Management Science

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Closer to Home: A Structural Estimate-Then-Optimize Approach to Improve Access to Healthcare Services

Fernanda Bravo, Ashvin Gandhi, Jingyuan Hu, Elisa F. Long

To cite this article:
Fernanda Bravo, Ashvin Gandhi, Jingyuan Hu, Elisa F. Long (2026) Closer to Home: A Structural Estimate-Then-Optimize Approach to Improve Access to Healthcare Services. Management Science 72(3):2345-2363. https://doi.org/10.1287/mnsc.2024.06274

https://pubsonline.informs.org/journal/mnsc

MANAGEMENT SCIENCE
Vol. 72, No. 3, March 2026, pp. 2345–2363
ISSN 0025-1909 (print), ISSN 1526-5501 (online)

Closer to Home: A Structural Estimate-Then-Optimize Approach to Improve Access to Healthcare Services

Fernanda Bravo,[a] Ashvin Gandhi,[a] Jingyuan Hu,[a,b] Elisa F. Long[a,*]

[a] Anderson School of Management, University of California, Los Angeles, Los Angeles, California 90095; [b] Paul Merage School of Business, University of California, Irvine, Irvine, California 92697

*Corresponding author

Contact: fernanda.bravo@anderson.ucla.edu, https://orcid.org/0000-0002-4625-7894 (FB); ashvin.gandhi@anderson.ucla.edu, https://orcid.org/0000-0002-7875-8773 (AG); jingyuan.hu@uci.edu, https://orcid.org/0009-0006-9741-8703 (JH); elisa.long@anderson.ucla.edu, https://orcid.org/0000-0001-9553-1151 (EFL)

Received: May 30, 2024
Revised: December 12, 2024; March 4, 2025
Accepted: March 10, 2025
Published Online in Articles in Advance: July 23, 2025

https://doi.org/10.1287/mnsc.2024.06274

Abstract. Geographic inequalities in healthcare access extend beyond rural–urban divides to include socioeconomic, racial, and other disparities. Proximity to hospitals, clinics, healthcare providers, and pharmacies varies widely, posing a challenge in deciding where to strategically locate such facilities. Demand for each service depends on local population health, individual preferences, provider capacity, and other factors. This study introduces a novel structural estimate-then-optimize (SETO) framework, combining structural demand estimation using a modified Berry–Levinsohn–Pakes approach that accounts for provider capacity with a choice-based optimal facility location model to maximize health service utilization. Our methodology is illustrated with a case study on the Federal Retail Pharmacy Program in California, a public–private partnership that administered millions of COVID-19 vaccinations. Demand estimates indicate that residents of socioeconomically vulnerable communities are more sensitive to travel distances to pharmacy-based vaccination sites. Strategically adding 500 retail stores serving lower-income communities increases predicted vaccinations by 2.9% overall (770,000 additional vaccinations statewide) and by 5.3% in the least healthy neighborhoods. Our integrative SETO approach outperforms heuristics that allocate resources based on current vaccination rates, existing service gaps, population density, or predicted demand. The case study demonstrates the importance of (1) accounting for heterogeneity in estimating demand and (2) selecting partnerships to complement existing networks with spatially heterogeneous supply and efficiently fill service gaps. Our study provides a systematic approach to optimize healthcare delivery networks, using publicly available aggregate data while accounting for individuals' preferences, highlighting the value of combining a structural demand model with prescriptive analytics.

History: Accepted by Stefan Sholtes, healthcare management.

Supplemental Material: The online appendix and data files are available at https://doi.org/10.1287/mnsc.2024.06274.

Keywords: structural estimation • BLP • choice model • facility location • healthcare access

1. Introduction

Across the United States, access to essential healthcare resources, such as acute care hospital beds, maternity units, primary care physicians, or pharmacy services, often depends on location (Bronner et al. 2021). Numerous studies show that proximity to a health service affects utilization, expenditures, and patient health outcomes. More than half of the geographic variation in healthcare utilization are explained by supply-side differences, with the rest attributed to patient-related demand (Finkelstein et al. 2016). Proximity to a hospital affects utilization of emergency care, such as treatment following a heart attack (Tay 2003), and routine inpatient services (Gowrisankaran et al. 2015, Shepard 2022). Similarly, nursing home facility choice is highly sensitive to distance, affecting both patient selection (Gandhi 2023) and local competition (Hackmann 2019).

Long distances to health service providers pose greater challenges for some communities. Economically and socially vulnerable individuals face additional hurdles

such as limited time off work, lower vehicle ownership, and limited public transit access (Syed et al. 2013). Rural Americans generally have worse health outcomes and shorter life expectancy (Deryugina and Molitor 2021), owing in part to fewer and lower quality healthcare providers (Finkelstein et al. 2021). Although cities have higher concentrations of physicians, many urban areas face shortages of primary care and specialist providers (Brown et al. 2016).

Pharmacists play a crucial role in dispensing medications, consulting with patients, and administering vaccinations. Five percent of Americans live in *pharmacy deserts*—areas without a pharmacy in a one-mile radius (10 miles in rural areas)—and lack access to essential services (Wittenauer et al. 2024). Pharmacy deserts increasingly arise as retail pharmacies close and the industry consolidates (Salako et al. 2018). At low-volume pharmacies, Medicaid reimbursements are often insufficient to cover operating costs, leading to further closures (Ippolito et al. 2020). Whereas pharmacy deserts are common in rural areas, some densely populated regions, particularly Black and Hispanic neighborhoods, also face pharmacist shortages (Guadamuz et al. 2021). In California, 2.5 million residents (6%) live in pharmacy deserts, including 400,000 in urban Los Angeles (LA; Wittenauer et al. 2024). Figure 1 highlights the inequality in pharmacy access by the Healthy Places Index (HPI), a composite measure of community well-being (Public Health Alliance of Southern California 2022).

Mitigating spatial inequalities in healthcare access has urgent policy importance. One potential mechanism is better positioning of medical providers, given a well-defined objective (e.g., maximize service utilization, minimize average distance traveled). However, this requires an understanding of individuals' preferences in selecting a provider. One empirical challenge is the lack of individual outcome data due to privacy concerns or constraints in obtaining patient health records. For example, disease prevalence is reported at an aggregate geographic level, and policymakers typically lack provider-level outcome data or information on which patients visited each healthcare facility.

Figure 1. (Color online) Average Distance to a Pharmacy Offering COVID-19 Vaccinations in California by HPI Quartile

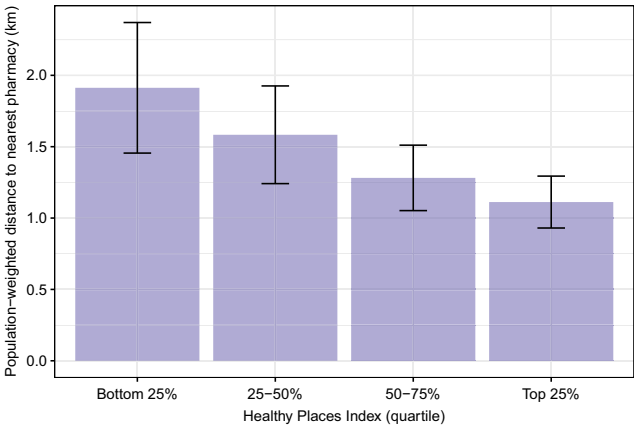

Note. Bars correspond to the population-weighted average distances by zip code within each HPI quartile, and error bars are 95% confidence intervals.

Consider an undifferentiated service, such as routine mammography, where service needs are *heterogeneous* across patient groups and individual preferences for a service provider vary. Provider availability differs by geographic area, contributing to inequities in access. Table 1 summarizes the key features often present in the healthcare contexts that motivate our work. Given these challenges, our study addresses an important question in healthcare delivery: *How to optimally design a network of providers offering a standard service, to improve access to care?* This design could involve adding, closing, or replacing facilities. We propose a data-driven approach to optimally design the healthcare delivery network, with an aim to realign a mismatch between a supply of service providers and patient demand for such service (Figure 2). Our contributions are summarized as follows:

- *Structural estimate-then-optimize (SETO) approach.* We develop a novel two-stage approach where the first stage corresponds to *structural demand estimation* that relies only on aggregated, publicly available healthcare utilization data. We employ a modified Berry–Levinsohn–Pakes (BLP) approach that, beyond allowing for spatial demand heterogeneity, also incorporates facility capacity limits (Berry et al. 1995). The second stage is an *optimization approach* that integrates the demand estimates into a choice-based facility location problem to identify the optimal set of locations that maximize predicted uptake. Our framework links an established technique from economics—structural demand estimation—with optimization, harnessing observational aggregated data to inform prescriptive modeling.
- *Case study on vaccine distribution.* Our study demonstrates the potential value of integrating structural demand estimation with prescriptive analytics to design more equitable and efficient vaccination networks. In this context, heterogeneity exists both on the demand

Table 1. Healthcare Features That Support a Structural Estimate-Then-Optimize Framework

Stakeholder	Healthcare characteristic
Demand side (by patients)	Only aggregate data on service utilization (no individual-level outcomes)
	Heterogeneous sensitivity to provider's proximity
Supply side (by providers)	Spatial inequality in the distribution of service providers
	Unobserved utilization rates by provider
	Limited service capacity per provider
	Undifferentiated health service

Figure 2. (Color online) Proposed Structural Estimate-Then Optimize (SETO) Modeling Framework

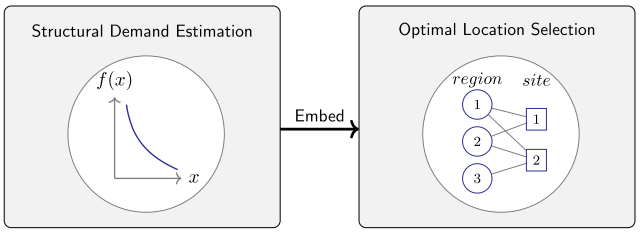

side as socioeconomically vulnerable communities show greater sensitivity to travel distances, and on the supply side as existing pharmacy sites fail to reach underserved populations. Using only aggregated data, our first stage estimates these heterogeneous demand elasticities by HPI, and our second stage optimally selects locations that complement the existing network, reducing service gaps and improving access. Our integrated SETO approach delivers more predicted gains than heuristic approaches that rely on demand estimation or optimization, but not both.

2. Related Literature

Our study builds on extensive prior work in healthcare accessibility, optimal facility location, and demand estimation. Each topic spans multiple fields, and we highlight the most relevant studies.

2.1. Accessibility of Health Services

Healthcare access has been a topic of discussion for the past half-century (Aday and Andersen 1974). Penchansky and Thomas (1981) outline five dimensions that define the relationship between patients and health systems: *accessibility* concerns the geographical proximity of healthcare providers; *availability* relates to services offered meeting patient needs; *accommodation* concerns health system features such as appointment scheduling; *acceptability* considers whether patients find providers' attributes (e.g., gender, race) acceptable; and *affordability* addresses whether costs are within the financial means of patients or their insurers. Ayer et al. (2014) summarize research within operations management that examine several of these dimensions.

Socioeconomic disparities in healthcare access are evident in various healthcare contexts including organ transplantation (Wang et al. 2022), decision support systems (Ganju et al. 2020), and appointment scheduling (Samorani et al. 2022). For patients needing treatment for chronic diseases, distance to a provider can be a major barrier to accessing care (Kelly et al. 2016). For instance, use of radiation therapy for breast cancer is notably lower for women residing farther from a provider, even after adjusting for socioeconomic status. Survey data confirm that lack of transportation correlates with lower healthcare utilization, particularly in rural communities (Arcury et al. 2005).

Prior studies show a link between proximity to influenza vaccination providers and uptake (Beshears et al. 2016). During the COVID pandemic, access to a retail pharmacy was associated with higher vaccination rates (Brownstein et al. 2022), including for residents of California, regardless of political background (Mazar et al. 2023). Although these studies report a correlation between distance and vaccine uptake, our demand estimation is more formal and economically microfounded. Despite a national vaccination campaign launching in 2021, vaccine availability was limited in rural and socioeconomically disadvantaged areas (Zhong et al. 2021, Rader et al. 2022). Chevalier et al. (2022) find that offering COVID vaccinations at Dollar General stores could substantially benefit low-income and minority households. Improving vaccination access has been studied under various objectives (Enayati and Özaltın 2020, Rastegar et al. 2021). Bennouna et al. (2023) use a machine learning–based epidemic model and optimization to efficiently allocate vaccines to different populations. Dey et al. (2024) provide a recent literature review.

Our case study extends these efforts by precisely estimating demand for a COVID vaccine under service provider capacity constraints and formulating an optimization model to strategically choose vaccination delivery sites. Our study is the first to *structurally estimate* different vaccination elasticities to travel distance—a heterogeneous demand curve—using only aggregate data. Our methodology could extend to other contexts with varying distance sensitivities and unequal provider availability, such as cancer screenings, kidney dialysis, and substance abuse treatment.

