff. Privacy**	**Alt: Single-Stop**	**Alt: Short-Term Set**	**Alt: No Weights**
First Stage						
Travel Costs	-0.039***	-0.052***	-0.042***	-0.040***	-0.022***	-0.042***
	(0.006)	(0.006)	(0.008)	(0.006)	(0.002)	(0.006)
Observations	232,585	293,559	112,592	232,585	95,649	232,585
Second Stage						
Advisory	-0.348*	-0.198	-0.131	-0.350*	-0.215**	-0.378**
	(0.182)	(0.203)	(0.095)	(0.183)	(0.100)	(0.146)
Observations	2,874	2,925	2,807	2,874	1,791	2,874
MWTP	8.92	3.81	3.12	8.75	9.77	9.00
T-value of Diff.	-	(-73.68)	(-54.39)	(-9.30)	(11.02)	(-0.50)

Notes: MWTP represents the negative value of the ratio of the advisory coefficient to the travel cost coefficient. We generate 1000 bootstrapped values of MWTP, and run t-tests to examine the difference between alternative practices and proposed practices. Significance levels: *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Evaluating Precision, Privacy, and Representation with Cell Phone Data: Evidence from Cape Cod Beaches

Xibo Wan, Mike Vo, Cristina Connolly, and Sandro Steinbach

Appendices for Online Publication

These appendices supplement our article "Evaluating Precision, Privacy, and Representation with Cell Phone Data: Evidence from Cape Cod Beaches" with the following material:

- Online Appendix A includes additional details on the beach matching process, and Veraset data.
- Online Appendix B describes details of empirical Bayesian estimator for market shares.
- Online Appendix C reports the results of a set of robustness checks.

Appendix

Appendix A Additional details and evidence

Appendix A.1 Matching Advisories with Beach POIs

This study integrates multiple datasets to analyze how beach advisories and water quality influence visitation patterns. The primary data sources include BEACON 2.0, which provides information on beach attributes, advisories, and water quality monitoring, and Advan's POI data, which contains geospatial identifiers for locations such as beaches. By linking these datasets, we construct a comprehensive dataset that allows for a detailed examination of the relationship between environmental conditions and beach visitation.

To match BEACON 2.0 beach locations with Advan POIs, we first perform a spatial join between BEACON's beach location data and Advan's beach polygons. The BEACON dataset provides both starting and ending coordinates for each beach, which are converted into geospatial data using GeoPandas. These locations are then projected into EPSG:3857 (Web Mercator) to ensure consistency with Advan's dataset. Using a spatial join, each beach point is assigned to the nearest beach polygon from Advan's POI dataset. This step ensures that each beach is accurately linked to its corresponding POI, allowing us to associate beach names and IDs with Advan's standardized location identifiers.

Once the beaches are matched to Advan POIs, we link BEACON's advisory results to these locations. BEACON's water monitoring dataset contains test results recorded at specific beaches, identified by beach name or ID. Since we have already established a mapping between beach names and Advan POIs, we use this relationship to connect contamination advisory to the corresponding POI locations. To facilitate time-based analysis, we extract the month and year of each water quality measurement and aggregate the results at the monthly level. The processed dataset, which includes contamination advisory matched to POIs, is then saved for further analysis.

Appendix A.2 Distance Decay in Veraset Data

A likely reason for the pattern shown in Figure 2a is that it displays raw visit counts without normalizing by the number of residing devices in each CBG. Areas with larger observed device populations naturally register more visits, making distant CBGs appear comparably active. To address this, Appendix Figure 5 presents an adjusted analysis that estimates the distance-decay relationship between visitation probability and travel distance under alternative functional forms, log-log, log-linear, polynomial, and cubic B-spline, while controlling for site, CBG, and week fixed effects. These controls account for baseline differences in site popularity, home-market size, and temporal variation, isolating the underlying functional shape of distance decay. Across all specifications, the estimated relationships remain smooth and monotonically decreasing, confirming that the Spectus data exhibit realistic spatial gradients consistent with standard recreation-demand theory.

