

Darkening the Mood: Are People Willing to Pay for Dark Skies?*

Mo Van de Sompel

October 15, 2026

Abstract

The invention of the lightbulb is widely credited as the primary driver of widespread residential electrification; to this day, when new communities are connected to power grids, the first purchases are often lighting devices. Indeed, light is so correlated with wealth and productivity that economists regularly use remote-sensed night sky brightness as a proxy for GDP in hard-to-reach locations. Yet the popular use of the term “light pollution” implies some unease with the trend of lighting up the night. This paper uses nationwide cell phone mobility data to quantify willingness to pay for sky darkness through a recreation demand study of campground visitation. Exploiting spatial variation in sky brightness from remote-sensing data and temporal variation from lunar phase and position, modeled with a full astronomical brightness model, I find a precisely estimated willingness to pay of roughly \$10–15 for a one magnitude per square arcsecond increase in darkness on an ordinary night, rising to approximately \$25 during a week containing a meteor shower. I find that the increases in light pollution since 2019 incur \$833 million in recreation damages annually.

*I thank David Lorenz for his help in providing light pollution data

1 Introduction

The loss of natural darkness due to artificial light at night is an increasingly common environmental issue across the developed world. Worldwide, sky brightness increased by nearly 10% annually between 2011 and 2022 (Kyba et al., 2023), and for the majority of Americans and Europeans, the Milky Way is now entirely invisible (Falchi et al., 2016). A growing body of research in various disciplines has documented significant ecological and health-related harms associated with nighttime luminosity, yet relatively little effort has been made to study the economic consequences, much less reduce the impacts, of so-called light pollution at scale.

Light pollution is perhaps most salient to the general public through its impact on recreational stargazing and amateur astronomy. Despite this, economic analysis of recreational damages from light pollution remains sparse: most of the literature that uses night sky brightness at all treats it as a proxy for economic activity in hard-to-measure places, rather than as an object of study in its own right (Henderson et al., 2012; Gibson et al., 2021). This paper estimates the marginal willingness to pay for dark skies using a revealed-preference recreation demand model of camping trips, estimated on a nationally representative sample of United States cell phone mobility data.

I utilize a set of cell-phone mobility data, focusing on all sites in the continental United States between 2019 and 2022. My demand model ultimately spans 215,000 origins and 6,508 unique sites, across 204 weeks, combining to roughly half a billion unique site-origin-week observations. I utilize annual remote-sensing light pollution data from Lorenz (2025) alongside the model of moon brightness from Krisciunas and Schaefer (1991) to create a spatially and temporally-varying measure of sky darkness for each observation. Implementing multiple specifications of recreation demand study, I find a willingness to pay for a one unit increase in sky darkness¹ around \$10-\$15 on regular nights, rising to \$25 on nights with meteor showers. I estimate a second suite of specifications allowing valuation to vary by

¹The scale generally runs between 16 and 22, I go into greater detail in Section 2

home sky brightness and income, finding a somewhat counterintuitive result that agents living under dark skies have high willingness to pay for darkness in their campsites, while people living under bright skies actively seek to avoid darkness. I conclude with a policy counterfactual, estimating how much consumer surplus would have been lost had 2019's recreation decisions been made under modern levels of site brightness. I find that the median block group incurs roughly \$500 in recreation damages, for a total \$833 million in damages annually. I ultimately conclude that while these welfare estimates are modest, the cost of abatement is quite low as well.

Among papers that do economically value light pollution itself, they generally fall into two broad categories. The first uses stated-preference surveys to elicit a contingent valuation of willingness to pay (WTP) for darker skies (Simpson and Hanna, 2010; Mitchell et al., 2017; Eiswerth et al., 2026). Simpson and Hanna (2010) finds an asymmetry in this WTP: respondents are willing to pay substantially more to prevent further brightening than to actively reduce brightness that already exists. Mitchell et al. (2017) finds that most visitors to the national parks they survey are willing to pay something to reduce light pollution, but lament that responses are too noisy to pin down a precise value, a pattern the authors attribute partly to unfamiliarity with pricing a nonmarket good. Eiswerth et al. (2026) link visitors' own home sky brightness to their contingent valuation responses, estimating the consumer-surplus and local-spending losses that would follow if a site's sky brightened enough to resemble what visitors already see at home. These studies are informative about the specific sites and populations they survey, but inherit some concerns of the contingent valuation literature: incentive compatibility, external validity, and sensitivity to survey design (Simpson and Hanna, 2010), compounded, as Mitchell et al. (2017) note, by the possibility that even willing respondents may struggle to value an unfamiliar good competently.

The second category takes a wider lens, estimating the systemic causes and aggregate costs of light pollution rather than an individual's marginal WTP. Mitchell and Gallaway (2019) use input-output modeling to estimate that dark-sky tourism generated \$5.8 billion

in visitor spending over ten years on the Colorado Plateau, illustrating the potential local economic stakes of preserving dark skies. Olsen et al. (2014) and Gallaway et al. (2010) instead model light pollution’s causes, tying its severity to local economic, industrial, land-use, and demographic factors and arguing that policy has meaningful leverage over it. Together, these papers document that the aggregate stakes of light pollution are large and that policy can plausibly change it, but neither type of study recovers a marginal willingness to pay for abatement – the object of interest here. The one exception is a contemporaneous working paper, Smith et al. (2026), which surveys visitors to established dark-sky zones on the Colorado Plateau and estimates a WTP of \$45 for a one-unit improvement on the Bortle scale, the first revealed-preference estimate of light pollution’s recreational value that I am aware of. I build on this contribution by moving to a nationally representative sample using cell phone mobility data, which improves external validity beyond a single region’s dark-sky visitors and, in Section 3.1, allows preferences to be broken out by the type of traveler.

Finally, this paper contributes to a growing literature using cell phone mobility data to estimate recreation demand models. Athey et al. (2018), Newbold et al. (2022), and Wan et al. (2026) each confront and resolve technical obstacles specific to this kind of data – among them sparse or censored visitation counts, ambiguous market and choice-set definitions, and the sheer scale of a nationwide panel – in order to recover credible demand parameters where a conventional survey could never reach enough sites or visitors. Dark skies are a natural candidate for this approach: the population of people who value them is thin and geographically dispersed in a way that would be prohibitively expensive to survey directly, but leaves a real footprint in nationwide travel behavior. I develop and estimate a nationwide model of camping trips using spatial variation in sky brightness across sites and temporal variation from lunar phase and position to identify willingness to pay for darkness, adopting the mobility-data literature’s proposed solutions where they apply directly to this setting and departing from them where necessary.

