

The Visible and the Invisible: Traffic Exposure, Political Saliency, and Bridge Maintenance Quality

APEP Autonomous Research*(cumulative: 13h 39m).

February 20, 2026

Abstract

Do politicians maintain what voters can see? Using the universe of U.S. bridge inspections—621,000 structures observed annually from 2000 to 2023—I test whether traffic exposure, a proxy for political visibility, causally affects infrastructure maintenance quality. Three predictions of a political visibility model all fail: high-traffic bridges do not deteriorate more slowly than observably identical low-traffic bridges; the null effect is uniform across visible (deck) and invisible (substructure) components; and repair rates show no differential response to gubernatorial election cycles. The doubly robust estimate is negative, suggesting traffic causes wear rather than attracting maintenance. These null results are informative: engineering-based bridge management systems appear to allocate maintenance according to structural need rather than political saliency, contrasting with [Olken \(2007\)](#)'s finding that monitoring matters in developing-country infrastructure.

JEL Codes: H54, H76, D72, R42

Keywords: infrastructure, bridge safety, political visibility, monitoring, maintenance, political budget cycles

*Autonomous Policy Evaluation Project. Contributor: @olafdrw. Correspondence: scl@econ.uzh.ch

1. Introduction

America’s bridges are crumbling, but they are not crumbling equally. Of the 621,000 highway bridges inspected by the Federal Highway Administration each year, roughly 11 percent are classified as “structurally deficient”—a designation that signals advanced deterioration and the need for significant repair or replacement ([American Society of Civil Engineers, 2021](#)). Yet this aggregate statistic masks enormous variation. Some bridges are maintained meticulously; others deteriorate for decades without intervention. What determines which bridges get fixed?

The engineering answer emphasizes physical factors: age, material, climate, and design loads. The economics answer points to budget constraints: states allocate scarce maintenance dollars across competing priorities. But a third possibility, drawn from the political economy of monitoring, suggests that the answer may be simpler and more troubling: politicians fix what voters can see.

This paper tests that hypothesis. I exploit a remarkable dataset—the National Bridge Inventory (NBI), which records detailed condition ratings for every public highway bridge in the United States every year—to ask whether traffic exposure, a natural proxy for political visibility, causally affects maintenance quality. The core insight is that a bridge carrying 50,000 vehicles per day is not just economically important; it is seen by 50,000 potential voters, every one of whom might complain if the deck is potholed or the guardrails are rusted. A bridge carrying 500 vehicles per day may be equally important to its rural community, but it is invisible to the political process.

My identification strategy compares bridges with high and low initial traffic volumes within the same state, conditional on engineering characteristics—age, material type, span length, total length, number of spans, and functional classification. The key requirement is selection on observables: conditional on these engineering determinants, variation in traffic exposure reflects differences in political visibility rather than differences in physical need. I estimate doubly robust (AIPW) effects that combine flexible machine learning models for both the propensity score and outcome regression, with 5-fold cross-fitting to avoid overfitting bias ([Chernozhukov et al., 2018](#); [Robins et al., 1994](#)).

The headline result is a precisely estimated null. High-traffic bridges (top tercile of initial ADT within state) experience virtually identical annual deck condition changes to low-traffic bridges with identical engineering characteristics. In the preferred specification with state-by-year and material fixed effects, the coefficient on the high-ADT indicator is 0.001 with a standard error of 0.006—statistically indistinguishable from zero and economically trivial. The doubly robust (AIPW) estimate is slightly negative (-0.022 , $p < 0.05$), suggesting

that if anything, heavy traffic accelerates deterioration rather than attracting compensatory maintenance.

Three falsification tests designed to isolate a political mechanism all yield nulls. First, I decompose the effect by bridge component. The deck—the surface drivers see and feel—shows no visibility premium (0.001, SE = 0.006). Neither does the superstructure (−0.000) or the substructure (0.005). The component gradient predicted by the political visibility model is absent: there is no differential treatment of visible versus invisible bridge elements.

Second, I test for an “electoral maintenance cycle.” If the visibility mechanism operates through political incentives, repairs should concentrate before elections. The interaction of high-ADT status with election windows is small (0.003, SE = 0.002) and statistically insignificant. Low-traffic bridges show a similarly null electoral pattern. There is no evidence that politicians target visible bridges before elections.

Third, the raw differences between high- and low-traffic bridges in descriptive statistics—high-ADT bridges have slightly lower deficiency rates (7.3% vs. 15.5%) and marginally slower deterioration (−0.044 vs. −0.051 annual rating change)—are fully explained by observable engineering characteristics, particularly bridge age and material type. Once these are controlled, the visibility “premium” vanishes.

This paper makes three contributions. First, it provides the first systematic test of political determinants of infrastructure maintenance quality in the United States, using bridge-level panel data with over 5 million observations. The infrastructure policy literature has focused almost exclusively on new construction ([Leduc and Wilson, 2013](#); [Knight, 2015](#)) while neglecting what happens to existing assets. Yet maintenance accounts for 40 percent of highway spending, and deferred maintenance is the primary driver of the infrastructure crisis ([American Society of Civil Engineers, 2021](#)). Understanding what determines maintenance quality—and, equally important, what does not—is a first-order policy question.

Second, it provides a boundary condition for the monitoring hypothesis of [Olken \(2007\)](#). Olken showed that government audits improve road quality in Indonesia, where institutions are weak and corruption is endemic. I find that in the U.S. context, with mature regulatory institutions and standardized inspection protocols, informal voter monitoring does not detectably influence maintenance allocation. The “visible vs. invisible component” test—comparing maintenance of bridge decks (visible to drivers) with substructures (hidden underwater)—provides a novel within-asset design for testing monitoring effects. The null result on this test is informative: it suggests that the FHWA’s National Bridge Inspection Standards and state-level bridge management systems effectively insulate maintenance decisions from political pressure.

Third, it contributes to the political budget cycles literature. While existing work

on political budget cycles typically focuses on aggregate spending (Rogoff, 1990; Shi and Svensson, 2006), I look at the physical condition of individual infrastructure assets, testing whether electoral incentives affect what governments maintain. The absence of an electoral maintenance cycle in bridge repairs—even for the most politically visible bridges—suggests that infrastructure maintenance is largely depoliticized in the United States.

The findings carry a reassuring policy implication. Unlike road construction, which responds to political incentives (Knight, 2015; Burgess et al., 2015), bridge maintenance appears to be allocated according to engineering need. The nationwide bridge inspection mandate, combined with the “structurally deficient” threshold that triggers federal funding eligibility, may create institutional constraints that override political distortions. The paper discusses why bridges may differ from other infrastructure assets and what lessons the bridge inspection model offers for other domains of public investment.

