

Beyond the Neighborhood: Regional Amenities in Sorting Models*

Mo Van de Sompel[†]

October 15th, 2026

Abstract

Equilibrium sorting models treat amenities as localized; they are enjoyed equally within a geographic unit and not at all outside of it. However, many environmental amenities generate diffuse benefits across broad regions. Incorporating such regional amenities into sorting models requires solving a fundamental problem: at what rate do benefits attenuate with distance and when are they negligible? I demonstrate empirically that arbitrary decisions made by the researcher about an amenity's market extent can cause estimates of its value to vary by orders of magnitude. I address this problem by endogenizing value attenuation with the results of a recreation demand study, microfounding a spatial index that estimates the exposure of any given location to a complete network of amenities. To demonstrate, I estimate a model of residential sorting across Iowa and find not only that people have significant willingness to pay for improvements in lake access (\$5,945 for a one standard deviation increase), but that failing to account for them threatens the exclusion restriction of the standard sorting instrument; failing to account for them attenuates the coefficient on price by 9%. In counterfactuals I simulate the creation of a lake, and my estimated model shows that interregional variation in supply elasticity allows the resulting demand shift to be absorbed by higher-elasticity locations, softening the price response everywhere and raising concerns for the ability of hedonic models to fully capitalize the value of regional goods even in low-elasticity places. Finally, I show the importance of spatially modeling environmental goods by simulating cleanup projects at two lakes that differ sharply in substitute availability: the isolated lake's cleanup is worth roughly 20 times more than the identical water-quality improvement at the lake with substitutes nearby.

*I thank my advisors, Jonathan Colmer, Lee Lockwood, Jay Shimshack, Ahnaf Rafi, and Julie Holland Mortimer for their help and contributions to this paper. I also thank Leora Friedberg, Amalia Miller, Kerem Cosar, Simon Anderson, Francesco Ruggieri, Colin Williams, Thomas Willingham, Camille Wixon, and Yiren Ding for their feedback at myriad presentations at UVA's workshops. Thanks also to the attendees of the 2025 and 2026 Southern Economics Association conference for their valuable questions and suggestions. Finally, I am hugely indebted to Yongjie Ji of the Iowa Lakes Project for his help in the estimation of my recreation demand model when data were inaccessible.

[†]University of Virginia. mv8zr@virginia.edu

1 Introduction

When asked about why they chose to live where they do, the public often cites nature and greenspace as one of the most important factors (Thomas et al., 2015; Urban Land Institute, 2015; Scott et al., 2020). This implies that environmental amenities may be important contributors to household sorting decisions, a phenomenon which has been evidenced by various papers studying the importance of local environmental amenities such as open space, air quality, and urban heat. But many environmental amenities are not local, features like beaches, mountains, lakes, or national parks are often well outside of one’s immediate neighborhood, but access to them may still be a significant factor in sorting decisions. These goods are characterized by what I call regionality: rather than having their benefits concentrated tightly within some geographic boundary, they impact a broad region, with the value at any point being determined by the amenity’s accessibility therefrom.

Regionality introduces problems for economists studying the value of these goods and their effects on sorting. A common tool for research in this space is the equilibrium sorting model, where households make a discrete choice decision over neighborhoods which are modeled as bundles of pure characteristics. This approach works well for localized goods: since people mostly care only about the air quality in their own neighborhood, we can measure PM2.5 there and use that as a proxy for the neighborhood’s exposure to air pollution, and the estimated coefficient tells us the willingness to pay (WTP) to mitigate that exposure. The problem with regional goods is precisely that they are *not* in the neighborhood, we require a way of measuring the location’s exposure to an amenity outside of it.

The core difficulty is in measuring value and accessibility. Realistically a location’s exposure to an amenity is a continuous function in space, peaking at the amenity itself and attenuating with distance. However, modeling this accurately would require knowledge of the amenity’s valuation, which is often what is being studied in the first place. Inspiration for heuristics can be drawn from the hedonic literature: one approach is an indicator for whether the amenity is present within the neighborhood itself or some distance therefrom

(Muehlenbachs et al., 2015; Keiser and Shapiro, 2019; Greenstone and Gallagher, 2008), another is simple distance to the amenity in question (Walsh et al., 2011; Wolf and Klaiber, 2017; Gamper-Rabindran and Timmins, 2013). These differ in their strengths and weaknesses, and may be chosen to measure exposure based on the perceived importance of size, numerosity, substitutability, and willingness to travel for access. A particular weakness of both methods is market extent, i.e., which neighborhoods count as having any exposure at all? Allowing market extent to be determined by municipal boundaries or distance thresholds is arbitrary, assuming no one outside gets any value and everyone inside gets equal value. Using distance raises the inverse problem: everyone is in the market no matter how far away. This is particularly problematic in the environmental space because, as Smith (1993) argues, the marginal valuation of a public good is likely to be less important than the number of people we attribute that value to. Leaving quantity to be determined arbitrarily risks misspecifying the value calculation entirely.

Rather than using these heuristics, this paper implements a methodology for including regional goods in equilibrium sorting models by estimating the exposure function directly. By using data on household travel decisions to amenities, I model a system of amenity demand in which households have access to a network of sites and I estimate parameters dictating the relative valuation of each one’s characteristics. This system is then summarized by a spatial index, in which each site in the network contributes value relative to its quality and location. The index is microfounded and carries an economic interpretation as the expected value of an opportunity to take a trip to an amenity from any given location. Once parameters are estimated, the index is easily calculated for every neighborhood in the sorting model and proxies as a measure for each location’s exposure to the full network of amenities.

This approach has several key benefits. First, market extent and value attenuation are determined by data. The parameter that affects how quickly the amenity loses value the further away it is determined by revealed-preference data showing how far people are actually willing to travel to access it. After a certain distance, costs are so high that its

contribution to the index is zero. Secondly, the approach naturally allows us to introduce multiple measures of amenity quality as well as accessibility. While this would be possible with heuristic methods using interactions between accessibility and quality, the more dimensions a researcher adds the more they are stretching housing sales data to finely disentangle the value of each dimension. Finally, this approach naturally deals with issues that arise due to substitutability and returns to variety when analyzing multiple similar amenities. When simulating counterfactuals with regional goods it's necessary to keep in mind that more than one may be accessible, and the costs of making a low-value option worse or unavailable are lower than doing the same for a high-value one.

To demonstrate, I implement my approach in an equilibrium sorting model of Iowa, with its lakes as the regional amenity of interest. Lakes exhibit many of the more problematic aspects of regional amenities: they are numerous, many locations have access to several nearby, but also geographically variant, some parts of the state have essentially no lakes. They vary along many dimensions of quality, at minimum there is the question of size, but cleanliness and boat access features are also relevant. Finally, they are plausibly important to utility, with 65% of Iowans visiting a lake in a given year and an average of 7.8 single-day trips among those going (Wan et al., 2021). The model is estimated on 3,500 Iowa neighborhoods across 3 year-blocks spanning 2010-2022, $\sim 650,000$ observed housing sales are used to estimate each neighborhood's representative price per year-block, drawing other demographic and tax data from various sources. To construct the exposure index, I utilize the travel cost data from the Iowa Lakes Project, a repeated multiyear survey of Iowans' recreation habits as estimated in Ji et al. (2020, 2022).

I start by demonstrating the problem of arbitrary market definition by taking a distance threshold approach to measuring lake availability: I estimate the sorting model separately for 15 different distance thresholds for a neighborhood to be considered having access to a lake. Within these specifications, the estimated willingness to pay to be within the threshold varies between \$13,000 and \$35,000. I further show that the lost consumer surplus attributed

to removing any given lake varies sharply across these fifteen specifications, with a median ratio of $12.7\times$ between a lake’s largest and smallest implied value, and some lakes exceeding $100\times$. This demonstrates the difficulties regional goods bring to public goods valuation.

I then implement my index approach. Estimating the same model using the index instead of distance thresholds, I find a WTP of \$5,945 for a one standard deviation increase in lake quality-accessibility. The median lake raises per-household consumer surplus by around \$5.51, with a few very high-value lakes over \$100 raising the mean to \$21.12. By supplementing the model with a supply side drawn from Baum-Snow and Han (2024), I am able to answer novel policy questions about sorting responses to amenity changes. To illustrate, I simulate a counterfactual of creating Gray’s Lake at its central location in Des Moines. Because of equilibrium sorting dynamics, demand-side substitution across the region surrounding the lake blunts part of the effect on prices. This implies that in addition to the known issue of incomplete capitalization of amenities into home price due to supply elasticity (Hilber and Mayer, 2009), regional goods may be incompletely capitalized even in inelastic markets due to competition from more-elastic neighbors. I further estimate two twin counterfactuals assessing the benefits of potential cleanup projects for Lakes Rathbun and Little Spirit. I model the welfare effects of equal improvements in Secchi depth, one of my dimensions of lake quality. I find that a cleanup in Rathbun, a relatively isolated lake with few nearby substitutes, delivers roughly 20 times higher welfare gains than the same improvement in Little Spirit, a lake with many nearby substitutes. This highlights the importance of taking seriously the spatial nature of the world our models approximate.