2 Background

2.1 Measuring Sky Brightness

Astronomers measure sky brightness in magnitudes per square arcsecond ($\text{mag}/\text{arcsec}^2$), a logarithmic unit inherited from the stellar magnitude scale and therefore *inverted*: a larger number means a fainter, darker sky, and a one-unit increase corresponds to a roughly twofold increase in the number of visible stars. This scale runs from roughly 16 $\text{mag}/\text{arcsec}^2$ at its most light-polluted (a bright urban core) a ceiling near 22 $\text{mag}/\text{arcsec}^2$, the brightness of a moonless, pollution-free sky lit only by natural sources. Because of this inversion, I refer throughout the paper’s results to sky darkness rather than brightness, to maintain the intuition that a higher measure is darker. A related but distinct convention used elsewhere in the recreational-astronomy and dark-sky-tourism literature, including Smith et al. (2026), is the Bortle scale (Bortle, 2001), a nine-point ordinal ranking of the naked-eye sky based on how faint an object remains visible. The two scales are highly correlated but not interchangeable, since Bortle class is subjective and requires measurement by an on-site human observer. A final unit is the nanoLambert (nL), similar to $\text{mag}/\text{arcsec}^2$ in being a measure of light intensity per area, but critically being linearly additive. This paper estimates preference for dark skies using in $\text{mag}/\text{arcsec}^2$ at zenith², since it is a continuous physical quantity rather than an ordinal perceptual one, it is the unit in which the underlying remote-sensing data are natively reported, and its logarithmic scale is more representative of human light sensitivity, and thus, perceived quality of the sky.

2.2 Sources of Sky Brightness

Night sky brightness is not stable across locations or time because of the multiple sources of light that contribute. Natural background skyglow comes from airglow³, zodiacal light⁴, and

²The point in the night sky directly overhead the observer

³UV radiation from the sun excites gas molecules in the upper atmosphere which emit light

⁴Natural sunlight being reflected off space dust back to the far side of the Earth

integrated starlight⁵, and is present roughly equally at all locations and times. It measures just over 22 mag/arcsec² (Falchi et al., 2016) and represents the effective cap of how dark the sky can be. Artificial skyglow is light pollution in the colloquial sense, it results from artificial light cast upward which collides with atmospheric particles and scatters across the sky. There is large variation in artificial skyglow across locations, not only due to degree of presence of artificial light, but also due to variance in density of particles to scatter across. Manhattan, the brightest site modeled in this paper, averages around 16.8 mag/arcsec², representing the rough floor of skyglow at which point few to no stars are visible. Moonlight is the last and most variable of the three: unlike the other two components, it changes substantially night to night as the moon moves through its phase cycle, and it also varies meaningfully across location as the moon appears higher or lower in the sky depending on the observer’s line of latitude. The moon naturally varies between 18 mag/arcsec² at its brightest to being a nonfactor while in its new phase, meaning it generates within-location variation for dark places but not for bright ones.

A core assumption of this paper is that these three sources of light are interchangeable. Because I include site fixed effects, variation in skyglow due to moon phase provides the majority of identifying variation. In other words, all welfare estimates are downstream of an assumption that an increase in skyglow on a given night due to moonlight would cause the same damages as if that increase were instead due to artificial light. On one hand, this is largely true from a physics perspective: increased levels of light in any form or color brighten the sky and outshine darker stars, thus obscuring the constellations and the Milky Way. However, it is worth addressing some ways in which they vary. Light pollution generally blocks out the lower parts of the horizon it emanates from: a site with a large city on one side will have more stars visible on the opposite horizon than the near one. The moon, on the other hand, moves across the horizon over the course of the night and blocks out the stars behind and near it at any given moment. Moonlight is generally white in color and feels

⁵The combined light of all the background stars reaching earth

intuitively natural to people, while light from light pollution is generally a dark orange to brownish shade when emanating from traditional streetlamps, and more of a blue-grey hue when cast by modern LEDs. In this paper, I measure skyglow at zenith, looking straight up, during the hours campers are most likely to be stargazing. I assume any effects of sites being more bright in certain parts of the sky are essentially noise, and that the likely preference for moonlight to light pollution makes all my welfare estimates lower bounds.

As mentioned in Section 2 these three components combine additively only when expressed in the linear radiance unit of nanoLamberts. Garstang (1989) proposes the following formulae for converting between the two:

$$L \text{ (cd/m}^2\text{)} = 1.085 \times 10^5 \times 10^{-0.4V}, \quad V = -2.5 \log_{10} \left(\frac{L}{1.085 \times 10^5} \right),$$

where V is the brightness in mag/arcsec² and L is nanoLamberts. Section 3.3 constructs this paper's darkness measure for each location and time by converting every component to nL, summing, and converting the total back to mag/arcsec².

3 Model

I propose a random utility model of recreation demand for campsites. Each period, a resident of Census Block Group (CBG) c chooses among the campsites j available to them or the outside option. I assume the indirect utility individual i from origin c receives from visiting site j in week w is

$$U_{i,c,j,w} = \lambda_{j,y} + \beta D_{j,w} + \rho M_w + \theta (D_{j,w} \times M_w) - \alpha t c_{c,j,y} + \mu_{c,j,w} + \epsilon_{i,c,j,w}$$

where $\lambda_{j,y}$ is a site-by-year fixed effect absorbing any time-invariant-within-year site characteristics, along with year-specific shocks common to that site; $D_{j,w}$ is sky darkness at site

j in week w , M_w is an indicator for weeks containing the peak of a major meteor shower; $tc_{c,j,y}$ is the round-trip travel cost of reaching j from c , varying by year due to changes in per-mile travel cost; and $\mu_{c,j,w}$ is an origin-side, time-varying unobservable. The coefficients of primary interest are β and θ , which capture the value of sky darkness to a camper.

Each origin CBG also has an outside option, taking a value of

$$U_{i,c,0,w} = \gamma_c + \eta_m + \epsilon_{i,c,0,w}$$

where γ_c is an origin-specific fixed effect, and η_m is a calendar-month fixed effect common across all origins, allowing the value of the outside option – and hence the propensity to travel anywhere at all – to move seasonally. I model γ_c and η_m as additively separable rather than allowing a fully flexible $\gamma_{c,m,y}$: while the fully interacted version is a more flexible object, observing visitations only at a weekly level leaves too little variation for estimation upon its inclusion.