2.2. Optimal Facility Location

A wide literature exists on optimal facility location, particularly the maximal covering location problem (MCLP; Church and ReVelle 1974), including extensive applications in healthcare (Daskin and Dean 2005, Jia et al. 2007, Ahmadi-Javid et al. 2017). Under stochastic demand, MCLP models allow for customer queuing at each location. Ekici et al. (2014) formulate an MCLP with underlying epidemic model to project the spread of an infectious disease. Deo and Sohoni (2015) and Jónasson et al. (2017) assign HIV testing capacity across geographically distant testing facilities, while allowing for congestion and the test collection rate to decline in turnaround time. Whereas these studies focus on optimal capacity allocation, our study relates to *strategic network design* (i.e., service locations) for equitable healthcare access. Prior studies have integrated distance sensitivity into MCLPs in healthcare settings including primary care (Nobles et al. 2014), influenza vaccination (Heier Stamm et al. 2017),

and HIV care (McCoy and Johnson 2014). These studies assume a homogeneous linear relationship between distance and service utilization, whereas our model incorporates empirically estimated heterogeneous demand curves, resulting in more equitable resource allocation than under a homogeneous demand model.

In contexts with uncertain travel distances and costly service delays, robust optimization can be used. Boutilier and Chan (2020) use predict-then-optimize to predict daily demand for emergency medical services and then optimally locate and route ambulances using robust optimization. Whereas they consider demand to be exogenous, we estimate a demand curve that is *endogenously* determined based on the set of service providers. In other words, we model demand for service as changing in response to the underlying cost (e.g., travel distance).

Most healthcare facility location studies assume a central planner assigns individuals to facilities. Lee et al. (2013) optimally select mass medical dispensing sites during an outbreak and assign households to facilities. Bertsimas et al. (2022) combine a dynamic compartmental model of COVID-19 with optimization to forecast the epidemic's trajectory, optimally select mass vaccination sites, assign populations to each site, and allocate vaccines by age group. Whereas their resulting strategy is driven by epidemic conditions, we embed individual preferences into a facility location problem via a choice model. *Choice-based* location models examine how individual preferences, influenced by factors such as distance or service quality, shape facility location strategies. Such preferences may be captured using a multinomial logit (MNL) discrete choice model, and then subsequently incorporated in an optimization model. For example, in choosing a healthcare provider, studies vary in modeling patient utility based on price and quality (Denoyel et al. 2017), time sensitivity (Zhang et al. 2012), waiting time (Krohn et al. 2021), or travel time (Hwang et al. 2022). Notably, the first three studies do not include empirical demand estimation but instead perform calibration exercises. In a study of maternity care, Hwang et al. (2022) use aggregate geographic covariates to estimate value function parameters. In contrast, our utility function explicitly depends on proximity to a provider and allows for heterogeneous distance elasticities. Without individual-level choice data, we rely on the BLP approach to structurally estimate demand.

Our approach is essentially an assortment optimization problem under a mixture of MNL models, with customer segments having different distance elasticities, along with capacity and cardinality constraints. Bront et al. (2009) show that the problem is NP-hard if the number of customer types reaches the number of products. Others have proposed integer programming (Méndez-Díaz et al. 2014) and conic programming (Şen et al. 2018) solution techniques. Similarly, we provide a mixed-integer problem that reformulates nonlinear terms to achieve a tractable facility location problem.

2.3. Structural Demand Estimation

The primary input into a choice-based facility location model is the demand curve for obtaining service, as a function of travel distance. We draw on the established literature from economics on structural demand estimation, which follows a random-coefficient logit model using aggregate choice data. This approach, known as BLP (Berry et al. 1995), has been widely used because of its ability to capture individual-level preferences using only aggregate data, and to address endogeneity using instruments.

More advanced implementations of BLP incorporate techniques such as utilizing moments from disaggregated data (Petrin 2002, Berry et al. 2004), minimizing the impact of parametric error (Berry and Pakes 2007), or including spatial preferences (Davis 2006), among other variations. Many applications of BLP use structural models to measure industry behavior (Nevo 2001), or to quantify specific counterfactual scenarios (e.g., Duch-Brown et al. (2023) studies online market integration). The operations management literature has gradually adopted BLP estimation for modeling demand, with applications in the fast-food industry (Allon et al. 2011), automobile sector (Guajardo et al. 2016), and online physician services (Xu et al. 2021). We expand on the classic BLP method by (1) incorporating capacity constraints via a fixed-point equation approach and (2) embedding the estimated demand model into a prescriptive optimization model.

2.4. Predict-Then-Optimize

This approach is widely used in decision-making problems with uncertain parameters. First, machine learning is used to estimate these parameters by minimizing prediction error through a loss function (e.g., mean squared error in regression). Then, the estimated parameters are incorporated into an optimization model to determine the best decision based on an objective function and relevant constraints.

One potential concern of this method is that small prediction errors can lead to suboptimal decisions, as noted by Tan and Frazier (2022), who analyze regret bounds for unconstrained optimization problems with smooth objectives. Robust optimization can mitigate sensitivity to estimation errors, leading to more conservative yet stable decisions (Bertsimas and Sim 2004). Recent integrative approaches guide parameter estimation using a decision-aware loss function, optimizing decision quality rather than predictive accuracy (Chung et al. 2022, Elmachtoub and Grigas 2022).

Our framework shares similarities with predict-then-optimize but differs in a key aspect: we employ an econometric approach to estimate demand using only aggregated service utilization data. Notably, our aim is not to maximize predictive accuracy, but to obtain unbiased estimates that rationalize the observed utilization rates.

3. Structural Estimate-Then-Optimize Framework

Our framework consists of a microfounded model of consumer demand for a healthcare service that relies only on aggregated data on service utilization. Demand curves represent the sum of decisions made by individuals with varying preferences and distances to a healthcare provider. Crucially, aggregating these individual-level decisions offers a structured way to estimate demand for a service given the locations of service providers. Last, we incorporate the estimated demand curve into an optimization model to select facility locations to maximize predicted service utilization.

3.1. Demand Model

Consider a set of locations $\mathcal{L}$ that provide an undifferentiated health service, such as eye exams, vaccinations, blood collections, etc. Individual i decides whether to receive service at a location $\ell_i \in L_i$ (referred to as "provider location" or "service location") from among i's feasible set of provider locations, $L_i \subset \mathcal{L}$. The location represents the provider that i most plausibly considers when choosing whether to obtain service and can vary across settings. For example, ℓ_i could be the most convenient location with available capacity, or ℓ_i might be the closest provider covered by i's insurance. In other contexts, such as a centrally organized blood drive or vaccination campaign, ℓ_i might be a service location assigned by a central planner.

A key feature of demand that we aim to capture is how consumers' decisions to obtain the health service depend on the proximity of available providers. Let $a_i \in \mathcal{A}$ denote the geographic area (e.g., city or zip code) and $s_i \in a_i$ denote a smaller subarea within a_i (e.g., census tract or block) where individual i resides. Let $d(s_i, \ell_i)$ denote a measure of travel distance (e.g., geodesic distance, travel time, public transit time) between the centroid of s_i and ℓ_i. Figure 3 illustrates an area a, with arrows representing the distance between each subarea s and the closest service location, ℓ or ℓ'.

3.1.1. Individual Preferences. We introduce a utility-based binary choice model. Individual i may choose to obtain service at location ℓ_i and receive additional utility u_i relative to no service:[1]

$$u_i = D_i + X_{a_i}\boldsymbol{\beta} + \xi_{a_i} + \epsilon_i. \tag{1}$$

A key term in this model is D_i, the disutility that individual i experiences from traveling from home subarea s_i to service location ℓ_i. We define $D_i := d(s_i, \ell_i)\alpha_{a_i}$, where α_{a_i} represents the individual's area-specific sensitivity to travel distance. The variation in α_{a_i} captures across-area differences in the travel cost to a service location (e.g., due to the availability of public transit).[2] We denote the parameters characterizing distance sensitivity by the vector $\boldsymbol{\alpha}$. The distance function $d()$ measures individual i's distance from the centroid of s_i to location ℓ_i, a reasonable assumption given the small geographic size of subareas. The vector $\boldsymbol{\beta}$ captures the systematic relationship between the area's observable characteristics X_{a_i} (e.g., demographics) and average preferences for obtaining the service. The term ξ_{a_i} represents the idiosyncratic preferences of area a_i due to unobservable factors, and ϵ_i is a logistic preference specific to i.

Under the utility model in Equation (1), individual i's demand follows the classic binary logit discrete choice model where the probability that she obtains service is

$$P_\epsilon(u_i \geq 0) = \frac{\exp(d(s_i, \ell_i)\alpha_{a_i} + X_{a_i}\boldsymbol{\beta} + \xi_{a_i})}{1 + \exp(d(s_i, \ell_i)\alpha_{a_i} + X_{a_i}\boldsymbol{\beta} + \xi_{a_i})}. \tag{2}$$

3.1.2. Demand Estimation. With individual-level choice data and exogenous variation in s_i, l_i, and X_{a_i}, it is straightforward to estimate the demand model from Section 3.1.1 via maximum likelihood. However, such favorable settings are exceedingly uncommon. When consumer data are highly sensitive or proprietary, a frequent occurrence in healthcare, we may observe only aggregate data (Table 1). We instead estimate demand using the BLP approach (Berry et al. 1995), which captures individual-level factors influencing decisions without directly observing individual choices. We consider the case where service utilization is observed by

Figure 3. (Color online) Conceptual Model of BLP Demand Estimation Given Aggregate Outcome Data

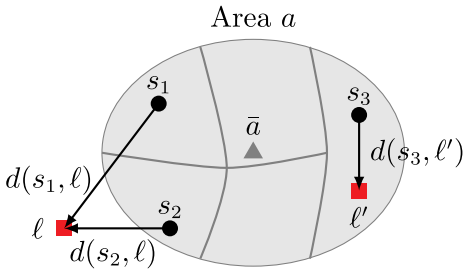

Note. Area a has population centroid $\bar{a}$ and subarea-level (s_1, s_2, s_3) variation in distance to a healthcare provider (ℓ and ℓ')

geographic area. BLP incorporates factors affecting demand that vary across individuals within a geographic area by treating these factors as *distributions* (e.g., the distribution of income or household locations within an area). These factors are commonly known as "random coefficients." Area-level demand curves are obtained by integrating individual-level demand curves over the distributions of random coefficients.

In our setting, the distance that each individual travels to obtain service, $d(s_i, \ell_i)$, varies across individuals within a given area. Correspondingly, travel disutility $D_i := d(s_i, \ell_i)\alpha_{a_i}$ is a "random coefficient" within the BLP framework. The share of individuals in area $a \in \mathcal{A}$ receiving the service is computed by aggregating demand from all individuals living in a:

$$\begin{aligned}\rho_a &:= \sum_{s_i \in a} \left[\int P_\epsilon(u_i \geq 0 | D_i, s_i) dF(D_i | s_i) \right] f(s_i | a) \\ &= \sum_{s_i \in a} \left[\int \frac{\exp(D_i + X_a \boldsymbol{\beta} + \xi_a)}{1 + \exp(D_i + X_a \boldsymbol{\beta} + \xi_a)} dF(D_i | s_i) \right] f(s_i | a) \\ &= \sum_{s_i \in a} \left[\sum_{\ell_i \in L_i} \frac{\exp(d(s_i, \ell_i)\alpha_a + X_a \boldsymbol{\beta} + \xi_a)}{1 + \exp(d(s_i, \ell_i)\alpha_a + X_a \boldsymbol{\beta} + \xi_a)} f(\ell_i | s_i) \right] f(s_i | a) \\ &= \sum_{s_i \in a} \sum_{\ell_i \in L_i} \frac{\exp(d(s_i, \ell_i)\alpha_a + X_a \boldsymbol{\beta} + \xi_a)}{1 + \exp(d(s_i, \ell_i)\alpha_a + X_a \boldsymbol{\beta} + \xi_a)} f(s_i, \ell_i | a), \end{aligned} \tag{3}$$

where $F(\cdot)$ and $f(\cdot)$ denote cumulative distribution functions and probability density functions, respectively. The third equality follows from the definition of D_i, and the distance measure $d(s_i, \ell_i)$ is defined by the set of discrete locations in L_i.