I contribute to two main literatures. The first set of contributions is in the valuation of environmental/public goods, where I increase the range of goods which can be valued, document and mitigate a potential threat to identification in sorting models, and document a novel dynamic of regional goods in hedonic models.

Equilibrium sorting and hedonic models have become workhorses for valuing, non-market goods, with the former especially commonly used in estimating policy counterfactuals. En-

environmental applications have been made in measuring the value of air quality (Sieg et al., 2004; Bayer et al., 2009; Tra, 2010, 2013; Hamilton and Phaneuf, 2015), open space (Walsh, 2007; Klaiber and Phaneuf, 2010), urban heat (Klaiber et al., 2017; Borsky et al., 2024), and others. Of particular interest to me are Hamilton and Phaneuf (2015), who find evidence that people sort at multiple layers, treating certain amenities as local to a region and others as local to a neighborhood. Within this space, a substantial methodological literature has developed asking how much the housing price function can tell us about preferences for the amenity in question (Kuminoff et al., 2010; Abbott and Klaiber, 2011; Klaiber and Smith, 2013; Kuminoff and Pope, 2014; Banzhaf, 2021; Zhang et al., 2015). I show that the prior step—deciding what counts as the amenity and who has access to it—is at least as consequential. By replacing researcher discretion over market extent with a parameter recovered from revealed preference, I increase the range of environmental goods which can be evaluated using these models, albeit at the cost of additional data.

I also show two econometrically relevant findings for the valuation of public goods. First, in the equilibrium sorting space: the existence of regional amenities threatens the standard identification strategy for sorting models. The ring-based instrument used to purge price of unobserved quality as introduced by Bayer et al. (2004) and Bayer and Timmins (2007) relies on distant housing characteristics affecting own price only through equilibrium sorting behavior, not by affecting utility directly. But accessibility of regional amenities is highly informed by land use outside the typical 3-5 mile extent of the ring, which itself may be correlated with housing characteristics outside the ring. Thus, if accessibility is not explicitly controlled for, it will correlate local unobserved utility with distant characteristics and violate the exclusion restriction. Second, in the hedonic space, it is well established that amenities are incompletely capitalized into house prices where housing supply is elastic (Hilber and Mayer, 2009; Hilber, 2017). Through my model, I show that regional goods are subject to a second, spatial version of this problem: because demand induced by a regional amenity is spread across an entire region, higher-elasticity locations absorb part of the re-

sponse and soften price effects everywhere, including in locations that are themselves supply constrained. Hedonic estimates of regional goods are therefore attenuated regardless of local supply elasticity—a concern that cannot be diagnosed by conditioning on the elasticity of the treated market.

Finally, the index approach also expands the set of answerable policy questions as regards environmental goods. As initially implemented by Phaneuf et al. (2008) and Kuwayama et al. (2022) in hedonic models, this index allows researchers to disentangle the effects of changes in an amenity locally (e.g., the immediate effects of proximity) vs regionally. By embedding it into an equilibrium sorting model, we can alleviate some of the strong assumptions required for general equilibrium valuations. In particular, hedonic work typically assumes that changes are marginal enough to not induce a new equilibrium, which has been reasonable in the chosen contexts, but precludes the valuation of, for example, a spatially isolated high-quality site. Further, the index allows for plausible counterfactuals in situations where substitution is a major concern. Previous work has been forced to treat additional availability crudely: a counterfactual where a new park opens simply increases total parkspace but does not account for the fact that the new park may pull users away from what was available before, or may be lower quality and never used, thus adding almost nothing. Under this framework, changes in an amenity enter through a second model which estimates, taking into account quality, substitutes, and accessibility, how much the change matters.

My second set of contributions are to the growing spatial environmental literature. As argued in Balboni and Shapiro (2025), the environmental and urban/spatial literatures have been largely separate to the detriment of both. Many environmental papers have established that environmental goods have spatial impacts on their surroundings (Walsh et al., 2011; Wolf and Klaiber, 2017; Gamper-Rabindran and Timmins, 2013), but often stop short of building a fully spatial model. My twin lake counterfactual shows the importance of doing so by finding that the same policy can have huge differences in welfare implications simply because of the spatial distribution of lakes. On the reverse side, work in urban economics using sorting

models has found that amenities and sorting affect urban form, inequality, and incumbent welfare (Couture et al., 2024; Almagro and Domínguez-Iino, 2024), but rarely direct their interest towards environmental goods. I model the spatial distribution of amenity exposure directly, using an index in the spirit of Hansen (1959) and that underpins the market-access measures now standard in quantitative spatial economics (Donaldson and Hornbeck, 2016; Ahlfeldt et al., 2015; Redding and Rossi-Hansberg, 2017). The relative contribution here is to estimate the index from microdata rather than calibrate it from macro moments: the rate at which an amenity loses value with distance is estimated from choices over the amenity itself, and the quality arguments of the index are those that demonstrably shift trip decisions. Environmental goods are a particularly natural setting for this, since trip-level revealed preference data are routinely collected for the sites whose value we wish to measure.

I also make a contribution in choice of geography. Most equilibrium sorting models use administrative boundaries as the base unit across which households make decisions. This is often a natural choice: access to many goods varies at the municipal boundary so it represents a meaningful bundle of goods. However, in some cases, the presence of the amenity in question varies significantly within a boundary. Homes in the same Census tract or block group may be more or less exposed to air pollution depending on proximity to a highway, and taking a single average metric as representative would cover meaningful heterogeneity. Further, boundaries such as Census geographies vary significantly in size, meaning representative metrics are typically more representative for small, urban geographies than larger rural and suburban ones. This is the first equilibrium sorting paper to use the Gridded Environmental Impacts Frame (EIF), a $.01^\circ \times .01^\circ$ grid that provides estimates of population demographics per cell for the entire United States. Using these cells as the sorting geography of choice allows spatial metrics to vary finely in the model without concern that strong heterogeneity is being masked.

The paper continues as follows, section 2 lays out the equilibrium sorting model, section 3 introduces the data, section 4 shows results and policy counterfactuals.

2 Model

The empirical strategy has two layers. The first is a residential sorting model in which agents make a discrete choice decision over different neighborhoods, each modeled as a bundle of characteristics. The second is a recreation-demand model of trips to lakes, which delivers an amenity exposure index for every location in the model, used as one of those neighborhood characteristics. I present the sorting model first, since it is the object ultimately estimated and disciplines the interpretation of everything else, followed by the housing-supply side used for general-equilibrium counterfactuals. I then show two alternative ways of measuring site exposure that do not require the travel-cost layer — indicator thresholds and continuous distance — since these are the natural benchmarks against which the index is judged in section 4; finally I present the travel-cost model itself and show how its output becomes the index.

2.1 Sorting Model

Neighborhoods are the products in a static discrete-choice demand system over residential location. A household i chooses to buy in neighborhood j during time period t if doing so gives it the highest utility among all $L = J \times T$ available location-time-periods. Utility from choosing $l = \{j, t\}$ is

$$u_{i,l} = \alpha p_l + X_l \beta + \xi_l + \epsilon_{i,l}$$

where p_l is the prevailing home price, X_l is a vector of observed neighborhood characteristics, ξ_l is a location-time-period-specific unobservable capturing everything about j 's desirability in period t that is not in X_l , and $\epsilon_{i,l}$ is a Type-I extreme value taste shock idiosyncratic to household i . Aggregating over households with i.i.d. $\epsilon_{i,l}$ gives the standard logit market

share for l ,

$$s_l = \frac{\exp(\alpha p_l + X_l \beta + \xi_l)}{\sum_{c=1}^L \exp(\alpha p_c + X_c \beta + \xi_c)}.$$

Inversion In order to estimate the model I follow the standard two-step approach of Berry (1994) using the instrument as implemented for equilibrium sorting by Bayer et al. (2004). First, observed market shares $s_l = n_l / \sum_c n_c$ are inverted via the contraction mapping¹

$$\delta_l^{(k+1)} = \delta_l^{(k)} + \ln s_l - \ln \hat{s}_l(\delta^{(k)}), \quad \hat{s}_l(\delta) = \frac{e^{\delta_l}}{\sum_c e^{\delta_c}},$$

which converges to the unique mean utility level $\delta_l = \alpha p_l + X_l \beta + \xi_l$ that makes the model exactly match observed shares in every location.