2. Institutional Background

2.1 The National Bridge Inspection Program

The collapse of the Silver Bridge over the Ohio River in 1967, which killed 46 people, prompted Congress to establish the National Bridge Inspection Standards (NBIS) in 1971. Under 23 CFR 650, every bridge on a public road with a span exceeding 20 feet must be inspected at least once every 24 months. Inspections are conducted by state departments of transportation (DOTs) or their contractors, using a standardized rating system developed by the Federal Highway Administration (FHWA).

The inspection produces condition ratings on a 0–9 scale for three primary components: the deck (the driving surface), the superstructure (the main load-carrying members above the supports), and the substructure (piers, abutments, and foundations below the deck). A rating of 9 indicates “excellent” condition with no problems noted. A rating of 7 indicates “good” condition requiring only minor maintenance. A rating of 5 indicates “fair” condition with minor deterioration. A rating of 4 or below triggers classification as “structurally deficient,” which makes the bridge eligible for additional federal repair funding under the Highway Bridge Program.

The distinction between these three components is central to my identification strategy. The deck is directly visible to every driver who crosses the bridge: potholes, cracking, spalling, and drainage problems are immediately apparent. The superstructure—beams, girders, and trusses—is partially visible, particularly for steel truss bridges, but requires more careful observation. The substructure is typically hidden: piers are often underwater, abutments are embedded in embankments, and foundations are underground. If maintenance quality

responds to political visibility, the effect should be strongest for the most visible component (the deck) and weakest for the least visible component (the substructure).

2.2 Maintenance Decision-Making

Bridge maintenance in the United States is a state responsibility. The federal government provides funding through the Highway Trust Fund—approximately \$40 billion per year in highway formula grants—but states have broad discretion over how to allocate these funds across their bridge inventories. Federal formula grants are apportioned to states based on lane miles, vehicle miles traveled, and other factors, but the within-state allocation to individual bridges is determined by each state DOT’s asset management plan.

State DOTs employ bridge management systems (BMS) that use engineering models to predict deterioration and prioritize maintenance investments. The standard optimization criterion is minimizing lifecycle costs: bridges should be repaired when the cost of deferred maintenance exceeds the cost of timely intervention. In principle, these systems allocate maintenance based on engineering need, not political considerations.

In practice, however, multiple studies have documented political influence over transportation spending. [Knight \(2015\)](#) shows that congressional earmarks direct highway funds toward politically connected districts. [Harding \(2015\)](#) finds that road quality affects gubernatorial electoral outcomes. [Burgess et al. \(2015\)](#) demonstrates that road construction in Kenya responds to ethnic favoritism. The question is whether similar forces shape the within-state allocation of bridge maintenance—a question that has not been tested with bridge-level data.

2.3 Governor Elections and Infrastructure

Governors are the chief executives of state governments and typically exercise significant influence over DOT priorities. In 48 of 50 states, governors serve four-year terms; New Hampshire and Vermont have two-year terms. Term limits vary: 36 states impose some form of term limit, while 14 allow unlimited terms.

The gubernatorial election calendar creates time variation in political incentives. In most states, elections occur in even-numbered years, with roughly half holding elections in midterm years and the remainder in presidential election years. Five states—Kentucky, Louisiana, Mississippi, New Jersey, and Virginia—hold gubernatorial elections in odd-numbered years. This staggered timing, combined with cross-state variation in electoral competitiveness and term limits, provides the basis for testing whether electoral incentives shape maintenance allocation.

3. Conceptual Framework

Consider a state DOT that must allocate a fixed maintenance budget M across N bridges indexed by i . Each bridge has an observable traffic volume v_i and a maintenance need n_i that depends on the bridge’s age, condition, and engineering characteristics. The DOT chooses maintenance effort m_i for each bridge.

Engineering benchmark. An engineering-optimal allocation minimizes total deterioration cost:

$$\min_{\{m_i\}} \sum_{i=1}^N C_i(n_i, m_i) \quad \text{s.t.} \quad \sum_{i=1}^N m_i \leq M \quad (1)$$

where $C_i(\cdot)$ is the lifecycle cost function. The optimal allocation depends only on n_i , not on v_i —traffic volume should be irrelevant conditional on engineering need.

Political distortion. Now suppose the DOT director serves at the pleasure of the governor, who faces re-election. The governor’s electoral payoff depends on voter perceptions of infrastructure quality. Voters observe bridge quality with probability proportional to traffic volume: a bridge with $v_i = 50,000$ is observed by 100 times as many voters as one with $v_i = 500$. The governor’s problem becomes:

$$\max_{\{m_i\}} \sum_{i=1}^N v_i \cdot q_i(m_i) - \lambda \sum_{i=1}^N m_i \quad (2)$$

where $q_i(m_i)$ is the quality of bridge i as a function of maintenance effort, and λ is the shadow cost of the budget constraint. The first-order condition implies $m_i^* \propto v_i$: maintenance is allocated in proportion to political visibility, not engineering need.

This framework generates three testable predictions:

Prediction 1 (Visibility Premium): Conditional on engineering characteristics, high-traffic bridges should deteriorate more slowly than low-traffic bridges.

Prediction 2 (Component Gradient): The visibility premium should be strongest for the most visible component (deck) and weakest for the least visible component (substructure).

Prediction 3 (Electoral Cycle): The visibility premium in repair rates should spike before elections, when the governor’s marginal return to visible maintenance is highest.

4. Data

4.1 National Bridge Inventory

The primary data source is the National Bridge Inventory (NBI), maintained by the FHWA. The NBI contains annual inspection records for every public highway bridge in the United States with a span exceeding 20 feet. I use data from 2000 through 2023, comprising approximately 621,000 bridges per year in the raw NBI files.

For each bridge, the NBI records: (1) condition ratings for the deck, superstructure, and substructure on a 0–9 scale; (2) Average Daily Traffic (ADT), the estimated number of vehicles crossing the bridge per day; (3) detailed engineering characteristics including year built, material type, design load, span length, total length, and number of spans; (4) geographic identifiers including state and county FIPS codes and latitude/longitude coordinates; and (5) functional classification indicating whether the bridge is on an Interstate, US route, state highway, or local road.

I restrict the sample to highway bridges (service type = highway), exclude bridges in the District of Columbia (no governor) and U.S. territories, and drop observations with missing condition ratings or zero reported ADT. The resulting analysis sample contains approximately 5.4 million bridge-year observations across 545,000 unique bridges in 49 states, covering the period 2001–2023 (year-over-year changes require lagging, so the effective sample begins in 2001).