In computing the aggregate demand for area a in Equation (3), a key input is the distribution over travel costs $f(s_i, \ell_i | a)$, which can be decomposed as follows:

$$f(s_i, \ell_i | a) := \underbrace{f(s_i | a)}_{\substack{\text{Spatial} \\ \text{population}}} \underbrace{f(\ell_i | s_i)}_{\substack{\text{Service} \\ \text{location}}}. \tag{4}$$

The first term, $f(s_i | a)$, is the observed spatial distribution of the population across subareas within area a. This is a reasonable assumption given that public data typically include an extremely granular spatial distribution of the population (e.g., census blocks). The second term, $f(\ell_i | s_i)$,[3] is a distribution over service locations, characterized by a set of parameters, $\boldsymbol{\lambda}$, to be estimated. In some cases, formulating this term is straightforward (e.g., if a central planner assigns individuals to the closest facility). In our case study, $f(\ell_i | s_i)$ is determined sequentially (discussed in Section 3.1.3), allowing each individual to receive service at her *preferred* location with remaining capacity.

The parameters to be estimated are $\boldsymbol{\theta} := (\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\lambda})$. BLP exploits the idea that for any given $\boldsymbol{\theta}$, exactly one set of idiosyncratic shocks $(\xi_a)_{a \in \mathcal{A}}$ will precisely rationalize the observed utilization rates (Berry 1994). Thus, in enforcing that these area-level utilization rates implied by the demand model in Equation (3) match the observed data, the unobserved area-level idiosyncratic preferences are an implicit function of $\boldsymbol{\theta}$. We indicate this relationship by denoting ξ_a as a function of $\boldsymbol{\theta}$. Following Berry et al. (1995), we estimate $\boldsymbol{\theta}$ using orthogonality conditions:

$$E_a[Z_a(\boldsymbol{\theta}) \xi_a(\boldsymbol{\theta})] = \vec{0}\,. \tag{5}$$

This condition indicates that, at the true parameter value $\boldsymbol{\theta}$, the implied idiosyncratic preferences for each area, $\xi_a(\boldsymbol{\theta})$, are uncorrelated with the area-level instruments $Z_a(\boldsymbol{\theta})$. Because $\xi_a(\boldsymbol{\theta})$ is a utility residual from the demand model, Equation (5) enforces that the residual variation in the utility model evaluated at the true $\boldsymbol{\theta}$ not be correlated with the instruments. The instruments Z_a must be chosen so that the moment conditions are both informative and plausible. In our case study, Z_a includes the observable area-level characteristics X_a, and the average distance to service providers in the area interacted with area-level demographics. Because service locations vary with $\boldsymbol{\theta}$, these latter instruments vary with $\boldsymbol{\theta}$ as well. We therefore denote the instruments Z_a as a function of $\boldsymbol{\theta}$.

We leverage the moments in Equation (5) via the generalized method of moments (GMM):[4]

$$\hat{\boldsymbol{\theta}} = \arg\min_{\boldsymbol{\theta}} \left(\frac{1}{|\mathcal{A}|} \sum_a Z_a(\boldsymbol{\theta}) \xi_a(\boldsymbol{\theta}) \right)' \Phi \left(\frac{1}{|\mathcal{A}|} \sum_a Z_a(\boldsymbol{\theta}) \xi_a(\boldsymbol{\theta}) \right), \tag{6}$$

where optimal weighting matrix Φ is chosen via an efficient two-step GMM (Hansen 1982). A description of the solution approach to obtain $\hat{\boldsymbol{\theta}}$ is provided in Online Appendix B, Section B.1.

3.1.3. Service Location Distribution.

We have presented the demand model and estimation quite generally to emphasize its flexibility. Here, we describe a specific model of sequential location selection used in estimation, connecting service provider availability to individual choices.

We assume that an individual's choice of service location $\ell_i \in L_i$ is a function of its proximity $d(s_i, \ell_i)$ and her sensitivity to travel distance α_{a_i}. Under this assumption, no additional parameters are introduced in the model. In other words, to determine the probability distribution over service locations, $f(\ell_i | s_i)$, the only parameters needed are the sensitivity to distance $\boldsymbol{\alpha}$. We assume a type I extreme value (i.e., Gumbel) distribution of idiosyncratic preferences over locations, resulting in an MNL discrete choice model over the set of service locations L_i:

$$f(\ell_i | s_i, L_i) := \frac{\exp(d(s_i, \ell_i)\alpha_{a_i})}{\sum_{\ell \in L_i} \exp(d(s_i, \ell)\alpha_{a_i})}. \tag{7}$$

If an individual dislikes traveling longer distances (i.e., $\alpha_{a_i} < 0$), then Equation (7) implies that i's likelihood of choosing location ℓ is decreasing in distance. A reasonable alternative model that takes this to an extreme is to assume individual i chooses the closest site within L_i.[5] A key advantage of using an MNL model instead is that it may better capture individuals' idiosyncratic preferences in their preferred location, such as preferring a site close to their workplace rather than home.

If all providers have infinite capacity, individuals could choose their service location from among all sites (i.e., that $L_i = \mathcal{L}$). In reality, healthcare providers have finite capacity, limiting an individual's set of available providers. To capture this, we assume that individuals' choice of service provider is fulfilled sequentially in a random order, where individual i's set of possible service locations ($L_i \subset \mathcal{L}$) is restricted to those *with remaining capacity*.[6] A sequential approach could reflect a random arrival process of identical individual types requesting service slots, such as an online booking of vaccination appointments. To implement this, we first randomly generate an ordering of all individuals in the population. Then, we iterate over the individuals and update each choice set L_i according to the location choices (Equation (7)) and service choices (Equation (2)) made by individuals $1,\dots,i-1$, so that only locations with remaining capacity are included in L_i. Thus, if location ℓ reaches capacity at individual j's turn, then ℓ is removed from the choice sets of all individuals following j. Finally, $f(\ell_i|s_i)$ in Equation (4) is obtained by integrating $f(\ell_i|s_i, L_i)$ in Equation (7) over the L_i for individuals who reside in subarea $s_i \in a$.

3.2. Optimization Model: Location Selection

We next formulate a facility location model to determine the optimal set of healthcare service locations to maximize the predicted demand served across all areas.

Consider $\mathcal{L}^+$ (with $\mathcal{L} \subseteq \mathcal{L}^+$) an augmented set of locations in the consideration set. Let $L \subset \mathcal{L}^+$ be the decision variable corresponding to the set of locations where service is offered, termed *open locations*. Each candidate location $\ell \in \mathcal{L}^+$ has a capacity K_ℓ, and we assume that candidate locations are otherwise identical. We allow for a limited budget on the number of locations offering service and include a cardinality constraint on the number of open locations, that is, $|L| \le N$.

We consider demand at the subarea level $s \in \mathcal{S}$ (with parent area a). Let m_s be the service-eligible population in subarea s, and let $\rho_{s\ell}$ be the share of the population in subarea s that obtains service at location ℓ. Using Equations (2) and (7), we derive an expression for $\rho_{s\ell}$ as follows:

$$\begin{aligned}\rho_{s\ell} &:= \frac{1}{m_s}\sum_{\{i\,|\,s_i=s\}} P_\epsilon(u_i > 0\,|\,\ell, s) f(\ell\,|\,s), \text{ if } \ell \in L \\ &= \frac{1}{m_s}\sum_{\{i\,|\,s_i=s\}} \left[\frac{\exp(d(s,\ell)\alpha_a + X_a\boldsymbol{\beta} + \xi_a)}{1+\exp(d(s,\ell)\alpha_a + X_a\boldsymbol{\beta} + \xi_a)}\right] \\ &\quad \left[\frac{\exp(d(s,\ell)\alpha_a)}{\sum_{k\in L}\exp(d(s,k)\alpha_a)}\right], \text{ if } \ell \in L \\ &= \left[\frac{\exp(d(s,\ell)\alpha_a + X_a\boldsymbol{\beta} + \xi_a)}{1+\exp(d(s,\ell)\alpha_a + X_a\boldsymbol{\beta} + \xi_a)}\right] \\ &\quad \left[\frac{\exp(d(s,\ell)\alpha_a)}{\sum_{k\in L}\exp(d(s,k)\alpha_a)}\right], \text{ if } \ell \in L \end{aligned} \tag{8}$$

and zero otherwise. Note that the second term in brackets assumes that individuals can choose any of the L open service locations without considering available capacity. This is different than in Section 3.1.3, where sequential service fulfillment tracks remaining capacity at open locations. Here, we incorporate a constraint to ensure that the predicted number of individuals served at each open location ℓ does not exceed its capacity K_ℓ.

For all subarea–location pairs (s,ℓ), we compute $\rho_{s\ell}$ using Equation (8) and the estimated distance sensitivities $\hat{\alpha}_a$, demographic-based preferences $\hat{\boldsymbol{\beta}}$, and idiosyncratic preferences $\hat{\xi}_a$. To ease notation, let $\bar{u}(s,\ell) := d(s,\ell)\hat{\alpha}_a + X_a\hat{\boldsymbol{\beta}} + \hat{\xi}_a$. Equation (8) simplifies to

$$\begin{aligned}\rho_{s\ell} &= \left[\frac{\exp(\bar{u}(s,\ell))}{1+\exp(\bar{u}(s,\ell))}\right]\left[\frac{\exp(\bar{u}(s,\ell) - X_a\hat{\boldsymbol{\beta}} - \hat{\xi}_a)}{\sum_{k\in L}\exp(\bar{u}(s,k) - X_a\hat{\boldsymbol{\beta}} - \hat{\xi}_a)}\right], \text{ if } \ell \in L \\ &= \frac{\frac{\exp(2\bar{u}(s,\ell))}{1+\exp(\bar{u}(s,\ell))}}{\sum_{k\in L}\exp(\bar{u}(s,k))}, \text{ if } \ell \in L \end{aligned} \tag{9}$$

and zero otherwise. In Equation (9), the term $\exp(-X_a\hat{\boldsymbol{\beta}} - \hat{\xi}_a)$ in the right bracket cancels out. Thus, the predicted demand from subarea s served at location ℓ is given by $m_s\rho_{s\ell}(L)$.

To formulate the optimization problem, we encode the main decision variable, the set of open locations $L \subset \mathcal{L}^+$, as a vector of binary decision variables $\mathbf{x} \in \{0,1\}^{|\mathcal{L}^+|}$: $x_\ell = 1$ if $\ell \in L$, and $x_\ell = 0$ otherwise. Thus, we can rewrite $\rho_{s\ell}$ as a function of $\mathbf{x}$ as follows:

$$\rho_{s\ell}(\mathbf{x}) = \frac{\frac{\exp(2\bar{u}(s,\ell))}{1+\exp(\bar{u}(s,\ell))}x_\ell}{\sum_{k\in\mathcal{L}^+}\exp(\bar{u}(s,k))x_k}, \text{ if } \sum_{k\in\mathcal{L}^+} x_k \ge 1,$$

and zero otherwise. To account for partial demand fulfillment, we introduce a decision variable $y_{s\ell} \le \rho_{s\ell}(\mathbf{x})$ as the share of the eligible population from subarea s that is served at location ℓ. Hence, a location ℓ may be opened to serve a fraction of the overall demand from

s to ℓ. The resulting optimization problem is

$$
\begin{aligned}
(\mathsf{P}_{\mathsf{loc}}) \quad \max_{\mathbf{x},\mathbf{y}} \ & \sum_{s\in\mathcal{S}}\sum_{\ell\in\mathcal{L}} m_s y_{s\ell} && && \text{(10a)}\\
\text{s.t.} \ & \sum_{s\in\mathcal{S}} m_s y_{s\ell} \le x_\ell K_\ell, && \forall \ell\in\mathcal{L}^+ && \text{(10b)}\\
& \sum_{\ell\in\mathcal{L}} x_\ell \le N && && \text{(10c)}\\
& y_{s\ell} \le \rho_{s\ell}(\mathbf{x}), && \forall \ell\in\mathcal{L}^+, s\in\mathcal{S} && \text{(10d)}\\
& x_\ell\in\{0,1\},\ y_{s\ell}\ge 0, && \forall \ell\in\mathcal{L}^+, s\in\mathcal{S}. && \text{(10e)}
\end{aligned}
$$

The objective (10a) maximizes the total expected demand served, where each vaccination is accounted for equally. Constraints (10b) ensure that the predicted demand served by an open location does not exceed its capacity. Note that multiple locations may serve the demand of a subarea, as determined by individuals' preferences reflected in $\rho_{s\ell}(\mathbf{x})$ through $y_{s\ell}$. Constraint (10c) imposes a total budget of N open locations. Constraints (10d) ensure that partial fulfillment of a subarea's demand is bounded by the total induced demand $\rho_{s\ell}(\mathbf{x})$. This formulation can be adapted to include other considerations, such as linear costs for selecting less desirable locations (e.g., stores more distant from public transit), limiting travel distances, or imposing equity constraints regarding distance traveled or demand served across groups.