Instrument Because ξ_l is unobserved to the econometrician but is correlated with price, α cannot be recovered by ordinary least squares. Following Bayer et al. (2004), I construct an instrument that takes advantage of the sorting behavior implied by the model to satisfy relevance, using the influence of distant substitute sites for which exclusion plausibly holds. The assumption is that the availability of distant sites only affects price through sorting decisions. In particular, for each location l I construct concentric annuli² around the neighborhood and define Z_l as the mean of every characteristic of X_l for each neighborhood in the annuli. I further use land-use characteristics L_l which inform the difficulty of construction in other nearby neighborhoods. The above are all assumed to be correlated with ξ_l , so we proceed as follows: given a guess of α , α_0^* , I partial price out of the mean utility,

$$\delta_l - \alpha_0^* p_l = X_l \psi + Z_l \phi + L_l \zeta + \tilde{\xi}_l, \tag{1}$$

¹With homogeneous preferences, as used throughout this paper, the standard log-inversion is equivalent to the fixed-point iteration.

²0–5, 5–10, 10–20, and 20–30 miles

and estimate (ψ, ϕ, ζ) by OLS. The fitted value $X_l\hat{\psi} + Z_l\hat{\phi} + L_l\hat{\zeta}$ is an estimate of the portion of mean utility explained by l 's own observed and nearby characteristics; the residual, $\tilde{\xi}_l$, is thus representative of influence originating beyond the annuli we set. Setting $\tilde{\xi}_l = 0$ and adding $\alpha_0^* p_l$ back, we have

$$\hat{\delta}_l = \alpha_0^* p_l + X_l\hat{\psi} + Z_l\hat{\phi} + L_l\hat{\zeta} \quad (2)$$

Which we substitute into the market clearing condition

$$\frac{\exp\left(\alpha_0^* p_l^{IV} + X_l\hat{\psi} + Z_l\hat{\phi} + L_l\hat{\zeta}\right)}{\sum_{c=1}^L \exp\left(\alpha_0^* p_c^{IV} + X_c\hat{\psi} + Z_c\hat{\phi} + L_c\hat{\zeta}\right)} = s_l \quad \forall l, \quad (3)$$

And solve for the price vector p^{IV} for which predicted logit shares equal the observed shares s_l . This new price vector is then used as the instrument on price to estimate

$$\delta_l = \alpha p_l + X_l\beta + \xi_l \quad (4)$$

which recovers a new $\hat{\alpha}_1^*$. In order to mitigate the effect of the initial guess of α_0^* , the above process is repeated until convergence³.

The instrument works because it is informed by the residual variation in $\tilde{\xi}$, which by construction can only originate from sites beyond our annuli. It is thus representative of the influence of distant site characteristics through the sorting process, distant enough that we assume no correlation between their characteristics and ξ_l . Note that this assumption is precisely what a regional amenity threatens: if a regional amenity itself is omitted or mismeasured, the resulting error is, by construction, similar for not just nearby locations but anything in the region, contaminating the instrument meant to be orthogonal to it. While including a measure of regional amenity exposure directly in X_l can keep this source of spatial correlation out of ξ_l , the cruder the exposure measure, the more of the true spatial

³This usually happens within 5 loops

correlation in lake access is left for ξ_l to absorb, and the more this identifying assumption is strained. I further strengthen the plausibility of this assumption in my work by including much wider annuli than is common in the literature, so as to capture only the influence on neighborhoods so far away that even regional amenities likely do not correlate.

Welfare. Given estimates $(\hat{\alpha}, \hat{\beta})$ and the recovered $\hat{\xi}_l$, the compensating variation from a change in characteristics that moves the choice set from $\{X_l\}$ to $\{X'_l\}$, holding $\hat{\xi}_l$ fixed at each location, is the standard logit consumer-surplus formula

$$\Delta CS = \frac{1}{\hat{\alpha}} \left(\ln \sum_{l=1}^L e^{\hat{\alpha} p_l + X_l \hat{\beta} + \hat{\xi}_l} - \ln \sum_{l=1}^L e^{\hat{\alpha} p_l + X'_l \hat{\beta} + \hat{\xi}_l} \right). \quad (5)$$

This is a partial equilibrium estimate measuring how much you would have to compensate the representative household to be indifferent to the new choice set. Because this welfare measure aggregates over the entire choice set, it prices in substitution: if a household's preferred location becomes less attractive relative to a close substitute as a result of a counterfactual equilibrium, ΔCS will reflect that the household can simply move.

Some counterfactuals in section 4 instead ask what happens if households are *not* free to re-sort. Particularly, when we are interested in the partial equilibrium spatial impact of a certain amenity, we want to know how much value different neighborhoods lose from changes in their characteristics due to change in the amenity. In that case, the relevant choice set for location l collapses to a single alternative, and the compensating variation for the household(s) already living in l is simply

$$\Delta CS_l^{\text{no re-sort}} = \frac{(X'_l - X_l) \hat{\beta}}{\hat{\alpha}},$$

I use this measure whenever I want to isolate how a change in local characteristics is distributed across the spatial model without also mixing in the (separate) question of how much substitution mutes it.

2.2 Counterfactuals and the Supply Side

In order to estimate general equilibrium counterfactuals I also implement a housing supply side. This allows prices to realistically fluctuate in response to changes in amenity value and accessibility. Particularly in the case of regional amenities, having both relatively elastic and inelastic neighborhoods available causes interesting interactions in sorting because amenities are partially shared across locations.

I assume constant elasticity of supply in each Census tract, $Q_l^s = A_l p_l^{\eta_l}$, with tract-level elasticities η_l taken from Baum-Snow and Han (2024)⁴ and the scale parameter A_l calibrated so that observed price and quantity satisfy the supply curve exactly at baseline. Given a change in characteristics, I solve for a new market-clearing price vector by iterating between (i) recomputing logit demand shares given the current price vector and (ii) adjusting each location’s price in proportion to the gap between demand and supply at that price, in the spirit of the tâtonnement process in Sieg et al. (2004), until prices converge. Because demand for a regional amenity spreads across every location within reach of it rather than concentrating on a single treated neighborhood, this setup lets locations with different supply elasticities absorb different shares of the same regional demand shock — the central mechanism behind the price-capitalization results in subsection 4.5.

2.3 Threshold and Distance Approaches

Before introducing the exposure index, I briefly introduce two measures of regional amenity exposure that have been common in the hedonic literature.

Indicator thresholds. A simple approach lets a location “have” the amenity if the nearest site falls within a given distance d of it. Defining dist_l as the minimum distance to a site, the threshold enters the sorting model as $\mathbb{1}[\text{dist}_l \leq d]$. In order to investigate the stability of estimates to d , I estimate the full sorting model separately for fifteen values of d from

⁴Where estimates are unavailable, I simply assume perfect elasticity, since tracts without estimates are typically low-population and remote. They generally do not affect counterfactuals significantly

1 to 15 kilometers. This measure’s biggest weakness is its arbitrary definition of market extent, allowing for no one outside the radius exposure and everyone inside equal exposure when realistically the function is continuous. Ideally, these issues would cancel out in such a way that the larger the radius, the lower the average value from exposure per-resident but the more residents, such that the marginal valuation of any given site is similar across specifications, thus making the arbitrary threshold innocuous.

Continuous distance. This approach is most commonly used for unique amenities such as distance to a city’s central business district. Still, it can be adapted to a context with multiple sites as distance to the nearest site, along with the square to allow for nonlinearity. These two measures simply enter as

$$X_l \ni (\text{dist}_l, \text{dist}_l^2) .$$

This avoids the arbitrary cutoff but introduces the concern that there is no limit to market extent: everyone gets some exposure to some amenity. Critically, both the threshold and distance measures are purely geometric: they encode where a site is but nothing about it. A site’s estimated value under either measure is therefore driven entirely by a single common valuation among all sites, and how many households happen to live near it, with no way to discount low-quality sites from high-quality ones. This is a particularly costly omission in a setting with real variance in quality, since it means removing a low-quality site is treated identically to removing a high-quality one at the same location.

2.4 Travel Cost Model and the Amenity Exposure Index

In my attempt to address the above concerns over market extent, value attenuation, and differences in quality, I attempt to measure amenity exposure more directly. I accomplish this through a quality-accessibility index built from a recreation demand model of trips to the amenity in question, following the logic of Phaneuf et al. (2008) and Kuwayama et al.

