Fueling Access: Gasoline Prices, Remoteness, and the Travel-Cost Elasticity of National-Park Demand

Tandem Young*

June 2026

Abstract

How sensitive is demand for public-land recreation to the cost of getting there? In a monthly panel of 371 contiguous-state national-park units spanning 1993–2025, I interact regional gasoline prices with each park's remoteness from population. Because gasoline prices are procyclical, the level of the elasticity is not identified; its remoteness gradient is, from cross-region price variation under a parallel-cyclical-response assumption that a non-fuel placebo supports. A one-standard-deviation increase in remoteness makes a park's gas-price elasticity roughly 0.07 log points more negative ($p \approx 0.04$ excluding the pandemic; $p = 0.003$ before it)—a point estimate that is stable across remoteness measures, price geographies, and sub-periods, and that reappears in pre-2010 data, which predate the 2014 and 2022 price cycles and the pandemic. The adjustment is extensive (fewer trips, unchanged visit duration), is concentrated at lower-volume parks, and shows no detectable diversion to nearer parks. Applied to the fuel-price changes implied by carbon pricing, the gradient yields a spatially regressive incidence: under the EPA social cost of carbon, parks at the 90th percentile of remoteness lose roughly seven percentage points more access than parks at the 10th—a recreation-access co-cost concentrated in the rural West and rarely counted.

Keywords: travel-cost method; recreation demand; national parks; gasoline prices; carbon pricing; distributional incidence.

JEL codes: Q26, Q51, Q54, R41, L83.

*University of Arkansas, Department of Economics. Email: tyoung@walton.uark.edu. This is a working draft; comments are welcome. All estimates and exhibits are reproducible from the replication code (`code/43–62`); the underlying data are public (NPS Visitor Use Statistics, EIA, PRISM, BEA, FHWA), and a replication archive will be deposited upon acceptance. The author declares no competing interests.

1 Introduction

Access to public lands depends on the cost of reaching them, and for the geographically dispersed United States national-park system that cost is dominated by the price of driving. The price responsiveness of national-park demand nonetheless remains largely unmeasured in a credibly identified way. Research on the system's long monthly visitation record has concentrated on non-price demand shifters—social media exposure (Wichman, 2024), local air quality (Keiser et al., 2018), and wildfire smoke—while Walls (2022), in a recent review of the economics of the park system, identifies the responsiveness of visitor demand to pricing as a pressing and underdeveloped question, one that bears directly on both congestion management (Carr and Newbold, 2025) and the system's chronic funding shortfalls. This paper estimates a travel-cost elasticity of park demand from regional fuel-price variation and uses it to quantify the distributional incidence of fuel-cost and carbon-pricing shocks on recreational access.

The central empirical obstacle is that gasoline prices are endogenous to the business cycle. Fuel is expensive when the economy, and with it discretionary travel demand, is strong, which biases the naive own-price relationship upward. I address this with a triple-difference design, in the spirit of difference-in-differences with a continuous treatment (Callaway et al., 2024), that exploits a basic feature of travel cost: because fuel raises the marginal cost of a trip in proportion to the distance driven, a uniform regional increase in the gas price should depress visitation more at parks far from population than at accessible ones. Interacting the regional gas price with a park's remoteness, within park-by-calendar-month and national year-month fixed effects, differences out the common procyclical confound and isolates the travel-cost channel. A more demanding variant of the design replaces the national time effects with region-by-time effects, so that identification rests solely on differences in remoteness among parks observed in the same region and month.

The main result is a robust negative gradient of approximately −0.07 per standard deviation of remoteness: each standard-deviation increase in a park's distance from population makes its demand about 0.07 log points more sensitive to the gas price. A methodological lesson accompanies the estimate. A sample beginning in 2010, the period over which the regional price series is most easily assembled, yields a much larger gradient that proves to be an artifact of the 2020 fuel-price collapse coinciding with the pandemic surge in outdoor recreation. The estimate becomes credible only after the price series is extended back to 1993 to span several independent oil-price cycles; it then survives the exclusion of the pandemic years and reappears in pre-2010 data alone.

Two further findings sharpen the interpretation. First, decomposing the response by margin shows that it is entirely extensive: remote parks lose trips and, proportionally, visitor-hours, but exhibit no change in the duration of a typical visit. The competing "trip-consolidation" margin, under which higher travel costs induce fewer but longer trips, appears only in the pandemic-contaminated sample and does not survive on the full panel. Second, I find no evidence that the demand lost at remote parks reappears at nearby ones, although this spatial-lag test for within-system substitution is underpowered. To the extent that the displaced trips leave the park system rather than reallocating within it, the welfare loss is a reduction in recreational access rather than a reshuffling of visits. Building on these results, I show that the welfare incidence of higher fuel prices is spatially regressive, falling disproportionately on access to the remote, rural-western public lands that are least substitutable.

I am explicit about the limits of what the design identifies. The level of the elasticity is confounded by procyclicality and cannot be credibly recovered with the available instruments: a state gasoline-tax instrument has a strong first stage once state-level prices are used, but is too underpowered on the short window of available tax variation to pin the level down (Appendix A). The claims of the paper rest entirely on the gradient, which is identified from cross-region gas-price variation interacted with remoteness under a parallel-cyclical-response assumption (Section 4.2) and does not require the level. The most demanding variant, which absorbs all region-by-time variation and identifies the gradient only from remoteness differences within a region, is same-signed but underpowered; I report it, and the decomposition behind it, alongside the headline estimate.

My contribution is not to establish that fuel matters for park demand—Stevens et al. (2014) and Oh and Hammitt (2011) already show that, and I place them among the related literature in Section 2—but to provide a credibly identified travel-cost *gradient*, estimated from regional gas-price variation interacted with remoteness across the full national-park panel, together with randomization-inference placebos and an application to the distributional incidence of carbon pricing. The monthly panel extends the month-of-year visitation averages of Fisichelli et al. (2015) into a longitudinal record.

The remainder of the paper is organized as follows. Section 2 reviews the related literature. Section 3 describes the visitation, price, and remoteness data. Section 4 sets out the travel-cost framework, the estimating equation, and the approach to inference. Section 5 presents the gradient, its margins, the substitution null, the robustness analysis, and where in the system the gradient operates. Section 6 develops the distributional incidence of fuel prices and carbon pricing, and Section 7 concludes.

2 Related literature

This paper sits at the intersection of three literatures: the demand for outdoor recreation, the price responsiveness of driving, and the distributional incidence of energy and carbon pricing.

The first is the travel-cost tradition in recreation demand, which infers the value of a site from the trips taken to it and treats the cost of travel as the implicit price of a visit (Phaneuf and Smith, 2005; Parsons, 2017; Lupi et al., 2020). Because visitation is a count, the workhorse estimators are count-data and random-utility models (Englin and Shonkwiler, 1995; Parsons et al., 2021), and a recurring practical question is how to build the travel-cost variable from distance, fuel, and the opportunity cost of time (McConnell, 1992; Larson, 1993). I borrow the field's core insight—that the gas price enters the cost of a visit in proportion to distance—but invert its usual use: rather than value a single site or recover a full demand system, I use regional fuel-price variation to trace how the gas-price elasticity of visitation varies with distance, and read that gradient as the object of interest. The closest park-demand antecedents are Stevens et al. (2014), who find national-park attendance more sensitive to fuel and travel costs than to entrance fees over an annual panel of thirty parks, and Oh and Hammitt (2011), who document gasoline-price-driven substitution away from a single state park that emerges only at high prices and is tempered by destination loyalty.

The second literature estimates how driving responds to the price of fuel. Aggregate studies find the short-run elasticity of gasoline demand small, and falling over time (Hughes et al., 2008), while higher-frequency designs recover larger responses (Levin et al., 2017); the medium-run elasticity of driving with respect to the gas price—the margin closest to a recreational trip—is around -0.2 (Gillingham, 2014), in line with the modest rebound estimates in the fuel-efficiency literature (Small and Van Dender, 2007; Gillingham et al., 2016). Two themes from this work shape the design here. One is that naive price regressions are badly confounded—by the business cycle, by anticipation, and by the endogeneity of prices to quantities—so credible estimates exploit sharp, plausibly exogenous variation; correcting for anticipation around tax changes, for instance, pulls implausibly large tax-instrument estimates back toward roughly -0.37 (Coglianese et al., 2017). The other, more specific, is that the driving response is itself heterogeneous in trip distance: Gillingham and Munk-Nielsen (2019) find a U-shaped response in commute length, with the longest-distance drivers among the most price-sensitive. A recreational analogue—that fuel-price sensitivity rises with the distance to a destination—is exactly the gradient this paper estimates.

A smaller body of work connects fuel prices directly to tourism and recreation, generally finding that dearer fuel depresses leisure travel, the more so for distant and car-dependent destinations (Becken and Lennox, 2012; Cho et al., 2014). Cho et al. (2014) in particular report that the travel-cost response varies widely across sites, foreshadowing the per-park heterogeneity documented below.

The third literature studies the distributional incidence of energy and carbon pricing. A carbon price raises the cost of gasoline, and a long line of work asks who bears that burden across the income distribution (Hassett et al., 2009; Williams et al., 2015; Davis and Knittel, 2019; Goulder et al., 2019; Cronin et al., 2019), with incidence depending heavily on whether households are ranked by annual or lifetime resources and on how revenue is returned. This work measures incidence through household energy *budgets*; it does not price the loss of *access* to the public lands that higher driving costs put further out of reach. Yet that access is itself unequally distributed—distance and travel cost are leading barriers to national-park visitation, and fall disproportionately on lower-income and minority populations who tend to live farther from the parks (Weber and Sultana, 2013; Scott and Lee, 2018; Xiao et al., 2022). The welfare analysis below quantifies one such omitted co-cost: a spatially regressive loss of recreation access that compounds the inequities the recreation-equity literature documents.

3 Setting and data

3.1 The national-park system and access to it

The National Park System is large, heavily used, and geographically dispersed. Its more than four hundred units—ranging from iconic western parks to small historic sites, monuments, seashores, and parkways—drew over 330 million recreation visits in 2024, a record (National Park Service, 2025). Visitation is extremely concentrated: a handful of flagship parks account for a large share of the total while hundreds of smaller units each receive a tiny fraction, a skew that turns out to shape how the travel-cost margin operates (Section 5.5). Most units lie far from major population centers—much of the system sits in the rural interior West—and the overwhelming majority of visitors arrive by private automobile, so for the typical park the cost of a visit is dominated by the cost of getting there.

That access is only partly within the agency's control, and the price levers it does hold are modest. Under the Federal Lands Recreation Enhancement Act of 2004 (U.S. Congress, 2004), only about a hundred of the system's units charge an entrance fee at all—commonly \$20 to \$35 per vehicle—while the rest are free, and an \$80 annual "America the Beautiful" pass covers entrance at every one of them. More recently the Service has layered timed-entry reservation systems onto a small number of congested parks (Carr and Newbold, 2025). These instruments, together with the system's persistent funding gap and a deferred-maintenance backlog in the tens of billions of dollars, are the focus of most policy discussion of park access and pricing (Walls, 2022); the responsiveness of demand to the cost of a visit, by contrast, is comparatively underdeveloped.

The fuel-cost margin sits outside all of it. For a household driving to a remote park, the round-trip fuel bill commonly exceeds the entrance fee by a wide margin, and unlike the fee it rises with the price of gasoline and with the distance to be covered—neither of which the agency sets or can waive. A policy that raises the price of driving, such as a carbon price, therefore operates on a cost of access that is larger than the agency's own price instrument at most units and that falls hardest on the remote parks this paper studies. The empirical questions are how much park demand responds to that margin, and how the response varies with remoteness.

3.2 Visitation

The visitation data are the National Park Service Visitor Use Statistics (VUStats), a balanced monthly panel of 406 administrative units observed from 1979 through 2025. The Service has tabulated visitation by month and unit for decades under a consistent set of counting definitions, which makes VUStats one of the longest continuous administrative records of outdoor-recreation demand in the United States. The primary outcome is recreation visits—the number of persons entering a unit for recreational purposes in a month, the Service's headline measure and the one used in nearly all prior economic work on park demand. The analysis of margins additionally uses recreation visitor-hours (total person-hours of recreational use), the implied hours per visit (visitor-hours divided by visits), overnight stays, and non-recreation visits (through-traffic and administrative entries), each of which the Service reports for every unit-month. The units span the full breadth of the system—large flagship parks, national monuments, historic sites, recreation areas, seashores, and parkways. Monthly visitation ranges from a few thousand at small units to several million at the busiest parks, a dispersion the analysis exploits when it distinguishes a per-park response from a visit-weighted one (Section 5.5).

3.3 Gasoline prices

Fuel prices are the Energy Information Administration's (EIA) monthly retail price of regular gasoline, taxes included, by region. The EIA reports prices for the five Petroleum Administration for Defense Districts (PADDs)—broad multi-state regions originally drawn for wartime fuel logistics and still the standard geography for US petroleum statistics—and, separately, for nine large states. Fourteen of these series are matched to parks through the EIA area crosswalk, so each park is assigned the monthly price of the smallest EIA region that contains it. The nine state series begin later than the PADD series—in 2000 for California, Colorado, Minnesota, New York, and Texas, and in 2003 for Florida, Massachusetts, Ohio, and Washington—so the 123 parks matched to a state series enter the panel only from those dates: the 1993–1999 panel contains roughly 200 parks priced at the five PADDs, and the full fourteen-region price geography is in place from 2003 (parks established or first reporting later enter the 371-park panel as they appear). Section 5.4 shows that this structured entry does not drive the results: pricing every park at its PADD from 1993 onward, which removes the entry, yields a slightly larger gradient (-0.075, $p = 0.009$, pandemic years excluded), as does its 1993–2009 sub-period alone (-0.079, $p = 0.010$; Table 6). Regional rather than national prices are essential to the design: the national price is absorbed by the time effects, and identification comes from the cross-region deviations that remain. Retail prices move closely together across regions because they share a common crude-oil component, but persistent wedges arise from differences in refining capacity, pipeline access, state taxes, and boutique-fuel requirements; it is these wedges, interacted with remoteness, that identify the gradient. One measurement caveat: the price is measured in the park's region, while visitors buy much of their fuel at the origin and en route, so for parks that draw visitors from other regions the treatment is measured with error, which works toward attenuation. I extend the series back to 1993 so that the panel spans the late-1990s lull, the early-2000s run-up, the 2008 spike, the 2014–2016 collapse, and the 2022 spike. These episodes supply several distinct oil-price cycles whose demand- and supply-driven components differ in origin (Kilian, 2009)—which matters because, as Section 5.1 shows, a gradient estimated from a single episode can be an artifact of that episode's particular coincidences. Figure 1 plots the fourteen series alongside the cross-region residual variation that survives the time fixed effects.