Assuming $\epsilon_{i,c,j,w}$ is distributed Type I extreme value and i.i.d. across alternatives, the market share of CBG c 's trips going to site j in week w is

$$s_{c,j,w} = \frac{\exp(\lambda_{j,y} + \beta D_{j,w} + \rho M_w + \theta(D_{j,w} \times M_w) - \alpha tc_{c,j,y} + \mu_{c,j,w})}{\exp(\gamma_c + \eta_m) + \sum_{k \in \mathcal{J}_c} \exp(\lambda_{k,y} + \beta D_{k,w} + \rho M_w + \theta(D_{k,w} \times M_w) - \alpha tc_{c,k,y} + \mu_{c,k,w})}$$

$$s_{c,0,w} = \frac{\exp(\gamma_c + \eta_m)}{\exp(\gamma_c + \eta_m) + \sum_{k \in \mathcal{J}_c} \exp(\lambda_{k,y} + \beta D_{k,w} + \rho M_w + \theta(D_{k,w} \times M_w) - \alpha tc_{c,k,y} + \mu_{c,k,w})}$$

where $\mathcal{J}_c$ is CBG c 's own choice set, defined here as every site it has ever been observed visiting. Taking logs and differencing eliminates the market-specific denominator entirely, in the spirit of Berry (1994):

$$\log \left(\frac{s_{c,j,w}}{s_{c,0,w}} \right) = \lambda_{j,y} + \beta D_{j,w} + \rho M_w + \theta (D_{j,w} \times M_w) - \alpha tc_{c,j,y} - \gamma_c - \eta_m + \mu_{c,j,w}$$

which is estimable by OLS with site-by-year, origin, and month fixed effects, and treating $\mu_{c,j,w}$ as the regression error. I cluster standard errors by origin CBG because $\mu_{c,j,w}$ is likely correlated within CBG due to constant choice set and localized shocks.

3.1 Heterogeneity by Home-Sky Brightness

A natural hypothesis is that preferences for dark skies are reference-dependent: a traveler accustomed to an already-dark home sky may value additional site darkness differently than a traveler accustomed to a bright one. I test this directly using each origin CBG's own home-sky brightness, H_c , interacting it continuously with travel cost and every darkness-related term:

$$\begin{aligned} \log \left(\frac{s_{c,j,w}}{s_{c,0,w}} \right) = & \lambda_{j,y} + (\beta + \beta_H H_c) D_{j,w} + \rho M_w + (\theta + \theta_H H_c) (D_{j,w} \times M_w) \\ & - (\alpha + \alpha_H H_c) tc_{c,j,y} - \gamma_c - \eta_m + \mu_{c,j,w} \end{aligned}$$

I prefer this approach to the random-coefficients route of Berry et al. (1995), as I view it as an unnecessary complexity. The theorized source of heterogeneity is directly observable rather than needing to be recovered from unobserved taste variation, and as argued by Klaiber and von Haefen (2019), including these random coefficients can significantly degrade in-sample predictions. I prefer interacting continuously with H_c over a discrete split (e.g., an urban/rural divide) because it lets the data reveal the shape of the relationship directly rather than imposing a single jump between two groups.

3.2 Zero Shares and the Epsilon Adjustment

As pointed out by Dubé et al. (2021) and Wan et al. (2026), among others, a pervasive issue in recreation demand modeling with cell phone mobility data is that most (origin, site, week) cells report zero visits. This may be due to legitimate non-use or censoring due to privacy protection. A zero share is not directly usable inside a log-share-ratio, so some adjustment is required. Ideally, one would adopt an Empirical Bayes approach in the spirit of Li et al. (2019), later applied by Wan et al. (2026), in which a prior for a cell’s visit count is drawn from other nearby CBGs visiting the same site in the same period and updated with the observed count. This is infeasible here: most sites are unvisited by *any* CBG in a given week, and among sites visited at all in a given week, the overwhelming majority are visited by exactly one CBG, leaving no donor pool to draw a prior from. Cheng and Wan (2026) confirm in simulation that attempting to manufacture a donor pool by expanding the time window instead introduces more bias than the simpler alternative: adding a small constant to every share. Following this second-best approach, for every site in CBG c ’s choice set in week w I construct

$$\tilde{s}_{c,j,w} = \hat{s}_{c,j,w} + \epsilon, \quad \tilde{s}_{c,0,w} = 1 - \sum_{k \in \mathcal{J}_c} \tilde{s}_{c,k,w}$$

with $\epsilon = 10^{-6}$ as the primary specification. Because ϵ is set by convention rather than tuned to the data, every specification in this paper is re-estimated at $\epsilon = 10^{-8}$ as a robustness check (Section A). Generally, the raw regression coefficients themselves rescale with ϵ in a way that has no economic meaning, but implied willingness-to-pay figures are stable across specifications.

3.3 Defining Darkness

Sky darkness $D_{j,w}$ is constructed as the sum, in linear brightness units, of three physically distinct components introduced in Section 2.2: artificial light pollution, natural background

skyglow, and moonlight, converted back to mag/arcsec² in the final step. While background skyglow is set by convention at 54nL, and annual site and origin light pollution is estimated by Lorenz (2025), the brightness of the moon at any given site and week is more difficult.

The standard astronomical model used for measuring the brightness of the moon is Krisciunas and Schaefer (1991), who estimate scattered moonlight brightness at a target point in the sky as:

$$B_{\text{moon}} = f(\rho) I^*(\alpha) 10^{-0.4kX_{\text{ext}}(Z_m)} [1 - 10^{-0.4k}]$$

In which brightness depends on the moon’s phase angle α (degrees from full), its distance from zenith Z_m , the target’s own distance from zenith Z , the moon-to-target angular separation ρ , and an atmospheric extinction coefficient k , which determines the extent of light scattering. Since this paper’s target is always brightness directly above a site, two useful simplifications follow immediately: $\rho = Z_m$, and the scattering airmass to the target, typically estimated as $X_{\text{ext}}(Z) = (1 - 0.96 \sin^2 Z)^{-1/2}$, is evaluated at $Z = 0$ to be exactly 1. The remaining equations are then

$$I^* = 10^{-0.4(3.84+0.026|\alpha|+4\times 10^{-9}\alpha^4)}$$

$$f(\rho) = 10^{5.36} [1.06 + \cos^2 \rho] + 10^{6.15-\rho/40}$$

$I^*(\alpha)$ is the moon’s illuminance outside the atmosphere as a function of phase and $f(\rho)$ is the total scattering function, combining Rayleigh scattering off air molecules with Mie scattering off aerosols. I assume an extinction coefficient of $k = .25$ for all sites, which ignores some degree of variation across region due to altitude and air pollution, but since this parameter is rarely estimated and even then only for single observatories, spatial data are unavailable. More problematic is that my visitation observations are weekly, such that I must estimate one site brightness value per week, despite significant change in moon brightness

within that span. To be as representative as possible, I compute B_{moon} at every hour from one hour after local sunset until local midnight, for every day of the calendar week. The week’s moonlight contribution, $B_{\text{moon},j,w}$, is the median of these hourly values, and is set to zero for any hour at which the moon is below the horizon. Generally, this still captures broad week by week trends in moon phase, and within week across site trends in moon altitude and position.