(2022). On any of a fixed number of choice occasions in a year, a household living in location l visits one of J sites. The (non-random) value of visiting site j from l in year y is

$$a_{j,l,y} = \kappa_{j,y} + \tau_y tc_{l,j,y},$$

where $\kappa_{j,y}$ is an alternative-specific constant (ASC) capturing site j 's average desirability in year y , later to be decomposed into site characteristics, and $tc_{l,j,y}$ is the round-trip travel cost between l and j ,

$$tc_{l,j,y} = 2 \left(c_y \cdot d_{l,j} + \frac{1}{3} \text{time}_{l,j} \cdot \text{wage}_{l,y} \right),$$

where c_y ⁵ is per-mile travel costs, $d_{l,j}$, the driving distance between l and j , and the opportunity cost of travel time is valued at one-third of the local wage multiplied by $\text{time}_{l,j}$, the driving time, following the standard travel-cost convention (Ji et al., 2020, 2022; Zhang et al., 2024; Cheng and Wan, 2026). τ_y is estimated, and reported throughout, as a raw coefficient rather than a magnitude to be subtracted; it is negative (Table 5), so higher travel cost lowers utility as expected. Adding a Type-I extreme value taste shock gives visitation probabilities

$$P(\text{visit } j \mid l) = \frac{e^{a_{j,l,y}}}{\sum_{k=1}^J e^{a_{k,l,y}}}$$

Writing $n_{i,j,y}$ for the number of trips household i took to site j in year y , out of the fixed number of choice occasions, the log-likelihood is

$$\ell(\{\kappa_{j,y}\}, \tau_y) = \sum_i \sum_{j=1}^J n_{i,j,y} \ln P(\text{visit } j \mid l_i),$$

I maximize this separately for each survey year, recovering one ASC per site-year along

⁵from AAA's published year-specific driving-cost estimates

with an annual travel-cost parameter τ_y . The ASCs are themselves the object of interest for recovering preferences over site characteristics: in a second stage, I regress $\kappa_{j,y}$ on observed site characteristics X_j ,

$$\kappa_{j,y} = \gamma X_j + \varepsilon_{j,y},$$

by OLS. While this could also be done separately by year, I find no evidence that preferences for these site characteristics change over time, so to simplify I use a pooled result⁶. The fitted $\hat{\gamma}$ (reported in Table 6) is what lets the index be recomputed for a site whose characteristics change, or for a hypothetical site that does not exist in the estimation sample, by predicting $\hat{\kappa}_j = \hat{\gamma} X_j$ rather than requiring κ_j to have been separately estimated for it. For the same reason, I use a single $\hat{\tau}$, averaged across all seven survey years, rather than the year-specific $\hat{\tau}_y$: the index needs one consistent travel-cost sensitivity applicable to every grid cell and year-block in the sorting model, not seven different ones tied to the recreation survey's own calendar.

The quality-accessibility index for location l is the expected value of the best available recreation opportunity from l ,

$$\hat{A}_l = E[\max(u_{i,1,l}, \dots, u_{i,J,l})] = \ln \left(\sum_{j=1}^J e^{\hat{\kappa}_j + \hat{\tau} t c_{l,j}} \right),$$

the standard logit inclusive value, evaluated at the pooled $(\hat{\tau}, \hat{\gamma})$ rather than at any single survey year's estimates. $\hat{A}_l$ rises in the quality of nearby sites and falls in the travel cost of reaching them. Removing or changing X_j for one site changes $a_{j,l}$ for that site at every location simultaneously, in proportion to each location's travel cost to it. Thus, a counterfactual index recomputes $\hat{A}_l$ for every neighborhood in the model, and the size of $\hat{\tau}$ determines how far away that change will be meaningfully felt. Figure 1 shows how the

⁶Ji et al. (2020) finds that preferences do shift for certain characteristics over time. Since the eventual welfare object here is net present value of all future amenity flows, it doesn't make sense to use different preference parameters each year

index varies for the same set of amenities given higher or lower τ , it is the key parameter dictating spatial attenuation of amenity exposure.




Figure 1: The value of A_l across space when $\gamma X_k = k$ and amenities are located at 0, .5, and 1. Note that the index captures both the value of being near any amenity and the value of being closer to higher-quality amenities

3 Context and Data

I implement the model using Iowa's lakes as the regional amenity of interest. Lakes are a natural setting for this approach because they fail every simplifying assumption that makes localized-amenity sorting models tractable: their recreation benefits are not confined to a single neighborhood, they vary along several dimensions of quality, many locations can reach more than one of them, and access is highly uneven across the state.

3.1 Recreation Demand Data

The travel-cost model requires origin-destination trip data together with lake characteristics; the construction of travel cost itself, and its estimation, are described in subsection 2.4. Trips come from the Iowa Lakes Project, a repeated survey of Iowa households' recreation behavior conducted by Iowa State University in 2002–2005, 2009, 2014, and 2019. Respondents report the number of single-day trips taken to each of Iowa's significant lakes, pictured in Figure 2,

over the preceding season; following Ji et al. (2020) I treat each household as facing 52 choice occasions per year. Driving distance and time are measured between the household’s (or, when calculating the index at the grid-cell level, the grid cell’s) centroid and each lake; the per-mile driving cost is taken from AAA’s *Your Driving Costs* series, and the opportunity cost of time from the household’s self-reported income (or, for the grid cell, its median income), imputed to an hourly wage. Lake characteristics are drawn from the Iowa Department of Natural Resources and the Iowa State Limnology laboratory. Of particular interest is Secchi depth, a measure of visual lake clarity that estimates how far into the water one can still see. Various impurities such as phytoplankton from algal blooms or Phosphorus from sediment runoff cause lower Secchi depth and give water an unclear look. Summary statistics for lake characteristics are in Table 1.

Statistic	N	Mean	SD	Min	Max
Log area	892	4.795	1.687	2.303	9.852
Secchi depth (m)	892	1.175	0.941	0.087	8.100
Boat ramp (0/1)	892	0.857	0.348	0	1
State park (0/1)	892	0.380	0.486	0	1

Table 1: Lake characteristics, pooled across the 7 travel-cost survey years (892 site-years, 123–139 distinct lakes per year)

3.2 The Gridded EIF

A core concern with choosing how to define neighborhoods in the sorting model is spatial granularity. My variable of interest, lake accessibility, should not meaningfully vary within the chosen geography. Census tracts and even block groups are a poor fit for this requirement: their boundaries are drawn so as to keep population per unit roughly similar, so their physical size varies enormously. A rural Iowa block group can span tens of square miles, while a Des Moines block group can be a few blocks. A single lake index value assigned to an entire rural block group could easily average over travel costs that differ by an hour’s drive from one edge of the unit to the other, which is precisely the within-unit heterogeneity the index




Figure 2: The Lakes Considered in the Recreation Demand Study

is meant to resolve.

I instead use the Census Bureau’s Gridded Environmental Impacts Frame (Gridded EIF), a privacy-protected data product that reports population counts by race, income, age, and sex on a uniform $0.01^\circ \times 0.01^\circ$ latitude-longitude grid (roughly 0.7 by 0.9 kilometers in Iowa) covering the entire United States back to 1999. Quantitative spatial economics has increasingly turned to fine, population-independent grids rather than administrative units precisely because market-access and amenity-exposure measures are sensitive to the within-unit heterogeneity that tracts and block groups paper over (Donaldson and Hornbeck, 2016; Redding and Rossi-Hansberg, 2017), but that literature has generally built its own custom grids from scratch, raster population counts, or nighttime lights, each with its own coverage and privacy tradeoffs. The Gridded EIF offers a nationally standardized, privacy-protected alternative with demographic detail at the same resolution, assembled directly from Census and IRS administrative records rather than imputed from satellite imagery. To my knowl-




Figure 3: Map of all used grid cells, colored by index value

edge, this paper is the first to use it as the geography for a sorting model. Because grid cells are a fixed-size geography defined independently of population density, political boundaries, or the location of the amenity being studied, a cell's travel cost to any given lake is far more homogeneous within the cell than within a tract or block group. Iowa is covered by over 100,000 such cells, but we will restrict to the 3,557 with adequate housing sales data (subsection 3.3). These are disproportionately the state's more populated cells, containing roughly 75% of Iowa's population and 85% of its observed housing sales. In addition to providing geometry, the gridded EIF provides my measure of per-cell % white, median income, and PM2.5 exposure, summarized in Table 2.