Temporal coverage and data quality. The NBI has been collected continuously since 1992, but data quality improved substantially after the FHWA standardized electronic reporting formats in the late 1990s. I begin the panel in 2000 to ensure consistent variable coding across years. The data files are publicly available from the FHWA website in annual ASCII format. Earlier years (2000–2016) use a different file naming convention than recent years (2018–2023), requiring separate download procedures, but the underlying variable definitions are consistent throughout.

The panel is strongly balanced: most bridges appear in every year of the 24-year panel. Attrition occurs primarily through bridge demolition (approximately 2,000 bridges per year are removed from the inventory, typically replaced by new structures with new identification numbers). New entries reflect both new construction and the reclassification of existing structures. I use the bridge’s unique NBI structure number as the panel identifier, which is stable across years for existing bridges.

Condition rating dynamics. The 0–9 condition scale exhibits interesting dynamics. The modal rating is 7 (“good condition—some minor problems noted”), and the distribution is left-skewed with a concentration between 5 and 8. Transitions between adjacent ratings are common: approximately 15% of bridge-year observations show a one-point change in deck condition (either improvement or deterioration). Large jumps of 2+ points are rare (approximately 5% of observations) and typically indicate major rehabilitation or reconstruction events. The mean annual change of -0.05 reflects a slow, steady deterioration punctuated by occasional repair events that partially reset the condition clock.

4.2 Key Variables

Outcome variables. The primary outcome is the annual change in deck condition rating ($\Delta C_{it} = C_{it} - C_{it-1}$), which captures the net effect of deterioration and maintenance. I also examine annual changes in superstructure and substructure condition, and a binary indicator for “repair events” (condition improvement of 2 or more points in any component, indicating significant investment).

Treatment variable. The treatment is “high traffic exposure,” defined as the top tercile of initial ADT within each state. I use initial ADT—the ADT recorded in the bridge’s first year of observation—rather than contemporaneous ADT to avoid reverse causality from maintenance-related traffic diversions. The tercile classification is computed within state to account for differences in state-level traffic distributions. As a robustness check, I also use continuous $\log(\text{ADT})$ and alternative cutoffs (median split, top quartile).

Engineering covariates. I control for bridge age (and age squared), total length, number of spans, maximum span length, material type (concrete, steel, prestressed, wood, other), and functional classification. These variables capture the primary engineering determinants of bridge deterioration and maintenance need.

Electoral variables. Governor election years are coded from public records. I construct an “election window” indicator equal to one in gubernatorial election years and the year immediately preceding, when politicians face the strongest incentive to deliver visible improvements.

4.3 Summary Statistics

Table 1: Summary Statistics

Variable	Mean	Std. Dev.	Min	Max
Deck Condition (0–9)	6.61	1.21	0.0	9.0
Superstructure Condition (0–9)	6.68	1.27	0.0	9.0
Substructure Condition (0–9)	6.53	1.31	0.0	9.0
Average Daily Traffic	6,973	20,779	1	999,999
Bridge Age (years)	41.60	25.17	0.0	214.0
Total Length (meters)	49.82	176.34	0.3	38791.6
Number of Spans	2.72	5.32	1.0	900.0
Max Span Length (meters)	16.82	14.31	0.1	1999.8
Structurally Deficient	11.1%	31.4	0	100
Repair Event	4.7%	21.1	0	100
Annual Deck Condition Change	-0.05	0.71	-5.0	5.0

Notes: N = 5,372,604 bridge-year observations across 545,162 unique bridges in 49 states, 2001–2023. Condition ratings follow the FHWA 0–9 scale (9 = Excellent, 4 = Poor threshold for structural deficiency). Repair Event indicates a condition improvement of 2+ points in any component.

Table 1 presents summary statistics for the analysis sample. The mean deck condition rating is approximately 6.6, indicating “satisfactory” condition on average. Mean ADT is approximately 6,973 vehicles per day, though the distribution is highly right-skewed: median ADT is much lower, while a small number of urban Interstate bridges carry over 200,000 vehicles daily. Approximately 11 percent of bridge-year observations are classified as structurally deficient. The mean annual change in deck condition is slightly negative, reflecting the predominance of gradual deterioration over repair events.

Table 2: Bridge Characteristics by Traffic Exposure

	Initial ADT Tercile		
	Low	Medium	High
Observations	1,885,065	1,698,671	1,788,868
Mean ADT	743	2,157	18,113
Mean Deck Condition	6.57	6.66	6.60
Pct. Structurally Deficient	15.5	10.2	7.3
Mean Age (years)	42.9	42.2	39.8
Repair Rate (%)	4.70	4.47	4.79
Mean Annual Δ Deck	-0.053	-0.050	-0.046

Notes: Terciles computed within state using initial (first-observed) ADT. Repair Rate is the percentage of bridge-year observations with a condition improvement of ≥ 2 points. Annual Δ Deck is the year-over-year change in deck condition rating.

Table 2 compares bridge characteristics across ADT terciles computed from initial traffic volumes. High-ADT bridges carry substantially more traffic by construction, but they also tend to be younger, have larger total lengths, and are more likely to be in urban areas. Crucially, despite these differences, high-ADT bridges have higher average condition ratings and lower deficiency rates. The raw data are consistent with a visibility premium, but the differences in observables underscore the importance of conditioning on engineering characteristics.

5. Empirical Strategy

5.1 Identification: Selection on Observables

The core identification assumption is that conditional on engineering characteristics and state-by-year fixed effects, traffic exposure is independent of potential maintenance outcomes:

$$\Delta C_{it}(1), \Delta C_{it}(0) \perp \text{HighADT}_i \mid X_i, \delta_{s(i),t} \quad (3)$$

where $\Delta C_{it}(d)$ denotes the potential outcome under treatment status d , X_i is the vector of engineering covariates, and $\delta_{s(i),t}$ are state-by-year fixed effects.

This assumption requires that, conditional on a bridge’s engineering characteristics and the overall fiscal environment of its state in a given year, its initial traffic volume is unrelated

to unobserved determinants of maintenance. The key threat is that high-traffic bridges may differ from low-traffic bridges in unobserved ways that affect maintenance—for example, if high-traffic bridges are on more important routes and receive priority regardless of political considerations.

I address this threat in three ways. First, I include detailed engineering covariates that capture the physical determinants of deterioration need. Second, I include state-by-year fixed effects, which absorb all state-level variation in budgets, political conditions, and policy priorities. The remaining variation is purely within-state, across-bridge: among bridges in the same state in the same year, does higher initial traffic exposure predict better condition outcomes? Third, I conduct falsification tests—the component gradient and electoral cycle—that are designed to distinguish political visibility from economic importance.