The problem $(\mathsf{P}_{\mathsf{loc}})$ is essentially an assortment optimization problem with cardinality, capacity, and partial fulfillment considerations, but is computationally challenging because of its nonlinear constraints, which rely on the choice probabilities $\rho_{s\ell}(\mathbf{x})$. Using techniques from assortment optimization under an MNL choice model (Bront et al. 2009), we propose an equivalent and tractable formulation.

We introduce an auxiliary variable $v_s \ge 0$ defined as

$$
v_s = \begin{cases} \dfrac{1}{\sum_{k\in\mathcal{L}} \exp(\bar{u}(s,k))x_k} & \text{if } \displaystyle\sum_{k\in\mathcal{L}^+} x_k \ge 1,\\ 0 & \text{otherwise,}\end{cases}
$$

and a binary decision variable $b_s \in \{0,1\}$, which together satisfy the following set of constraints:

$$
\begin{aligned}
v_s \sum_{k\in\mathcal{L}^+} \exp(\bar{u}(s,k))x_k = b_s \quad & \forall s\in\mathcal{S}, && \text{(11a)}\\
v_s \le M_s b_s \quad & \forall s\in\mathcal{S}, && \text{(11b)}\\
b_s \ge x_\ell \quad & \forall s\in\mathcal{S}, \ell\in\mathcal{L}^+. && \text{(11c)}
\end{aligned}
$$

Here, $M_s = \overline{v}_s := \max_{k\in\mathcal{L}^+}[\exp(\bar{u}(s,k))^{-1}]$. Even after incorporating these auxiliary variables, the constraints (10d) and (11a) remain nonlinear. Following the linearization approach from Wu (1997), we define $w_{s\ell} := v_s x_\ell$ and introduce a set of linear constraints:

$$
\begin{aligned}
v_s - w_{s\ell} \le M_s - M_s x_\ell \quad & \forall s\in\mathcal{S}, \ell\in\mathcal{L}^+, && \text{(12a)}\\
w_{s\ell} \le v_s \quad & \forall s\in\mathcal{S}, \ell\in\mathcal{L}^+, && \text{(12b)}\\
w_{s\ell} \le M_s x_\ell, \quad & \forall s\in\mathcal{S}, \ell\in\mathcal{L}^+. && \text{(12c)}
\end{aligned}
$$

The resulting (equivalent) optimization problem to $(\mathsf{P}_{\mathsf{loc}})$ is

$$
\begin{aligned}
(\tilde{\mathsf{P}}_{\mathsf{loc}}) \quad \max_{\mathbf{x},\mathbf{y},\mathbf{b},\mathbf{v},\mathbf{w}} \ & \sum_{s\in\mathcal{S}}\sum_{\ell\in\mathcal{L}^+} m_s y_{s\ell}\\
\text{s.t.} \ & \sum_{s\in\mathcal{S}} m_s y_{s\ell} \le x_\ell K_\ell, && \forall \ell\in\mathcal{L}^+\\
& \sum_{\ell\in\mathcal{L}^+} x_\ell \le N\\
& \text{Constraints (11a)–(11c) and (12a)–(12c)}\\
& y_{s\ell} \le \frac{\exp(2\bar{u}(s,\ell))}{1+\exp(\bar{u}(s,\ell))} w_{s\ell}\\
& x_\ell, b_s\in\{0,1\},\ v_s, w_{s\ell}, y_{s\ell}\ge 0, \quad \forall s\in\mathcal{S}, \ell\in\mathcal{L}^+.
\end{aligned}
$$

Here, $\mathbf{b}$, $\mathbf{v}$, and $\mathbf{w}$ are auxiliary variables used to linearize the original problem $(\mathsf{P}_{\mathsf{loc}})$. The resulting formulation $(\tilde{\mathsf{P}}_{\mathsf{loc}})$ is a mixed-integer linear program amenable to commercial optimization solvers.

3.3. Policy Evaluation

The primary outcome of $(\tilde{\mathsf{P}}_{\mathsf{loc}})$ is the set of optimally chosen locations, captured by $\tilde{\mathbf{x}}$ and the corresponding $\tilde{L} := \{\ell\in\mathcal{L}^+ \mid \tilde{x}_\ell = 1\}$. This optimization problem, however, assumes individuals select service locations simultaneously while satisfying capacity constraints only in expectation. This contrasts with the sequential decision making often observed in reality, where individuals take into account locations' remaining capacity when making their choices. Although direct optimization of this sequential decision-making process is computationally intractable, we can evaluate the performance of the chosen locations, $\tilde{L}$, more realistically by using the sequential appointment assignment process used in the demand estimation (Section 3.1.3). To do so, we first randomly order all individuals in the population. Each individual selects their appointment location and whether to obtain service, among the open locations $\tilde{L}$ with remaining capacity, according to the estimated MNL probabilities. Once a site reaches its capacity, it is removed from the set of open locations for all subsequent individuals. Hence, for a given $\tilde{L}$, we report the predicted demand served under the sequential appointment fulfillment process as the main outcome of our approach.[7] A detailed description of the entire estimate-then-optimize approach is provided in Online Appendix B, Section B.3.

4. Case Study: COVID-19 Vaccine Distribution in California

In early 2021, the U.S. government launched the Federal Retail Pharmacy Program (FRPP) to expand COVID vaccination distribution points through a partnership with 21 national chains and independent pharmacy networks. By August 2023, more than 300 million doses were administered through the FRPP, across 41,000 locations

nationwide (U.S. Centers for Disease Control and Prevention 2023). Nearly half of all Americans live within one mile of a pharmacy offering COVID vaccinations and 89% live within five miles. However, travel distances vary widely, particularly for Black communities, who often face distances exceeding ten miles (Berenbrok et al. 2021). This disparity extends beyond rural areas, with limited vaccination options among communities of color and those with chronic health conditions (Rader et al. 2022). In California, residents in the bottom HPI quartile face average travel distances 72% greater than those in the top quartile (Figure 1). Median vaccination rates are 19 percentage points lower in the bottom HPI quartile (Figure 4), in part because of these structural barriers to vaccination.

Using our SETO approach, we evaluate an FRPP expansion to include other partnerships to help alleviate a major barrier to vaccination access. Using an estimated spatial demand model (stage 1), we predict COVID vaccination uptake under counterfactual scenarios of offering vaccines at these additional locations, which are optimally selected by the facility location model (stage 2). Our case study offers insights into how a national vaccine campaign might be strategically improved, underscoring the framework's utility in various healthcare settings.

Our main analysis focuses on dollar stores, but we also examine partnering with Starbucks or public high schools. After the FRPP's 2021 launch, the U.S. Centers for Disease Control and Prevention director and Dollar General confirmed discussions about administering COVID vaccines in stores (Roberts 2021). Dollar General operates over 17,000 locations in 46 states—nearly double the next largest private vaccine provider—with 75% serving rural communities (Bomey 2021). Two other chains, Dollar Tree and Family Dollar, each operate over 7,000 stores nationwide. Collectively, these discount retailers offer a promising way to expand vaccination access in low-income communities (Chevalier et al. 2022).

Figure 4. (Color online) Zip-Code-Level COVID-19 Vaccination Rates in California, by HPI Quartile (as of March 1, 2022)

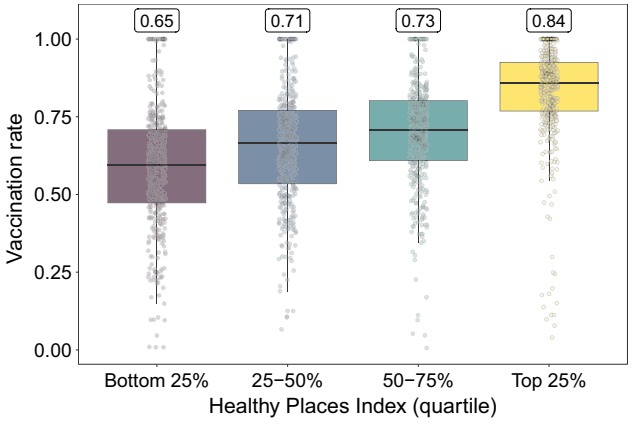

Note. Numbers above each boxplot are aggregated vaccination rates across all zip codes within each HPI quartile.

4.1. Data

Our case study integrates data on vaccination rates, spatial demographics, and store locations.

4.1.1. Vaccination Rates. We obtain the most granular population-level COVID vaccination data available in California: vaccinations administered by zip code (California Department of Public Health 2024). Our main outcome variable is the proportion of the population aged 5+ who are fully vaccinated as of March 1, 2022, although we consider alternative time points in Online Appendix C, Section C.5.

4.1.2. Vaccination Locations. Geographic coordinates of the 4,035 retail pharmacies in California participating in the FRPP are from https://www.vaccines.gov/ as of June 2021. Geodesic distances are computed between each vaccination site and the population centroid of each census tract or block. We exclude other dispensing sites such as mass vaccination centers or physician offices because pharmacies accounted for 90% of COVID vaccinations administered nationwide (IQVIA Institute for Human Data Science 2023).

Alternative vaccination sites include the three largest dollar store chains (Dollar General, Dollar Tree, and Family Dollar), which collectively operate 1,016 stores across California. Dollar store locations are obtained from ScrapeHero (2024), Starbucks locations are from Safegraph (2022), and public school locations are from the California Department of Education (2024).

4.1.3. Demographics. Our main demographic variable of interest, the Healthy Places Index, is a composite score constructed by the Public Health Alliance of Southern California (2022) that measures the well-being of a community. An area's HPI index is weighted across eight domains: socioeconomics, education, healthcare access, housing, neighborhood conditions, pollution/clean air, social factors, and transportation access (Maizlish et al. 2019). We group each region into HPI quartiles, where the top 25% are the healthiest, and the bottom 25% are the most vulnerable.

Population estimates of California's residents aged 5+ for each of the 376,762 census blocks are used, totaling 37,336,716.[8] The geographic coordinates of each block's population centroid characterize the spatial distribution of the vaccine-eligible population (age 5+; U.S. Census Bureau 2021). Demographic data are obtained from the 2019 American Community Survey for 1,751 zip codes in the state.[9] Key variables include race/ethnicity, unemployment rate, poverty

rate, college graduation rate, median household income, median home value, population, population density, and health insurance status. Summary statistics are available in Online Appendix A, Table A1.

4.2. Implementing the Estimate-Then-Optimize Framework

We describe the model parameters and include specific details on how to operationalize each step of our framework. We also provide a small working example at https://github.com/juanhu96/Demo.

4.2.1. Estimating Vaccination Demand. Areas ($\mathcal{A}$) correspond to zip codes, the most granular geographic region that reports our primary outcome, COVID vaccination rates (ρ_a^{obs}). Subareas ($\mathcal{S}$) correspond to census blocks, which are typically less than one acre in size, allowing us to account for variations in travel distance at an extremely granular level. The set of locations $\mathcal{L}$ includes pharmacies participating in the FRPP. We compute the geodesic distance $d(s,\ell)$ from the centroid of each block s to every location ℓ. This implies that all residents of a block have identical travel distance to any given site, a reasonable assumption given the small size of census blocks.

We allow sensitivity to travel distance α_{a_i} to vary by HPI quartile $\mathcal{H} := \{1,2,3,4\}$; that is, $\alpha_{a_i} = \alpha_{h(a_i)}$, where $h(a_i)$ denotes the HPI quartile of individual i's zip code, resulting in a vector $\boldsymbol{\alpha}$ consisting of the distance sensitivities for each HPI quartile. This permits heterogeneity in travel distance disutility across socioeconomic status (e.g., long travel distances may be a larger deterrent for lower-HPI individuals who have limited transportation access or time off work). Online Appendix C, Section C.6, examines *household income* as a single dimension to measure heterogeneity in distance sensitivity.

An individual's utility from vaccination (Equation (1)) incorporates zip-code-level demographic variables, X_{a_i}, including race/ethnicity composition (Black, Asian, Hispanic, and other), composition of health insurance (employer provided, Medicare, Medicaid, and others), fraction of college graduates, unemployment rate, poverty rate, median household income, median home value, and population density. Including these factors in the utility model helps capture the influence that community demographics exert on latent preferences for vaccination.

For each subarea, we consider the set of possible vaccination locations to be the M closest sites, to avoid diluting the choice probabilities with excessively distant locations and to ease computation. In both the estimation and optimization stages, individuals can obtain service at one of these M locations.[10] We initially set $M = 5$, but we vary this in sensitivity analysis. Capacity per vaccination location is assumed to be constant, $K_\ell = 10{,}000\ \forall \ell \in \mathcal{L}$, but we consider $K_\ell \in \{8{,}000, 12{,}000\}$ in sensitivity analysis. We compute K_ℓ such that total capacity (40.35 million vaccines) sufficiently covers California's vaccine-eligible population.