3.4 Measuring remoteness

A park's remoteness is the great-circle distance from its boundary-polygon centroid to the nearest metropolitan county, defined as a county of at least 250,000 residents, computed from the 2023 Census county cartographic file and the most recent (2024) Bureau of Economic Analysis county population and held fixed over the panel—the moderator is deliberately time-invariant, since the park-by-calendar-month effects absorb its level. Because the relevant cost is driving distance, this great-circle measure is a proxy; I therefore also compute the actual road distance and driving time from each park to its nearest metropolis by routing through the road network. Road circuity is indeed higher in the rural West—its correlation with longitude is -0.22—but because that inflation is roughly proportional, the road and great-circle measures rank parks almost identically (correlation 0.99), so the great-circle proxy is not biased by circuity in the cross-park ordering that identifies the gradient. The gradient itself is the same sign and magnitude under the road measures, though estimated slightly less precisely (Section 5.4). Counties that crossed the population threshold during the sample are metros for the whole panel under this definition; recomputing remoteness

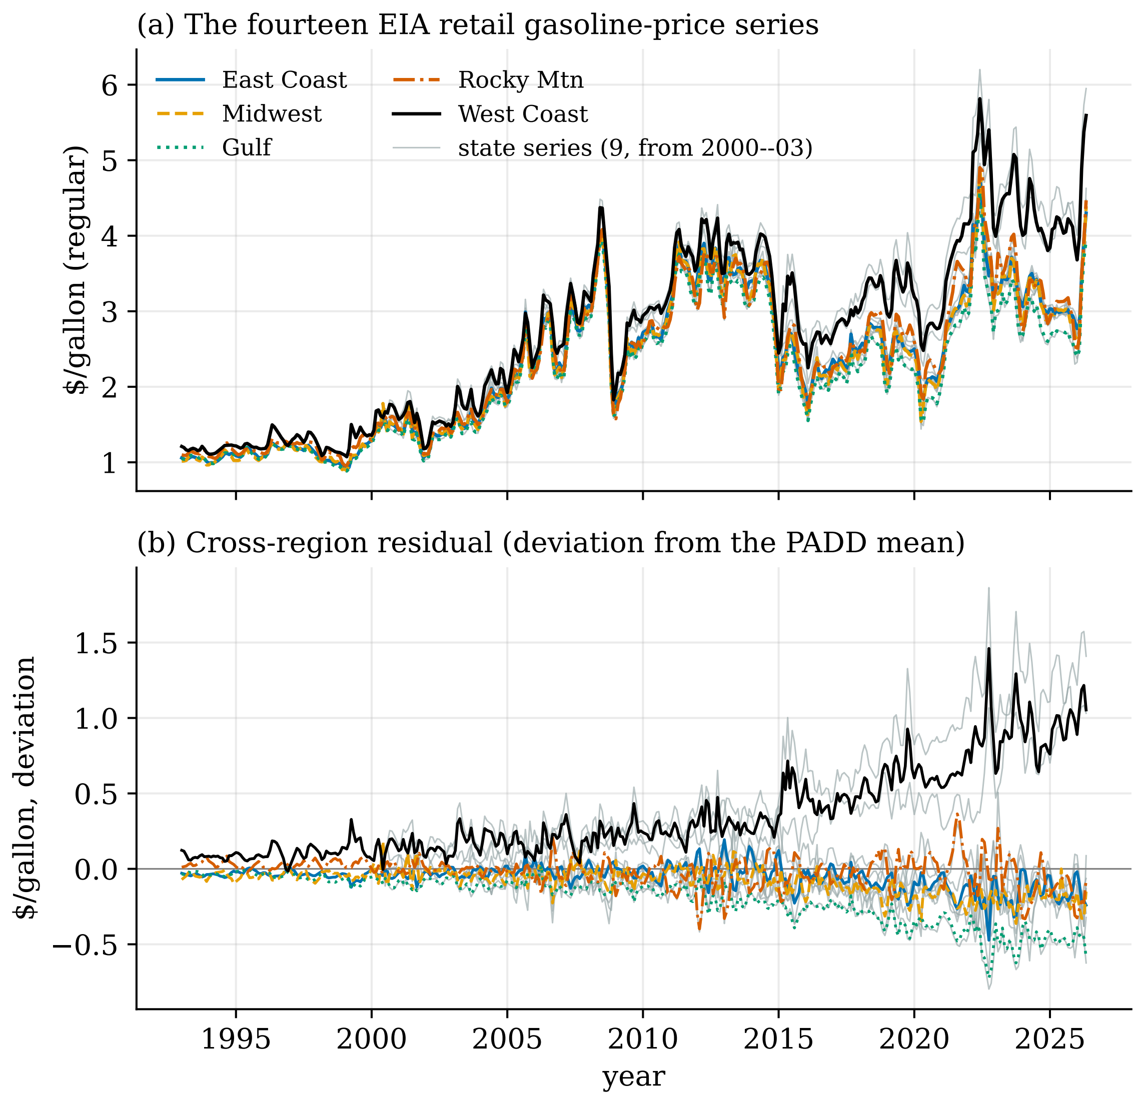

Figure 1: The fourteen EIA gasoline-price series, 1993–2025, and the cross-region residual variation that survives the time fixed effects. The five PADD aggregates are highlighted; the nine state series (grey) enter in 2000–2003.

from 1993 population instead yields a similar but less precise gradient (-0.056, $p = 0.09$; Section 5.4). As a robustness alternative I use a gravity-weighted index of population access, which sums regional population discounted by distance and so rewards proximity to many people rather than to the single nearest metropolitan area; the two measures rank parks similarly and yield similar gradients (Section 5.4). Remoteness is standardized to a z-score over the unit-months of the estimation sample, with a 90th-to-10th-percentile spread across parks of 2.7 standard deviations, and is mapped in Figure 2. The most accessible parks lie within commuting distance of major metropolitan areas, while the most remote sit hundreds of kilometers from any large population center: the mean distance to a metropolitan county rises from about 12 kilometers in the accessible tercile to 257 in the remote one (Table 1). Because fuel cost enters trip cost through distance, this cross-park variation in remoteness is the source of identification. The analysis is restricted to parks in the contiguous states, excluding Alaska, Hawaii, and the island territories, for which driving from population centers is not the relevant margin; a handful of contiguous-state units reachable only by boat or seaplane remain in the baseline, and excluding them as well leaves the point estimate essentially unchanged (-0.064, park-clustered $p = 0.054$; Section 5.4). All specifications control for temperature, entered quadratically, and precipitation—PRISM 4-km gridded monthly means aggregated to each park (Daly et al., 2008)—because weather is a first-order driver of monthly visitation and is correlated with both region and season; the quadratic in temperature accommodates the inverted-U relationship between warmth and outdoor recreation, with visitation falling off in both cold and hot months.

3.5 Estimation sample

The full 1993–2025 panel of 371 contiguous-state parks contains 116,634 park-months with price, weather, and visitation data, of which 108,494 remain after the pandemic years are excluded. Because the outcome enters in logs, the park-months with zero recorded recreation visits—1.7% of the full panel, 1.3% of the pandemic-excluded one—are dropped, leaving

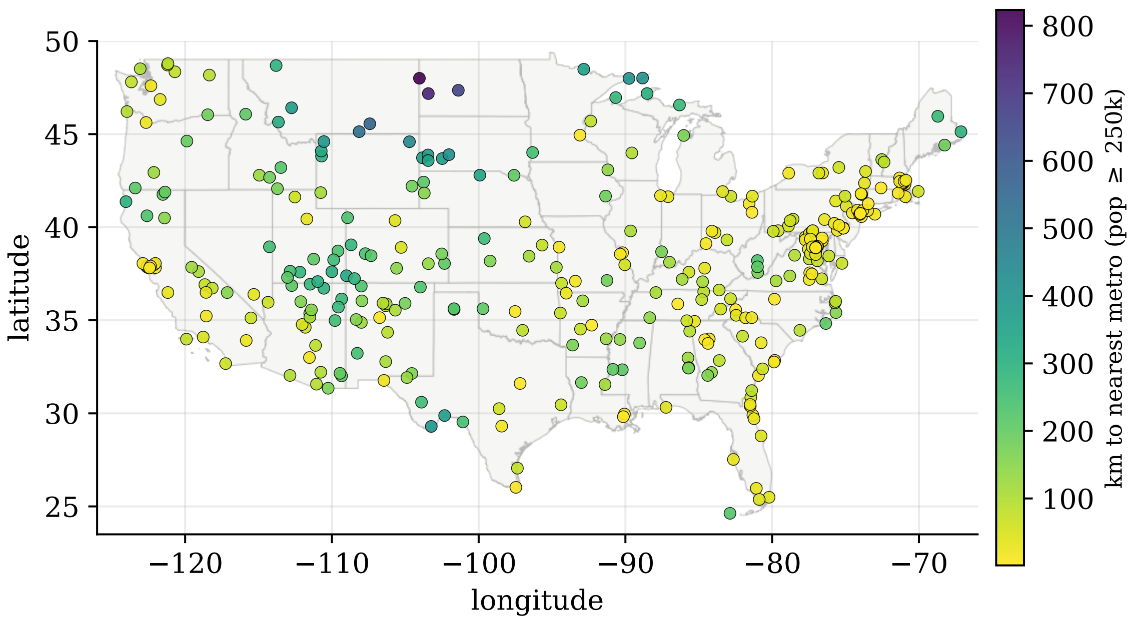

Figure 2: Park remoteness: the cross-sectional identifying variation.

114,611 and 107,131 park-months; the baseline regressions retain 114,579 and 107,098 of these after dropping the small number of singleton park-by-calendar-month cells absorbed by the fixed effects. The zeros are more common at accessible parks (2.5% of the accessible tercile versus 0.6% at remote parks), but Section 5.5 shows that this differential dropping does not drive the gradient: the zeros are concentrated in a handful of chronically closed micro-units, and zero-inclusive estimation within the stratum of parks where the gradient operates leaves it intact. Table 1 reports summary statistics by remoteness tercile for the pandemic-excluded sample. Remote parks are smaller in visitation—a mean of about 51,000 monthly recreation visits, against 100,000 at accessible parks—and they sit in colder, drier locations, so the weather controls and the park-by-calendar-month effects do real work. The gas price itself is nearly flat across terciles (a tercile mean of $2.43 to $2.52 per gallon), which confirms that remoteness is not a proxy for facing systematically cheaper or dearer fuel; the identifying variation is the *interaction* of the common regional price swings with how far each park sits from its visitors.

Table 1: Summary statistics for the travel-cost estimation sample (contiguous-state parks, 1993–2025, pandemic years excluded). The last three columns report variable means within remoteness terciles (accessible, middle, remote).

	Mean	SD	Accessible	Mid	Remote
Recreation visits	70,856	170,355	99,568	62,134	50,772
Gas price ($/gal)	2.48	0.96	2.52	2.49	2.43
Dist. to metro (km)	115.1	126.6	12.2	76.0	257.4
Remoteness (z)	0.00	1.00	−1.21	0.18	1.04
Temperature (°C)	12.8	9.4	14.3	13.3	10.8
Precip (mm)	78.6	68.1	94.4	86.8	54.6
N (unit-months)	107,131		35,788	35,672	35,671
Parks	371		132	123	116

4 Empirical strategy

4.1 A travel-cost framework

The travel-cost method values a recreation site from the trips people take to it, treating the cost of getting there as the implicit price of a visit (Phaneuf and Smith, 2005). I adapt that logic to a comparative-statics prediction: how a change in the price of fuel shifts demand, and how that shift varies with a site's distance from population.

Consider a representative visitor deciding how many trips n_i to take in a period to park i, located a distance d_i from her origin. The cost of a round trip is approximately

$$c_i \;=\; 2d_i\,\frac{p^{\text{gas}}}{\mu} \;+\; \omega\,\tau_i \;+\; f_i, \tag{1}$$

where p^{gas} is the price of gasoline, μ vehicle fuel economy, ω the value of travel time, τ_i round-trip travel time, and f_i the entrance fee and other on-site outlays. Trip demand slopes down in this full price, $n_i = n(c_i)$ with $n'(\cdot) < 0$.

Because only the fuel term in (1) depends on the gas price, $\partial c_i/\partial p^{\text{gas}} = 2d_i/\mu$, and the gas-price elasticity of visitation is

$$\varepsilon_i \;\equiv\; \frac{\partial \ln n_i}{\partial \ln p^{\text{gas}}} \;=\; \frac{n'(c_i)}{n_i}\,\frac{2d_i\,p^{\text{gas}}}{\mu}. \tag{2}$$

Two features of (2) organize the analysis. First, $\varepsilon_i < 0$: dearer fuel means fewer trips. Second, and central here, the elasticity grows more negative with distance, $\partial\varepsilon_i/\partial d_i < 0$, provided trip demand is not too convex in the full price. Log-concavity of $n(\cdot)$ in the full price is sufficient, as with linear demand; constant-elasticity demand, though log-convex, satisfies the weaker requirement whenever non-fuel trip costs are positive.[1] Intuitively, fuel is a larger share of the trip cost at a remote park, so a given proportional increase in the gas price raises that park's full trip price by more. Equivalently, under the same condition the cross-partial $\partial^2 \ln n_i/(\partial \ln p^{\text{gas}}\,\partial d_i) < 0$. This is the gradient prediction, and its sample analogue is the interaction coefficient γ in the estimating equation below.