Figure A1: Distance-Decay Relationship Between Visitation and Travel Distance

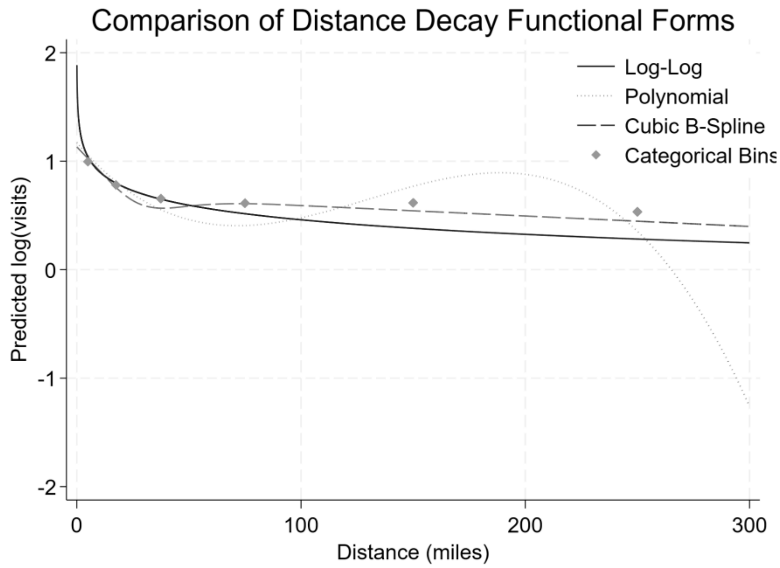

Notes: This figure compares alternative functional forms for modeling the distance-decay relationship between visitation and travel distance. The log-log specification captures a smooth, monotonic decline in visitation with increasing distance, while the polynomial and cubic B-spline allow for greater flexibility. The categorical-bin estimates (orange dots) provide nonparametric benchmarks for model comparison.

Appendix A.3 Data Coverage of Vareset Data

To evaluate the stability of the Veraset mobility data over time, we examined two complementary indicators of data coverage: (1) the sampling ratio, defined as the share of observed mobile devices relative to the census population within each CBG, and (2) POI coverage, reflecting the completeness of the beach visitation panel across years.

The Veraset dataset maintains consistent POI coverage for all monitored Cape Cod beaches from 2019 through 2022, with no entries or exits of major sites during the study period. This consistency ensures that temporal differences in visitation patterns are not driven by changing observation units.

Meanwhile, in Appendix Figure 6, the sampling ratio shows a steady and monotonic improvement in coverage over time, from approximately 12 percent in 2019 to 25 percent in 2022, based on mean estimates with non-overlapping 95 percent confidence intervals. This upward trend indicates a progressive increase in the share of devices contributing geolocation data, reflecting broader smartphone adoption and enhanced data vendor aggregation, rather than fluctuations in measurement quality.

Figure A2: Sampling Ratio Across CBGs Over Years

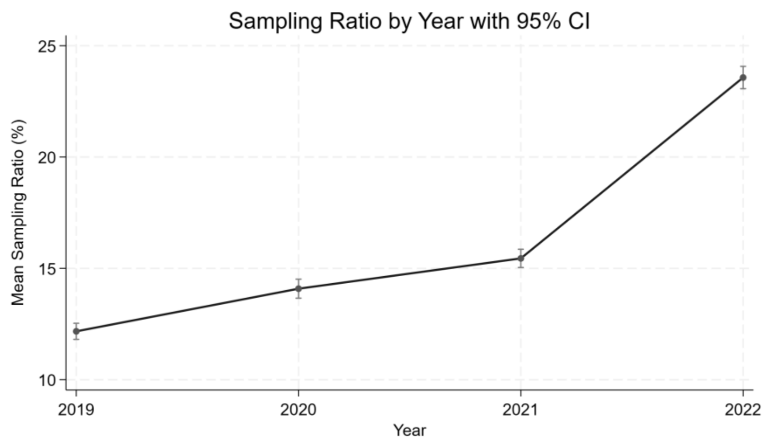

Notes: This figure shows the mean sampling ratio of mobile devices across CBGs from 2019 to 2022, with 95% confidence intervals. The sampling ratio gradually increased over time, reflecting both data provider expansion and improvements in device coverage.