Statistic	N	Mean	SD	Min	Max
% White	10,671	87.19	12.72	5.71	100.00
Median income (\$10,000s)	10,671	8.07	5.96	-1.28	59.33
PM 2.5 ($\mu g/m^3$)	10,671	7.87	1.15	5.66	11.57

Table 2: Gridded-EIF-derived characteristics across the 10,671 modeled grid-cell-by-year-block locations. A handful of very low-income-decile, low-population cells produce the negative minimum for median income; this reflects the underlying synthetic income imputation in sparsely-populated cells rather than a data error.

3.3 Housing Sales and Prices

Home prices are built from CoreLogic’s deed-transfer records, which cover Iowa property sales as far back as the early 20th century. I keep arm’s-length transactions with a recorded price of at least \$5,000, dropping split-parcel sales and same-day duplicate records at the same parcel, and restrict to the three year-blocks used in estimation (2010–2013, 2014–2017, 2018–2022), following Klaiber and Phaneuf (2010) and Bayer et al. (2009) in splitting the sample this way to balance transaction density against within-block price appreciation. This leaves 643,011 individual sales matched to a parcel- or block-level coordinate.

A grid cell’s representative price for a year-block should not simply be the average sale price observed there for all the years in the year-block, since sale prices drift within a year-block due to appreciation over the years a block spans. I first deflate every sale to 2010 dollars using BLS CPI-U annual averages, then, for each cell-by-year-block product h , recover a representative log-price π_h from

$$\ln P_h^i = \pi_h + y_t^i \theta_t + \varepsilon_h^i,$$

where P_h^i is the (CPI-deflated) price of the i th sale assigned to product h and y_t^i are calendar-year indicators. This is estimated as a single two-way fixed-effects regression across all products and all sales at once, with π_h recovered as the product fixed effect and θ_t a common year effect shared across every cell, following the price-adjustment approach of Bayer et al.

(2009). $\rho_h = \exp(\widehat{\pi}_h)$ is the representing price entering the sorting model.

A cell-by-year-block product is only as reliable as the number of sales backing it. Rather than impute a price for thinly-transacted cells (these cells often have almost no homes in them to start with), I require at least 5 sales in a given year-block for a grid cell to be assigned its own product in that block. Because the model is already static, I do not want certain locations to only be available at certain year-blocks, so any grid cell that cannot clear this bar in all three year-blocks is dropped from the choice set entirely. This leaves 3,557 grid cells, or 10,671 total grid-cell-by-year-block locations, with the sales-per-product counts summarized in Table 3.

Statistic	N	Mean	SD	Min	Max
Price (ρ_h , 2010\$)	10,671	130,065	78,950	11,471	853,604
Sales per product	10,671	49.95	49.37	5	380

Table 3: Recovered grid-cell-by-year-block prices and the number of sales backing each. The minimum of 5 is the binding cutoff described above, not an approximation of it.

3.4 Other Neighborhood Characteristics

As discussed in subsection 2.1, we require data on land use across the state of Iowa. Typically, studies such as Klaiber and Phaneuf (2010) and Bayer et al. (2004) draw on local data about land use that has detailed description on whether land is used for particular densities of residential, intensities of commercial use, and more. With my area of study being the full state of Iowa, no such locally-assembled land use data exist, and combining multiple local sources may leave me with inconsistent definitions of land use. Instead, I draw on the National Land Cover Database (NLCD), which classifies every 30-meter pixel in the country into one of several land cover categories, including four separate intensities of developed land (open, low, medium, and high development) alongside open water, forest, shrub/scrub, grassland, pasture/hay, cultivated crops, and wetlands. This is less-detailed than is ideal, but given over 80% of Iowa’s land is used for agriculture, it is likely sufficient.

I also compute median structural housing characteristics — square footage, bedrooms,

bathrooms, and lot size — per grid-cell-by-year-block product from the same CoreLogic transaction records used for price, taking the within-product median for all sales. Finally, I construct an indicator for whether a grid cell’s centroid lies within 10 miles of Iowa’s state border, motivated by the concern that the index, built only from Iowa’s own lakes, cannot see any recreational lakes just across the border which households near the border plausibly have access to.

Statistic	N	Mean	SD	Min	Max
Median sqft (100s)	10,671	17.08	5.89	4.60	56.23
Median bedrooms	10,671	3.06	0.57	1.00	20.00
Median bathrooms	10,671	2.19	0.87	1.00	8.00
Median lot size (acres)	10,671	0.33	0.51	0.00	18.62
On border (0/1)	10,671	0.227	0.419	0	1
% open water (0–1mi)	10,671	3.20	7.86	0.00	71.65
% developed, open (0–1mi)	10,671	11.16	6.61	0.00	38.98
% developed, low (0–1mi)	10,671	21.97	11.56	0.00	68.12
% developed, medium (0–1mi)	10,671	13.42	11.68	0.00	66.74
% developed, high (0–1mi)	10,671	3.20	4.71	0.00	49.42
% forest (0–1mi)	10,671	6.49	9.59	0.00	69.39
% shrub/scrub (0–1mi)	10,671	0.00	0.03	0.00	1.20
% grassland (0–1mi)	10,671	0.45	2.43	0.00	51.81
% pasture/hay (0–1mi)	10,671	8.01	9.37	0.00	64.96
% cultivated crops (0–1mi)	10,671	29.47	26.01	0.00	95.35
% wetlands (0–1mi)	10,671	2.41	4.60	0.00	51.49

Table 4: Structural housing characteristics, border proximity, and nearest-ring (0–1 mile) land use shares across the 10,671 modeled locations. Land use percentages are shares of land area, 0–100.

4 Results

4.1 The Trouble with Thresholds

I first estimate the sorting model with lake access measured as an indicator for being within d kilometers of the nearest lake, separately for every integer d from 1 to 15. Each threshold requires a full re-estimation of the model because these are fifteen different definitions of what counts as “having a lake nearby.” The implied WTP for lake access is estimated as




Figure 4: Willingness to pay for lake access, defined with 15 different distance cutoffs

$\beta_{\mathbb{1}(d_{l,lake} < d)}/\alpha$ for each of the models. An ideal measure of lake access would be invariant to this choice in one specific sense: widening the threshold waters down the definition of “accessible,” so per-person WTP should fall smoothly as more, more-distant households are included, with the two effects roughly offsetting so that the *total* value attributed to lakes is stable. Figure 4 shows this benchmark fails. Per-person WTP does fall as expected out to 4km, but then it rises again and stabilizes after 7km, even as the treated population continues to grow from 20% to 65%. Total value derived from lakes is therefore not stable across specifications — it increases purely because of how the researcher chose to define accessibility, which is the opposite of what the benchmark requires. This is concerning, but to see the problem more clearly we must consider the exercise of valuing individual lakes.

For a given threshold d and a given lake, I compute the consumer surplus lost from removing that lake. This is accomplished by removing the lake, re-estimating whether each neighborhood falls within the threshold under the resulting counterfactual distances, and using Equation 5 to price the change for a representative agent. If the threshold measure is accurately capturing the value of lake accessibility, this quantity would ideally be roughly *constant* across the fifteen specifications: the watered down definition of access roughly cancels out with the additional number of people who have access such that a researcher’s arbitrary choice of cutoff does not change how valuable a given lake is. Unfortunately, this is not the case. At many low thresholds, lakes are trivially worthless simply because they are

not within the distance cutoff of any of my grid cells, this is an uninteresting artifact of the model, so Figure 5 therefore plots the minimum non-zero value against the maximum across the fifteen specifications for each lake. Even restricting to non-zero comparisons, the ratio between the largest and smallest implied value has a median of $12.7\times$ and can exceed $100\times$. In other words, arbitrary decisions made by the researcher can generate orders of magnitude different estimations of value. Figure 6 illustrates the same point not just using minimum and maximum, tracing out the full fifteen-threshold profile for six representative lakes spanning the value distribution: none of the six profiles is flat, most aren't even monotonic.




Figure 5: Minimum non-zero and maximum lost consumer surplus from removing each lake, across fifteen separately-estimated distance-threshold specifications (1–15km), ranked by each lake's maximum implied value. 11 lakes worth \$0 at every threshold are omitted.




Figure 6: Threshold-by-threshold implied value for six representative lakes spanning the value distribution. Most profiles are neither flat nor smoothly monotonic, and no two share the same shape.