The level of the elasticity and its slope in distance are separately informative, and only the slope is identified here. The level mixes the travel-cost response with anything else that moves visitation when fuel is expensive—most importantly the business cycle, since fuel is dear precisely when incomes and discretionary travel are high. The slope nets that confound out: a procyclical demand shock common to all parks shifts the level of ε_i but not its slope in distance, which arises only because fuel cost scales with distance. A *region-level* demand shock can still masquerade as a gradient, but only through cross-region differences in average remoteness—a channel the region-by-time variant of the design closes by construction. Section 4.2 implements this prediction as a gas-price–remoteness interaction and reports both variants.

The framework also predicts *which* margin adjusts. A visitor chooses both how many

[1] Differentiating (2), holding travel time fixed, gives $\partial\varepsilon_i/\partial d_i = (2p^{\text{gas}}/\mu)\big[\partial \ln n_i/\partial c_i + c_i^{\text{fuel}}\,\partial^2 \ln n_i/\partial c_i^2\big]$, where $c_i^{\text{fuel}} = 2d_i p^{\text{gas}}/\mu$ is the fuel cost. The first term is negative, so log-concavity ($\partial^2 \ln n_i/\partial c_i^2 \le 0$) is sufficient but not necessary. With constant-elasticity demand $n_i = Ac_i^{-\eta}$, the bracket equals $-\eta(\omega\tau_i + f_i)/c_i^2$, which is negative whenever non-fuel trip costs are positive. Sufficiently convex demand—a loyal core of visitors insensitive to cost—can reverse the sign, which is precisely the remoteness-inelastic behavior of the highest-volume flagship parks documented in Section 5.5.

trips to take and how long to stay on each, and the two are distinct choices (McConnell, 1992; Larson, 1993). A fuel shock raises the fixed cost of *taking* a trip—the round-trip drive—without changing the marginal value of time once a visitor is on site, so it should reduce the number of trips while leaving the duration of a typical visit unchanged. The response should thus be extensive (fewer trips, and proportionally fewer visitor-hours) rather than intensive (longer stays). The competing "trip-consolidation" hypothesis, under which higher travel costs induce fewer but *longer* trips, predicts instead that visit duration rises. Section 5.2 tests these against each other.

Finally, the same elasticity delivers the welfare object. Because visitor origins are unobserved, I cannot measure each park's d_i; I proxy a park's representative trip distance d_i by its standardized remoteness Remote_i (the z-score defined in Section 3) and identify γ from how the gas-price response varies across parks of differing remoteness. A fuel-price change $\Delta \ln p^{\text{gas}}$—including one induced by a carbon price—then changes visitation at park i by approximately $\varepsilon_i \, \Delta \ln p^{\text{gas}}$, and the *differential* access loss between a remote and an accessible park (evaluated at the 90th versus 10th percentiles of remoteness, a 2.7-standard-deviation spread),

$$\left(\varepsilon_{\text{remote}} - \varepsilon_{\text{accessible}}\right) \Delta \ln p^{\text{gas}} \;=\; \gamma \left(\text{Remote}_{\text{remote}} - \text{Remote}_{\text{accessible}}\right) \Delta \ln p^{\text{gas}}, \tag{3}$$

depends only on the identified gradient γ, not on the confounded level. Section 6 evaluates (3) for the fuel-price changes implied by carbon pricing.

4.2 Specification and identification

The estimating equation, on log recreation visits, is

$$\ln \text{TRV}_{it} \;=\; \beta \ln \text{gas}_{a(i),t} \;+\; \gamma\big(\ln \text{gas}_{a(i),t} \times \text{Remote}_i\big) \;+\; X_{it}'\theta \;+\; \delta_{i,m(t)} \;+\; \lambda_t \;+\; u_{it}. \tag{4}$$

Here $a(i)$ denotes the EIA price region containing park i (Section 3) and $m(t)$ the calendar month of year-month t. The park-by-calendar-month effects $\delta_{i,m(t)}$ absorb each park's seasonal profile and its time-invariant level of remoteness; the national year-month effects λ_t absorb the common path of oil prices together with aggregate macroeconomic and income shocks; X_{it} collects the weather controls; and standard errors are clustered by park.

The coefficient of interest is γ. With the time effects in place, the level coefficient β is identified only from cross-region deviations in the gas price and remains procyclically confounded; it is positive and is not interpreted. The interaction γ, by contrast, is identified from the differential response of demand to the gas price across parks of differing remoteness. In the most demanding specification I replace the national time effects with region-by-year-month effects, which absorb all regional gas variation and any region-level demand shock, so that γ is identified purely from differences in remoteness among parks within the same region and month. I report the implied elasticity across the range of remoteness in Figure 5 and the gradient across the progression of fixed-effects specifications in Table 2.

Identification of γ requires that, conditional on the fixed effects and weather, a park's remoteness be uncorrelated with any other determinant of how strongly its demand responds to the regional business cycle. In the baseline, the fixed effects remove *additive* shocks at the park-season and national-time level, but two channels survive them. First, an additive region-by-time demand shock correlated with the region's gas-price deviation would load onto γ, because average remoteness differs across regions; the region-by-year-month variant removes this channel by construction, and Section 5.1 decomposes the gradient into its between- and within-region components. Second, a remoteness-correlated *slope* on the cycle—were remote parks systematically different in their non-fuel cyclical sensitivity, through visitor income, seasonality, or amenity mix—would generate a gas-by-remoteness gradient under either set of time effects. I assess this parallel-cyclical-response assumption with three checks reported in Section 5.4: in a horse race with a procyclical regressor that is not fuel—gateway-county income interacted with remoteness—the gas interaction is undisturbed while the income interaction collapses to zero; the gradient is essentially unchanged when weather is allowed to affect remote and accessible parks differently; and it is present for recreation visits but weaker and insignificant for non-recreation through-traffic.

Table 2: Main results: the gas-price by remoteness gradient across the fixed-effects progression. Column 1 is the full 1993–2025 panel with national year-month effects; column 2 replaces them with region-by-year-month effects; column 3 excludes the pandemic; column 4 restricts to the pre-COVID period (≤2019); and column 5 to 2010 onward. Standard errors clustered by park in parentheses.

	Log recreation visits				
	1993–2025	1993–2025	excl. 2020–21	≤2019	2010+ †
	(1)	(2)	(3)	(4)	(5)
ln(Gas price)	0.187		0.109	0.006	0.281
	(0.230)		(0.213)	(0.184)	(0.201)
ln(Gas price) × Remoteness (z)	−0.065*	−0.037	−0.068**	−0.089***	−0.101***
	(0.033)	(0.040)	(0.033)	(0.030)	(0.031)
Region × year-month	-	x	-	-	-
Year-month	x	-	x	x	x
Park × calendar-month	x	x	x	x	x
Observations	114,579	114,579	107,098	90,836	62,612
R^2	0.937	0.941	0.946	0.956	0.941

Significance levels: * p < 0.1, ** p < 0.05, *** p < 0.01. Cells report the coefficient with the standard error in parentheses.
Column 5 (2010 onward, †) is reported for comparison only: the paper shows this estimate is inflated by the 2020 pandemic fuel-price collapse and it is not a preferred specification.

4.3 Inference and placebo

Gasoline prices are highly persistent, with a month-to-month autocorrelation near 0.95, so a conventional lead–lag event study is uninformative: the leads and lags of a near–random-walk treatment are collinear, and the implied "pre-trends" are mechanical rather than diagnostic. The appropriate placebo for a persistent, continuous treatment of this kind is randomization inference (Young, 2019), which builds a null distribution by repeatedly reassigning the treatment moderator and re-estimating. I implement two permutation schemes, each with 500 draws, holding the panel and the gas series—and thus the price's persistence and region structure—fixed. The first permutes the cross-sectional remoteness assignment across all 371 parks; it treats remoteness as exchangeable across the whole system and so tests the full, blended gradient. The observed gradient lies well in the tail of this null ($p = 0.018$; Figure 3). The second permutes remoteness only among parks within the same gas region, preserving each region's remoteness composition; because every draw retains the between-region component, this null is centered away from zero, and it tests only the within-region claim. Against this stricter null the observed gradient is not in the tail ($p = 0.12$), so the within-region component is not separately established—the randomization analogue of the underpowered region-by-time specification in column 2 of Table 2, and a result I return to in Section 5.1.

Analytic inference must also confront the geography of the treatment. The gas price varies across only fourteen EIA regions, so although the interaction is defined at the park level, its identifying variation is effectively assigned at the region-by-month level. Clustering by park, as in the tables, treats each park as an independent draw and could understate uncertainty if shocks are correlated within a region. I therefore also cluster the headline interaction at the gas-region level. With only fourteen clusters the conventional cluster-robust variance is downward biased, so I report the cluster-jackknife (CRV3) variant, which is better behaved with few clusters; it yields $p = 0.015$ (Table 6). A restricted wild-cluster bootstrap over the fourteen regions—the resampling complement to the jackknife, with 999 Rademacher draws imposing the null (MacKinnon et al., 2023)—gives $p = 0.030$, the most conservative of the region-level tests but still rejecting at conventional levels.[2]

[2]The conventional CRV1 region-clustered standard error is smaller still (0.020, $p = 0.006$) but is expected to be downward biased with fourteen unbalanced clusters. With so few clusters I read the jackknife and bootstrap p-values as evidence that region-level clustering does not overturn the gradient, not as a gain in precision over the park-clustered baseline.

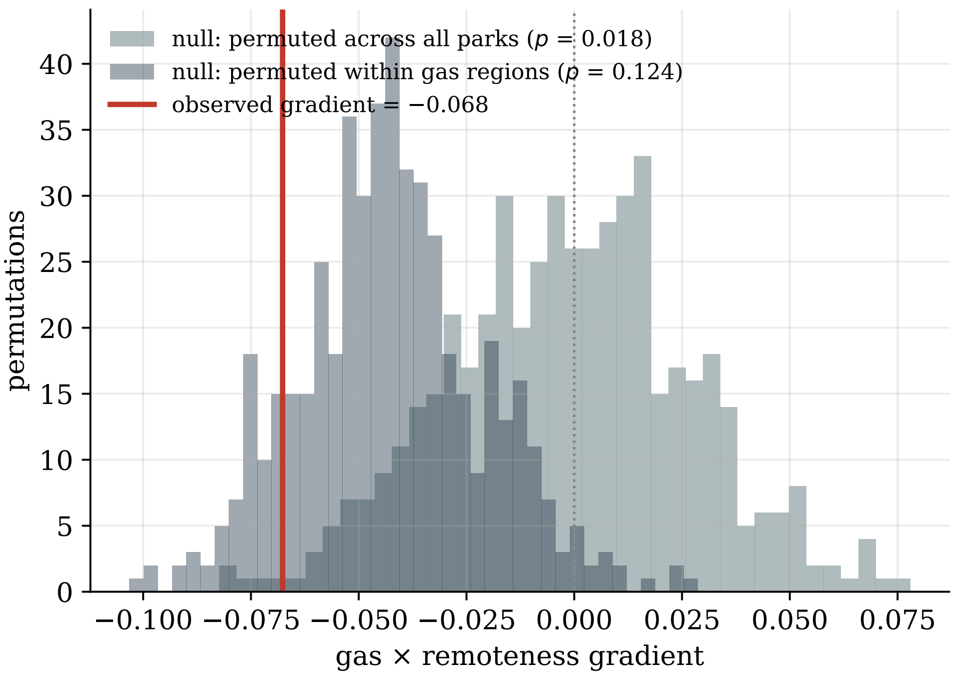

Figure 3: Randomization-inference placebo on the pandemic-excluded 1993–2025 panel: the observed gradient (−0.068) against two permutation nulls (500 draws each). Permuting remoteness across all parks tests the blended gradient (p = 0.018); permuting only within gas regions preserves each region's remoteness composition—so its null centers on the retained between-region component—and tests the within-region claim, against which the observed gradient is not in the tail ($p = 0.12$).

5 Results

5.1 The travel-cost gradient

Figure 4 shows the identifying variation before the formal estimates. Following the logic of the estimator itself, both log recreation visits and the gas-price-by-remoteness interaction are residualized on the park-by-calendar-month and year-month fixed effects, the weather controls, and the gas-price main effect; by the Frisch–Waugh theorem, the slope of the resulting scatter is numerically the gradient γ of the estimating equation. The relationship is visibly negative and close to linear across the distribution of the residualized interaction: when the gas price is high relative to its region and season, visitation falls at remote parks relative to accessible ones.