4 Data

The primary data source is Advan’s Weekly Patterns+ product, a cell phone mobility dataset scaled to estimate population-level visitation to points of interest (POI). I define the universe of sites as every Advan POI carrying NAICS code 721211 (RV parks and campgrounds) located in the continental United States, restricted further to sites for which a sky-darkness measurement is available (Section 3.3); this leaves 6,508 candidate sites. For each site-week, Advan reports estimated visits by visitor home Census Block Group (CBG); I treat this as the raw count of trips from origin c to site j in week w . Following Wan et al. (2026), I treat all non-campground visits reported in each CBG-week as the outside option.

I supplement this mobility data with site and origin characteristics from several sources. Temperature and precipitation at each site-week come from the PRISM Climate Group’s gridded daily weather product, aggregated to weekly site-level means. Sky darkness comes from Lorenz (2025), who uses remote sensing data to estimate specifically the degree of skyglow generated by light pollution across the world for 2020 and 2022. I impute linearly for annual valuations from 2019 and 2021 but they are near-identical regardless. Meteor shower peak dates come from the American Meteor Society’s shower calendars, restricted to each year’s nine to ten officially recognized major showers. Median household income at the CBG level, used to construct the opportunity cost of travel time, is drawn from the Census Bureau; AAA’s annual “Your Driving Costs” reports supply the marginal per-mile

cost of driving. I do not include a measure of historical cloud cover despite its relevance to star visibility both because it varies so much within a week as to be too noisy and because campers plausibly may not anticipate the presence of clouds prior to their decision to make a trip, as they often clear up or appear unpredictably.

4.1 Cleaning and Sample Construction

While Advan’s coverage is excellent, changes in data availability are present. In particular, during November 2022, data ownership was transferred from Safegraph to Advan, and stable trends in total visitation become murky. Rather than risk data construction practices becoming a driver in identification, I include only data before the transfer, resulting in a range covering January 2019 to November 2022.

Additional problems arise due to coverage: certain CBGs are not well-represented either due to low population or below-average phone ownership. Table 1 reports the sample lost at each stage of construction, on the origin (CBG) side. Starting from the full panel of 295,049 distinct origin CBGs, 217,304 are ever observed visiting a campground site; of these, 998 are dropped for having 180 or fewer of the 204 possible weeks of total-mobility coverage (e.g. zero trips are reported for those weeks). A final restriction drops origin CBGs located in Alaska or Hawaii, since every site in those states is already excluded from the site universe (below), and any origin CBGs with an estimated resident population below 100. These CBGs, overwhelmingly airports, generate implausibly large “home” visitation volumes relative to their population, almost certainly reflecting errors in home-location estimation. Among the remaining CBGs, any missing weeks of total mobility are imputed using CBG-by-month, CBG-by-year, and tract-by-month-by-year fixed effects. This leaves a final estimation universe of 214,724 origin CBGs, each entering the panel as its own market as defined in Section 3.

Table 1: Sample construction

Stage	Origin CBGs remaining
Raw national mobility panel	295,049
Observed visiting a campground at least once	217,304
Retain: >180 of 204 weeks of total-mobility coverage	216,306
Drop: Alaska/Hawaii origin	−787
Drop: estimated resident population <100	−795
Final estimation universe	214,724

I flag one identification concern these restrictions raise: dropping low-total-visit CBGs disproportionately removes lower-population and lower-income origins, and the direction of any resulting bias in the darkness and shower coefficients is a priori ambiguous. Lower-income CBGs may disproportionately value stargazing as a cheap recreational opportunity and camp more, or their on-average darker home skies may dissuade them from making a trip at all. As noted by Zhang et al. (2024), these issues are always present in work using cell phone mobility data because of sampling error induced by differences in cell phone ownership and use intensity. Further, on the site side, 48.9% of sites have unreported weeks, not weeks with 0 visitors, which also appear, but simply unreported. I interpret this as the site closing, as it is seasonally correlated and more common at higher latitudes. Because a middle-ground assumption about how to treat these site closures would be arbitrary, I estimate two variants side by side throughout: Variant A restricts the eligible site pool to the 3,328 sites with zero missing weeks over the full window, while Variant B treats every one of the 6,508 continental U.S. sites as eligible and treats a missing week identically to an observed zero-visit week. The two variants return relatively similar results throughout, which I take as evidence that little is lost by not resolving this choice.

The final choice set for each origin CBG is defined, for both variants, as every site that CBG’s residents are ever observed visiting at least once over the full sample window, this realizes 2,312,498 origin-CBG-by-site pairs, each observed for 204 weeks.

4.2 Summary Statistics

Table 2 reports summary statistics for the estimation sample. Sky darkness, constructed as in Section 3.3, ranges from 16.59 to 21.99 mag/arcsec² across site-weeks; the ceiling near 21.99 reflects the natural-background floor discussed in Section 2. Of the 204 weekly analysis windows spanning January 2019 through November 2022, 33 contain the calendar peak of at least one recognized meteor shower.

Table 2: Summary statistics

Candidate sites, Variant A (full-coverage)	3,328
Candidate sites, Variant B (all continental U.S.)	6,508
Final origin CBG universe	214,724
Realized origin-CBG $\times$ site pairs	2,312,498
Sky darkness, mag/arcsec ² (min / mean / max)	16.59 / 20.04 / 21.99
Weeks in sample	204
Weeks containing a meteor shower peak	33
Panel rows, Variant A (full-coverage sites only)	387,359,076
Panel rows, Variant B (all continental U.S. sites)	469,469,892

4.3 Travel Cost

Round-trip travel cost from origin c to site j combines real driving distance and duration, pulled for every realized origin-CBG-by-site pair from the OSRM routing engine, with a per-mile and per-hour dollar cost. It is estimated as

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

The per-mile rate c_y is AAA’s sales-weighted operating cost per mile (fuel, maintenance, tires) plus the marginal cost of depreciation from additional annual mileage, both taken directly from AAA’s published annual reports rather than reconstructed from category averages. The per-hour rate is one-third of the origin CBG’s implied hourly wage (median household income divided by 2,080 annual hours), the standard convention for the opportunity cost of travel time. A floor of one mile and three minutes is applied to every realized

pair before computing cost, to avoid attributing an implausible cost to a site which happens to be near the CBG centroid. All nominal dollar figures are deflated to real 2019 dollars using the BLS Consumer Price Index.