4.2 The Index Approach

I begin by maximizing the site-choice log-likelihood of subsection 2.4 separately for each of the seven survey years, recovering one travel-cost parameter τ_y and one ASC $\kappa_{j,y}$ per site. Table 5 reports τ_y together with its asymptotic standard error. τ drifts substantially across years; since the welfare object estimated here is meant to be the net-present value of all future amenity flows, rather than a snapshot tied to any one survey wave, I use the average of τ_y across all seven years (as described in subsection 2.4) for the index calculation in every year-block.

Year	τ_y (SE)	Mean κ_j	SD κ_j	Sites
2002	-1.780 (0.010)	-3.772	1.124	123
2003	-1.559 (0.008)	-3.937	1.139	123
2004	-1.449 (0.008)	-3.964	1.065	123
2005	-1.697 (0.009)	-3.762	1.125	123
2009	-1.869 (0.009)	-3.776	1.034	123
2014	-2.771 (0.015)	-3.760	1.054	138
2019	-2.841 (0.021)	-3.977	1.244	139

Table 5: First-stage travel-cost model results by survey year. τ_y and its standard error come from the inverse Hessian of the maximum-likelihood site-choice model at convergence; mean and SD of κ_j summarize the recovered site ASCs for that year.

The $\kappa_{j,y}$ recovered above are not used directly; the index needs to be recomputable for a site whose characteristics change or that does not exist in the estimation sample at all, which requires decomposing $\kappa_{j,y}$ into the underlying site characteristics via the second-stage regression of subsection 2.4, $\kappa_{j,y} = \gamma X_j + \varepsilon_{j,y}$. I pool all 892 site-year observations into a single regression rather than fit γ separately by year, on the same premise as earlier that while preferences may shift over time, we would prefer to model them as constant. Table 6 tests directly whether $\hat{\gamma}$ meaningfully shifts over time by adding a full set of year fixed effects and finds that regression coefficients are essentially unchanged. I move forward using the pooled model without fixed effects as $\hat{\gamma}$.

	(1)	(2)
Secchi depth	0.225 (0.021)	0.227 (0.021)
Log area	0.436 (0.017)	0.436 (0.016)
State park	0.238 (0.058)	0.238 (0.058)
Boat ramp	0.289 (0.088)	0.287 (0.087)
Year FE	No	Yes
R^2	0.574	0.581
N	892	892

Table 6: Second-stage decomposition of $\kappa_{j,y}$ into site characteristics, pooled across all 892 site-year observations, with and without a full set of year fixed effects. Robust standard errors in parentheses. The year fixed effects are jointly significant ($F = 2.53$, $p = 0.020$) but leave every characteristic coefficient essentially unchanged, which is what licenses pooling them into a single $\hat{\gamma}$.

With these simplifications we ensure that any change in $\hat{A}_t$ across year-blocks is not due

to varying preferences over time. I do allow variance within cell across year-blocks due to travel cost itself: the per-mile driving cost c_y ⁷ and the local wage used to value travel time⁸, both of which enter $\hat{A}_l$ directly. This is done with the idea that there will be consistent changes in travel cost over time due to rising wages and changing transportation prices that are worth modeling. With this, $\hat{A}_l$ is computed as in subsection 2.4, using $\hat{\tau}$, the fitted $\hat{\kappa}_j = \hat{\gamma}X_j$ for every one of the state's 143 lakes, and each cell's own block-specific travel cost to each of them. Their distribution is pictured in Figure 7, and the map of values by grid cell is pictured in Figure 3. The resulting index is standardized to mean zero and unit standard deviation⁹ before entering the sorting model as a single characteristic in X_l . Table 7 reports the full set of estimated coefficients.




Figure 7: Distribution of the accessibility index $\hat{A}_l$ across all 10,671 modeled grid-cell-by-year-block locations.

⁷The median of AAA's published rate over the years making up that block

⁸Also the median for each cell among the years in the block

⁹These are affine transformations of a term measured in utility and thus do not affect the economic interpretation

Variable	Coefficient	SE
Price (\$10,000s)	−0.1487***	(0.0049)
Access index (SD)	0.0884***	(0.0152)
Median income (\$10,000s)	0.1050***	(0.0056)
PM 2.5	−0.1852***	(0.0086)
% White	0.0004	(0.0009)
Sqft (100s)	−0.0045	(0.0038)
Bathrooms	0.5501***	(0.0276)
Bedrooms	0.0757***	(0.0286)
On border	0.0494*	(0.0275)

Table 7: Second-stage sorting model results, access index specification. *, **, *** denote significance at the 90%, 95%, and 99% level. Land use coefficients are suppressed.

Dividing the index coefficient by the price coefficient gives a willingness to pay of \$5,945 for a one standard deviation increase in the index. Using this coefficient, I can price the lost consumer surplus to the representative household under the counterfactual removal of each individual lake. As in subsection 4.1, these are partial equilibrium estimates: the loss in surplus to a representative household allowed to re-sort across the full choice set in response to the lake’s removal, with prices held fixed¹⁰. This gives a right-skewed distribution of lake values across the state’s 143 lakes, with a median of \$5.51 and a mean of \$21.12 per household, the gap driven by a small number of high-value sites against a long tail of lakes valued near zero.

A \$5,945 willingness to pay for a one-standard-deviation move might seem hard to square with a \$5.51 median lake value, but the two are consistent once the skew of $\hat{A}_l$ is taken into account: a one standard deviation increase can move a neighborhood from the 44th to the 95th percentile of exposure, so most of the distribution sits well within a single standard deviation of the mean and only a small number of locations are actually being moved by that much. This magnitude is also not dissimilar from the results Kuwayama et al. (2022) found when including recreation access in their hedonic regression in the context of Tampa Bay.

¹⁰In full equilibrium, prices also adjust, and — as subsection 4.5 and subsection 2.2 show — can absorb enough of the shock that the resulting change in consumer surplus is substantially smaller.

We can now benchmark the index approach against the threshold and distance approaches. Figure 8 compares the index specification to the continuous-distance specification, with each lake's per-household ΔCS ranked by its value under the index model. Figure 9 repeats this comparison against the threshold models, plotting each of the fifteen threshold specifications' per-household values. Overall, the distance model tends to value sites more highly than the index approach does, though not uniformly so. The threshold models show a similar tendency (setting aside the many lakes they mechanically set to zero value), especially for the lower-value lakes.




Figure 8: Per-household ΔCS for each lake under the index and continuous-distance approaches, allowing re-sorting at fixed prices, ranked by the index approach's valuation. The index values decline smoothly because of ranking.




Figure 9: Per-household ΔCS for each lake under the index and 15 different threshold models, allowing re-sorting at fixed prices, ranked by the index approach's valuation. The index values decline smoothly because of ranking.

What is driving these disparities is the subject of the next section: the approaches that only take distance into account ignore the other things that drive a lake's value to a given location.

4.3 How Site Value Varies with Distance

The clearest illustration of why modeling exposure using simple heuristics is so difficult is to look directly at how a lake's value at a given location changes with distance. For each lake, and for every grid cell in the state, I compute that cell's change in partial equilibrium value upon lake removal. Figure 10 plots this per-cell value against distance from the lake for four representative lakes, together with the binned mean trend. None of the four patterns is well approximated by a simple decay function, nor by a binary cutoff for inside and outside a threshold. In fact, value is not even monotonic as a function of distance. This non-monotonicity is not noise, it is the direct result of substitution through the underlying index. Two cells at equal distances but opposite directions from the same lake can have different

valuations of it, because one may have immediate access to a second lake nearby while the other has only the first. Non-monotonicity arises because a cell can be further from a lake but still value it more due to isolation. Fitting any single cutoff or continuous, monotonic function of distance to this pattern — across 143 lakes whose value-by-distance profiles take none of the same shape — cannot recover it, no matter how flexibly the function is specified, because the object being approximated is simply not a function of distance alone.




Figure 10: Per-household site value (no re-sorting, not weighted by population) against distance from the lake, for four lakes with continuous grid-cell coverage out to 30km. Blue points are individual grid cells; the red line is a binned mean.

4.4 Omitted Regional Amenities and Instrument Validity

The ring-based instrument of subsection 2.1, following Bayer et al. (2004), relies on an exclusion restriction: the characteristics of distant locations used to construct p_{iv} must be

uncorrelated with a given location’s own unobserved utility ξ_i . This restriction is harder to defend in the presence of a regional amenity. If lake access is not measured directly, it goes into ξ_i , but the lakes themselves are plausibly located outside the 5 mile rings standard in the literature. This means if any neighborhood characteristics correlate with being near lakes, including the controls for land use themselves, the exclusion restriction is violated. The problem here is precisely that the good is regional: it is far enough away to not be controlled for as contributing to local utility, but close enough that it still affects utility.