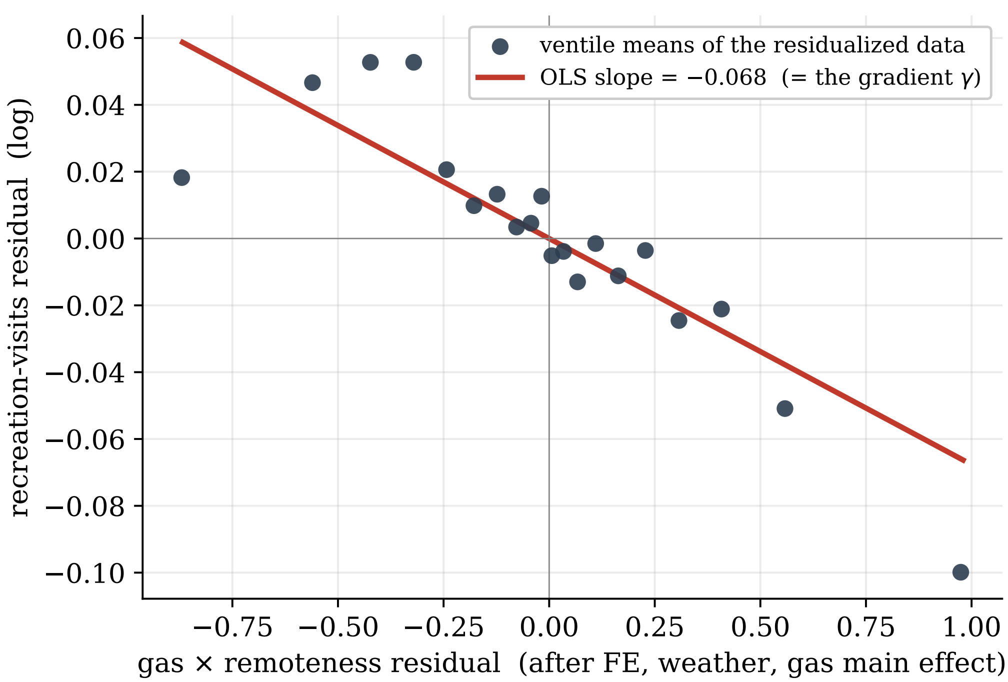

Figure 4: The estimand as a binscatter. Ventile means of log recreation visits against the gas-price-by-remoteness interaction, both residualized on the park-by-calendar-month and national year-month fixed effects, the weather controls, and the gas-price main effect (pandemic-excluded 1993–2025 sample). By the Frisch–Waugh theorem the OLS slope through the underlying observations (-0.068) is exactly the gradient of Table 2, column 3.

Table 2 and Figure 5 report the central estimate. In the baseline specification on the full 1993–2025 panel, a one-standard-deviation increase in remoteness makes the gas-price elasticity of demand 0.065 log points more negative ($p = 0.050$). The estimate is essentially unchanged when the pandemic years are excluded (-0.068, $p = 0.039$) and is larger and more precise in the pre-COVID sample (-0.089, $p = 0.003$). The most demanding specification, which absorbs all regional gas variation through region-by-year-month effects and identifies the gradient only from differences in remoteness within a region and month, yields a same-signed but statistically insignificant estimate of -0.037 ($p = 0.36$). This attenuation partly reflects a loss of identifying variation: because remote parks tend to cluster within regions, conditioning on region-by-time effects removes roughly two-fifths of the cross-park variation in remoteness (about 57% survives within region), and the standard error rises modestly (from 0.033 to 0.040). The point estimate also falls by about half (from -0.065 to -0.037); but the two specifications are not statistically distinguishable—their confidence intervals overlap substantially—so the cleanest specification is best read as same-signed but underpowered rather than as a contradiction. A decomposition states plainly where the identifying variation lives: splitting the interaction into a between-region component (the gas price times each region's mean remoteness) and the within-region remainder yields -0.108 ($p = 0.02$) and -0.039 ($p = 0.33$), respectively, on the full pandemic-excluded sample, with the headline gradient their variance-weighted blend. The cross-region alignment of remoteness with regional gas behavior thus carries most of the statistical significance—and since the between component varies only at the region level, even its park-clustered p-value is generous—while the within-region component is same-signed but underpowered, consistent with column 2 and with the within-region randomization test of Section 4.3. I make no claim to identify the gradient within region; the object I defend is the between-region gradient, under the assumption that, conditional on the fixed effects, a region's average remoteness is uncorrelated with how strongly its non-fuel demand co-moves with the regional gas price.

Because the gradient's significance rests on these cross-region comparisons, the central concern is that some non-fuel determinant of demand that varies with regional remoteness also moves with the regional gas price. I probe it by subjecting a non-fuel procyclical regressor—gateway-county income—to the same between/within decomposition (Table 3). The result is informative precisely because it is not a clean null: on its own, income's between-region interaction is negative and only about a fifth smaller than the gas term, though imprecise (-0.079, $p = 0.26$, on the subsample with income observed, where the gas between-gradient is -0.098). Were the gas gradient merely generic regional cyclicality, fuel and income would be hard to separate. They separate cleanly: in a horse race that includes both between-region interactions, the gas gradient *strengthens* to -0.128 ($p = 0.001$) while income's turns positive and insignificant ($+0.050$, $p = 0.53$). A regional confound common to fuel and income would leave income's interaction intact and attenuate the gas one; the data show the reverse, which points to a fuel-specific channel. This is the sharpest test the design admits without visitor origins, and it supports—without proving—the identifying assumption. Figure 7 collects the gradient across every sample window, robustness perturbation, and identification cut in a single view.

Figure 5 traces the implied elasticity across the range of remoteness. It is procyclically positive at the most accessible parks and declines monotonically toward zero at the most remote; the level at any point reflects the procyclical confound, while the slope is the identified object. The tercile-specific elasticities plotted alongside the fitted line are individually imprecise and not monotone, each being estimated from the thin residual gas variation within a tercile, but they are consistent with the fitted gradient.

The estimate is not an artifact of any single episode. A sample beginning in 2010 yields a gradient of -0.10 (Table 2, column 5), but this estimate collapses to roughly zero—a small, statistically insignificant positive estimate—once the pandemic years are removed, because the 2020 gas-price crash coincided with the pandemic surge in outdoor recreation. Extending the price series to 1993 and excluding the pandemic restores a robust negative gradient, and the estimate reappears in the 1993–2009 sub-period alone (-0.069, $p = 0.036$), a window that contains neither the 2014 price cycle nor the pandemic. A pandemic-driven artifact

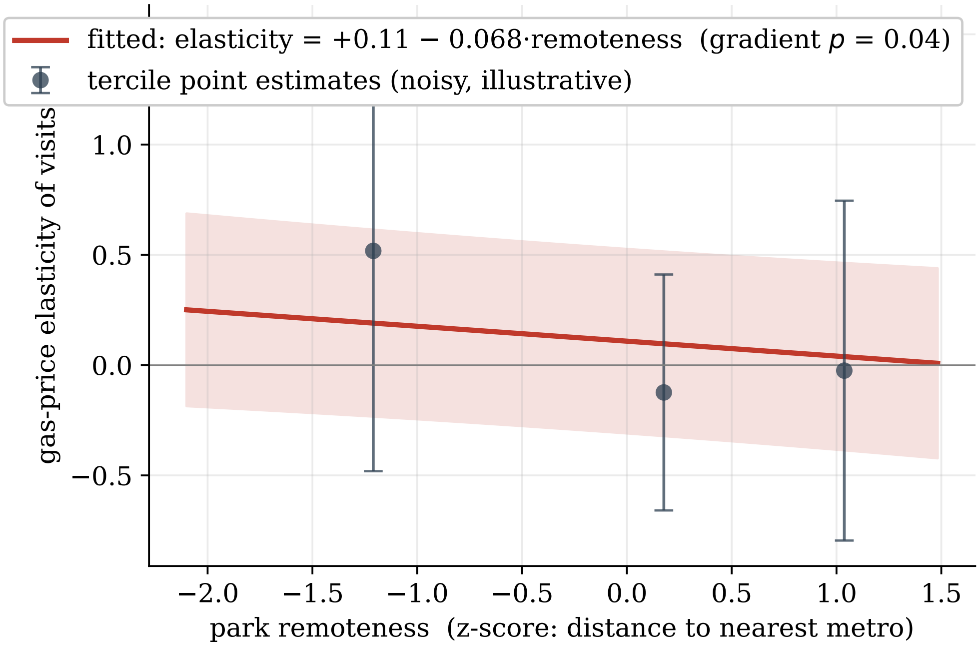

Figure 5: The travel-cost gradient: predicted gas-price elasticity by remoteness, estimated on the pandemic-excluded 1993–2025 sample (gradient −0.068). The fitted line is the identified object and its level reflects the procyclical confound; the tercile points are illustrative and imprecise. The shaded band and tercile bars are 95% confidence intervals (standard errors clustered by park).

Table 3: Where the identifying variation lives, and a non-fuel placebo. Columns split the gas-price-by-remoteness interaction (and an income-by-remoteness placebo) into a between-region component (the regressor × each region's mean remoteness) and a within-region component. Gateway-county income is a non-fuel procyclical regressor: if the gradient were generic regional cyclicality rather than a fuel-cost effect, income would reproduce the gas pattern. Pandemic-excluded 1993–2025 panel with non-missing income; park-clustered standard errors.

	Between-region	Within-region
Gas price × remoteness	−0.098	−0.039
	(0.046)	(0.039)
Income × remoteness (placebo)	−0.079	−0.048
	(0.069)	(0.057)

cannot reproduce itself in data that predate the pandemic by a decade. Figure 6 collects these estimates by sample window: on the short 2010-onward panel the significant gradient disappears once the pandemic years are removed, whereas on the extended 1993–2025 panel it is negative and stable across every sub-period, including the pre-2010 window.

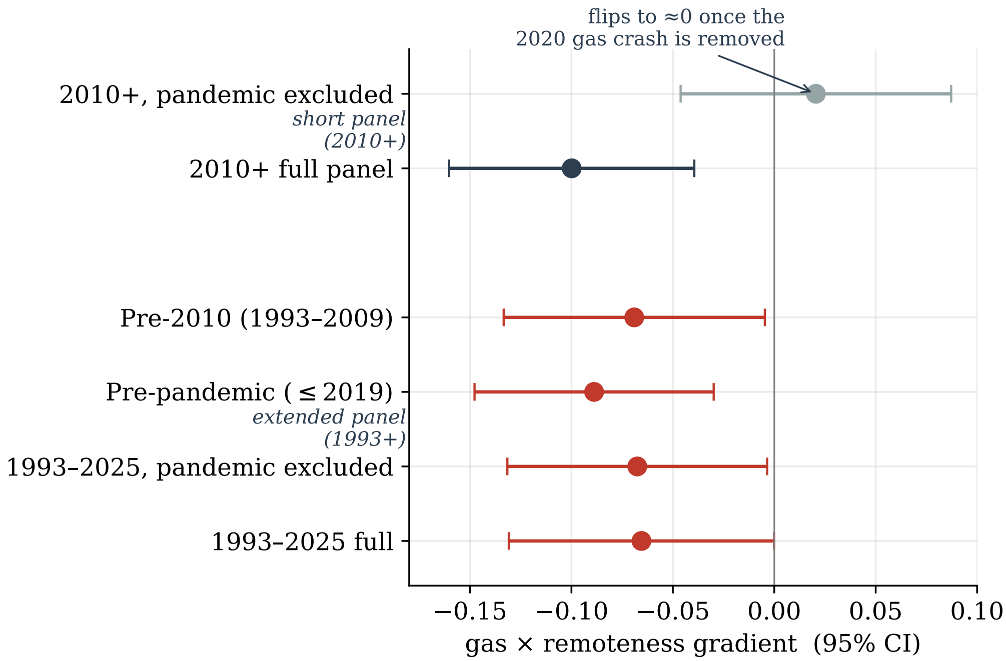

Figure 6: The gas-price by remoteness gradient by sample window. On the short 2010-onward panel (top), the gradient is large and significant on the full sample but collapses toward zero once the pandemic years are excluded, because the 2020 fuel-price crash coincided with the pandemic-era boom in outdoor recreation. On the extended 1993–2025 panel (bottom), which spans several independent oil-price cycles, the gradient is negative and stable across the full sample, the pandemic exclusion, the pre-pandemic window, and the pre-2010 window alone. Bars are 95% confidence intervals.

5.2 Margins of adjustment

When fuel becomes more expensive, do visitors to remote parks take fewer trips, or do they take the same trips and stay longer? Table 4 and Figure 8 decompose the response across outcomes; the hours-per-visit outcome drops a small number of park-months with implausible implied durations, a trim disclosed and shown immaterial in the table note. The gradient is -0.068 for trips and -0.070 for total visitor-hours, but only -0.002 ($p = 0.87$) for hours per visit and statistically zero for overnight stays. The adjustment is therefore entirely extensive: visitors take fewer trips, and visitor-hours fall in proportion, while the duration of a typical visit is unchanged. The trip-consolidation margin, under which travelers respond to higher travel costs with fewer but longer trips, appears in the pandemic-contaminated 2010-onward sample but does not replicate on the full panel.

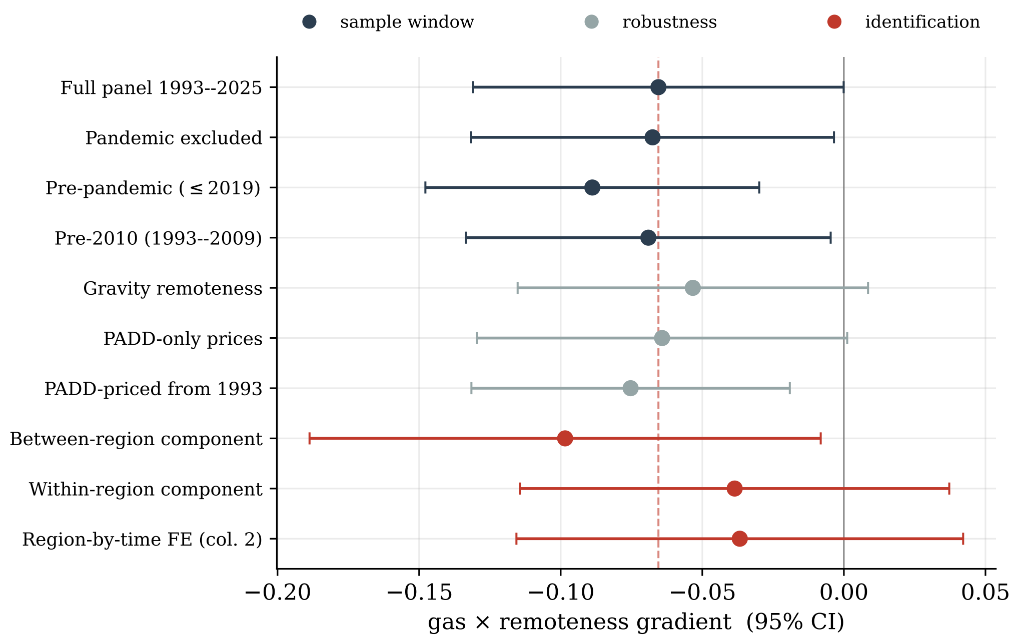

Figure 7: The gradient across every cut, in one view: sample windows (top), robustness perturbations (middle), and the identification decomposition (bottom). The dashed line marks the baseline (−0.065). The point estimate is negative throughout; the between-region component carries the significance, while the within-region component and the region-by-time specification are same-signed but underpowered. Bars are 95% confidence intervals (clustered by park).