Figure A3: Distribution of Dwell Time (mins)

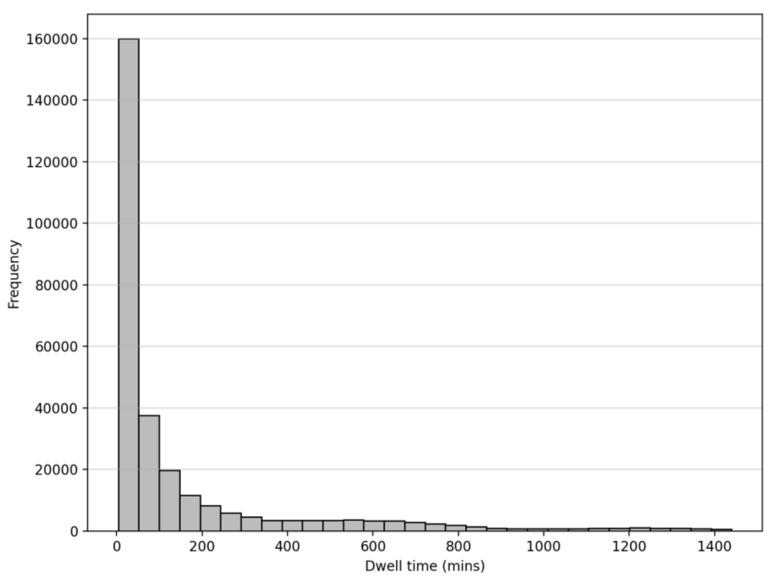

Notes: This figure displays the distribution of dwell times across all observed visits. Most visits are short, with a median duration under 30 minutes, while a small number of long stays generate a right-skewed tail.

Appendix B Empirical Bayes Estimator for Market Shares

This section examines how demand estimates vary depending on the choice of empirical Bayes priors. Appendix Table 4 presents estimation results from a logit demand model using both observed market shares and mean empirical Bayes posterior market shares constructed with different prior sample sizes. Column (1) is estimated using observed market shares, where observations with zero market share are dropped from the estimation. Columns (2) through (6) present results using empirical Bayes posterior means, each constructed with a different number of similar markets included 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 in the empirical Bayes prior. Across specifications, the empirical Bayes estimates differ from those based on observed market shares in Column (1). The most noticeable differences are found in key parameters such as the travel cost coefficient and advisory coefficient, suggesting that adjustments for zero shares meaningfully influence demand estimates. Results across Columns (4) through (6) remain relatively stable, with improved precision as the number of markets in the prior increases. This consistency indicates that the main findings of the paper are robust to different choices of large prior sample size. The counterfactual and welfare results in the main analysis use priors constructed from 50 markets.

Appendix Table 4, 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 maximum likelihood estimates (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 5 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.

Table B1. The Logit Demand Estimates with Empirical Bayes Share (All Data)

Panel A: 1st Stage	Observed	EB20	EB30	EB50	EB70	EB90
Travel Costs	-0.007***	-0.119***	-0.089***	-0.039***	-0.030***	-0.031***
	(0.001)	(0.014)	(0.013)	(0.006)	(0.003)	(0.002)
Observation	31918	232585	232585	232585	232585	232585
Panel B: 2nd Stage	Observed	EB20	EB30	EB50	EB70	EB90
Advisory	0.084	-0.684	-0.286	-0.348*	-0.220**	-0.129*
	(0.103)	(0.421)	(0.344)	(0.182)	(0.090)	(0.070)
Observations	1190	2874	2874	2874	2874	2874

Notes: This table presents logit IV demand estimates for different market share construction approaches. Column (1) uses observed market shares and has fewer observations because zeros are dropped in estimation. Columns (2) through (6) use mean empirical Bayes (denoted "EB") posteriors as market shares, with different numbers of markets in the prior. Column (2) has 20 CBGs closest in travel distance in the prior, Column (3) has 30 CBGs, and so on, until Column (6) has 90 markets in the prior. Standard errors in parentheses are clustered by CBG. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table B2. Summary statistics of Observed Shares and Empirical Bayes Posterior Shares