Following Section 3.1.3, we model individuals' choice of service locations under capacity constraints. We first randomly order all 37.34 million individuals in California. Then, individuals choose a location and decide to vaccinate (or not) in that order. An individual's set of possible vaccination locations L_i is determined by the remaining capacity at nearby sites after all previous individuals have made their location and vaccination decisions. We test the robustness of our demand estimates to the assignment mechanism in Online Appendix C, Sections C.3 and C.7.

Our main demand estimates measure travel costs using the log-distance between each subarea and vaccination location. This functional form can result in the optimization model seeking to reduce small travel distances, while neglecting reductions in longer travel distances. We also consider an alternative distance function that treats all distances below a threshold $\overline{d}$ as constant. We explore various thresholds $\overline{d} = \{0.5$ miles, 1 mile$\}$ in Online Appendix C, Section C.4.

With the need to assign more than 376,000 census blocks in California to one or more vaccination locations, our mixed-integer facility location problem is computationally intractable. We therefore formulate this portion (the second stage of our framework) at the census tract level as follows. First, we compute block-level utility estimates using the estimated demand model parameters. Then, we aggregate demand up to the census tract level and compute approximate utility estimates at this level, which serve as inputs to the optimization model (see Online Appendix B, Section B.2, for details).

4.2.2. Optimizing Location Selection. We let $\mathcal{L}^+$ include the existing FRPP pharmacies ($N = 4{,}035$) and any new candidate locations. The model ($\tilde{\mathsf{P}}_{\mathsf{loc}}$) selects the optimal set of new sites assuming all FRPP pharmacies remain operational, referred to as *network expansion*. We consider partnering with dollar stores, assuming network expansion to 100 to 1,000 locations (out of 1,016 stores statewide). We also consider partnering with public high schools (out of 1,225 locations statewide) or Starbucks (out of 3,172 locations statewide), as detailed in Online Appendix E.[11]

Our primary outcome measure is the predicted gain in vaccinations compared with a *pharmacy-only* policy, which utilizes only the $N = 4{,}035$ pharmacies participating in the FRPP in California as of June 2021. To maintain consistency, we employ the same parameter values throughout the estimation and optimization stages. We assume that partner locations have equal capacity to existing pharmacies, but we test a scenario

where the capacity is reduced by one-half (Online Appendix C, Section C.2).

We solve ($\tilde{P}_{loc}$) using Python 3.9.12 and the Gurobi 10.0.1 solver. The algorithm is configured to halt execution once the optimality gap falls below 0.1% or run time reaches six hours.

4.3. Results

We first present empirical demand estimates under the pharmacy-only policy, and we validate the choice-based model's predictions through policy evaluation. We then conduct counterfactual analyses to assess network expansion and compare results to various benchmarks. We explore site closures and replacement strategies in Online Appendix G.

4.3.1. Demand Estimates. Table 2 gives our main estimates and key demographic predictors of vaccine demand. HPI strongly predicts vaccination uptake, with less healthy areas showing conditionally lower vaccination rates, suggesting that the observed spatial inequalities in vaccinations mirror those seen in other dimensions of health. Regions with higher employment rates and more college-educated residents also show greater vaccine demand. Although these estimates are suggestive, individual coefficients should be interpreted with some caution because of high correlations among demographic covariates.[12]

Table 2. Structural Demand Estimates

	Model outcome	
	Fraction Fully Vaccinated	
Independent variable	Coefficient	Standard error
HPI Quartile 4 (most healthy)	Ref.	
HPI Quartile 3	−0.304***	(0.084)
HPI Quartile 2	−0.451***	(0.112)
HPI Quartile 1 (least healthy)	−0.605***	(0.149)
Log-distance × HPI Quartile 4	−0.063	(0.070)
Log-distance × HPI Quartile 3	−0.142***	(0.050)
Log-distance × HPI Quartile 2	−0.124***	(0.043)
Log-distance × HPI Quartile 1	−0.161***	(0.043)
Race White	Ref.	
Race Black	−0.052	(0.342)
Race Asian	2.001***	(0.179)
Race Hispanic	0.987***	(0.167)
Race Other	7.077**	(3.049)
Health Insurance: Employer	−1.202*	(0.696)
Health Insurance: Medicare	0.197	(0.739)
Health Insurance: Medicaid	−0.973	(0.729)
Health Insurance: Other	−2.319***	(0.760)
College Graduation Rate	1.727***	(0.282)
Unemployment Rate	−1.777***	(0.686)
Poverty Level	−0.011	(0.472)
Log(*Median Household Income*)	−0.110*	(0.060)
Log(*Median Home Value*)	0.114**	(0.050)
Log(*Population Density*)	−0.078***	(0.022)
Constant	1.813**	(0.721)

Notes. Estimates represent a demand model of the share of eligible population fully vaccinated for COVID as of March 1, 2022. Parameter estimation is described in Section 3.1.2. HPI is a composite measure of a community's health and well-being. Log-distance is the distance between an individual's census block centroid and the selected vaccination location.

Significance levels: *$p < 0.05$; **$p < 0.01$; ***$p < 0.001$.

The estimates of interest are the distance sensitivities by HPI quartile, α_h, which reveal a clear and economically meaningful pattern: α_h is negative and statistically significant for all HPI quartiles except the top one, whose coefficient is statistically indistinguishable from zero. Less healthy communities experience greater declines in vaccination rates as distance increases. Doubling the average distance from 1 km to 2 km equates to a 2.4 percentage-point decrease in vaccinations for the bottom HPI quartile, compared with a 0.5 percentage-point reduction for the top HPI quartile. This aligns with the observation that longer travel distances are more burdensome for individuals with lower socioeconomic status. The coefficients remain significant after controlling for observable demographics, suggesting that the results cannot be explained by spatial demographics alone, but indicate a genuine sensitivity to distance. Figure 5 shows predicted vaccination rates by HPI quartile, highlighting the unequal uptake and greater sensitivities for the bottom quartiles.

The preceding evidence underscores the strong correlation between distance to a vaccination site and uptake. One potential threat to this model specification is selection bias among pharmacies in the FRPP, which might arise if the selection process is ex ante correlated with both vaccination propensity and HPI—beyond what is captured by observable demographics. However, this is unlikely because most participating locations are national chains (e.g., CVS, Walgreens) with established store locations before the COVID vaccine

Figure 5. (Color online) Predicted COVID-19 Vaccination Rates

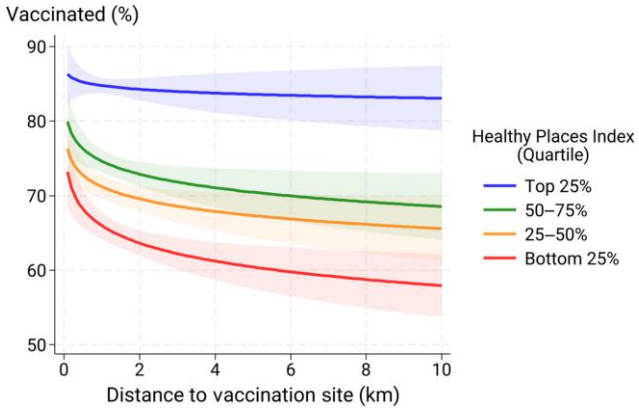

Note. Within each HPI quartile, we compute the mean vaccination rate, assuming that all individuals are assigned to a location with the given distance (x-axis).

gained U.S. Food and Drug Administration approval. Including additional demographics helps control for the tendency of retail pharmacies to locate in higher-income, urban areas.

We conduct extensive robustness checks, as detailed in Online Appendix C. Demand estimates remain qualitatively similar across various specifications, including alternative choice sets (Table C11), capacity constraints (Table C12), functional forms for distance (Table C12), cross-sectional time periods (Table C13), pandemic controls (Table C13), and demand fulfillment (Table C15). In all cases, the estimates consistently reflect significant vaccination inequalities by HPI and greater distance sensitivities for low-HPI individuals.

4.3.2. Baseline Network Validation. Under the baseline pharmacy-only scenario, predicted vaccinations vary widely between the top and bottom HPI quartiles (Table 3). This gap arises because lower-HPI residents live farther from FRPP pharmacies (Figure 1) and have steeper travel distance elasticities (Figure 5). A map visualizing the uneven distribution of pharmacies within Los Angeles is shown in Figure 6. Of the 26.4 million predicted vaccinations, three-quarters (19.5 million) occur at pharmacies within a one-mile walking distance, though fewer are walkable for residents of the lowest HPI quartile. Our model's predicted vaccinations broadly align with observed zip-code-level vaccination rates (Figure 4). As expected with BLP estimation, goodness-of-fit measures confirm strong model performance (Online Appendix A, Table A4). At the county level, the model slightly underpredicts vaccination rates in some highly vaccinated counties (Figure A2 in Online Appendix A), likely because of more restrictive choice sets and capacities imposed during the policy evaluation stage.

4.3.3. Network Expansion. Augmenting the FRPP network to include 500 *optimally selected* dollar stores statewide generates 0.77 million additional predicted vaccinations (a 2.9% increase), a marginal gain of approximately 1,500 vaccinations per added dollar store. Although residents of the bottom HPI quartile accrue much of the vaccination benefits (a 5.3% increase), adding dollar stores does benefit high-HPI communities (a 2.6% increase) by expanding capacity and allowing individuals across the HPI spectrum to obtain a vaccination closer to home. Notably, in the lower-HPI communities that have lower automobile ownership and access to public transit, *walkable* vaccinations increase by 18% (Table 3, bottom panel). We observe diminishing returns as the number of partner locations increases, as shown by the SETO curve in Figure 7. With just 100 optimally selected dollar stores, a gain of 540,000 vaccinations is predicted, representing 70% of the gains with 500 dollar stores. With 200 dollar stores, more than 700,000 added vaccinations are predicted, with over 260,000 vaccinations accruing to the bottom HPI quartile.

Dollar stores often locate in lower-income areas: 431 out of 1,016 stores statewide are in the bottom HPI quartile. The optimal expansion policy indeed prioritizes these locations: out of 500 stores selected, 226 are in the bottom HPI quartile (Table 4, heterogeneous demand). The selected dollar stores tend to be located farther from existing FRPP pharmacies (average distance of 2.8 km), and a sizable proportion of dollar stores inside pharmacy deserts are selected (56 out of 71).

We illustrate the geographic variation in pharmacies and potential partners for Los Angeles County in Figure 6, with darker regions indicating lower-HPI tracts. Out of 500 optimally selected dollar stores statewide, 102 are located in LA County, primarily in lower-HPI areas. Alternatively, an optimal network expansion policy using public high schools selects 151 schools in LA County, providing more access points in underserved communities, and boosting predicted vaccinations by approximately one million statewide. Finally, a Starbucks partnership selects 117 locations in LA County and also generates about one million additional vaccinations statewide, although the gain primarily accrues to the top HPI quartile. Additional details are given in Online Appendix E.

We conduct several robustness checks using the revised demand estimates obtained previously. We perform the subsequent optimization and policy evaluation for alternative distance forms, choice sets, capacity constraints, demand fulfillment, cross-sectional time periods, and pandemic controls. All results are reported in Online Appendix C.

Table 3. Predicted Vaccinations Under Network Expansion with 500 Additional Dollar Stores Statewide

			HPI quartile			
Strategy		All	Bottom 25%	25%–50%	50%–75%	Top 25%
Pharmacy only	Total	26.44	5.62	6.68	6.83	7.31
	Walkable	19.46	3.98	5.07	4.96	5.45
Pharmacy + dollar	Total	+0.77	+0.30	+0.14	+0.14	+0.19
	Walkable	+1.54	+0.73	+0.29	+0.29	+0.24

Notes. All values are reported in millions of vaccinations. The "pharmacy + dollar" strategy is the change in vaccinations compared with pharmacy only. Walkable vaccinations refer to travel distances of less than one mile.