4.4 Robustness Checks

I report every combination of three specification choices: the epsilon convention used to adjust zero shares, $\epsilon \in \{10^{-6}, 10^{-8}\}$ (Section 3.2); the meteor shower definition, either every officially recognized major shower or only the two most prominent annual showers, the Perseids and Geminids; and the site-quality variant, A or B (Section 4.1). Each of these eight combinations is estimated for both specifications discussed in Section 3 for sixteen regressions in total, reported in the appendix.

5 Results

I report two specifications side by side. The baseline specification (Section 5.1) pools all origin CBGs and estimates a single set of preference parameters for the full population; it is the specification I take as the primary result, since it does not rest on any origin-level classification's finer distinctions and is correspondingly the more conservative claim. Section 5.2 then estimates the continuous home-sky-brightness heterogeneity specification of Section 3.1. I report it as a secondary, more exploratory result with that scope caveat attached throughout.

5.1 Baseline Specification

Table 3: Baseline specification: willingness to pay (\$)

	Variant A (full-coverage sites)	Variant B (all continental U.S. sites)
Darkness, $D_{j,w}$ (β)	10.26*** (0.57)	14.79*** (0.54)
Darkness $\times$ shower (θ)	13.33*** (0.97)	10.85*** (0.95)
Meteor shower, M_w (ρ)	-299.52*** (19.17)	-251.46*** (18.81)
Avg. temperature ($^{\circ}\text{C}$)	28.37*** (0.31)	34.70*** (0.35)
Total precipitation (mm)	-0.74*** (0.02)	-0.85*** (0.02)
Site-year, origin, month FE	Yes	Yes
$\hat{\alpha}$ (travel cost)	-0.000154*** (0.0000014)	-0.000136*** (0.0000012)
Panel rows	387,359,076	469,469,892

*** $p < 0.001$. Every WTP entry above is $-\hat{\beta}/\hat{\alpha}$ for that row's own coefficient, with standard errors computed via the delta method from the model's coefficient covariance matrix; $\hat{\alpha}$, the travel-cost coefficient each ratio divides by, is reported directly above

The headline result is that campers are willing to pay an additional \$10-\$15 for an additional unit of darkness, coming up to around \$25 during meteor showers. Values vary slightly between the two variants but are close enough to be economically similar. As reported in the appendix, changing the definition of ϵ shifts the estimates by less than \$2, but using the more conservative definition of meteor showers raises WTP for darkness during them to \$38.53 and \$55.15. The standalone effect of a meteor shower, independent of darkness, is large and negative at around -\$300, but this is partially an artifact of the scale of the darkness measure, which runs between 16 and 22. Accounting for this, WTP for a meteor shower ranges between \$86 and \$6, a surprising but robust result across all specifications, though the range does reach positive when using the conservative definition of meteor showers. Taken together, these estimates evidence that dark skies have value even on a typical night, and that the value of darkness rises further when there is a salient reason to look up.


5.2 Heterogeneity by Home-Sky Brightness and Income

Because I estimate separate values of α by origin income and darkness, reporting results in WTP values would require evaluating $\hat{\alpha}$ at a representative value of both, but in reality, WTP varies significantly across the range. As a result, I report the raw coefficients in Table 4 and a contour plot of WTP for darkness in Figure 1. The variation in WTP is driven almost exclusively by home brightness, counterintuitively, with darker home skies resulting in higher valuation of darker campsites. I theorize three possible causes. First, sorting behavior may pull individuals with greater value of dark skies to both live in dark places and visit darker skies. Second, relative improvement is important, agents with dark skies at home do not need to go camping to see stars, so a site with comparable skies is relatively less desirable. Finally, astronomical sophistication may vary by home darkness: remembering that most of the variation identifying the coefficient on darkness comes from changes in the lunar cycle, agents living under relatively dark skies see every month how much the moon blocks out the stars and may plan trips accordingly. In the reverse direction, people with bright skies at home might prefer camping during a full moon because it provides more similar lighting conditions to what they are used to; true darkness may be uncomfortable or at least inconvenient. Meteor showers improve WTP for darkness across the board because of the strong positive connection between darkness and income during showers. $\hat{\theta}_H$ is insignificant, indicating that home darkness does not meaningfully affect visitation trends during meteor showers, possibly indicating that agents with dark skies at home can observe them locally, while agents in cities are unwilling to pay the significant cost of getting to a dark site to view them.


Table 4: Home-sky brightness and home income, jointly: regression coefficients

	Variant A (full-coverage sites)	Variant B (all continental U.S. sites)
Travel cost (α)	−0.000836*** (0.0000300)	−0.000755*** (0.0000255)
Darkness, $D_{j,w}$ (β)	−0.3261*** (0.0056)	−0.2503*** (0.0047)
Meteor shower, M_w (ρ)	−0.0453*** (0.0029)	−0.0340*** (0.0025)
Avg. temperature ($^{\circ}\text{C}$)	0.00436*** (0.00003)	0.00471*** (0.00003)
Total precipitation (mm)	−0.000111*** (0.000003)	−0.000115*** (0.000002)
Darkness $\times$ shower (θ)	−0.00182*** (0.00033)	−0.00192*** (0.00029)
Travel cost $\times H_c$ (α_H)	−0.0000563*** (0.00000095)	−0.0000493*** (0.00000080)
Travel cost $\times \log(\text{income})$ (α_{li})	0.000156*** (0.000002)	0.000139*** (0.0000017)
Darkness $\times H_c$ (β_H)	0.01735*** (0.00028)	0.01343*** (0.00023)
Darkness $\times \log(\text{income})$ (β_{li})	−0.00163*** (0.00008)	−0.00141*** (0.00007)
Darkness $\times$ shower $\times H_c$ (θ_H)	0.0000047 (0.0000085)	−0.0000173* (0.0000077)
Darkness $\times$ shower $\times \log(\text{income})$ (θ_{li})	0.000335*** (0.000019)	0.000334*** (0.000017)
Site-year, origin, month FE	Yes	Yes
Panel rows	386,356,620	468,278,328

*** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$. Raw regression coefficients and standard errors, clustered by origin CBG – not converted to willingness-to-pay dollars (see Figure 1 for the WTP surface over (H_c, income) instead).