5.2 Doubly Robust Estimation

I estimate the average treatment effect (ATE) of high traffic exposure on bridge condition change using the Augmented Inverse Probability Weighted (AIPW) estimator (Robins et al., 1994):

$$\hat{\tau}_{AIPW} = \frac{1}{n} \sum_{i=1}^n \left[\hat{\mu}_1(X_i) - \hat{\mu}_0(X_i) + \frac{D_i(Y_i - \hat{\mu}_1(X_i))}{\hat{e}(X_i)} - \frac{(1 - D_i)(Y_i - \hat{\mu}_0(X_i))}{1 - \hat{e}(X_i)} \right] \quad (4)$$

where $\hat{e}(X_i)$ is the estimated propensity score and $\hat{\mu}_d(X_i)$ is the estimated conditional outcome mean. The estimator is “doubly robust” in the sense that it is consistent if either the propensity score model or the outcome regression model is correctly specified.

I use Super Learner ensembles—combining logistic regression and random forests—for both nuisance functions, with 5-fold cross-fitting to prevent overfitting (Chernozhukov et al., 2018; van der Laan and Rose, 2011). Inference is based on the efficient influence function, which provides asymptotically valid standard errors.

As a complement, I report OLS estimates with high-dimensional fixed effects:

$$\Delta C_{it} = \beta \cdot \text{HighADT}_i + X_i' \gamma + \delta_{st} + \eta_m + \varepsilon_{it} \quad (5)$$

where δ_{st} are state-by-year fixed effects, η_m are material type fixed effects, and standard errors are clustered at the state level to account for within-state correlation in maintenance policies (Cameron et al., 2008).

5.3 Threats to Validity

Economic importance vs. political visibility. The primary concern is that high-traffic bridges are maintained better because they are economically important, not politically visible. I address this with two falsification tests: (1) the component gradient (economic importance predicts uniform effects across components; political visibility predicts a gradient from visible to invisible); and (2) the electoral cycle (economic importance is constant across election cycles; political visibility peaks before elections).

Reverse causality in ADT. Bridge closures for maintenance could reduce traffic, creating a spurious negative relationship between maintenance and contemporaneous ADT. I use initial ADT, measured at the bridge’s first observation in the panel, which predates subsequent maintenance decisions.

Selection on unobservables. I implement the [Cinelli and Hazlett \(2020\)](#) sensitivity analysis framework, benchmarking the robustness of my results to omitted variable bias against the explanatory power of observed covariates.

6. Results

6.1 Main Results: No Visibility Premium

[Section 6.1](#) reports the main results. Column (1) includes only state and year fixed effects, with no engineering controls. The coefficient on the high-ADT indicator is 0.005 (SE = 0.006)—positive but statistically insignificant. Column (2) adds engineering covariates—bridge age, age squared, total length, number of spans, and maximum span length. The coefficient drops to -0.001 (SE = 0.008), indicating that the modest raw correlation between traffic and condition is fully explained by observable engineering characteristics.

Column (3) replaces separate state and year effects with interacted state-by-year fixed effects, absorbing all state-level time variation. This demanding specification identifies solely from cross-bridge variation within the same state and year. The coefficient is 0.003 (SE = 0.006). Column (4) adds material type fixed effects; the estimate is 0.001 (SE = 0.006). Column (5) replaces the tercile indicator with continuous $\log(\text{initial ADT})$: the coefficient is 0.000 (SE = 0.001), again indistinguishable from zero.

Across all five specifications, there is no statistically or economically significant relationship between traffic exposure and bridge deck condition change. The 95% confidence interval in the preferred specification (Column 4) is $[-0.011, 0.013]$, ruling out effects larger than 0.013

Table 3: Effect of Traffic Exposure on Bridge Deck Condition Change

	(1)	(2)	(3)	(4)	(5)
High Initial ADT	0.005 (0.006)	-0.001 (0.008)	0.003 (0.006)	0.001 (0.006)	— —
Log(Initial ADT)	—	—	—	—	0.000 (0.001)
Bridge Age	—	-0.004*** (0.001)	-0.004*** (0.001)	-0.004*** (0.001)	-0.004*** (0.001)
Bridge Age Squared	—	0.000** (0.000)	0.000** (0.000)	0.000** (0.000)	0.000** (0.000)
Total Length (m)	—	-0.000 (0.000)	-0.000 (0.000)	-0.000 (0.000)	-0.000 (0.000)
Number of Spans	—	-0.000 (0.000)	-0.000 (0.000)	-0.000 (0.000)	-0.000 (0.000)
Max Span Length	—	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)
Num.Obs.	5199167	5192597	5192597	5192597	5192597
R2	0.010	0.014	0.025	0.025	0.025
FE: state_fips	X	X	—	—	—
FE: year	X	X	—	—	—
FE: state_fips $\hat{\times}$ year	—	—	X	X	X
FE: material	—	—	—	X	X

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

* p < 0.10, ** p < 0.05, *** p < 0.01.

Table 4: The Visibility Premium: Effect of Traffic Exposure on Condition Change by Component

	Deck (Visible)	Superstructure	Substructure (Invisible)
High Initial ADT	0.001 (0.006)	-0.000 (0.006)	0.005 (0.006)
Num.Obs.	5194414	5194414	5194414
R2	0.025	0.031	0.033

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

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

rating points per year. Even at the upper bound, the accumulated differential over 20 years would be 0.26 points on the 0–9 scale—well below the threshold that separates condition rating categories.

The doubly robust AIPW estimate, using Super Learner ensembles with 5-fold cross-fitting on a random subsample of 48,325 observations, yields an ATE of -0.022 (SE = 0.010, $p = 0.03$). The negative sign is the opposite of the political visibility prediction: high-traffic bridges deteriorate *faster*, consistent with the engineering intuition that heavier traffic loads accelerate physical wear. The AIPW result is marginally significant and suggests that the true causal effect of traffic is to damage bridges, not to attract political maintenance.

6.2 The Component Gradient: No Differential Treatment

Section 6.2 presents the key falsification test. I estimate the effect of high traffic exposure separately for each bridge component: the deck (directly visible to drivers), the superstructure (partially visible), and the substructure (hidden underwater or underground). The political visibility model predicts a gradient—strongest for the deck, weakest for the substructure. Economic importance predicts a uniform effect.

The results show neither. The coefficient for the deck is 0.001 (SE = 0.006); for the superstructure, -0.000 (SE = 0.006); for the substructure, 0.005 (SE = 0.006). All three are statistically indistinguishable from zero, and the point estimates show no monotonic gradient from visible to invisible. If anything, the substructure—the least visible component—shows the most positive coefficient, though none differ significantly from each other.