Figure 6. (Color online) Maps of FRPP Pharmacies in Los Angeles County and Optimal Locations of Partner Stores Assuming Network Expansion with 500 Locations Statewide

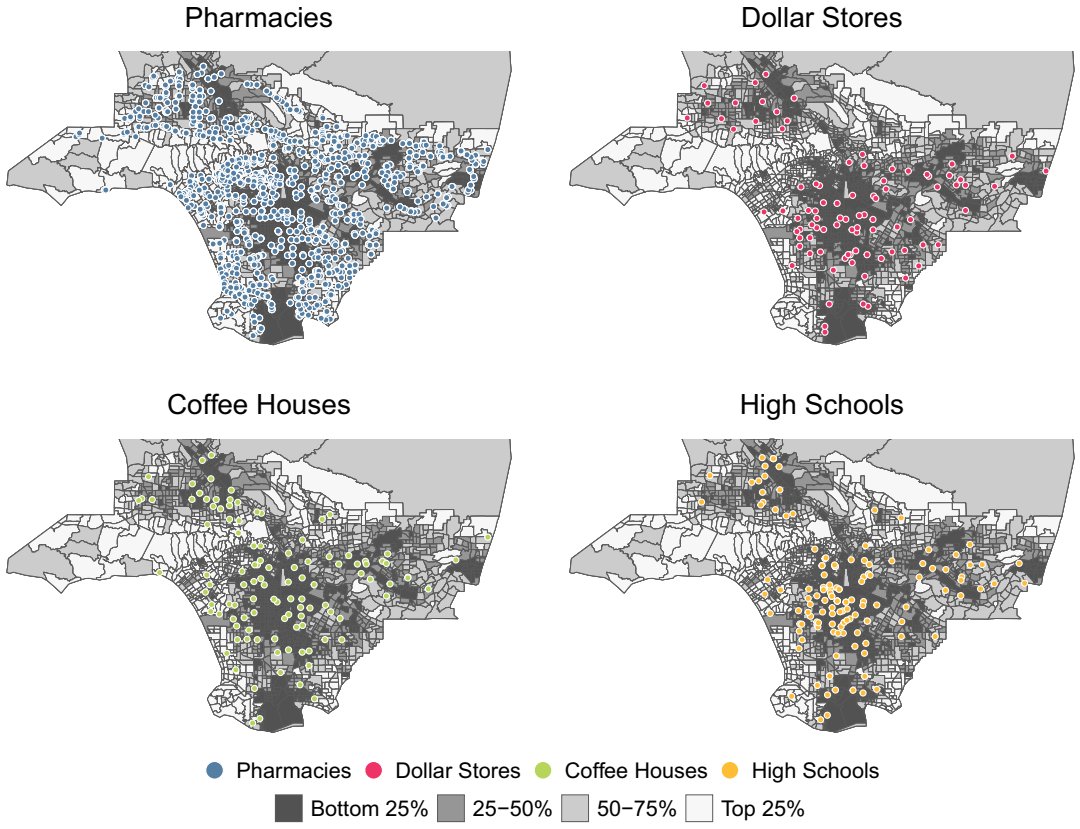

Notes. LA County has 982 FRPP pharmacies, 162 dollar stores, 788 Starbucks, and 307 high schools. Network expansion with 500 locations statewide selects 102 dollar stores, 117 Starbucks, or 151 high schools in LA County. Tracts with <1,500 residents or >50% in institutional settings (e.g., dorms, nursing homes, prisons) are undefined.

We also examine to what extent misspecification of the demand model affects its performance. If the estimation step ignores demand *heterogeneity* (Online Appendix C, Table C16), then the optimization model prioritizes different populations in selecting dollar stores (Table 4). More dollar stores are selected in the top HPI quartile (74 versus 56 with heterogeneous demand), with fewer in the bottom quartile (181 versus 226 with heterogeneous demand). We also observe a shift in selecting dollar stores closer to existing pharmacies (average distance of 1.4 km versus 2.8 km with heterogeneous demand), with fewer selected inside pharmacy deserts (7% versus 11% under heterogeneous demand). Altogether, these results highlight the importance of flexibly accounting for heterogeneity while estimating demand, and failing to do so results in a more regressive

Table 4. Characteristics of Selected Dollar Stores Under Heterogeneous and Homogeneous Demand Models, Assuming Network Expansion with 500 Locations Statewide

		Selected dollar stores	
	All dollar stores	Heterogeneous demand	Homogeneous demand
No. of stores (%)	1,016	500	500
in HPI Quartile 4 (most healthy)	89 (9%)	56 (11%)	74 (15%)
in HPI Quartile 3	195 (19%)	91 (18%)	104 (21%)
in HPI Quartile 2	301 (30%)	127 (26%)	141 (28%)
in HPI Quartile 1 (least healthy)	431 (42%)	226 (45%)	181 (36%)
No. of stores (%) in pharmacy deserts	71 (7%)	56 (11%)	36 (7%)
Avg. distance to the nearest pharmacy (km)	2.1	2.8	1.4

Notes. Heterogeneous demand results are our main model results. Homogeneous results are based on the model with constant preferences across locations within the choice set.

Figure 7. (Color online) Predicted Added Vaccinations with Network Expansion of 100 to 1,000 Dollar Stores Statewide

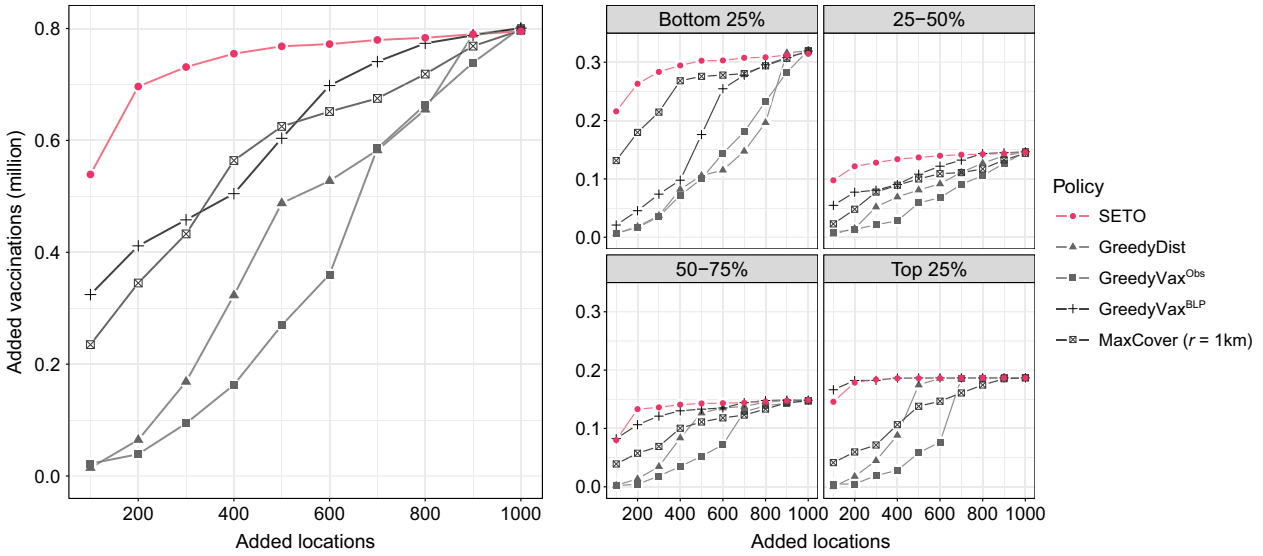

Note. Added vaccinations are relative to the pharmacy-only scenario.

allocation of resources. We further discuss the effect of ignoring demand heterogeneity on predicted vaccinations in Online Appendix D, Section D.1.

We also assess the effect of ignoring provider *capacity* in demand estimation, and we observe steeper distance elasticities, suggesting a lower willingness to travel for vaccination (Table C12 in Online Appendix C). This occurs because, with unlimited capacity, most individuals can be served at their closest pharmacy. To match the observed vaccination rates, the demand model estimates distance elasticities that are greater in magnitude than under the original capacitated model. Intuitively, failing to account for capacity results in the BLP demand model incorrectly rationalizing individuals as highly vaccine hesitant or distance sensitive, when in reality, they simply could not find capacity at nearby sites. The effect on predicted vaccinations is detailed in Online Appendix D, Section D.2.

4.3.4. Alternative Benchmarks. We compare the performance of our SETO approach to benchmark policies. The maximal covering location problem (MaxCover) does not rely on demand estimation and simply aims to maximize the population within radius r of a service location. GreedyVax^{Obs} prioritizes areas with the lowest observed vaccination rates, whereas GreedyDist focuses on areas farthest from existing pharmacies. We also include a heuristic policy, GreedyVax^{BLP}, based on the BLP-based demand estimates that prioritizes new locations that result in the most predicted vaccinations, assuming an area's demand is served at the closest new location. Figure 7 shows predicted vaccination gains under each benchmark with 100 to 1,000 dollar stores, using the policy evaluation approach from Section 3.3. Additional benchmark details are in Online Appendix F.

SETO delivers superior performance in predicted vaccinations, especially when fewer new locations are added. For instance, if 200 new locations are added, MaxCover ($r = 1$ km)—one of the best-performing benchmarks—generates only half the vaccination gains as SETO, amounting to 350,000 fewer vaccinations. This performance gap persists across HPI quartiles, with our approach yielding nearly 100,000 more vaccinations in the bottom HPI quartile compared with MaxCover. Unaware of demand sensitivities, MaxCover prioritizes densely populated areas that often have pharmacies nearby and overlooks potential sites in suburban and rural neighborhoods. Interestingly, GreedyVax^{BLP} performs similarly to MaxCover, despite not relying on optimization. This heuristic is demand aware, favoring areas with high predicted vaccination uptake, which tend to be high HPI and densely populated. However, this policy fails to account for the existence of other vaccination sites, and it ultimately selects new locations that overlap with existing pharmacies. Simpler heuristics like GreedyDist and GreedyVax^{Obs} perform poorly as they ignore distance sensitivities and the spillover effects of opening a new location. To illustrate the limitations of a greedy policy, we examine three California counties (Online Appendix F, Section F.2). In areas with more homogeneous demand (e.g., Orange County), GreedyVax^{Obs} performs competitively. However, in regions with heterogeneous demand (e.g., Los Angeles County), this myopic policy fails to prioritize communities with average vaccination rates but highly distance-sensitive populations.

Notably, SETO outperforms both the optimization-only and BLP-estimation-only policies (Figure 8). For example, under network expansion with 500 new locations, the best heuristics provide 77%–80% of the predicted benefits of our SETO approach. This underscores that SETO's advantage lies in integrating the estimated demand model within an optimization framework. By doing so, it ensures that resources are allocated where they are most likely to be used, while also considering the distribution of existing service locations to effectively address spatial gaps in service providers.

4.4. Feasibility and Operational Challenges

Despite vaccinating millions of Americans, the FRPP falls short in reaching vulnerable communities. We propose a data-driven, transparent approach to expand vaccination access by partnering with nonhealthcare organizations and strategically select new locations to complement the existing network. Adding 100 optimally selected dollar stores in California could boost vaccinations by over half a million—a 2% gain, comparable to prior interventions. Unlike vaccine education, which showed no impact among healthcare staff (Berry et al. 2022), or nudges (Dai et al. 2021) and persuasive advertising (Larsen et al. 2023), which yield modest gains, our model offers a scalable, cost-efficient solution that reduces geographic inequities.

Beyond only vaccination gains, the cost and feasibility of nonpharmacy partnerships are key considerations. Each option has trade-offs as summarized in Online Appendix E, Table E20. Starbucks attracts repeat visitors, enabling frequent vaccinations, but physical space and patient privacy are concerns. High schools offer ample indoor space but face organizational challenges due to decentralized districts. Dollar stores previously hosted pop-up COVID vaccination sites in 2021 (Online Appendix E, Table E20), but efforts were ad hoc and lacked coordination. Our approach helps identify the most effective partnerships, and policymakers can weigh these factors against potential vaccination gains.

Administering vaccinations in nonhealthcare settings presents staffing challenges. Some dollar store chains may soon offer pharmacy services themselves (Katje 2021), whereas another option could entail visiting nurses via agencies like SnapNurse.[13] Staffing costs can be approximated using SnapNurse's average California wage of $30/hour (ZipRecruiter 2024). A three-month campaign with one or two nurses per site for an 8-hour shift per week costs $2,900–$5,800 per location. Partnering with 500 optimally chosen dollar stores would cost $1.4–$2.9 million, yielding 0.77 million vaccinations at $2–$4 per dose. Even a six-month campaign would cost about $5 million—just 0.002% of California's $300 billion budget. Given the protective benefits, these additional vaccinations offer significant social value.

Our case study has several limitations. First, we perform a static analysis to determine vaccination locations, but time-varying vaccination uptake and inventory availability may warrant a dynamic policy for site selection. Second, operational considerations, including the costs of opening new locations and varying capacities across sites, are not incorporated into the baseline model. Last, our analysis does not account for immunizations administered outside of pharmacies nor the age- or occupation-based prioritization implemented early in the vaccination campaign.