(a) Variant A (full-coverage sites)



(b) Variant B (all continental U.S. sites)

Figure 1: Willingness to pay for site darkness as a function of origin home-sky brightness H_c and home median log-income. Left: the ordinary-week value, $-\left[\hat{\beta} + \hat{\beta}_H H_c + \hat{\beta}_{li} \log(Inc_c)\right] / \left[\hat{\alpha} + \hat{\alpha}_H H_c + \hat{\alpha}_{li} \log(Inc_c)\right]$. Right: the meteor-week value, same as left but including the $\hat{\theta}$ terms. Results are all significant outside a thin band where WTP crosses zero, in both variants.

5.3 Counterfactual Welfare Under Current Pollution Levels

A natural question regarding the slow creep of light pollution is how much damage has already been done. In this section, I estimate how much worse off my model predicts agents would have been had sites in 2019 had the pollution levels they have in 2025. I estimate

this using Variant A of the model in Section 5.1. While a similar estimate could be done using the heterogeneous preference model, the potential for agent sorting in response to new levels of home light pollution and the notion that preferences may themselves be informed by home light pollution makes the plausibility of such a counterfactual questionable.

Inclusive value. The core welfare object is the inclusive value term, economically interpreted as the expected value of the opportunity to make a trip from origin c in week w

$$IV_{c,w} = \frac{-1}{\alpha} E[\max(u_{i,c,0,w}, u_{i,c,1,w}, \dots, u_{i,c,N_c,w})] = \frac{-1}{\alpha} \ln \left(e^{\gamma_c + \eta_m} + \sum_{j \in \mathcal{J}_c} e^{V_{c,j,w}} \right).$$

In order to perform counterfactual analysis I simply construct the term using alternate values of $D_{j,w}$ and compute the difference.

The counterfactual. For each of the 3,328 Variant A sites, I take the most recent measure of light pollution available (2025), then combine that pollution level with the site’s actual 2019 background and week-by-week moonlight as in Section 3.3 to obtain a counterfactual $D_{j,w}^{cf}$. Between 2019-2025, the average site saw a decrease in darkness of $-0.102 \text{ mag/arcsec}^2$, with 2,922 of 3,328 brighter than they started.

From a representative trip to total welfare. $IV_{c,w}$ is the expected utility of a single representative trip-allocation decision for CBG c in week w . Aggregating to total welfare change thus requires us to sum across all opportunities to take a trip in every CBG and week for 2019. The dollar-denominated total annual welfare change is then

$$\Delta W = -\frac{1}{\hat{\alpha}} \sum_c \sum_{w \in 2019} N_{c,w} (IV_{c,w}^{cf} - IV_{c,w}),$$

with $N_{c,w}$ taken from the same total-mobility product used to construct the shares themselves. Results are shown in Table 5, the median CBG sees a loss of roughly \$500 in surplus, heavily right skewed with an average of \$4000. In aggregate, just the additional light pollu-

tion incurred between 2019 and 2025 causes annual recreation damages around \$830 million in 2019 dollars.

Table 5: Counterfactual welfare, 2019 trips under 2025-measured pollution

Mean site darkness change, 2019 $\rightarrow$ 2025 (mag/arcsec ²)	−0.102
Sites brighter / darker (of 3,328)	2,922 / 397
Median annual ΔW_c per CBG	−\$509
Mean annual ΔW_c per CBG	−\$3,949
CBGs worse off / better off (of 211,113)	190,551 / 20,338
Aggregate annual welfare change	−\$833,775,059

6 Conclusion

I find evidence of modest but significant willingness to pay for dark skies among camping recreation trips. While my estimates vary across specification, I generally find a willingness to pay between \$10 and \$15 for an additional unit of darkness, rising to roughly \$25 during meteor showers. I find further evidence of significant heterogeneity in this value, with people living in dark areas valuing dark skies in their campsites much more than those living with bright skies, who actively seek to avoid dark sites.

Using the results of this model, I find that light pollution incurred since 2019 costs \$830 million in recreation damages annually, despite being a relatively marginal increase in pollution at the typical site. Simple abatement procedures such as installing shielded fixtures typically pay for themselves within years by reducing municipal electrical costs (Buckley), so while the damages to recreation may be modest, they are also easy to limit.

As I stated in Section 2, the primary identifying variation in darkness is due to moon cycle, so a core assumption of my paper is that aversion to moonlight and light pollution are similar. In reality, I expect aversion to moonlight to be less, and thus treat the welfare estimates of this paper to be a lower bound on damages to recreation.

I stress that this paper does not represent a holistic valuation of the damages from light pollution: lost productivity due to poor sleep, lost ecosystem services from animals

threatened by bright lights, possible health risks, climate damages from electricity resulting in wasted light, all contribute to social costs. Nonetheless, recreational stargazing is likely the most salient of these to the general public.

Future research may gain insight into the valuation of sky brightness at the location of residence by taking a hedonic approach. As I document, there is possible evidence for sorting based on dark skies. This approach would require careful controls for possible confounders, as skyglow is closely correlated with many metrics such as productivity, remoteness, crime, neighborhood makeup, and so on. Another approach may stick with the recreation demand study but attempt a better solution to the zero-market share problem. As mentioned, improved approaches have been proposed and implemented, but in my nationwide setting the data is simply too sparse to form priors over the distribution of site visitation. Finally, future work may utilize the results of policy shocks from jurisdictional light pollution laws as identifying variation rather than relying on moonlight directly. This would require confirmation that such laws meaningfully impact the level of local light pollution, which would be feasible using time series remote-sensing data.

A Robustness to the Epsilon Convention

Table 6 and Table 7 re-estimate the baseline specification of Section 5 and the continuous home-brightness/income heterogeneity specification of Section 3.1, respectively, at $\epsilon = 10^{-8}$ in place of the $\epsilon = 10^{-6}$ used throughout the main text (Section 3.2), holding the full meteor-shower calendar fixed. In the heterogeneity specification, every coefficient rescales by roughly the same factor (around $1.55\text{--}1.65\times$) moving from $\epsilon = 10^{-6}$ to $\epsilon = 10^{-8}$, consistent with the baseline specification’s own ϵ -dependence and with the WTP ratio itself being invariant to this rescaling. Signs and significance are stable throughout with one exception: $\hat{\theta}_H$, the home-brightness moderation of the darkness $\times$ shower interaction, is statistically indistinguishable from zero at both ϵ values in Variant A ($t = 0.55$ and $t = 1.71$), but in Variant B it is

significantly negative at $\epsilon = 10^{-6}$ ($t = -2.24$) and indistinguishable from zero at $\epsilon = 10^{-8}$ ($t = -0.99$) – the only coefficient in either table whose significance is sensitive to this choice.