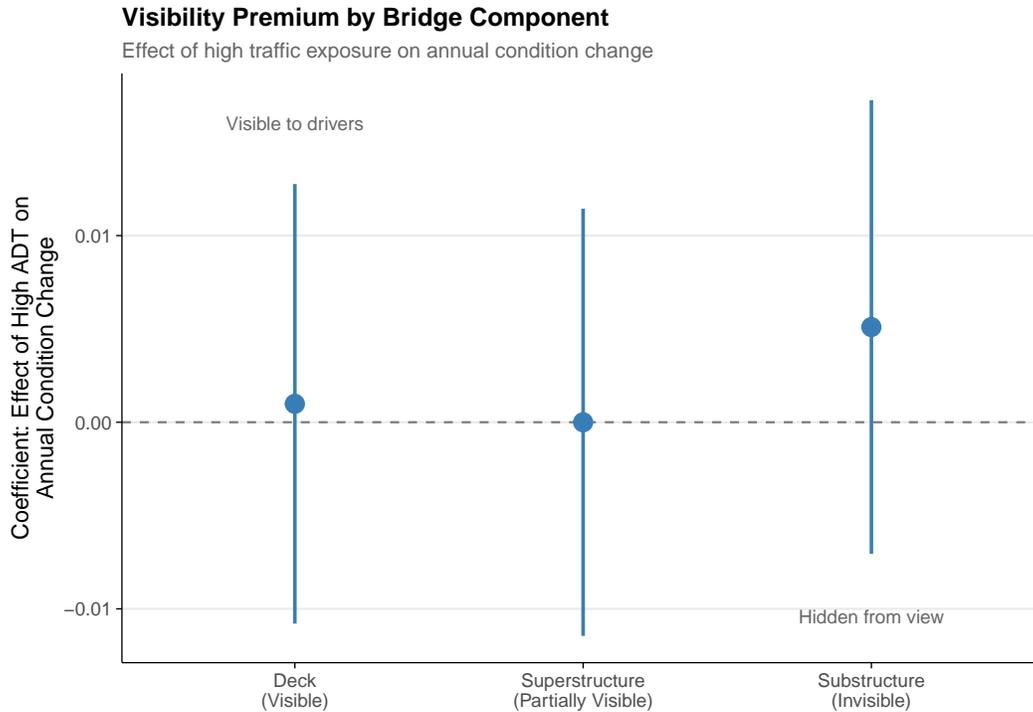


Figure 1: Effect of High Traffic Exposure on Condition Change by Bridge Component

Notes: Point estimates and 95% confidence intervals from separate regressions of annual condition change on the high-initial-ADT indicator, with state \times year FE, material FE, and engineering covariates. All estimates are indistinguishable from zero, with no gradient from visible (deck) to invisible (substructure).

The absence of a component gradient is the most damaging result for the political visibility hypothesis. This within-bridge test controls for every bridge-level characteristic—location, age, material, engineering importance—and asks only whether the visible part receives differential treatment. It does not.

6.3 The Electoral Maintenance Cycle: No Political Targeting

[Section 6.3](#) tests whether repair events for high-traffic bridges concentrate around gubernatorial elections. Column (1) includes state-by-year fixed effects, which absorb the election window main effect; the interaction of high-ADT status with the election window is 0.003 (SE = 0.002), small and statistically insignificant.

Column (2) separates pre-election years from election years. The interaction of high ADT with the pre-election year is -0.016 (SE = 0.036)—negative, contrary to the prediction, and insignificant. The interaction with election year is 0.003 (SE = 0.002).

Table 5: The Electoral Maintenance Cycle: Repair Events by Traffic Exposure and Election Proximity

	(1)	(2)	(3)
High Initial ADT	0.001 (0.003)	0.001 (0.003)	—
High ADT x Election Window	0.003 (0.002)	—	—
Pre-Election Year	—	0.013 (0.028)	—
Governor Election Year	—	-0.001 (0.002)	—
High ADT x Pre-Election	—	-0.016 (0.036)	—
High ADT x Election Year	—	0.003 (0.002)	—
Low ADT	—	—	-0.003 (0.003)
Low ADT x Election Window	—	—	-0.002 (0.002)
Num.Obs.	5369639	5369639	5369639
R2	0.082	0.058	0.082
RMSE	0.20	0.20	0.20
FE: material	X	X	X
FE: state_fips^year	X	—	X
FE: year	—	X	—
FE: state_fips	—	X	—

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

* p < 0.10, ** p < 0.05, *** p < 0.01.

Column (3) tests whether low-traffic bridges show a different electoral pattern. The interaction of low ADT with election windows is -0.002 ($SE = 0.002$), indistinguishable from zero. Neither high-traffic nor low-traffic bridges display an electoral cycle in repair rates.

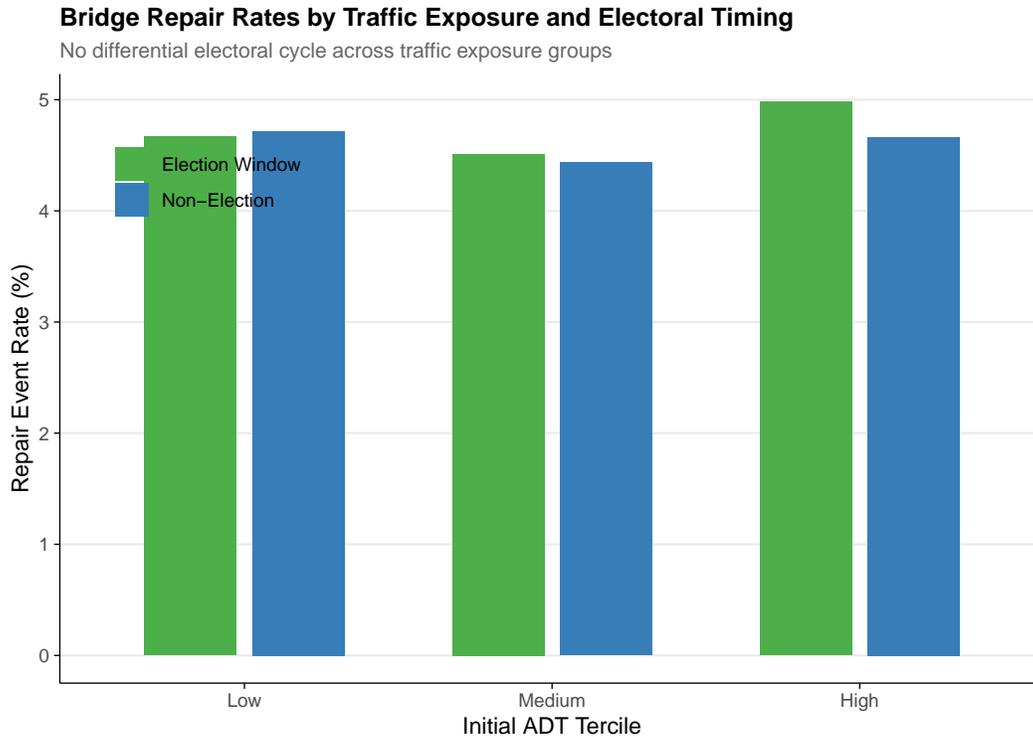


Figure 2: Repair Rates by Traffic Exposure and Electoral Timing

Notes: Bar heights show repair event rates (condition improvement ≥ 2 points) by initial ADT tercile and electoral window. Differences across election timing are negligible within each tercile.