4.5. Discussion

Integrating structural estimation and optimization into policy design can uncover and mitigate disparities in healthcare access, extending beyond COVID vaccination to other healthcare services. Our findings highlight the importance of accounting for *heterogeneity* in both demand and supply. First, we show that demand elasticities vary by socioeconomic status, with vulnerable communities being more sensitive to distance (Table 2), a pattern likely applicable to other health services. Other barriers like limited transportation may further restrict access to infrequent care (e.g., annual vaccinations, cancer screenings), whereas frequent care needs (e.g., dialysis, substance abuse treatment) pose even greater challenges. Ignoring demand heterogeneity can lead to inefficient or inequitable service networks (Table 4 and Online Appendix D, Section D.1). Second, given a spatially heterogeneous supply of service

Figure 8. Predicted Vaccination Gains Under Different Approaches

	Observational data	BLP estimation
No optimization	GreedyDist +0.49 million $\text{GreedyVax}^{\text{Obs}}$ +0.27 million	$\text{GreedyVax}^{\text{BLP}}$ +0.60 million
With optimization	MaxCover +0.62 million	SETO +0.77 million

Note. All scenarios assume network expansion with 500 additional dollar stores statewide.

providers, strategically selecting new sites to complement the existing network helps eliminate geographic service gaps (Figure 6). In contrast, when demand is more uniform, simpler policies like a greedy approach may perform well (Online Appendix F, Section F.2). Targeting distance-sensitive underserved populations yields greater benefits than prioritizing high-density areas (as with MaxCover) or low-utilization regions (as with GreedyVax^{Obs}). Finally, embedding demand estimation within an optimization model can adequately account for network spillover effects—something not achievable with demand estimation alone (Figure 8).

To illustrate the generalizability of our approach in a different healthcare setting, consider a policymaker deciding where to expand opioid use disorder (OUD) treatment centers. In 2021, the California Department of Health Care Services (2025) initiated the Behavioral Health Continuum Infrastructure Program (BHCIP) with a budget of $2.2 billion to "construct, acquire, and expand properties" for behavioral health services. BHCIP awarded grants to for-profit and nonprofit entities, including inpatient and outpatient treatment facilities for OUD. In one funding round, $518 million was awarded to 45 partners to build 1,292 inpatient beds and serve 130,000 patients (California Department of Health Care Services 2025). However, the geographic distribution of grants does not match populations or estimated needs. For example, a single facility near Sacramento (population 1.6 million) received funds for 36,000 outpatient treatments per year, whereas Los Angeles County (population 9.7 million) received funds for only 28,000 treatments, despite having significantly higher opioid prescriptions and overdose deaths (California Department of Public Health 2025). In 2024, California passed Proposition 1, a $6.4 billion behavioral health bond measure, including $4.4 billion for substance abuse treatment facilities, appropriated similarly to BHCIP.

Given the financial stakes and potential health impact of such programs, it is essential for state leadership to allocate funding effectively. This requires examining questions like: *Which communities are expected to face high rates of OUD? What is the existing infrastructure available to treat OUD? How far can individuals travel to receive outpatient treatment?* Our SETO framework can address these questions, relying on publicly available data on treatment facilities, opioid-related hospitalizations, overdose deaths, and treatment rates. A related study combined a predictive epidemic model of OUD with optimization for facility allocation (Luo and Stellato 2024). Our approach complements this study by endogenously linking individuals' decision making (i.e., whether and where to initiate OUD treatment) with policy decisions (i.e., locations of treatment facilities) to predict outcomes (i.e., OUD treatment completion). Our approach is not only *data driven* but also *demand driven*, better matching public funds to the needs of the most vulnerable patients.

5. Conclusions

Expanding access to essential health services is vital for addressing disparities in care, yet few studies rigorously examine policy interventions through prescriptive analytics. This study contributes to the literature by developing a data-driven framework combining structural econometrics and facility location optimization, offering an interdisciplinary approach to tackling health inequality. Heterogeneity in the demand for and supply of health services can pose significant challenges, but they also present opportunities to measure how utilization depends on several factors including distance and socioeconomics, a necessary component to determine where additional resources are most needed.

Optimizing healthcare delivery networks pose challenges to both the private and public sectors, and our structural estimate-then-optimize approach is applicable across diverse settings. In the private sector, major pharmacy chains (Walgreens, CVS, and Rite Aid) have announced over 1,000 store closures based on low utilization and local market dynamics (Roeloffs 2024). Many health systems pursue cross-market mergers and hospital acquisitions to expand their geographic footprint, affecting access to care and costs, particularly when hospitals are geographically close (Cooper et al. 2019). In the public sector, optimizing healthcare networks is equally critical. For example, the UK's National Health Service (NHS) launched the Community Diagnostic Centres program in 2021 (National Health Services England 2024), establishing over 160 diagnostic sites in shopping centers, universities, and stadiums to alleviate wait times and improve access to care. With a single policymaker overseeing service expansion, our framework can systematically compare potential testing sites based on patient needs, facility proximity, and space availability. By predicting utilization, our approach supports cost-effectiveness analyses and investment decisions, demonstrating its adaptability across both private and public healthcare systems.

We conclude our paper by suggesting avenues for future research. Investigating disparities and distance elasticities across various healthcare domains offers fertile ground for empirical and analytical research in the operations management community. Future work could extend the prescriptive aspect of our framework by relaxing some assumptions (e.g., heterogeneous capacity per facility), specifying alternative objectives (e.g., cost, equity), or considering context-specific constraints (e.g., allowing repeat visits), while also developing scalable and efficient solutions for the resulting

problem. Additionally, leveraging observational data in diverse formats to improve operational decision making remains an understudied area. Our study highlights the value of combining econometrics models with prescriptive analytics, underscoring their potential to inform and refine policymaking and operational strategies.

Acknowledgments

The authors thank Zhijian Li for his excellent research assistance and Peter Rossi for his valuable insights during the early stages of this work. The authors are also grateful to department editor Stefan Scholtes, the associate editor, and the three anonymous reviewers for their constructive feedback throughout the review process, which significantly improved this manuscript.

Endnotes

[1] Note that to remain as consistent as possible with our empirical case study, Equation (1) is not written in fullest generality. For example, it is straightforward to extend the model to allow β to vary with i (Berry et al. 1995).

[2] The model can be extended to allow variation across individuals to capture individual-level factors (i.e., using α_i), such as vehicle ownership or job flexibility.

[3] We omit the conditioning on a because knowing s_i uniquely determines an individual's area, so $f(\ell_i|s_i,a) = f(\ell_i|s_i)$.

[4] Note that ξ_a is an implicit function of θ through the constraint that the model-implied market shares are equal to the observed market shares, that is, $\rho_a(\boldsymbol{\theta}) = \rho_a^{obs}$. This implicit function is enforced at each step of the optimization either through a contraction mapping (Berry et al. 1995) or solver (Reynaerts et al. 2012). Alternatively, in treating the equality as an equilibrium constraint, the GMM objective can be optimized using an mathematical program with equilibrium constraints (MPEC) (Dubé et al. 2012).

[5] In our case study, we verify that our demand estimates are similar under this alternative model (see Online Appendix C, Section C.3).

[6] One could further limit an individual's choice set to locations within a maximum distance to avoid assigning appointments that are extremely far from home.

[7] The main results are reported for a single ordering of the population. In Online Appendix C, Section C.7, we perform sensitivity analysis for several population orderings.

[8] Census blocks are small geographic regions within a tract, with approximately 100 residents on average. We consider only census blocks with nonzero population.

[9] The analysis includes 1,659 zip codes with nonzero populations.

[10] To estimate demand via BLP, all individuals must be assigned a service location. In the unusual case where all M closest locations are at capacity, we assign the individual to the Mth location. Although this assumption does not strictly respect capacity constraints, we find that our demand estimates are similar as M increases (see Online Appendix C, Section C.1).

[11] The partner's locations refer to the geographic coordinates of the vaccination distribution points and the specific deployment strategy could be on-site, pop-up, or mobile clinic.

[12] For example, households with employer-provided health insurance have lower unemployment and poverty rates.

[13] Founded in 2017, SnapNurse deployed over 10,000 nurses during COVID, growing revenue from $3 million in 2019 to nearly $90 million in 2020 (Barber 2021).

References

Aday LA, Andersen R (1974) A framework for the study of access to medical care. *Health Services Res.* 9(3):208–220.

Ahmadi-Javid A, Seyedi P, Syam SS (2017) A survey of healthcare facility location. *Comput. Oper. Res.* 79(March):223–263.

Allon G, Federgruen A, Pierson M (2011) How much is a reduction of your customers' wait worth? An empirical study of the fast-food drive-thru industry based on structural estimation methods. *Manufacturing Service Oper. Management* 13(4):489–507.

Arcury TA, Gesler WM, Preisser JS, Sherman J, Spencer J, Perin J (2005) The effects of geography and spatial behavior on health care utilization among the residents of a rural region. *Health Services Res.* 40(1):135–156.

Ayer T, Keskinocak P, Swann J (2014) Research in public health for efficient, effective, and equitable outcomes. Newman A, Leung J, eds. *Bridging Data and Decisions*, INFORMS TutORials in Operations Research (INFORMS, Catonsville, MD), 216–239.

Barber P (2021) Mission of mercy: Meet the traveling nurses vaccinating Sonoma County. *Press Democrat* (May 7), https://www.pressdemocrat.com/article/news/mission-of-mercy-meet-the-traveling-nurses-vaccinating-sonoma-county/.

Bennouna A, Joseph J, Nze-Ndong D, Perakis G, Singhvi D, Lami OS, Spantidakis Y, Thayaparan L, Tsiourvas A (2023) COVID-19: Prediction, prevalence, and the operations of vaccine allocation. *Manufacturing Service Oper. Management* 25(3):1013–1032.

Berenbrok LA, Tang S, Coley KC, Boccuti C, Jingchuan G, Essien UR, Dickson S, Hernandez I (2021) Access to potential COVID-19 vaccine administration facilities: A geographic information systems analysis. Accessed February 2, 2021, https://westhealth.org/wp-content/uploads/2021/02/Access-to-Potential-COVID-19-Vaccine-Administration-Facilities-2-2-2021.pdf.

Berry ST (1994) Estimating discrete-choice models of product differentiation. *RAND J. Econom.* 25(2):242–262.

Berry S, Pakes A (2007) The pure characteristics demand model. *Internat. Econom. Rev.* 48(4):1193–1225.

Berry S, Levinsohn J, Pakes A (1995) Automobile prices in market equilibrium. *Econometrica* 63(4):841–890.

Berry S, Levinsohn J, Pakes A (2004) Differentiated products demand systems from a combination of micro and macro data: The new car market. *J. Political Econom.* 112(1):68–105.

Berry SD, Goldfeld KS, McConeghy K, Gifford D, Davidson HE, Han L, Syme M, et al. (2022) Evaluating the findings of the IMPACT-C randomized clinical trial to improve COVID-19 vaccine coverage in skilled nursing facilities. *JAMA Internal Medicine* 182(3):324–331.

Bertsimas D, Sim M (2004) The price of robustness. *Oper. Res.* 52(1):35–53.

Bertsimas D, Digalakis V Jr, Jacquillat A, Li ML, Previero A (2022) Where to locate COVID-19 mass vaccination facilities? *Naval Res. Logist.* 69(2):179–200.

Beshears J, Choi JJ, Laibson DI, Madrian BC, Reynolds GI (2016) Vaccination rates are associated with functional proximity but not base proximity of vaccination clinics. *Medical Care* 54(6):578–583.

Bomey N (2021) Dollar General, CDC exploring partnership to speed up COVID-19 vaccine rollout. *USA Today* (March 9), https://www.usatoday.com/story/money/2021/03/09/dollar-general-cdc-covid-vaccines/6925995002/.

Boutilier JJ, Chan TC (2020) Ambulance emergency response optimization in developing countries. *Oper. Res.* 68(5):1315–1334.

Bronner K, Eliassen SM, King A, Leggett C, Punjasthitkul S, Skinner J (2021) The Dartmouth atlas of health care: 2018 data update. Report, The Dartmouth Institute for Health Policy and Clinical Practice, Lebanon.

Bront JJM, Méndez-Díaz I, Vulcano G (2009) A column generation algorithm for choice-based network revenue management. *Oper. Res.* 57(3):769–784.

Brown EJ, Polsky D, Barbu CM, Seymour JW, Grande D (2016) Racial disparities in geographic access to primary care in Philadelphia. *Health Affairs* 35(8):1374–1381.

Brownstein J, Cantor JH, Rader B, Simon KI, Whaley CM (2022) If you build it, will they vaccinate? The impact of COVID-19 vaccine sites on vaccination rates and outcomes. NBER Working Paper No. 30429, National Bureau of Economic Research, Cambridge, MA.

California Department of Education (2024) Public schools and districts data files. Accessed May 2024, https://www.cde.ca.gov/ds/si/ds/pubschls.asp.

California Department of Health Care Services (2025) Behavioral health continuum infrastructure program. Accessed February 2025, https://www.infrastructure.buildingcalhhs.com/.

California Department of Public Health (2024) COVID-19 vaccine progress dashboard data by zip code. Accessed May 2024, https://data.chhs.ca.gov/dataset/covid-19-vaccine-progress-dashboard-data-by-zip-code.

California Department of Public Health (2025) Overdose surveillance dashboard. Accessed February 2025, https://skylab.cdph.ca.gov/ODdash/.