Table 6: Epsilon robustness, baseline specification: WTP (\$)

	Variant A		Variant B	
	$\epsilon = 10^{-6}$	$\epsilon = 10^{-8}$	$\epsilon = 10^{-6}$	$\epsilon = 10^{-8}$
Darkness	10.26*** (0.57)	9.40*** (0.55)	14.79*** (0.54)	13.87*** (0.52)
Darkness $\times$ shower	13.33*** (0.97)	14.17*** (0.95)	10.85*** (0.95)	11.73*** (0.93)
Meteor shower (standalone)	-299.52*** (19.17)	-316.28*** (18.81)	-251.46*** (18.81)	-269.20*** (18.46)

*** $p < 0.001$. Each entry is $-\hat{\beta}/\hat{\alpha}$ for that row's own coefficient, with the delta-method SE in parentheses

Table 7: Epsilon robustness, continuous home-brightness/income heterogeneity specification: regression coefficients

	Variant A		Variant B	
	$\epsilon = 10^{-6}$	$\epsilon = 10^{-8}$	$\epsilon = 10^{-6}$	$\epsilon = 10^{-8}$
Travel cost (α)	-0.000836*** (0.000030)	-0.001353*** (0.000047)	-0.000755*** (0.000026)	-0.001221*** (0.000040)
Darkness, $D_{j,w}$ (β)	-0.3261*** (0.0056)	-0.5071*** (0.0085)	-0.2503*** (0.0047)	-0.3901*** (0.0071)
Meteor shower, M_w (ρ)	-0.0453*** (0.0029)	-0.0741*** (0.0044)	-0.0340*** (0.0025)	-0.0561*** (0.0039)
Avg. temperature ($^{\circ}\text{C}$)	0.00436*** (0.00003)	0.00681*** (0.00005)	0.00471*** (0.00003)	0.00732*** (0.00004)
Total precipitation (mm)	-0.000111*** (0.000003)	-0.000174*** (0.000004)	-0.000115*** (0.000002)	-0.000179*** (0.000004)
Darkness $\times$ shower (θ)	-0.00182*** (0.00033)	-0.00282*** (0.00050)	-0.00192*** (0.00029)	-0.00303*** (0.00045)
Travel cost $\times H_c$ (α_H)	-0.00005632*** (0.00000095)	-0.00008451*** (0.00000146)	-0.00004933*** (0.00000080)	-0.00007402*** (0.00000122)
Travel cost $\times \log(\text{income})$ (α_{li})	0.000156*** (0.000002)	0.000242*** (0.000003)	0.000139*** (0.000002)	0.000215*** (0.000003)
Darkness $\times H_c$ (β_H)	0.01735*** (0.00028)	0.02695*** (0.00043)	0.01343*** (0.00023)	0.02090*** (0.00035)
Darkness $\times \log(\text{income})$ (β_{li})	-0.00163*** (0.00008)	-0.00249*** (0.00012)	-0.00141*** (0.00007)	-0.00215*** (0.00011)
Darkness $\times$ shower $\times H_c$ (θ_H)	0.0000047 (0.0000085)	0.0000216 (0.0000127)	-0.0000173* (0.0000077)	-0.0000113 (0.0000114)
Darkness $\times$ shower $\times \log(\text{income})$ (θ_{li})	0.000335*** (0.000019)	0.000510*** (0.000029)	0.000334*** (0.000017)	0.000509*** (0.000026)

*** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$. Raw regression coefficients, clustered by origin CBG, not converted to willingness-to-pay dollars; see Section 5.2

B Robustness to the Meteor Shower Definition

Table 8 and Table 9 re-estimate the same two specifications (the latter again the continuous home-brightness/income heterogeneity specification; see Appendix A) restricting the meteor-shower indicator to only the two most prominent annual showers, the Perseids and Geminids, in place of the full calendar of recognized major showers used throughout the main text (Section 4.4), holding $\epsilon = 10^{-6}$ fixed. As in the baseline specification, the darkness $\times$ shower interaction and the standalone shower effect move substantially under this redefinition (the former roughly sevenfold in both variants, the latter roughly doubling), while every other coefficient shifts only slightly. The one coefficient this redefinition destabilizes outright is $\hat{\theta}_H$: insignificant in Variant A ($t = 0.55$) and significantly negative in Variant B ($t = -2.24$) under the full shower calendar (Table 7), it becomes large, positive, and highly significant in both variants under the two-shower definition ($t = 27.06$ and $t = 28.45$) – a full sign flip in Variant B, and the only coefficient in either table whose sign is sensitive to this choice.

Table 8: Shower-definition robustness, baseline specification: WTP (\$)

	Variant A		Variant B	
	All showers	Major showers only	All showers	Major showers only
Darkness	10.26*** (0.57)	12.83*** (0.54)	14.79*** (0.54)	16.20*** (0.51)
Darkness $\times$ shower	13.33*** (0.97)	38.53*** (2.06)	10.85*** (0.95)	55.15*** (2.07)
Meteor shower (standalone)	-299.52*** (19.17)	-726.00*** (39.73)	-251.46*** (18.81)	-1048.29*** (40.01)

*** $p < 0.001$. Each entry is $-\hat{\beta}/\hat{\alpha}$ for that row's own coefficient, with the delta-method SE in parentheses; see Section 5.1

Table 9: Shower-definition robustness, continuous home-brightness/income heterogeneity specification: regression coefficients