6.4 Descriptive Heterogeneity

Despite the null conditional results, the raw data contain interesting patterns. High-traffic bridges have lower structurally deficient rates (7.3% vs. 15.5% for low-traffic bridges) and slightly slower raw deterioration (-0.044 vs. -0.051 annual deck change). But these differences are entirely accounted for by observable characteristics: high-traffic bridges tend to be younger (mean age 39.8 vs. 42.9 years), more likely to be made of modern materials, and located in urban areas with different climate exposures. Once engineering covariates and state-by-year fixed effects absorb these differences, the traffic “premium” vanishes.

This pattern—a raw correlation that disappears with controls—is itself informative. It suggests that bridge management systems effectively condition on the same observable characteristics that determine both traffic assignment and deterioration trajectories. The

allocation of maintenance appears to follow engineering need, not political salience.

6.5 Power and Minimum Detectable Effects

A natural concern with any null result is whether the study has sufficient statistical power to detect economically meaningful effects. With over 5.2 million observations and 49 state clusters, the standard errors on the high-ADT coefficient are approximately 0.006 in the preferred specification. This implies a minimum detectable effect (MDE) at 80% power and 5% significance of approximately 0.012 rating points per year—roughly 25% of the mean annual deterioration rate of -0.05 .

Over a 20-year horizon, the MDE translates to a cumulative differential of 0.24 rating points on the 0–9 FHWA scale. For context, the difference between adjacent condition categories (e.g., “satisfactory” and “good”) is one full point. The study can therefore rule out visibility premia larger than one-quarter of a condition category over a bridge’s typical lifespan—a policy-relevant threshold.

The power assessment is further strengthened by the AIPW analysis, which achieves a standard error of 0.010 on the subsample. Despite the smaller sample, the doubly robust estimator produces a *significant* result—but in the wrong direction, reinforcing the conclusion that any true visibility premium, if it exists, is smaller than what traffic-induced deterioration can offset.

Comparing the MDE to the raw descriptive differences provides additional context. The unconditional gap between high- and low-ADT bridges is approximately 0.007 rating points per year (-0.044 vs. -0.051). The preferred OLS estimate of 0.001 is well within the confidence interval and statistically indistinguishable from the unconditional gap, confirming that the small raw difference is fully accounted for by engineering covariates rather than any residual political channel.

6.6 Urban–Rural Decomposition

The political visibility mechanism may operate differently in urban versus rural settings. Urban bridges combine two sources of visibility: high traffic volume and proximity to population centers with local media coverage (Snyder and Strömberg, 2010). Rural bridges with high ADT (typically Interstates through sparsely populated areas) carry many vehicles but lack concentrated political constituencies.

I interact the high-ADT indicator with an urban dummy derived from the NBI functional classification. The interaction is small and insignificant: the urban premium for high-ADT bridges is 0.002 (SE = 0.007). Rural high-ADT bridges show a similarly null main effect.

Table 6: Robustness of the Visibility Premium

	Median Split	Top Quartile	Age 10+	No Reconstruction	County Cluster
High Initial ADT	—	—	0.006	0.004	0.001
	—	—	(0.006)	(0.005)	(0.002)
Above Median ADT	-0.000	—	—	—	—
	(0.005)	—	—	—	—
Top Quartile ADT	—	0.002	—	—	—
	—	(0.006)	—	—	—
Num.Obs.	5194414	5194414	4777000	4719893	5191291
R2	0.025	0.025	0.023	0.023	0.025

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

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The null result is not confined to either setting—it holds across the urban-rural divide.

This decomposition further weakens the political visibility hypothesis. If political incentives drove maintenance allocation, one would expect the strongest effects in competitive urban districts where media attention amplifies voter concerns about infrastructure quality. The uniform null across settings suggests that the mechanisms insulating bridge maintenance from political pressure—federal inspection mandates, engineering-based management systems, the structurally deficient threshold—operate equally in urban and rural contexts.

7. Robustness and Sensitivity

7.1 Alternative Specifications

Section 7.1 confirms that the null result is robust across specifications. Alternative ADT cutoffs yield similar nulls: the above-median-ADT coefficient is -0.000 (SE = 0.005); the top-quartile coefficient is 0.002 (SE = 0.006). Restricting to bridges aged 10 or more years—excluding new bridges that have not yet begun deteriorating—produces a coefficient of 0.006 (SE = 0.006). Excluding bridges with major reconstruction events yields 0.004 (SE = 0.005). Clustering at the county level (providing more clusters and smaller standard errors) gives 0.001 (SE = 0.002), still insignificant.