I test this directly by re-estimating the full instrument-construction and 2SLS procedure of subsection 2.1 with each lake-access measure removed from both the own-cell and ring-averaged characteristics used to build p_{iv} . Table 8 reports the resulting price coefficient $\hat{\alpha}$. Relative to a specification with no lake-access measure at all, including the index shifts $\hat{\alpha}$ by 9.2 percent; using the heuristic approaches shifts it by even more. In each case, $\hat{\alpha}$ is attenuated towards zero when the instrument doesn’t include lake access, the expected direction of when a valuable amenity is loaded into price. As a comparison, removing median number of bedrooms — a characteristic with its own significant, well-measured effect on price, but which does not plausibly inform utility for distant locations — from the index specification moves $\hat{\alpha}$ by only 0.35 percent. The instability is therefore not a generic consequence of dropping any one regressor; it is specific to measures of lake access. I stress that this occurs even under my rings which control for land use characteristics up to 30 miles out, well beyond what is common in the literature.

Specification	$\hat{\alpha}$	SE	% shift
No lake-access measure	−0.1361	(0.0047)	—
Access index	−0.1487	(0.0049)	+9.2%
Distance to nearest lake	−0.1586	(0.0053)	+16.5%
Threshold ($\leq 8\text{km}$)	−0.1552	(0.0052)	+14.0%

Table 8: Price coefficient $\hat{\alpha}$ under different lake-access specifications, each re-estimated with the full instrument-construction procedure. % shift is relative to the no-lake-access baseline.

This sensitivity is also not the result of instrument weakness leading to large swings. Table 9 reports the first-stage F -statistic and partial R^2 of p_{iv} with and without the access

index; both specifications easily clear conventional weak-instrument thresholds. Together these results show that regional amenities are a real threat to the exclusion restriction for the ring instrument in sorting models. In my case, not controlling for their presence would overstate all WTP estimates by roughly 10%.


Specification	First-stage F (robust)	Partial R^2
With access index	1,985.6	0.2247
Without access index	1,781.2	0.2249

Table 9: First-stage strength of p_{iv} , with and without the access index. F is the robust (HC1) Wald statistic on p_{iv} in the regression of price on p_{iv} and the remaining exogenous controls; partial R^2 is net of those same controls.

4.5 Counterfactuals: Creating Gray’s Lake

I estimate a general equilibrium response to the creation of a major amenity, illustrated with Gray’s Lake, an urban lake in downtown Des Moines and one of the highest-value lakes in the model. Creating it in partial equilibrium raises per-household surplus by \$95.66 (aggregate market welfare $\approx$ \$51.0 million). The shock to demand is concentrated around Des Moines and its immediate suburbs. Closing the model with the tract-level supply elasticities as per subsection 2.2 shows how this demand shift translates into prices; as is visible in Figure 11, prices rise sharply in Des Moines and its immediate suburbs, while falling modestly everywhere else.


More interesting is the role of supply elasticity in shaping these price changes. Because there is meaningful heterogeneity in supply elasticity within Des Moines, high-elasticity neighborhoods respond to the demand shock partially through increases in the quantity of housing supplied rather than through price alone. Because these neighborhoods are themselves substitutes for the less-elastic neighborhoods, they effectively dampen the price increase there by absorbing some of the excess demand that would otherwise fall on the more inelastic neighborhoods. Figure 12 zooms in on Polk County, containing Des Moines, and visualizes how the shock to the index does not translate directly into a change in price, with



(a) Change in GE Prices

(b) Binary change in GE Prices

Figure 11: Effect on Prices of Creating Gray's Lake Statewide



(a) Change in Index Value

(b) Supply Elasticity

(c) Change in Price

Figure 12: Effect of Gained Lake Exposure on Prices, Moderated by Elasticity

certain neighborhoods seeing far less of an increase than others.

While these results are simply summary statistics of the mechanical output of the estimated model, I demonstrate this relationship via simple regression, limited to cells in Polk County:

$$\Delta \text{Price}_l = \beta_0 + \beta_1 \Delta \text{Access}_l + \beta_2 \text{Elasticity}_l + \beta_3 (\Delta \text{Access}_l \times \text{Elasticity}_l) + \beta_4 \overline{\text{Elasticity}}_{-l} + \beta_5 (\Delta \text{Access}_l \times \overline{\text{Elasticity}}_{-l}) + \epsilon_l$$

where $\overline{\text{Elasticity}}_{-l}$ is the mean elasticity of all other cells within 5 miles of l . The results

are in Table 10 and show that β_5 is negative: the more elastic a cell's neighbors, even conditional on its own elasticity, the less it is affected by the demand shock.

Variable	Coef.	SE
Constant	-73.48	9.64***
ΔAccess	5,966.08	173.33***
Elasticity	-9.21	6.28
$\Delta\text{Access} \times \text{Elasticity}$	-444.06	205.39**
$\overline{\text{Elasticity}}_{-l}$	73.90	20.24***
$\Delta\text{Access}_l \times \overline{\text{Elasticity}}_{-l}$	-1,906.59	448.36***
R^2	0.991	
N	1,377	

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 10: Equilibrium price response to Gray's Lake's creation, Polk County cells only. Robust (HC1) standard errors.

This finding is more than the standard concern that amenities are incompletely capitalized into price where supply is elastic (Hilber and Mayer, 2009). That is present here and evidenced by the significant negative interaction on β_3 . The novel issue here is that the elasticity in *other* locations further prevents complete capitalization even in highly inelastic cells because the amenity shock is common across them. This is a further dampening effect unique to regional amenities, such that hedonic models may struggle to accurately value regional amenities even when accurately measured and in a low supply-elasticity context.

4.6 Counterfactuals: Lake Cleanups and Nearby Substitutes

Because water clarity, which I measure via Secchi depth, is one of the characteristics valued in the travel-cost model, the index approach can also price a quality improvement to an existing lake, not merely its removal. I use this to evaluate a federally-funded project to reduce Phosphorus loading in Lake Rathbun in southern Iowa, budgeted at \$12 million, and to demonstrate the importance of taking the full spatial distribution of amenities into account. Phosphorus reduction is not itself a characteristic in the travel-cost model, but it is one of the primary determinants of Secchi depth, so I simulate a range of potential

counterfactual Secchi depths the cleanup may result in, recomputing the quality-accessibility index for each, and compute the compensating variation for each, essentially the total value of the cleanup.


I compare Rathbun against Little Spirit Lake in northern Iowa, chosen to hold water quality and other characteristics somewhat similar while varying substitute availability sharply. Table 11 summarizes the two lakes: both are similarly turbid at baseline, but they differ starkly in isolation. Rathbun’s nearest lake is 29.4km away, and no lake lies within 15km of it at all. Little Spirit sits adjacent to Big Spirit Lake and has another four lakes within 15km, including one with a Secchi depth of 6.26 meters, one of the clearest in the state. Figure 13 maps both lakes and their county neighbors.

Statistic	Rathbun Lake	Little Spirit Lake
Secchi depth, mean (m)	0.78	0.82
Secchi depth, range (m)	0.48–1.23	0.29–1.25
Lake size (ln acres)	9.31	6.43
State park	Yes	No
Boat ramp	Yes	Yes

Table 11: Lake characteristics for the two cleanup counterfactuals. Secchi depth statistics are computed across the seven travel-cost survey years.

I simulate improvements over baseline up to 8 meters, which would make both lakes cleaner than any in the state. Figure 14 displays the results in terms of total social surplus¹¹. Roughly doubling Rathbun’s Secchi depth to 1.5 meters improves social welfare by \$30,000, while doing the same for Little Spirit only improves welfare by \$1,300. The larger the simulated improvement, the greater the level gap, at 3 meters Rathbun improves surplus by \$95,000, Little Spirit only \$5,000.

¹¹This multiplies the representative household’s compensating valuation by the number of households in the model.



(a) Rathbun Lake (Appanoose, Lucas, Monroe, and Wayne Counties) (b) Little Spirit Lake (Dickinson County)

Figure 13: The two cleanup lakes (red) and all other lakes in the depicted regions (blue).




Figure 14: Simulated total statewide welfare gain from improving Secchi depth at Rathbun Lake and Little Spirit Lake, relative to each lake's estimated baseline. Shaded bands are ± 1 SE.

This gap in cleanup value does not purely result from availability of substitutes. In addition to being larger, Rathbun is also in a State Park, meaning it is, by the model's estimates,

a more desirable destination. This demonstrates the importance of rigorously modeling the spatial and quality distribution of regional goods. Improving an already high-value lake leads to first-order gains from the many people that already visit it, while improving a low-value, rarely-visited lake with substitutes mostly increases surplus through second-order substitution behavior. The result is that the exact same cleanup causes orders of magnitude differences in welfare improvement. This exercise would not be feasible under the heuristic approaches as they do not account for quality, but even if they were, they could not capture the importance of this substitution behavior.