5.3 No within-system substitution

A natural question is whether the trips lost at remote parks reappear at nearer ones, the spatial-substitution channel documented by Oh and Hammitt (2011). To test for it, I add to the specification a spatial lag of neighboring parks' remoteness interacted with the gas price, estimated on the 2010-onward panel on which the substitute set is constructed (Table 5; Figure 9). In the within-region specification the neighbor interaction is small and statistically indistinguishable from zero (+0.018, $p = 0.79$; randomization-inference $p = 0.52$), and it remains so on the extended, pandemic-excluded 1993–2025 panel (−0.008, $p = 0.90$; Table 6). This is a weak test, however. Own and neighbor remoteness are correlated at 0.84 because remote parks cluster geographically; the own gradient is itself imprecise in this specification (−0.072, $p = 0.18$); and the neighbor estimate's confidence interval is wide enough to admit diversion as large as the own effect. I therefore read the result as no *detectable* within-

Table 4: Margins of adjustment: the gas-price by remoteness gradient by outcome—recreation visits, visitor-hours, hours per visit, and overnight stays. Standard errors clustered by park in parentheses.

	Log monthly outcome			
	Recreation visits (1)	Visitor-hours (2)	Hours per visit (3)	Overnight stays (4)
ln(Gas price)	0.109	0.089	−0.021	−0.756
	(0.213)	(0.211)	(0.103)	(0.477)
ln(Gas price) × Remoteness (z)	−0.068**	−0.070*	−0.002	0.018
	(0.033)	(0.038)	(0.015)	(0.101)
Year-month	x	x	x	x
Park × calendar-month	x	x	x	x
Observations	107,098	107,064	106,931	36,047
R^2	0.946	0.955	0.939	0.922

Significance levels: * p < 0.1, ** p < 0.05, *** p < 0.01. Cells report the coefficient with the standard error in parentheses.
Hours per visit drops the 133 park-months (at 3 units) with implied durations above 100 hours, which are reporting artifacts; without the trim the duration gradient is -0.002 ($p = 0.88$).

system substitution rather than as positive evidence of none—a reading also consistent with the broader absence of cross-park spillovers in the system. To the extent that the displaced trips do not reappear at nearby parks, the welfare loss developed in Section 6 is a reduction in access rather than a reshuffling of visits; the data cannot decisively rule out the latter.

5.4 Robustness and mechanism

Table 6 and Figure 10 show that the gradient's point estimate is stable across a wide battery of perturbations: an alternative gravity-based remoteness measure (-0.053, $p = 0.09$), the actual road distance and driving time to the nearest metropolis in place of great-circle distance (-0.058, $p = 0.08$, and -0.053, $p = 0.11$; the road and great-circle rankings correlate at 0.99, so circuity does not bias the measure even though it is higher in the West, and the road gradient carries the same between/within structure as the baseline—a between-region

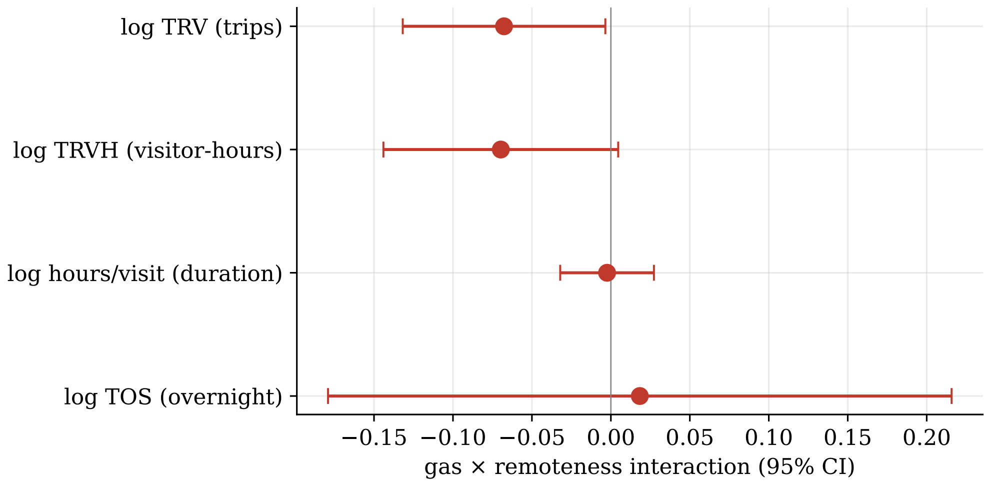

Figure 8: Margins of adjustment: the response is entirely extensive, with trips and hours falling and per-visit duration unchanged. Bars are 95% confidence intervals (clustered by park).

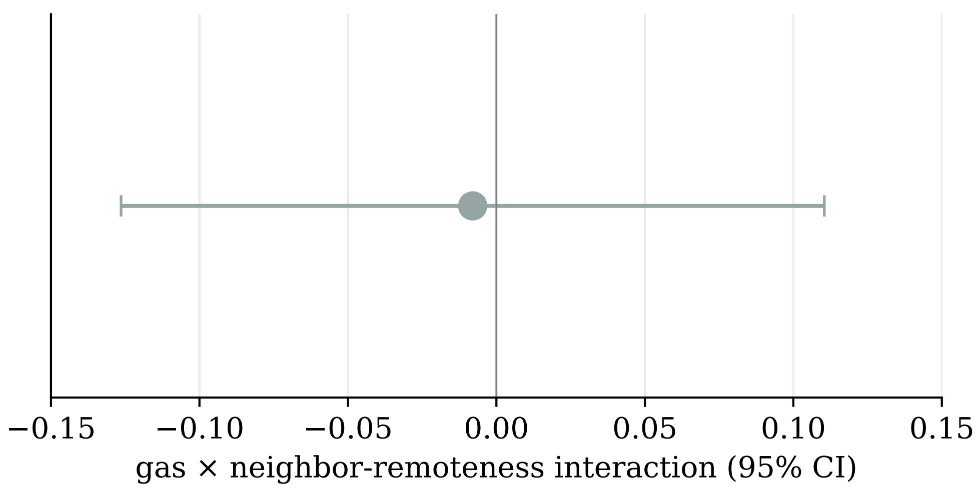

Figure 9: Spatial-substitution test on the extended pandemic-excluded panel: the neighbor-remoteness by gas-price interaction is indistinguishable from zero, though the test is underpowered. The bar is the 95% confidence interval (clustered by park).

Table 5: Spatial-substitution test: no detectable diversion of lost trips to nearby parks. All specifications use the 2010–2025 panel including the pandemic years, on which the substitute-set construction is available. Specification A is the within-region design (gas $\times$ own and neighbor remoteness, region-by-year-month effects); A2 re-estimates A with national year-month effects in place of the region-by-time effects (which also permits a gas-price main effect); B replaces neighbor remoteness with the count of nearby substitute parks; RI is the randomization-inference p-value for the neighbor interaction. The extended-panel (1993–2025, pandemic-excluded) neighbor test is reported in Table 6.

	Term	Coef.	SE	p
A: own remoteness	ln(Gas) $\times$ own remoteness	−0.072	0.054	0.185
A: neighbor remoteness	ln(Gas) $\times$ neighbor remoteness	0.018	0.065	0.787
A2: own remoteness	ln(Gas) $\times$ own remoteness	−0.050	0.052	0.335
A2: neighbor remoteness	ln(Gas) $\times$ neighbor remoteness	−0.063	0.049	0.206
B: own remoteness	ln(Gas) $\times$ own remoteness	−0.105	0.041	0.011
B: # substitutes	ln(Gas) $\times$ # substitutes	−0.007	0.036	0.843
RI: neighbor	ln(Gas) $\times$ neighbor remoteness	0.018		0.524

component of -0.103 ($p = 0.03$) with a null within-region part), a price series built only from the five broader PADD regions (-0.064, $p = 0.055$), the exclusion of the pandemic, and the exclusion of any single region one at a time (a range of -0.079 to -0.053). On the pandemic-excluded panel, the gradient likewise survives excluding each major oil-price episode one at a time—the 2007–08 spike (-0.069), the 2014–16 collapse (-0.066), and the 2022 spike (-0.078)—a one-month lag of the price (-0.069), restriction to the May–September driving season (-0.063), and the exclusion of the handful of boat- and seaplane-access units (-0.064, $p = 0.054$). Allowing remote parks their own linear trend leaves the gradient negative and, if anything, larger (-0.125, $p < 0.01$; the trend interaction itself is insignificant), and remoteness measured from 1993 county population gives a similar but less precise estimate (-0.056, $p = 0.09$), consistent with attenuation from the older vintage. The point estimates all lie well within the baseline confidence interval; precision varies, with the gravity, road-distance, PADD-only, vintage, and boat-exclusion variants significant at the 10% rather than 5% level, and the driving-time variant ($p = 0.11$) just outside it. It is also unchanged when the weather controls are interacted with remoteness, allowing temperature and precipitation to bear on remote and accessible parks differently (-0.066, $p = 0.04$), and a first-difference test finds no asymmetry: the gradient is statistically indistinguishable for month-over-month gas-price increases versus decreases (difference +0.02, $p = 0.75$). As a check on the mechanism (Figure 11), the gradient is negative and significant for recreation visits (-0.068, $p = 0.04$) but smaller and insignificant for non-recreation through-traffic (-0.048, $p = 0.49$). The pattern is consistent with a recreational travel-cost effect, but it is only suggestive: a stacked specification cannot reject equality of the two gradients ($\Delta = 0.030$, $p = 0.71$), so on its own this check does not rule out a remoteness-correlated regional confound common to all park traffic. A non-fuel placebo points the same way more sharply. Gateway-county income, itself procyclical, interacted with remoteness yields an imprecise interaction of similar magnitude on its own (-0.061, $p = 0.19$; its within-region level elasticity is a precise zero, Appendix B), so the standalone placebo is only weakly informative; the discriminating test is the horse race, in which the gas-by-remoteness gradient is undisturbed (-0.064, $p = 0.02$) while the income interaction collapses to zero (+0.002, $p = 0.97$). The same conclusion holds for the between-region component that carries the gradient's significance, where the gas term in fact strengthens once income is included (Section 5.1, Table 3). The differential response is thus specific to the fuel price rather than to generic remoteness-correlated cyclicality.

To guard against the concern that the headline rests on a single fortunate specification, Figure 12 reports a specification curve: the gradient estimated across all sixty-four combinations of the defensible analytical choices—two remoteness measures, two price geographies, national versus region-by-time fixed effects, four sample windows, and weather controls with or without a remoteness interaction. The gradient is negative in 94% of these specifications, and significant at the 5% level in one-third of them overall. That share is entirely a matter of identifying power, not of which fork is chosen: significance reaches two-thirds (66%) among the specifications that use the national time effects—the design that identifies the headline—and zero among the region-by-time specifications that the decomposition in Section 5.1 already shows are underpowered. The sign of the effect is thus not a product of specification search, and its statistical significance is concentrated exactly where the design has power.

Table 6: Robustness of the gas-price × remoteness gradient on the extended 1993–2025 panel. Rows are park-clustered except the gas-region clustering row, which uses the CRV3 cluster-jackknife over the 14 gas regions. Rows marked ^{e} are estimated on the pandemic-excluded panel. The leave-one-region-out rows report the most- and least-negative point estimate across the 14 single-region drops (standard errors and N do not apply). The neighbor-remoteness row (^{n}) reports a different estimand—the gas × neighbor-remoteness interaction from the within-system substitution check of Section 5.3—not the own gradient.

	Coef.	SE	p	N
Baseline (distance to nearest metro)	−0.065	0.033	0.050	114,579
Gravity population-access index	−0.053	0.032	0.092	114,579
PADD-only gas-price series	−0.064	0.033	0.055	114,445
Exclude pandemic (2020–21)	−0.068	0.033	0.039	107,098
Weather × remoteness controls ^{e}	−0.066	0.032	0.041	107,098
Remoteness-specific linear trend	−0.125	0.026	<0.001	114,579
Exclude 2007–08 price spike ^{e}	−0.069	0.034	0.044	99,916
Exclude 2014–16 price collapse ^{e}	−0.066	0.033	0.049	95,542
Exclude 2022 price spike ^{e}	−0.078	0.032	0.015	103,293
One-month-lagged gas price ^{e}	−0.069	0.033	0.035	102,970
Driving season only (May–Sep) ^{e}	−0.063	0.031	0.044	44,749
Exclude boat/seaplane-access units ^{e}	−0.064	0.033	0.054	104,446
All parks PADD-priced from 1993 ^{e}	−0.075	0.029	0.009	117,232
All parks PADD-priced, 1993–2009 only	−0.079	0.030	0.010	62,167
Remoteness from 1993 population ^{e}	−0.056	0.033	0.091	107,098
Road distance to nearest metro	−0.058	0.033	0.080	114,579
Driving time to nearest metro	−0.053	0.033	0.106	114,579
Leave-one-region-out (most negative)	−0.079			
Leave-one-region-out (least negative)	−0.053			
Neighbor-remoteness substitution ^{e,n}	−0.008	0.060	0.895	107,098
Gas-region clustering (CRV3, 14 clusters)	−0.065	0.023	0.015	114,579
Pre-pandemic only (≤2019)	−0.089	0.030	0.003	90,836
Pre-2010 only (1993–2009)	−0.069	0.033	0.036	51,965

5.5 Where the gradient operates

Two features locate the gradient within the park system. First, it is a *per-park* response—one that weights each park-month equally in logs—rather than a visit-weighted one. Re-estimated by Poisson pseudo-maximum-likelihood, which weights by visit counts and so is

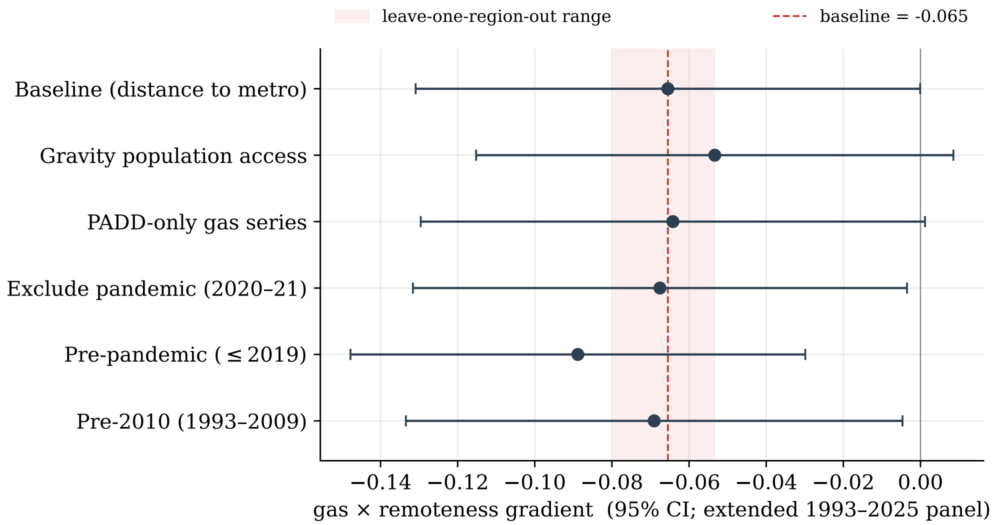

Figure 10: Robustness of the gradient across remoteness measures, the geography of the price series, sub-periods, and leave-one-region-out (95% confidence intervals, clustered by park; the full battery is in Table 6).