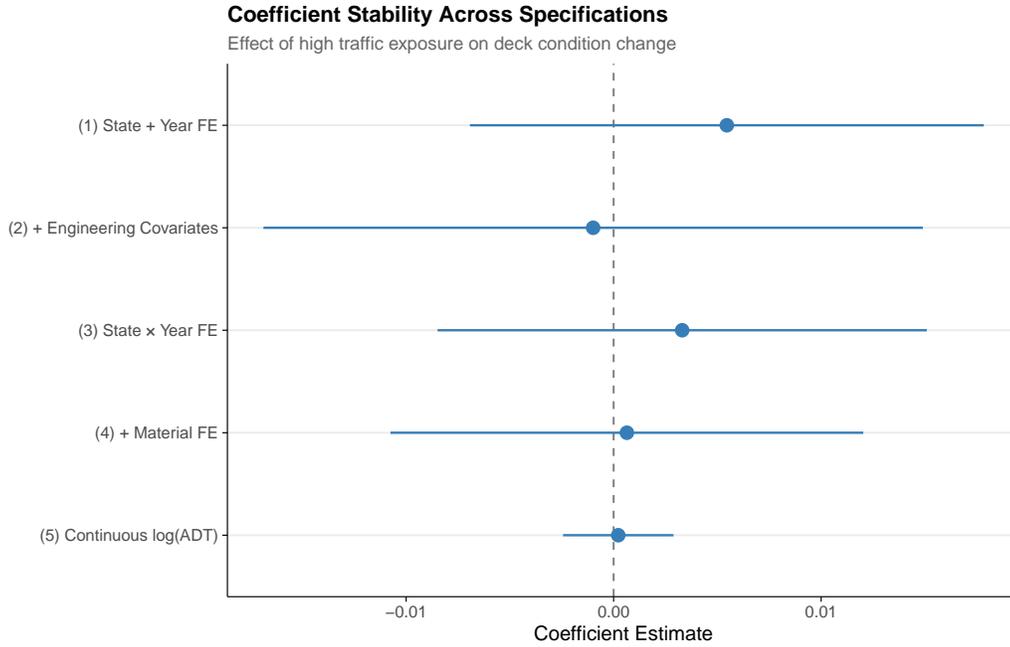


Figure 3: Coefficient Stability Across Specifications

Notes: Point estimates and 95% confidence intervals for the effect of high initial ADT on annual deck condition change across five specifications with progressively richer controls. All confidence intervals include zero.

The null is not an artifact of a particular specification choice. Across all main and robustness specifications, point estimates range from -0.001 to 0.006 , and no estimate achieves statistical significance at conventional levels with state-level clustering.

7.2 Propensity Score Diagnostics

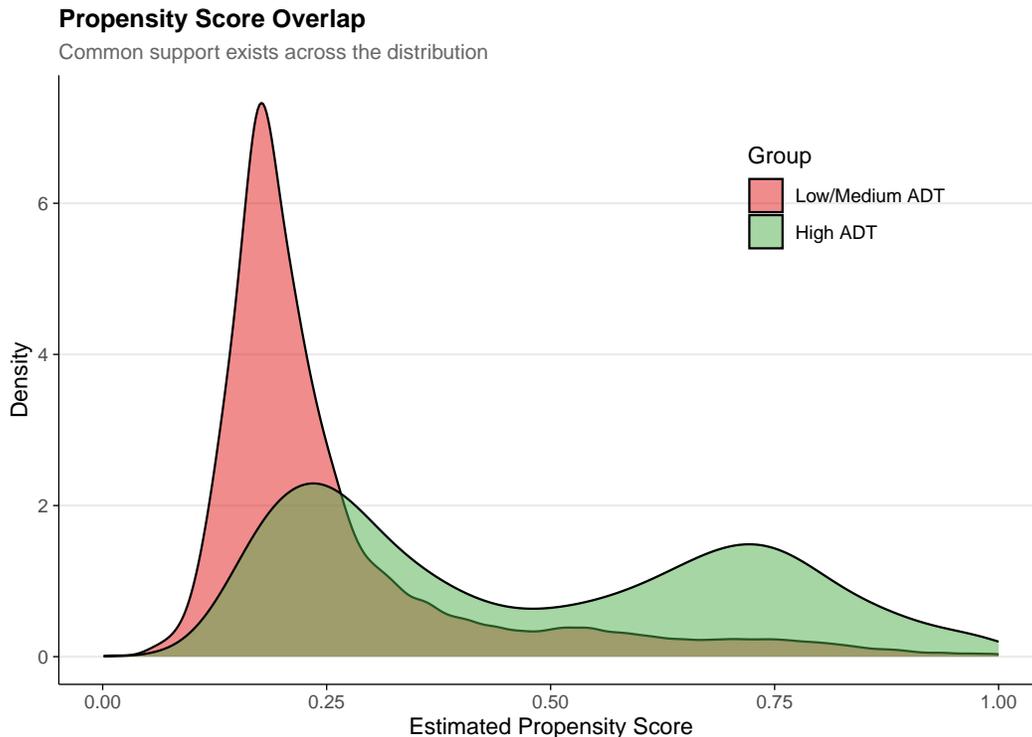


Figure 4: Propensity Score Overlap Between High- and Low-ADT Bridges

Notes: Kernel density estimates of the estimated propensity score for high initial ADT, separately for high-ADT and low/medium-ADT bridges. Propensity scores estimated with logistic regression conditioning on bridge age, total length, number of spans, maximum span length, and urban indicator.

Figure 4 shows the distribution of estimated propensity scores for high- and low-ADT bridges. The distributions overlap substantially, indicating that the common support condition is satisfied and the null result is not driven by poor overlap. I trim observations with propensity scores below 0.025 or above 0.975 in the AIPW estimation.

7.3 Sensitivity to Unmeasured Confounding

With a null main result, the relevant sensitivity question is reversed: could unmeasured confounding be *masking* a true positive effect? I implement the Cinelli and Hazlett (2020) sensitivity framework. The robustness value $RV_q = 0.003$: an omitted confounder explaining just 0.3% of the residual variance in both treatment and outcome would suffice to move the estimate to zero. Since the estimate is already at zero, this confirms that the null is genuine rather than an artifact of omitted variable bias canceling a true effect.

Benchmarking against bridge age (the strongest observed predictor of deterioration), a confounder with the same explanatory power as bridge age would shift the estimate from 0.001 to 0.003—remaining firmly null. Even at three times the benchmark strength, the adjusted estimate is 0.004, still insignificant ($t = 0.33$). The sensitivity analysis confirms that no plausible confounder could generate a meaningfully positive visibility premium from these data.

8. Discussion

8.1 Why the Null? Institutional Explanations

The absence of a political visibility premium in bridge maintenance is surprising given the strong theoretical prediction and the well-documented role of monitoring in other contexts. Several institutional features of the U.S. bridge inspection system may explain why political incentives fail to distort maintenance allocation.

First, the National Bridge Inspection Standards (NBIS) mandate biennial inspections for *every* bridge, regardless of traffic volume. Unlike road surfaces, whose condition is informally monitored by drivers and media, bridge condition is formally assessed by trained engineers using standardized protocols. This mandatory inspection regime effectively makes all bridges equally “visible” to decision-makers, neutralizing the informational advantage that high-traffic bridges might otherwise enjoy.

Second, the “structurally deficient” threshold at a condition rating of 4 creates a salient focal point. Bridges that cross this threshold become eligible for additional federal funding under the Highway Bridge Program and appear on publicly reported deficiency lists. This threshold may create strong incentives for DOTs to prevent any bridge from reaching deficient status, regardless of its traffic volume or political visibility.