To briefly address the policy itself, against a \$12 million project cost, even Rathbun's most generous welfare estimate falls well short of justifying the expenditure using this valuation alone. However, it is worth pointing out that the index will not capture all potential benefits of such an improvement. It may partially capitalize things that are correlated with lake exposure, such as sediments running off into local drinking water, but would miss the value of ecosystem services generated by a cleaner lake, local and nonlocal agricultural spillovers, or effects on downstream water quality.

References

- J. K. Abbott and H. A. Klaiber. An embarrassment of riches: Confronting omitted variable bias and multi-scale capitalization in hedonic price models. *Review of Economics and Statistics*, 93(4):1331–1342, 2011.
- G. M. Ahlfeldt, S. J. Redding, D. M. Sturm, and N. Wolf. The economics of density: Evidence from the Berlin Wall. *Econometrica*, 83(6):2127–2189, 2015.
- M. Almagro and T. Domínguez-Iino. Location sorting and endogenous amenities: Evidence from amsterdam. Working Paper 32304, National Bureau of Economic Research, 2024. URL <http://www.nber.org/papers/w32304>.
- C. Balboni and J. S. Shapiro. Spatial environmental economics. In *Handbook of Regional and Urban Economics*, volume 6, pages 585–652. Elsevier, 2025.
- H. S. Banzhaf. Difference-in-differences hedonics. *Journal of Political Economy*, 129(8):2385–2414, 2021.
- N. Baum-Snow and L. Han. The microgeography of housing supply. *Journal of Political Economy*, 132(6):1897–1946, 2024.
- P. Bayer and C. Timmins. Estimating equilibrium models of sorting across locations. *The Economic Journal*, 117(518):353–374, 2007.
- P. Bayer, R. McMillan, and K. Rueben. An equilibrium model of sorting in an urban housing market, 2004.
- P. Bayer, N. Keohane, and C. Timmins. Migration and hedonic valuation: The case of air quality. *Journal of Environmental Economics and Management*, 58(1):1–14, 2009.
- S. T. Berry. Estimating discrete-choice models of product differentiation. *The RAND Journal of Economics*, pages 242–262, 1994.

- S. Borsky, E. Fesselmeyer, and L. Vogelsang. Urban heat and within-city residential sorting. *Journal of Environmental Economics and Management*, 127:103014, 2024.
- N. Cheng and X. Wan. Causal recreation demand estimation with cell phone mobility data: Evidence from the 2021 huntington beach oil spill. *Land Economics*, 102(3):339–360, 2026.
- V. Couture, C. Gaubert, J. Handbury, and E. Hurst. Income growth and the distributional effects of urban spatial sorting. *Review of Economic Studies*, 91(2):858–898, 2024.
- D. Donaldson and R. Hornbeck. Railroads and american economic growth: A “market access” approach. *Quarterly Journal of Economics*, 131(2):799–858, 2016.
- S. Gamper-Rabindran and C. Timmins. Does cleanup of hazardous waste sites raise housing values? evidence of spatially localized benefits. *Journal of Environmental Economics and Management*, 65(3):345–360, 2013.
- M. Greenstone and J. Gallagher. Does hazardous waste matter? evidence from the housing market and the superfund program. *The Quarterly Journal of Economics*, 123(3):951–1003, 2008.
- T. L. Hamilton and D. J. Phaneuf. An integrated model of regional and local residential sorting with application to air quality. *Journal of Environmental Economics and Management*, 74:71–93, 2015.
- W. G. Hansen. How accessibility shapes land use. *Journal of the American Institute of Planners*, 25(2):73–76, 1959.
- C. A. Hilber and C. Mayer. Why do households without children support local public schools? linking house price capitalization to school spending. *Journal of Urban Economics*, 65(1):74–90, 2009.
- C. A. L. Hilber. The economic implications of house price capitalization: A synthesis. *Real Estate Economics*, 45(2):301–339, 2017.

- Y. Ji, D. A. Keiser, and C. L. Kling. Temporal reliability of welfare estimates from revealed preferences. *Journal of the Association of Environmental and Resource Economists*, 7(4): 659–686, 2020.
- Y. Ji, D. A. Keiser, C. L. Kling, and D. J. Phaneuf. Revenue and distributional consequences of alternative outdoor recreation pricing mechanisms: Evidence from a micropanel data set. *Land Economics*, 98(3):478–494, 2022.
- D. A. Keiser and J. S. Shapiro. Consequences of the clean water act and the demand for water quality. *The Quarterly Journal of Economics*, 134(1):349–396, 2019.
- H. A. Klaiber and D. J. Phaneuf. Valuing open space in a residential sorting model of the twin cities. *Journal of environmental economics and management*, 60(2):57–77, 2010.
- H. A. Klaiber and V. K. Smith. Quasi experiments, hedonic models, and estimating trade-offs for local amenities. *Land Economics*, 89(3):413–431, 2013.
- H. A. Klaiber, J. K. Abbott, and V. K. Smith. Some like it (less) hot: Extracting trade-off measures for physically coupled amenities. *Journal of the Association of Environmental and Resource Economists*, 4(4):1053–1079, 2017.
- N. V. Kuminoff and J. C. Pope. Do “capitalization effects” for public goods reveal the public’s willingness to pay? *International Economic Review*, 55(4):1227–1250, 2014.
- N. V. Kuminoff, C. F. Parmeter, and J. C. Pope. Which hedonic models can we trust to recover the marginal willingness to pay for environmental amenities? *Journal of Environmental Economics and Management*, 60(3):145–160, 2010.
- Y. Kuwayama, S. Olmstead, and J. Zheng. A more comprehensive estimate of the value of water quality. *Journal of Public Economics*, 207:104600, 2022.
- L. Muehlenbachs, E. Spiller, and C. Timmins. The housing market impacts of shale gas development. *American Economic Review*, 105(12):3633–3659, 2015.

- D. J. Phaneuf, V. K. Smith, R. B. Palmquist, and J. C. Pope. Integrating property value and local recreation models to value ecosystem services in urban watersheds. *Land Economics*, 84(3):361–381, 2008.
- S. J. Redding and E. Rossi-Hansberg. Quantitative spatial economics. *Annual Review of Economics*, 9(1):21–58, 2017.
- M. M. Scott, R. L. Santos, O. Arena, C. Hayes, and A. Simon. Community Ties: Understanding What Attaches People to the Place Where They Live, 2020.
- H. Sieg, V. K. Smith, H. S. Banzhaf, and R. Walsh. Estimating the general equilibrium benefits of large changes in spatially delineated public goods. *International economic review*, 45(4):1047–1077, 2004.
- V. K. Smith. Nonmarket valuation of environmental resources: an interpretive appraisal. *Land Economics*, pages 1–26, 1993.
- E. Thomas, I. Serwicka, and P. Swinney. Why do people live where they do? *Urban Demographics*, 2015.
- C. I. Tra. A discrete choice equilibrium approach to valuing large environmental changes. *Journal of Public Economics*, 94(1–2):183–196, 2010.
- C. I. Tra. Measuring the general equilibrium benefits of air quality regulation in small urban areas. *Land Economics*, 89(2):291–307, 2013.
- Urban Land Institute. *America in 2015: A ULI Survey of Views on Housing, Transportation, and Community*. Urban Land Institute, Washington, D.C., 2015.
- P. J. Walsh, J. W. Milon, and D. O. Scrogin. The spatial extent of water quality benefits in urban housing markets. *Land Economics*, 87(4):628–644, 2011.
- R. Walsh. Endogenous open space amenities in a locational equilibrium. *Journal of urban Economics*, 61(2):319–344, 2007.

- X. Wan, Y. Ji, and W. Zhang. A report to the iowa department of natural resources—the iowa lakes valuation project 2019: Summary and findings. 2021.
- D. Wolf and H. A. Klaiber. Bloom and bust: Toxic algae’s impact on nearby property values. *Ecological Economics*, 135:209–221, 2017.
- C. Zhang, K. J. Boyle, and N. V. Kuminoff. Partial identification of amenity demand functions. *Journal of Environmental Economics and Management*, 71:180–197, 2015.
- W. Zhang, X. Wan, W. Fan, and Y. Ji. Uncovering disparities in water-based outdoor recreation using cell phone mobility data. *Environmental Research Letters*, 19(11):114057, 2024.