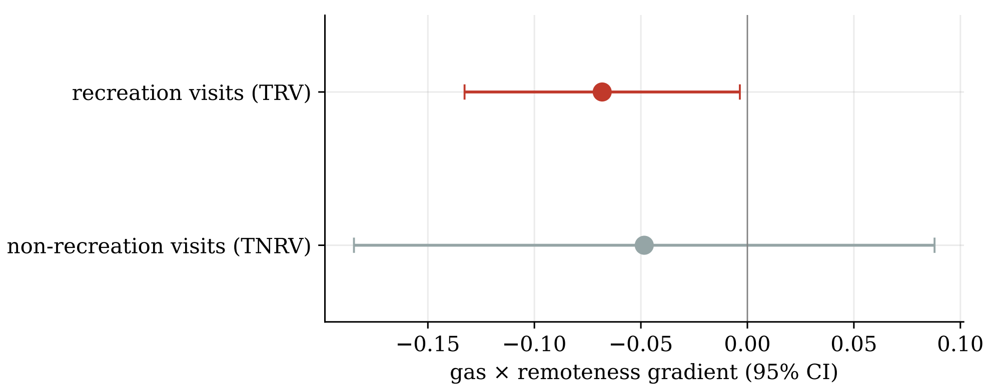

Figure 11: Recreation versus non-recreation gradient: present for recreation visits, weaker and insignificant for non-recreation through-traffic (the difference is not statistically significant). Bars are 95% confidence intervals (clustered by park).

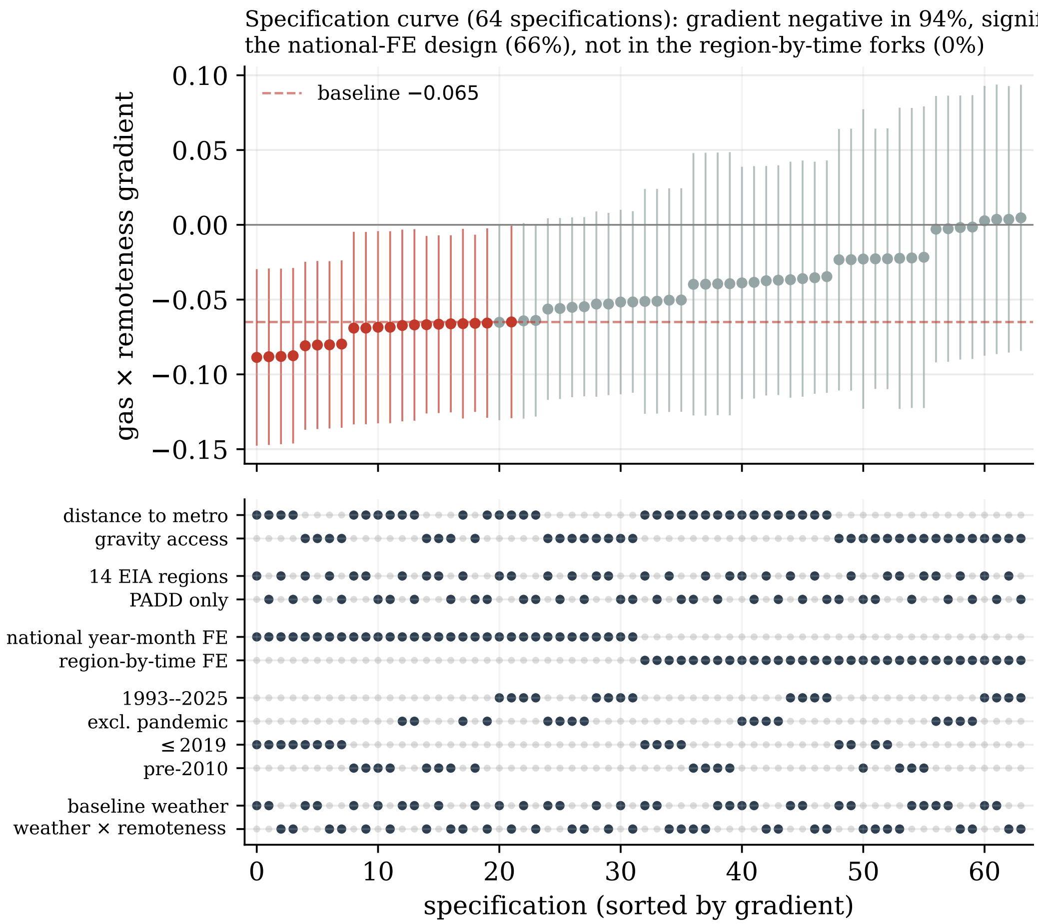

Figure 12: Specification curve: the gas-price by remoteness gradient across all 64 combinations of the defensible analytical choices (top, sorted by point estimate, with 95% confidence intervals; red where significant at the 5% level), with the choices active in each specification marked below. The gradient is negative in 94% of specifications; significance is concentrated in the national-time-FE design (66%) and absent from the underpowered region-by-time forks (0%). The per-park (log) estimand throughout; the count-weighted Poisson estimand, ≈ 0 by construction, is treated separately in Section 5.5.

dominated by the busiest parks, the gradient is essentially zero (-0.002, $p = 0.96$). The reason is heterogeneity across the size distribution: splitting the panel at the median of mean monthly visitation, the gradient is -0.187 ($p = 0.002$) among the smaller half of parks but a precise $+0.015$ ($p = 0.68$) among the larger half. Because the most-visited tenth of parks alone account for about three-fifths of all recreation visits, a count-weighted estimate loads almost entirely on parks at which remoteness does not bind. Figure 13 displays these estimates side by side.

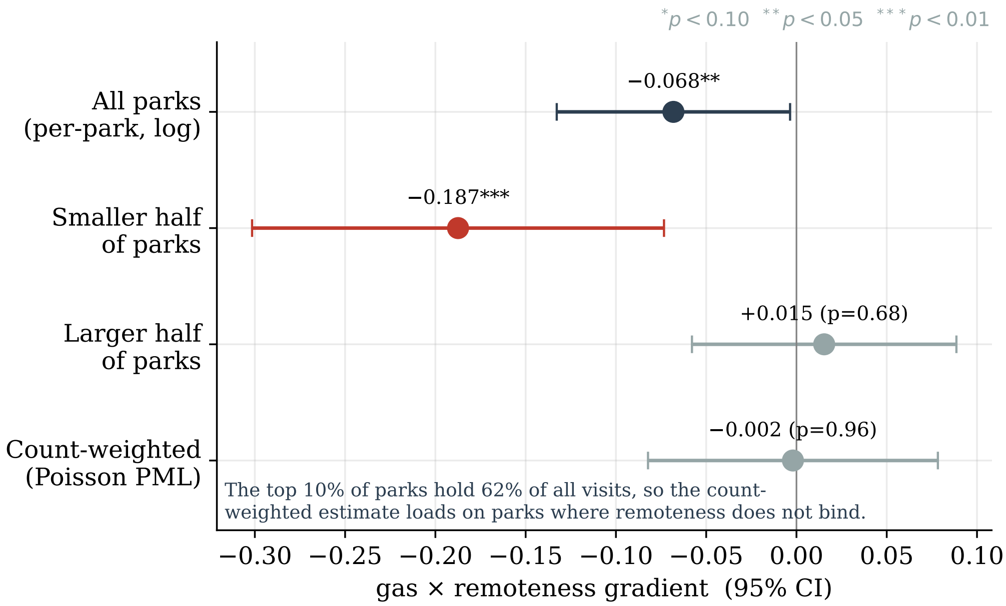

Figure 13: Where the gradient operates, estimated on the pandemic-excluded 1993–2025 panel. The per-park (log) gradient is concentrated in the smaller half of parks (−0.187) and absent in the larger half (+0.015); the count-weighted Poisson estimate, dominated by the highest-volume parks, is essentially zero. Because the top tenth of parks hold roughly three-fifths of all recreation visits, the system-wide count-weighted estimate loads on exactly the parks where remoteness does not bind. Bars are 95% confidence intervals; the smaller-park estimate is significant at the 1% level ($p = 0.002$).

This pattern is economically sensible. Demand at the highest-volume, marquee parks—the destinations that anchor planned, multi-day vacations—is remoteness-inelastic: visitors who have organized a trip around a flagship park do not abandon it when fuel rises, the destination loyalty Oh and Hammitt (2011) document at a single site. The travel-cost margin instead operates at the system's many smaller, more dispersed units, for which a visit is closer to a discretionary day-trip and the drive is a larger share of the total cost. The result is thus not in tension with the inelastic *aggregate* visitation other work reports (Stevens et al., 2014); it says that the fuel-cost margin, while nearly invisible in the system-wide count, is real and concentrated exactly where access is most fragile.

The distinction matters for what follows. The welfare question in Section 6 concerns *access*—how many parks, and which, become harder to reach when driving is dear—not the aggregate visit count, so the per-park (log) specification is the estimand that answers it, and it is the one I carry forward.

The per-park gradient is likewise not an artifact of dropping the zero-visit months noted in Section 3, but the zeros deserve more care than a full-panel Poisson comparison can provide, because a count-weighted estimator places almost no weight on the micro-units where the zeros occur. The zeros have structure: half of the 1,363 zero months in the pandemic-excluded panel sit at just ten micro-units, and seventeen units with more than two years of zeros are chronic closure or reporting regimes rather than demand outcomes. A naive $\log(1 + \text{visits})$ or inverse-hyperbolic-sine regression that treats those closure episodes as demand zeros yields a gradient near zero—the extensive-margin sensitivity of log-like estimators that Chen and Roth (2024) formalize—and the incidence of zeros itself moves with gas-by-remoteness (a linear-probability estimate of -0.008, $p = 0.01$), so the zeros are not ignorable in general. Handled where the gradient actually operates, the result is intact: among the smaller half of parks, the Poisson gradient is -0.202 with zeros kept and -0.205 without ($p = 0.001$ in both cases); excluding the seventeen chronic-closure units, the zero-inclusive $\log(1 + \text{visits})$ gradient is -0.067 ($p = 0.09$); and on the 239 parks that never record a zero, the log estimate is -0.082 ($p = 0.049$). The differential incidence of zeros across remoteness terciles therefore does not drive the result.

6 The distributional incidence of fuel prices and carbon pricing

Conventional analyses of carbon-pricing incidence ask how the policy's cost is distributed across the income distribution, measured through households' energy budgets (Hassett et al., 2009; Williams et al., 2015; Goulder et al., 2019). That accounting captures the price households pay at the pump but not what they forgo with the driving they give up. One omitted item is access to public lands: when the price of driving rises, the parks that become harder to reach are disproportionately the remote ones, so the policy carries a recreation-access co-cost that budget-based incidence studies do not price. Each forgone visit is not negligible—travel-cost estimates put the average willingness to pay for a national-park visit near $100 (Neher et al., 2013)—so a differential in access is a differential in realized recreation benefits. This section quantifies that co-cost with the identified gradient.

Because the design identifies the differential response rather than its level, I express the result as an incidence—the access loss of remote parks relative to accessible ones—rather than as an absolute change in visits. A carbon price raises the price of gasoline by its carbon content, 8.887 kilograms of carbon dioxide per gallon (U.S. Environmental Protection Agency, 2023a), and applying the gradient to the implied change in the log fuel price gives the differential access loss across the remoteness distribution (Table 7; Figure 14). Expressed relative to a recent pump price of $3.65 per gallon, a carbon price of $50 per metric ton of CO_2 (a 12% fuel-price increase) lowers access at parks at the 90th percentile of remoteness by 2.1 percentage points more than at the 10th; a price of $100 per ton (24%), by 4.0 points; and the Environmental Protection Agency's social cost of carbon, near $190 per ton (U.S. Environmental Protection Agency, 2023b)—denominated in 2020 dollars, so conservative against the 2022–2024 nominal price base—or an additional $1.69 per gallon (a 46% increase), by 6.9 points. Because the scenarios are linear in the gradient, their uncertainty is the gradient's: propagating the park-clustered standard error gives an exact 95% confidence interval of 0.4 to 13.5 points on the headline gap (Table 7), wide but bounded away from zero. To the extent that the displaced demand leaves the system rather than reallocating within it (Section 5.3), this differential is a loss of recreational access rather than a reshuffling of visits. Two caveats bound the magnitudes. The scenarios assume the carbon charge passes through fully to the pump price; pass-through of fuel taxes is often found to be high (Li et al., 2014) but can be incomplete (Harju et al., 2022), in which case the magnitudes scale down approximately proportionally. And the gradient is identified from market price variation at monthly frequency, whereas a carbon charge is a permanent, salient policy change: demand responses to persistent, legislated tax changes are systematically larger than responses to equally sized market price movements (Li et al., 2014; Coglianese et al., 2017), as are long-run responses relative to short-run ones (Small and Van Dender, 2007), so to the extent the differential response scales with the level response, the scenarios understate the long-run differential and are best read as conservative.