	Variant A		Variant B	
	All showers	Major showers only	All showers	Major showers only
Travel cost (α)	−0.000836*** (0.000030)	−0.000836*** (0.000030)	−0.000755*** (0.000026)	−0.000756*** (0.000026)
Darkness, $D_{j,w}$ (β)	−0.3261*** (0.0056)	−0.3271*** (0.0056)	−0.2503*** (0.0047)	−0.2516*** (0.0047)
Meteor shower, M_w (ρ)	−0.0453*** (0.0029)	−0.0888*** (0.0059)	−0.0340*** (0.0025)	−0.1200*** (0.0052)
Avg. temperature (°C)	0.00436*** (0.00003)	0.00441*** (0.00003)	0.00471*** (0.00003)	0.00476*** (0.00003)
Total precipitation (mm)	−0.000111*** (0.000003)	−0.000110*** (0.000003)	−0.000115*** (0.000002)	−0.000113*** (0.000002)
Darkness $\times$ shower (θ)	−0.00182*** (0.00033)	−0.01254*** (0.00068)	−0.00192*** (0.00029)	−0.01040*** (0.00062)
Travel cost $\times H_c$ (α_H)	−0.00005632*** (0.00000095)	−0.00005631*** (0.00000095)	−0.00004933*** (0.00000080)	−0.00004932*** (0.00000080)
Travel cost $\times \log(\text{income})$ (α_{li})	0.000156*** (0.000002)	0.000156*** (0.000002)	0.000139*** (0.000002)	0.000139*** (0.000002)
Darkness $\times H_c$ (β_H)	0.01735*** (0.00028)	0.01741*** (0.00028)	0.01343*** (0.00023)	0.01349*** (0.00023)
Darkness $\times \log(\text{income})$ (β_{li})	−0.00163*** (0.00008)	−0.00160*** (0.00008)	−0.00141*** (0.00007)	−0.00138*** (0.00007)
Darkness $\times$ shower $\times H_c$ (θ_H)	0.0000047 (0.0000085)	0.0004848*** (0.0000179)	−0.0000173* (0.0000077)	0.0004816*** (0.0000169)
Darkness $\times$ shower $\times \log(\text{income})$ (θ_{li})	0.000335*** (0.000019)	0.000684*** (0.000040)	0.000334*** (0.000017)	0.000640*** (0.000038)

*** p<0.001, ** p<0.01, * p<0.05. Raw regression coefficients, clustered by origin CBG, not converted to willingness-to-pay dollars; see Section 5.2

References

- Susan Athey, David Blei, Robert Donnelly, Francisco Ruiz, and Tobias Schmidt. Estimating heterogeneous consumer preferences for restaurants and travel time using mobile location data. In *AEA Papers and Proceedings*, volume 108, pages 64–67. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203, 2018.
- Steven Berry, James Levinsohn, and Ariel Pakes. Automobile prices in market equilibrium. *Econometrica*, 63(4):841–890, 1995.
- Steven T Berry. Estimating discrete-choice models of product differentiation. *The RAND Journal of Economics*, 25(2):242–262, 1994.

- John E Bortle. Introducing the bortle dark-sky scale. *Sky & Telescope*, 101(2):126, 2001.
- Cara Buckley. Pennsylvania, pittsburgh, and light pollution.
- Nieyan Cheng and Xibo Wan. Causal recreation demand estimation with cell phone mobility data: Evidence from the 2021 huntington beach oil spill. *Land Economics*, 102(3):339–360, 2026.
- Jean-Pierre Dubé, Ali Hortaçsu, and Joonhwi Joo. Random-coefficients logit demand estimation with zero-valued market shares. *Marketing Science*, 40(4):637–660, 2021.
- Mark Eiswerth, Liam Mullen, and Guanyi Yang. Contingent behavior modeling for dark skies valuation at great sand dunes national park. *American Journal of Agricultural Economics*, 2026.
- Fabio Falchi, Pierantonio Cinzano, Dan Duriscoe, Christopher CM Kyba, Christopher D Elvidge, Kimberly Baugh, Boris A Portnov, Nataliya A Rybnikova, and Riccardo Furgoni. The new world atlas of artificial night sky brightness. *Science advances*, 2(6):e1600377, 2016.
- Terrel Gallaway, Reed N Olsen, and David M Mitchell. The economics of global light pollution. *Ecological economics*, 69(3):658–665, 2010.
- RH Garstang. Night sky brightness at observatories and sites. *Publications of the Astronomical Society of the Pacific*, 101(637):306–329, 1989.
- John Gibson, Susan Olivia, Geua Boe-Gibson, and Chao Li. Which night lights data should we use in economics, and where? *Journal of Development Economics*, 149:102602, 2021.
- J. Vernon Henderson, Adam Storeygard, and David N. Weil. Measuring economic growth from outer space. *American Economic Review*, 102(2):994–1028, April 2012. doi: 10.1257/aer.102.2.994. URL <https://www.aeaweb.org/articles?id=10.1257/aer.102.2.994>.

- H. Allen Klaiber and Roger H von Haefen. Do random coefficients and alternative specific constants improve policy analysis? an empirical investigation of model fit and prediction. *Environmental and Resource Economics*, 73(1):75–91, 2019.
- Kevin Krisciunas and Bradley E Schaefer. A model of the brightness of moonlight. *Publications of the Astronomical Society of the Pacific*, 103:1033–1039, 1991. doi: 10.1086/132921.
- Christopher CM Kyba, Yiğit Öner Altıntaş, Constance E Walker, and Mark Newhouse. Citizen scientists report global rapid reductions in the visibility of stars from 2011 to 2022. *Science*, 379(6629):265–268, 2023.
- Jing Li et al. Compatibility and investment in the us electric vehicle market. *Unpublished manuscript, MIT*, 2019.
- David J. Lorenz. Light pollution atlas. <https://djlorenz.github.io/astronomy/lp/>, 2025. Accessed: 2026-10-09.
- David Mitchell and Terrel Gallaway. Dark sky tourism: economic impacts on the colorado plateau economy, usa. *Tourism Review*, 74(4):930–942, 2019.
- David M Mitchell, Terrel Gallaway, and Reed Olsen. Estimating the willingness to pay for dark skies. *International Journal of Research in Engineering and Technology*, 2017.
- Stephen C Newbold, Sarah Lindley, Shannon Albeke, Joshua Viers, George Parsons, and Robert Johnston. Valuing satellite data for harmful algal bloom early warning systems. *Resources for the Future*, (22-23), 2022.
- Reed N Olsen, Terrel Gallaway, and David Mitchell. Modelling us light pollution. *Journal of Environmental Planning and Management*, 57(6):883–903, 2014.
- Stephanie N Simpson and Brid Gleeson Hanna. Willingness to pay for a clear night sky: Use of the contingent valuation method. *Applied Economics Letters*, 17(11):1095–1103, 2010.

- Jordan W Smith, J Adam Beeco, Sharolyn J Anderson, Panagiotis Tsigaris, Zachary D Miller, Emily J Wilkins, Anna B Miller, Chase Lamborn, and John C Barentine. Artificial light at night erodes cultural ecosystem services at scale. *Available at SSRN 6642558*, 2026.
- Xibo Wan, Yongwang Ren, Ruohao Zhang, Jiameng Zheng, and Wendong Zhang. Quantifying recreational benefits from fish consumption advisories: Insights from cell phone mobility data. *Land Economics*, 102(3):313–338, 2026.
- Wendong Zhang, Xibo Wan, Wenran Fan, and Yongjie Ji. Uncovering disparities in water-based outdoor recreation using cell phone mobility data. *Environmental Research Letters*, 19(11):114057, 2024.