Variable	Obs	Mean	Std. Dev.	Min	Max
Panel A: All data					
Observed Share	236,464	0.00056	0.00407	0.00000	0.31379
Bayes Share (20 CBGs in prior)	236,464	0.00051	0.00377	0.00000	0.31177
Bayes Share (30 CBGs in prior)	236,464	0.00051	0.00380	0.00000	0.31177
Bayes Share (50 CBGs in prior)	236,464	0.00052	0.00387	0.00000	0.31177
Bayes Share (70 CBGs in prior)	236,464	0.00053	0.00390	0.00000	0.30571
Bayes Share (90 CBGs in prior)	236,464	0.00053	0.00391	0.00000	0.31236
Panel B: Nonzero subsample					
Observed Share	36,361	0.00365	0.00982	0.00003	0.31379
Bayes Share (20 CBGs in prior)	36,361	0.00296	0.00924	0.00001	0.31177
Bayes Share (30 CBGs in prior)	36,361	0.00300	0.00930	0.00000	0.31177
Bayes Share (50 CBGs in prior)	36,361	0.00309	0.00947	0.00000	0.31177
Bayes Share (70 CBGs in prior)	36,361	0.00314	0.00953	0.00000	0.30571
Bayes Share (90 CBGs in prior)	36,361	0.00316	0.00954	0.00000	0.31236
Panel C: Zeros subsample					
Observed Share	200,103	0.00000	0.00000	0.00000	0.00000
Bayes Share (20 CBGs in prior)	200,103	0.00007	0.00013	0.00000	0.00760
Bayes Share (30 CBGs in prior)	200,103	0.00006	0.00010	0.00000	0.00550
Bayes Share (50 CBGs in prior)	200,103	0.00006	0.00008	0.00000	0.00399
Bayes Share (70 CBGs in prior)	200,103	0.00005	0.00007	0.00000	0.00382
Bayes Share (90 CBGs in prior)	200,103	0.00005	0.00007	0.00000	0.00353

Notes: This table presents summary statistics for Observed Shares and Empirical Bayes Posterior Shares, categorized into three panels: (A) All data, (B) Nonzero subsample, and (C) Zero subsample. Observed Shares represent raw market shares, while Bayes Shares incorporate empirical Bayes priors with different sample sizes of CBGs with similar distance used in prior construction.

Appendix C Robustness Checks

Table C1. Robustness Checks: Alternative Minimum Dwell Time

	(1)	(2)	(3)	(4)
Panel A: First Stage				
Travel Costs	-0.039***	-0.039***	-0.023***	-0.028***
	(0.006)	(0.007)	(0.005)	(0.006)
Observations	232,585	150,097	103,742	79,102
Panel B: Second Stage				
Advisory	-0.348*	-0.485***	-0.430***	-0.151
	(0.182)	(0.136)	(0.129)	(0.130)
Observations	2,874	2,840	2,728	2,663
Min. Dwell Time (mins)	10	30	60	90

Notes: Each column corresponds to a different minimum dwell time for a recreational visit used to define the choice set. Panel A reports first-stage coefficients on travel costs; Panel B reports second-stage coefficients on the advisory dummy. Robust standard errors in parentheses. Significance levels: *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table C2. Robustness Checks: Alternative Distance Cutoffs

	(1)	(2)	(3)	(4)	(5)
Panel A: First Stage					
Travel Costs	-0.039***	-0.053***	-0.047***	-0.043***	-0.040***
	(0.006)	(0.008)	(0.007)	(0.007)	(0.006)
Observations	232,585	215,807	225,473	229,249	232,268
Panel B: Second Stage					
Advisory	-0.348*	-0.331*	-0.359*	-0.351*	-0.349*
	(0.182)	(0.182)	(0.184)	(0.183)	(0.182)
Observations	2,874	2,872	2,872	2,874	2,874
Distance Cutoff (mile)	300	100	150	200	250

Notes: Each column corresponds to a different maximum travel distance used to define the choice set. Panel A reports first-stage coefficients on travel costs; Panel B reports second-stage coefficients on the advisory dummy. Robust standard errors in parentheses. Significance levels: *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table C3. Robustness Checks: Lag of Advisories

	(1)	(2)
First Stage		
Travel Costs	-0.039***	-0.039***
	(0.006)	(0.006)
Observations	232,585	232,585
Second Stage		
Advisory	-0.348*	-0.327**
	(0.182)	(0.141)
Advisory (Lag)		-0.059
		(0.178)
Observations	2,874	2,815
MWTP	8.92	8.38

Notes: Column (1) reproduces the baseline estimates from Table 2 (Proposed Practices). Column (2) reports the updated results including a lagged advisory variable. Robust standard errors in parentheses. Significance levels: *** $p<0.01$, ** $p<0.05$, * $p<0.1$.