Three features make this co-cost worth counting. First, it is geographically concentrated: the parks that lose the most access are the remote, rural-western units far from metropolitan populations (Figure 14), so the burden falls on exactly the regions that energy-cost increases are already scrutinized for hitting hardest. Second, it compounds existing inequities in who reaches the parks. Distance and travel cost are among the leading barriers to national-park visitation, and they fall most heavily on lower-income and minority populations, who on average live farther from the system's units (Weber and Sultana, 2013; Scott and Lee, 2018; Xiao et al., 2022); a policy that raises the cost of distance widens that gap, echoing a broader finding that the benefits of public recreation, and the burden of the fees used to ration it, are themselves unequally distributed (More and Stevens, 2000; Huhtala and Pouta, 2009; Ji et al., 2022). Third, the magnitude is of policy-relevant size. At the EPA social cost of carbon the point estimate of the differential access loss between the most- and least-remote parks (the 90th versus 10th percentile of remoteness) is about seven percentage points—a spread across the remoteness distribution rather than a system-wide level change. Applied to the roughly 51,000 mean monthly visits of a remote-tercile park, that differential is on the order of 3,500 forgone visits per month relative to an accessible park, and it implies a difference of about 0.15 in the fuel-price elasticity of visitation between the most- and least-remote parks—the same modest order as the entrance-fee elasticities estimated for the park system (Schwartz and Lin, 2006; Melstrom and Vasarhelyi, 2019). The comparison is to a *differential* between park types, computed from observed baseline visitation rather than from the confounded elasticity level, and its confidence interval is wide with a small lower bound. The difference from a fee is that the fuel-cost margin is one the agency neither sets nor can waive. None of this argues against carbon pricing, whose climate benefits are first-order; it

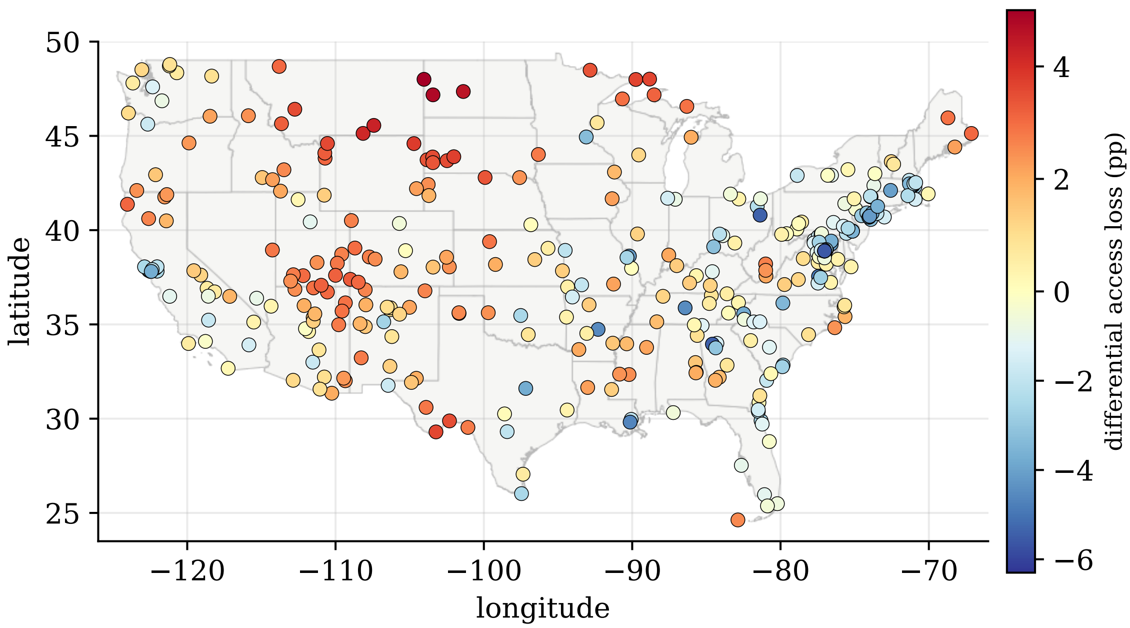

Figure 14: Distributional incidence: the differential access loss under the EPA social cost of carbon, concentrated in the rural West.

argues for entering the recreation-access co-cost into the ledger, and for pairing the policy with measures—fee relief at remote units, or transit and shuttle access—that blunt a burden it otherwise places on the least-accessible public lands.

Table 7: Distributional incidence of fuel-price and carbon-pricing shocks on remote-park access. Access losses apply the pandemic-excluded gradient of -0.068 per standard deviation of remoteness to $\Delta \ln(\text{gas}) = \ln(1+\Delta\text{gas}/p_0)$, relative to the 2022–2024 mean pump price $p_0 = \$3.65$/gal. The last two columns report the loss gap between parks at the 90th and 10th percentiles of remoteness (a 2.7-SD spread), in percentage points (larger = more access lost), with the exact 95% confidence interval from propagating the park-clustered standard error of the gradient (the scenarios are linear in the gradient). Carbon prices are per metric ton of CO_2; the EPA social cost of carbon is denominated in 2020 dollars, so against the 2022–2024 nominal price base the implied shock is conservative.

	Δ gas ($/gal)	% gas increase	Access loss / SD (pp)	Remote−accessible loss (pp)	95% CI
Carbon $50/t	0.44	12.2	0.78	2.10	[0.1, 4.1]
Carbon $100/t	0.89	24.4	1.47	3.98	[0.2, 7.7]
EPA SCC (∼$190/t)	1.69	46.3	2.57	6.94	[0.4, 13.5]
+$1.00/gal	1.00	27.4	1.64	4.42	[0.2, 8.6]

7 Discussion and conclusion

Fuel prices are a real and spatially uneven shifter of national-park demand. Using regional gasoline-price variation interacted with park remoteness, I estimate a robust travel-cost gradient of about -0.07 per standard deviation of remoteness: when fuel is expensive, remote parks lose visits relative to accessible ones. The adjustment is entirely extensive—fewer trips and proportionally fewer visitor-hours, with no change in the duration of a visit—and I find no evidence that the displaced demand diverts to nearer parks, though that test is underpowered. The response is a per-park one, borne by the system's many lower-volume units rather than its few high-traffic ones, whose demand is remoteness-inelastic. Applying the gradient to the fuel-price changes implied by carbon pricing yields a spatially regressive incidence, with the most-remote parks losing several percentage points more access than the most-accessible, a burden concentrated in the rural West.

A methodological lesson generalizes beyond this application. Estimated on the convenient post-2010 window alone, the design delivers a large, significant, and wrong answer, because a single dominant episode—the 2020 fuel-price crash arriving together with a once-in-a-century shock to outdoor recreation—masquerades as identifying variation. No control strategy diagnosed the problem; only widening the sample to span several independent price cycles did, after which the estimate stabilized across every sub-period, including the years before 2010. Continuous-treatment designs whose identifying variation comes from one macroeconomic episode are fragile in precisely this way, and the remedy is more cycles, not more covariates.

The results speak to two policy debates. The first is the design of recreation access and pricing. The agency's own price instruments—the roughly \$20 to \$35 entrance fee at the minority of units that charge one, the \$80 annual pass, and the timed-entry reservations recently added at a few congested parks—operate at the busy, accessible destinations where, as Section 5.5 shows, demand is least sensitive to the cost of reaching the park. The fuel-cost margin works in the opposite place: it binds at the remote, lower-volume units, it is typically larger than the entrance fee for a household driving any real distance, and it lies entirely outside the agency's control. Two implications follow. One is that the distributional footprint of access policy is incomplete if it counts only fees, because a uniform fee change and a fuel-price change fall on very different parks. The other is that instruments meant to broaden access—fee-free days, discounted or waived passes, or subsidized transit and shuttle service—deliver the most additional access where they offset the fuel margin, namely at remote units, and that fee relief there forgoes little revenue precisely because those units are lightly visited.

The second debate, less widely recognized, is the distributional design of fuel and carbon pricing. A carbon price is evaluated mainly through its effect on household energy budgets and through how its revenue is returned (Hassett et al., 2009; Goulder et al., 2019; Cronin et al., 2019). The recreation-access co-cost measured here is a distinct margin, and it points the same way as the budget incidence: it is spatially regressive, concentrated on the rural-western public lands that are least substitutable and on the lower-income and minority populations who already face the longest trips to reach them (Weber and Sultana, 2013; Scott and Lee, 2018). At the EPA social cost of carbon the gap reaches about seven percentage points of access between the most- and least-remote parks (95% confidence interval 0.4 to 13.5)—small beside the climate benefits that justify the policy, but not zero, and at present omitted from the ledger entirely. Recognizing it does not weigh against carbon pricing; it argues for entering the co-cost into the accounting and, where it matters, for pairing the policy with access measures of the kind just described. More broadly, the finding suggests that distance to amenities is an under-examined dimension of the incidence of any policy that raises the price of driving.

Two limitations bound these claims. The identified object is a gradient rather than a level: because gasoline prices are procyclical, the average elasticity is confounded and is not interpreted, and a state gasoline-tax instrument, while it has a strong first stage, is too underpowered on the available window of tax variation to recover the level (Appendix A). And because the data are monthly counts at the park level without visitor origins, remoteness proxies the representative trip distance rather than measuring it. Both limitations point to natural extensions: a longer and broader panel of state-level retail prices and gasoline taxes to recover the level; and origin-resolved or mobile-device trip data to measure trip distance directly (a road-network distance and driving time, which I do use, still locate the park rather than the visitor). Neither is required for the gradient, which is identified under the design's parallel-cyclical-response assumption, but each would refine the welfare magnitudes that follow from it.

References

Susanne Becken and James Lennox. Implications of a long-term increase in oil prices for tourism. *Tourism Management*, 33(1):133–142, 2012. doi: 10.1016/j.tourman.2011.02.012.

Brantly Callaway, Andrew Goodman-Bacon, and Pedro H. C. Sant'Anna. Difference-in-differences with a continuous treatment. NBER Working Paper 32117, National Bureau of Economic Research, 2024.

Taurey Carr and Stephen C. Newbold. More people would visit Yellowstone if it weren't so crowded: Optimal congestion pricing for high demand protected areas. *Journal of Travel Research*, 64(6):1499–1514, 2025. doi: 10.1177/00472875251316267.

Jiafeng Chen and Jonathan Roth. Logs with zeros? Some problems and solutions. *Quarterly Journal of Economics*, 139(2):891–936, 2024. doi: 10.1093/qje/qjad054.

Seong-Hoon Cho, J. M. Bowker, Donald B. K. English, Roland K. Roberts, and Taeyoung Kim. Effects of travel cost and participation in recreational activities on national forest visits. *Forest Policy and Economics*, 40:21–30, 2014. doi: 10.1016/j.forpol.2013.12.004.

John Coglianese, Lucas W. Davis, Lutz Kilian, and James H. Stock. Anticipation, tax avoidance, and the price elasticity of gasoline demand. *Journal of Applied Econometrics*, 32(1):1–15, 2017. doi: 10.1002/jae.2500.

Julie Anne Cronin, Don Fullerton, and Steven Sexton. Vertical and horizontal redistributions from a carbon tax and rebate. *Journal of the Association of Environmental and Resource Economists*, 6(S1):S169–S208, 2019. doi: 10.1086/701191.

Christopher Daly, Michael Halbleib, Joseph I. Smith, Wayne P. Gibson, Matthew K. Doggett, George H. Taylor, Jan Curtis, and Phillip P. Pasteris. Physiographically sensitive mapping of climatological temperature and precipitation across the conterminous United States. *International Journal of Climatology*, 28(15):2031–2064, 2008. doi: 10.1002/joc.1688.

Lucas W. Davis and Christopher R. Knittel. Are fuel economy standards regressive? *Journal of the Association of Environmental and Resource Economists*, 6(S1):S37–S63, 2019. doi: 10.1086/701187.

Jeffrey Englin and J. S. Shonkwiler. Estimating social welfare using count data models: An application to long-run recreation demand under conditions of endogenous stratification and truncation. *The Review of Economics and Statistics*, 77(1):104–112, 1995. doi: 10.2307/2109996.

Federal Highway Administration. Monthly motor fuel reported by states. U.S. Department of Transportation, Office of Highway Policy Information, 2024. URL `https://www.fhwa.dot.gov/policyinformation/motorfuelhwy_trustfund.cfm`. State gasoline excise-tax rates by state and year, 2013–2023.

Nicholas A. Fisichelli, Gregor W. Schuurman, William B. Monahan, and Pamela S. Ziesler. Protected area tourism in a changing climate: Will visitation at US national parks warm up or overheat? *PLOS ONE*, 10(6):e0128226, 2015. doi: 10.1371/journal.pone.0128226.

Kenneth Gillingham. Identifying the elasticity of driving: Evidence from a gasoline price shock in California. *Regional Science and Urban Economics*, 47:13–24, 2014. doi: 10.1016/j.regsciurbeco.2013.08.004.