Chevalier JA, Schwartz JL, Su Y, Williams KR (2022) Distributional impacts of retail vaccine availability. *J. Urban Econom.* 127(January): 103382.

Chung TH, Rostami V, Bastani H, Bastani O (2022) Decision-aware learning for optimizing health supply chains. Preprint, submitted November 15, https://arxiv.org/abs/2211.08507.

Church R, ReVelle C (1974) The maximal covering location problem. *Papers Regional Sci. Assoc.* 32(1):101–118.

Cooper Z, Craig SV, Gaynor M, Van Reenen J (2019) The price ain't right? Hospital prices and health spending on the privately insured. *Quart. J. Econom.* 134(1):51–107.

Dai H, Saccardo S, Han MA, Roh L, Raja N, Vangala S, Modi H, Pandya S, Sloyan M, Croymans DM (2021) Behavioural nudges increase COVID-19 vaccinations. *Nature* 597(7876):404–409.

Daskin MS, Dean LK (2005) Location of health care facilities. Brandeau ML, Sainfort F, Pierskalla WP, eds. *Operations Research and Health Care* (Springer, Boston), 43–76.

Davis P (2006) Spatial competition in retail markets: Movie theaters. *RAND J. Econom.* 37(4):964–982.

Denoyel V, Alfandari L, Thiele A (2017) Optimizing healthcare network design under reference pricing and parameter uncertainty. *Eur. J. Oper. Res.* 263(3):996–1006.

Deo S, Sohoni M (2015) Optimal decentralization of early infant diagnosis of HIV in resource-limited settings. *Manufacturing Service Oper. Management* 17(2):191–207.

Deryugina T, Molitor D (2021) The causal effects of place on health and longevity. *J. Econom. Perspect.* 35(4):147–170.

Dey S, Kurbanzade AK, Gel ES, Mihaljevic J, Mehrotra S (2024) Optimization modeling for pandemic vaccine supply chain management: A review and future research opportunities. *Naval Res. Logist.* 71(7):976–1016.

Dubé JP, Fox JT, Su CL (2012) Improving the numerical performance of static and dynamic aggregate discrete choice random coefficients demand estimation. *Econometrica* 80(5):2231–2267.

Duch-Brown N, Grzybowski L, Romahn A, Verboven F (2023) Evaluating the impact of online market integration—Evidence from the EU portable PC market. *Amer. Econom. J.: Microeconomics* 15(4):268–305.

Ekici A, Keskinocak P, Swann JL (2014) Modeling influenza pandemic and planning food distribution. *Manufacturing Service Oper. Management* 16(1):11–27.

Elmachtoub AN, Grigas P (2022) Smart "predict, then optimize." *Management Sci.* 68(1):9–26.

Enayati S, Özaltın OY (2020) Optimal influenza vaccine distribution with equity. *Eur. J. Oper. Res.* 283(2):714–725.

Finkelstein A, Gentzkow M, Williams H (2016) Sources of geographic variation in health care: Evidence from patient migration. *Quart. J. Econom.* 131(4):1681–1726.

Finkelstein A, Gentzkow M, Williams H (2021) Place-based drivers of mortality: Evidence from migration. *Amer. Econom. Rev.* 111(8): 2697–2735.

Gandhi A (2023) Picking your patients: Selective admissions in the nursing home industry. Preprint, submitted June 19, http://dx.doi.org/10.2139/ssrn.3613950.

Ganju KK, Atasoy H, McCullough J, Greenwood B (2020) The role of decision support systems in attenuating racial biases in healthcare delivery. *Management Sci.* 66(11):5171–5181.

Gowrisankaran G, Nevo A, Town R (2015) Mergers when prices are negotiated: Evidence from the hospital industry. *Amer. Econom. Rev.* 105(1):172–203.

Guadamuz JS, Wilder JR, Mouslim MC, Zenk SN, Alexander GC, Qato DM (2021) Fewer pharmacies in Black and Hispanic/Latino neighborhoods compared with white or diverse neighborhoods, 2007–15. *Health Affairs* 40(5):802–811.

Guajardo JA, Cohen MA, Netessine S (2016) Service competition and product quality in the US automobile industry. *Management Sci.* 62(7):1860–1877.

Hackmann MB (2019) Incentivizing better quality of care: The role of Medicaid and competition in the nursing home industry. *Amer. Econom. Rev.* 109(5):1684–1716.

Hansen LP (1982) Large sample properties of generalized method of moments estimators. *Econometrica* 50(4):1029–1054.

Heier Stamm JL, Serban N, Swann J, Wortley P (2017) Quantifying and explaining accessibility with application to the 2009 H1N1 vaccination campaign. *Health Care Management Sci.* 20(1):76–93.

Hwang K, Asif TB, Lee T (2022) Choice-driven location-allocation model for healthcare facility location problem. *Flexible Services Manufacturing J.* 34(4):1040–1065.

Ippolito B, Levy JF, Anderson GF (2020) Abandoning list prices in Medicaid drug reimbursement did not affect spending. *Health Affairs* 39(7):1202–1209.

IQVIA Institute for Human Data Science (2023) Trends in vaccine administration in the United States. Accessed January 13, 2023, https://www.iqvia.com/insights/the-iqvia-institute/reports-and-publications/reports/trends-in-vaccine-administration-in-the-united-states.

Jia H, Ordóñez F, Dessouky M (2007) A modeling framework for facility location of medical services for large-scale emergencies. *IIE Trans.* 39(1):41–55.

Jónasson JO, Deo S, Gallien J (2017) Improving HIV early infant diagnosis supply chains in sub-Saharan Africa: Models and application to Mozambique. *Oper. Res.* 65(6):1479–1493.

Katje C (2021) Dollar General makes moves to be the next corner pharmacy. *Business Insider* (July 7), https://markets.businessinsider.com/news/stocks/dollar-general-makes-moves-to-be-the-next-corner-pharmacy-1030586485.

Kelly C, Hulme C, Farragher T, Clarke G (2016) Are differences in travel time or distance to healthcare for adults in global north countries associated with an impact on health outcomes? A systematic review. *BMJ Open* 6(11):e013059.

Krohn R, Müller S, Haase K (2021) Preventive healthcare facility location planning with quality-conscious clients. *OR Spectrum* 43(1):59–87.

Larsen BJ, Ryan TJ, Greene S, Hetherington MJ, Maxwell R, Tadelis S (2023) Counter-stereotypical messaging and partisan cues: Moving the needle on vaccines in a polarized United States. *Sci. Adv.* 9(29):eadg9434.

Lee EK, Pietz F, Benecke B, Mason J, Burel G (2013) Advancing public health and medical preparedness with operations research. *Interfaces* 43(1):79–98.

Luo J, Stellato B (2024) Frontiers in operations: Equitable data-driven facility location and resource allocation to fight the opioid epidemic. *Manufacturing Service Oper. Management* 26(4): 1229–1244.

Maizlish N, Delaney T, Dowling H, Chapman DA, Sabo R, Woolf S, Orndahl C, Hill L, Snellings L (2019) California Healthy Places Index: Frames matter. *Public Health Rep.* 134(4):354–362.

Mazar A, Jaro D, Tomaino G, Carmon Z, Wood W (2023) Distance to vaccine sites is tied to decreased covid-19 vaccine uptake. *PNAS Nexus* 2(12):pgad411.

McCoy JH, Johnson ME (2014) Clinic capacity management: Planning treatment programs that incorporate adherence. *Production Oper. Management* 23(1):1–18.

Méndez-Díaz I, Miranda-Bront JJ, Vulcano G, Zabala P (2014) A branch-and-cut algorithm for the latent-class logit assortment problem. *Discrete Appl. Math.* 164(1):246–263.

National Health Services England (2024) Community diagnostic centres. Accessed September 2, 2024, https://www.england.nhs.uk/long-read/community-diagnostic-centres/.

Nevo A (2001) Measuring market power in the ready-to-eat cereal industry. *Econometrica* 69(2):307–342.

Nobles M, Serban N, Swann J (2014) Spatial accessibility of pediatric primary healthcare: Measurement and inference. *Ann. Appl. Statist.* 8(4):1922–1946.

Penchansky R, Thomas JW (1981) The concept of access: Definition and relationship to consumer satisfaction. *Medical Care* 19(2):127–140.

Petrin A (2002) Quantifying the benefits of new products: The case of the minivan. *J. Political Econom.* 110(4):705–729.

Public Health Alliance of Southern California (2022) Healthy Places Index 2022. Accessed May 2024, https://www.healthyplacesindex.org/.

Rader B, Astley CM, Sewalk K, Delamater PL, Cordiano K, Wronski L, Rivera JM, et al. (2022) Spatial modeling of vaccine deserts as barriers to controlling SARS-CoV-2. *Comm. Medicine* 2(1):1–11.

Rastegar M, Tavana M, Meraj A, Mina H (2021) An inventory-location optimization model for equitable influenza vaccine distribution in developing countries during the COVID-19 pandemic. *Vaccine* 39(3):495–504.

Reynaerts J, Varadha R, Nash JC (2012) Enhancing the convergence properties of the BLP (1995) contraction mapping. VIVES Discussion Paper, KU Leuven, Brussels.

Roberts B (2021) Dollar general stores could aid in administering vaccines. *Spectrum News 1* (March 16), https://spectrumnews1.com/ky/louisville/news/2021/03/16/dollar-general-could-assist-in-vaccine-distribution.

Roeloffs MW (2024) Here's why drug stores are closing in minority neighborhoods: Walgreens, CVS and Rite Aid shutter more than 1,000. Forbes (January 14), https://www.forbes.com/sites/maryroeloffs/2024/01/14/heres-why-drug-stores-are-closing-in-minority-neighborhoods-walgreens-cvs-and-rite-aid-shutter-more-than-1000/.

Safegraph (2022) Starbucks locations in the USA. Accessed May 2022, https://www.safegraph.com/.

Salako A, Ullrich F, Mueller KJ (2018) Update: Independently owned pharmacy closures in rural America, 2003–2018. *Rural Policy Brief* 2018(2):1–6.

Samorani M, Harris SL, Blount LG, Lu H, Santoro MA (2022) Overbooked and overlooked: Machine learning and racial bias in medical appointment scheduling. *Manufacturing Service Oper. Management* 24(6):2825–2842.

ScrapeHero (2024) Dollar store locations in the USA. Accessed May 2024, https://www.scrapehero.com/store/product/dollar-store-locations-in-the-usa/.

Şen A, Atamtürk A, Kaminsky P (2018) A conic integer optimization approach to the constrained assortment problem under the mixed multinomial logit model. *Oper. Res.* 66(4): 994–1003.

Shepard M (2022) Hospital network competition and adverse selection: Evidence from the Massachusetts health insurance exchange. *Amer. Econom. Rev.* 112(2):578–615.

Syed ST, Gerber BS, Sharp LK (2013) Traveling towards disease: Transportation barriers to health care access. *J. Community Health* 38(5):976–993.

Tan S, Frazier PI (2022) Regret bounds and experimental design for estimate-then-optimize. Preprint, submitted October 27, https://arxiv.org/abs/2210.15576.

Tay A (2003) Assessing competition in hospital care markets: The importance of accounting for quality differentiation. *RAND J. Econom.* 34(4):786–814.

U.S. Census Bureau (2021) American community survey, 2015–2019. https://www.census.gov/topics/research/guidance/planning-databases.2021.html.

U.S. Centers for Disease Control and Prevention (2023) Federal retail pharmacy program. Accessed May 2024, https://archive.cdc.gov/#/details?url=https://www.cdc.gov/vaccines/covid-19/retail-pharmacy-program/index.html.

Wang G, Zheng R, Dai T (2022) Does transportation mean transplantation? Impact of new airline routes on sharing of cadaveric kidneys. *Management Sci.* 68(5):3660–3679.

Wittenauer R, Shah PD, Bacci JL, Stergachis A (2024) Locations and characteristics of pharmacy deserts in the United States: A geospatial study. *Health Affairs Scholar* 2(4):qxae035.

Wu TH (1997) A note on a global approach for general 0–1 fractional programming. *Eur. J. Oper. Res.* 101(1):220–223.

Xu Y, Armony M, Ghose A (2021) The interplay between online reviews and physician demand: An empirical investigation. *Management Sci.* 67(12):7344–7361.

Zhang Y, Berman O, Verter V (2012) The impact of client choice on preventive healthcare facility network design. *OR Spectrum* 34(April):349–370.

Zhong H, Wang G, Dai T (2021) Wheels on the bus: Impact of vaccine rollouts on demand for public transportation. Preprint, submitted July 7, http://dx.doi.org/10.2139/ssrn.3874150.

ZipRecruiter (2024) What is the average Snap nurse salary by state. Accessed May 2024, https://www.ziprecruiter.com/Salaries/What-Is-the-Average-Snap-Nurse-Salary-by-State.