Kenneth Gillingham and Anders Munk-Nielsen. A tale of two tails: Commuting and the fuel price response in driving. *Journal of Urban Economics*, 109:27–40, 2019. doi: 10.1016/j.jue.2018.09.007.

Kenneth Gillingham, David Rapson, and Gernot Wagner. The rebound effect and energy efficiency policy. *Review of Environmental Economics and Policy*, 10(1):68–88, 2016. doi: 10.1093/reep/rev017.

Lawrence H. Goulder, Marc A. C. Hafstead, GyuRim Kim, and Xianling Long. Impacts of a carbon tax across US household income groups: What are the equity-efficiency trade-offs? *Journal of Public Economics*, 175:44–64, 2019. doi: 10.1016/j.jpubeco.2019.04.002.

Jarkko Harju, Tuomas Kosonen, Marita Laukkanen, and Kimmo Palanne. The heterogeneous incidence of fuel carbon taxes: Evidence from station-level data. *Journal of Environmental Economics and Management*, 112:102607, 2022. doi: 10.1016/j.jeem.2021.102607.

Kevin A. Hassett, Aparna Mathur, and Gilbert E. Metcalf. The incidence of a U.S. carbon tax: A lifetime and regional analysis. *The Energy Journal*, 30(2):155–177, 2009. doi: 10.5547/ISSN0195-6574-EJ-Vol30-No2-8.

Jonathan E. Hughes, Christopher R. Knittel, and Daniel Sperling. Evidence of a shift in the short-run price elasticity of gasoline demand. *The Energy Journal*, 29(1):113–134, 2008. doi: 10.5547/ISSN0195-6574-EJ-Vol29-No1-9.

Anni Huhtala and Eija Pouta. Benefit incidence of public recreation areas—have the winners taken almost all? *Environmental and Resource Economics*, 43(1):63–79, 2009. doi: 10.1007/s10640-008-9252-6.

Yongjie Ji, David A. Keiser, Catherine L. Kling, and Daniel J. Phaneuf. Revenue and distributional consequences of alternative outdoor recreation pricing mechanisms: Evidence from a micro panel data set. *Land Economics*, 98(3):478–494, 2022. doi: 10.3368/le.98.3.090721-0107.

David Keiser, Gabriel Lade, and Ivan Rudik. Air pollution and visitation at U.S. national parks. *Science Advances*, 4(7):eaat1613, 2018. doi: 10.1126/sciadv.aat1613.

Lutz Kilian. Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market. *American Economic Review*, 99(3):1053–1069, 2009. doi: 10.1257/aer.99.3.1053.

Douglas M. Larson. Joint recreation choices and implied values of time. *Land Economics*, 69(3):270–286, 1993. doi: 10.2307/3146593.

Laurence Levin, Matthew S. Lewis, and Frank A. Wolak. High frequency evidence on the demand for gasoline. *American Economic Journal: Economic Policy*, 9(3):314–347, 2017. doi: 10.1257/pol.20140093.

Shanjun Li, Joshua Linn, and Erich Muehlegger. Gasoline taxes and consumer behavior. *American Economic Journal: Economic Policy*, 6(4):302–342, 2014. doi: 10.1257/pol.6.4.302.

Frank Lupi, Daniel J. Phaneuf, and Roger H. von Haefen. Best practices for implementing recreation demand models. *Review of Environmental Economics and Policy*, 14(2):302–323, 2020. doi: 10.1093/reep/reaa007.

James G. MacKinnon, Morten Ørregaard Nielsen, and Matthew D. Webb. Cluster-robust inference: A guide to empirical practice. *Journal of Econometrics*, 232(2):272–299, 2023. doi: 10.1016/j.jeconom.2022.04.001.

Kenneth E. McConnell. On-site time in the demand for recreation. *American Journal of Agricultural Economics*, 74(4):918–925, 1992. doi: 10.2307/1243189.

Richard T. Melstrom and Leah Vasarhelyi. Modeling recreation demand and fees at national parks. *Annals of Tourism Research*, 77:175–178, 2019. doi: 10.1016/j.annals.2018.12.020.

Thomas A. More and Thomas Stevens. Do user fees exclude low-income people from resource-based recreation? *Journal of Leisure Research*, 32(3):341–357, 2000.

National Park Service. Visitation numbers. U.S. Department of the Interior, 2025. URL https://www.nps.gov/aboutus/visitation-numbers.htm. 2024 recreation visits: 331,863,358.

Christopher Neher, John Duffield, and David Patterson. Valuation of national park system visitation: The efficient use of count data models, meta-analysis, and secondary visitor survey data. *Environmental Management*, 52(3):683–698, 2013. doi: 10.1007/s00267-013-0080-2.

Chi-Ok Oh and William E. Hammitt. Impact of increasing gasoline prices on tourism travel patterns to a state park. *Tourism Economics*, 17(6):1311–1324, 2011. doi: 10.5367/te.2011.0093.

George R. Parsons. Travel cost models. In Patricia A. Champ, Kevin J. Boyle, and Thomas C. Brown, editors, *A Primer on Nonmarket Valuation*, volume 13 of *The Economics of Non-Market Goods and Resources*, chapter 6, pages 187–233. Springer, Dordrecht, 2 edition, 2017. doi: 10.1007/978-94-007-7104-8_6.

George R. Parsons, Christopher G. Leggett, Joseph A. Herriges, Kevin J. Boyle, Nancy E. Bockstael, and Zhe Chen. A site-portfolio model for multiple-destination recreation trips: Valuing trips to national parks in the southwestern United States. *Journal of the Association of Environmental and Resource Economists*, 8(1):1–25, 2021. doi: 10.1086/710714.

Daniel J. Phaneuf and V. Kerry Smith. Recreation demand models. In Karl-Göran Mäler and Jeffrey R. Vincent, editors, *Handbook of Environmental Economics*, volume 2, chapter 15, pages 671–761. Elsevier, 2005. doi: 10.1016/S1574-0099(05)02015-2.

Zvi Schwartz and Li-Chun Lin. The impact of fees on visitation of national parks. *Tourism Management*, 27(6):1386–1396, 2006. doi: 10.1016/j.tourman.2005.12.015.

David Scott and Kang Jae Jerry Lee. People of color and their constraints to national parks visitation. *The George Wright Forum*, 35(1):73–82, 2018. URL https://www.jstor.org/stable/26452993.

Kenneth A. Small and Kurt Van Dender. Fuel efficiency and motor vehicle travel: The declining rebound effect. *The Energy Journal*, 28(1):25–51, 2007. doi: 10.5547/ISSN0195-6574-EJ-Vol28-No1-2.

Thomas H. Stevens, Thomas A. More, and Marla Markowski-Lindsay. Declining national park visitation: An economic analysis. *Journal of Leisure Research*, 46(2):153–164, 2014.

U.S. Congress. Federal Lands Recreation Enhancement Act, 2004. Pub. L. No. 108-447, Div. J, Title VIII; codified at 16 U.S.C. §§ 6801–6814.

U.S. Environmental Protection Agency. Greenhouse gas emissions from a typical passenger vehicle. Technical Report EPA-420-F-23-014, U.S. Environmental Protection Agency, Office of Transportation and Air Quality, 2023a. URL https://www.epa.gov/greenvehicles/greenhouse-gas-emissions-typical-passenger-vehicle.

U.S. Environmental Protection Agency. EPA report on the social cost of greenhouse gases: Estimates incorporating recent scientific advances. Technical report, U.S. Environmental Protection Agency, 2023b. URL https://www.epa.gov/system/files/documents/2023-12/epa_scghg_2023_report_final.pdf.

Margaret Walls. Economics of the US national park system: Values, funding, and resource management challenges. *Annual Review of Resource Economics*, 14:579–596, 2022. doi: 10.1146/annurev-resource-111920-021401.

Joe Weber and Selima Sultana. Why do so few minority people visit national parks? Visitation and the accessibility of "America's Best Idea". *Annals of the Association of American Geographers*, 103(3):437–464, 2013. doi: 10.1080/00045608.2012.689240.

Casey J. Wichman. Social media influences national park visitation. *Proceedings of the National Academy of Sciences*, 121(15):e2310417121, 2024. doi: 10.1073/pnas.2310417121.

Roberton C. Williams, III, Hal Gordon, Dallas Burtraw, Jared C. Carbone, and Richard D. Morgenstern. The initial incidence of a carbon tax across income groups. *National Tax Journal*, 68(1):195–214, 2015. doi: 10.17310/ntj.2015.1.09.

Xiao Xiao, KangJae Jerry Lee, and Lincoln R. Larson. Who visits U.S. national parks (and who doesn't)? A national study of perceived constraints and vacation preferences across diverse populations. *Journal of Leisure Research*, 53(3):404–425, 2022. doi: 10.1080/00222216.2021.1899776. Published online 18 August 2021.

Alwyn Young. Channeling Fisher: Randomization tests and the statistical insignificance of seemingly significant experimental results. *Quarterly Journal of Economics*, 134(2): 557–598, 2019. doi: 10.1093/qje/qjy029.

A A state gasoline-tax instrument

The level of the elasticity is confounded by the procyclicality of fuel prices, which motivates an instrument. Instrumenting the retail price with state gasoline taxes (Li et al., 2014)—excise-tax rates by state and year from the Federal Highway Administration's Monthly Motor Fuel Reported by States (Federal Highway Administration, 2024)—fails when the price series is measured at the regional level: the geography of the regional series and the state-level tax do not align, and the first stage is essentially dead ($F = 0.2$). With a state-level price series the first stage is strong ($F = 45.6$), but the available window of tax variation is thin: nine states over 2013–2023, only six of which change their tax at all, with most of the movement small annual indexation and the few large legislated increases confined to a handful of state-years. The instrumented level elasticity is accordingly uninformative (+2.5, $p = 0.12$). The instrument is therefore valid but underpowered, and the level remains unidentified (Table A1). Figure A1 shows both halves of this conclusion: the tax moves the price strongly only when the price is measured at the state level, and even there the instrumented level elasticity's confidence interval comfortably spans zero.

Table A1: State gasoline-tax instrument (state-level gas prices, 9 states, 2013–2023). The first stage is strong, but the short window of tax variation leaves the instrumented level elasticity imprecise; the gradient is not separately identified here. The first-stage coefficient is the change in the log retail price per $1 of state tax; at the sample-mean price of $2.96/gal it implies dollar-on-dollar pass-through of roughly 0.85 (0.71 excluding the pandemic), consistent with the high pass-through found by Li et al. (2014).

	First-stage F	First stage ($\Delta \ln p$ per $1/gal tax)	IV level elast.	IV level p	N
Full 2013–2023	45.6	0.29	2.50	0.120	14,629
Exclude pandemic (2020–21)	37.5	0.24	1.47	0.314	12,101

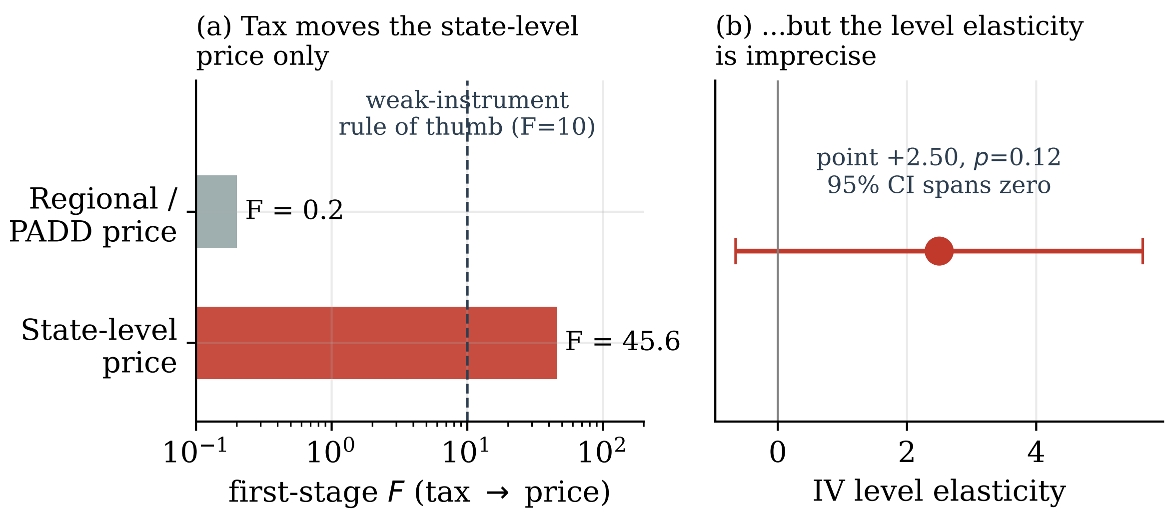

Figure A1: The state gasoline-tax instrument. Panel (a): the first-stage F from regressing the log gas price on the state tax is far below the weak-instrument rule of thumb when the price is measured at the regional (PADD) level, but strong once the price is measured at the state level. Panel (b): even with the strong state-level first stage, the instrumented level elasticity is imprecise—its 95% confidence interval spans zero—so the level remains unidentified on the available 2013–2023 tax variation.

B Income elasticity

A within-region income elasticity of visitation is indistinguishable from zero (Table A2). Gateway-county income—annual per-capita personal income of the park's gateway county, from the Bureau of Economic Analysis regional accounts—is not visitor income, and a national income elasticity would be collinear with the year effects, so the data cannot speak to the income elasticity of park demand.

Table A2: Income elasticity of park visitation (null). The within-region income elasticity is indistinguishable from zero; the significant national-trend row is collinear with the macroeconomic cycle and is not separately identified, and the National-Park interaction is an increment on (confounded) gateway-county income, not visitor income.

	Term	Coef.	SE	p
Region×year FE	ln(Income)	0.056	0.128	0.661
+ park trend	ln(Income)	0.056	0.128	0.661
National trend (confounded)	ln(Income)	0.059	0.028	0.035
Heterogeneity: base	ln(Income)	−0.033	0.125	0.792
Heterogeneity: NP extra	ln(Income) × National-Park designation	0.204	0.078	0.009