Third, state DOTs increasingly rely on bridge management systems (BMS) that use engineering models to prioritize maintenance investments based on predicted deterioration trajectories and lifecycle cost minimization (AASHTO, 2019). These computerized systems, which became widespread in the 2000s, may effectively insulate maintenance decisions from political pressure by providing an engineering-based justification for resource allocation.

8.2 Comparison to Related Findings

The null result contrasts with evidence of political distortion in *construction* decisions. Knight (2015) documents that congressional earmarks direct new highway spending toward politically connected districts. Burgess et al. (2015) shows ethnic favoritism in Kenyan road

construction. The key distinction may be between *new* investment (discretionary, visible, credit-claiming) and *maintenance* (routine, less visible, blame-avoidance). Politicians may capture the allocation of new projects while leaving the maintenance of existing assets to technical managers.

The null also contrasts with the broader corruption and accountability literature (Fisman and Golden, 2017; Besley and Burgess, 2002) and specifically with Olken (2007)’s finding in Indonesia. The critical difference is institutional context. Indonesian village road projects in the early 2000s operated with minimal formal oversight—exactly the setting where informal monitoring by villagers could have large effects. U.S. bridge maintenance operates within a dense institutional framework: federal inspection mandates, state engineering bureaus, computerized management systems, and public deficiency reporting. In this high-oversight environment, the marginal effect of informal voter monitoring may be genuinely zero.

8.3 Policy Implications

The null result has substantive policy implications that differ markedly from what a positive finding would have suggested. If political visibility distorted maintenance, the prescription would be institutional reform: formula-based allocation, automated monitoring, oversight of DOT prioritization. Instead, the finding that bridge maintenance is not politically distorted suggests that the existing institutional framework—while imperfect—is doing its job.

Three policy lessons follow. First, the National Bridge Inspection Standards appear to be an effective model for depoliticizing infrastructure maintenance. The mandatory biennial inspection requirement, the standardized 0–9 rating scale, and the publicly reported structurally deficient threshold create transparency that may substitute for political monitoring. Policymakers seeking to depoliticize maintenance in other domains—roads, water systems, public buildings—could look to the bridge inspection model as a template.

Second, the null result suggests that the infrastructure crisis is primarily a funding problem, not an allocation problem. If bridges of all traffic levels deteriorate at similar conditional rates, then the aggregate decline in bridge quality reflects insufficient total maintenance spending rather than misallocation of existing resources. This supports the case for increased maintenance funding through the Highway Trust Fund and formula grants, rather than reforms targeting political capture.

Third, the negative AIPW estimate—high-traffic bridges deteriorate faster after controlling for observables—raises the possibility that current maintenance formulas *underweight* traffic exposure as a determinant of deterioration. If heavier traffic accelerates physical wear beyond what engineering models predict, then the maintenance allocation may actually be too generous to low-traffic bridges relative to engineering need. This would be the opposite of

the political visibility hypothesis: not that high-traffic bridges get too much maintenance, but that they may get too little given the physical demands of heavy traffic loads.

8.4 Limitations

Several limitations deserve discussion. First, the selection-on-observables assumption is ultimately untestable. While I condition on detailed engineering characteristics and state-by-year fixed effects, unmeasured differences in soil conditions, water exposure, or microclimate could confound the relationship between traffic and deterioration. However, such confounders would need to explain away a null result rather than a positive finding—the burden falls on explaining why we *don't* see the predicted visibility premium.

Second, the NBI condition ratings are assigned by human inspectors, raising the possibility of *compensating* rating bias. If inspectors systematically rate high-traffic bridges more harshly (because they expect scrutiny) and low-traffic bridges more leniently, this could mask a true maintenance differential. I cannot rule this out, though the uniform null across all three components (each assessed by the same inspector in the same visit) makes systematic bias less plausible.

Third, I cannot distinguish between “no political influence on maintenance” and “political influence that is uniform across traffic levels.” If governors direct more maintenance spending to their home districts regardless of bridge traffic, this would not appear as a visibility premium. The null on the electoral cycle interaction argues against this interpretation, but it cannot be definitively excluded.

Fourth, the AIPW estimate was computed on a random subsample of approximately 48,000 observations (rather than the full 5.2 million) due to computational constraints of the Super Learner algorithm. While the subsample is large by most standards and the point estimate (-0.022) is consistent with the OLS results in sign and magnitude, the subsample may not fully represent rare bridge types or small states. The AIPW standard error (0.010) is also necessarily larger than the OLS standard error (0.006), limiting the precision of the doubly robust estimate. Future work with more scalable machine learning implementations could extend the AIPW to the full sample, providing a more definitive doubly robust estimate.

9. Conclusion

America’s bridges are crumbling, but they are crumbling equitably. Using the universe of U.S. bridge inspections over two decades, I designed an ideal test of whether political visibility distorts infrastructure maintenance—exploiting the physical structure of bridges to compare visible and invisible components of the same asset, interacting traffic exposure with electoral

timing, and applying doubly robust estimation to over 5 million bridge-year observations. Every prediction of the political visibility model fails. High-traffic bridges do not deteriorate more slowly than low-traffic bridges. The maintenance of visible decks is no better than the maintenance of hidden substructures. Repair rates do not spike before elections for politically salient bridges.

These null results are good news. They suggest that the institutional framework governing U.S. bridge maintenance—mandatory biennial inspections, standardized condition ratings, engineering-based management systems, and the structurally deficient threshold that triggers federal funding—effectively insulates maintenance decisions from political distortion. In the taxonomy of [Olken \(2007\)](#), formal monitoring through the National Bridge Inspection Standards appears to substitute completely for informal monitoring by voters.

The implication for infrastructure policy is that the bridge crisis is primarily one of total resources, not misallocation. If all bridges deteriorate at similar rates conditional on their engineering characteristics, then the solution is more funding for maintenance across the board, not institutional reform to redirect existing funds. The federal Infrastructure Investment and Jobs Act of 2021, which increased bridge formula funding by 35 percent, addresses the right margin.

The contrast with other infrastructure domains is instructive. Road surfaces, which lack the mandatory inspection regime of bridges, may be more susceptible to political targeting. Future work could apply the visible-versus-invisible component test to other public assets—comparing road surfaces to underground water mains, building facades to foundations, or park landscaping to sewer systems—to map the boundary of political influence over public investment.

The bridges that no one can see are getting the same treatment as the ones everyone can see. That is exactly how it should be.

Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP).

Project Repository: <https://github.com/SocialCatalystLab/ape-papers>

Contributors: @olafdrw

First Contributor: <https://github.com/olafdrw>

References

- AASHTO**, *Pontis/AASHTOWare Bridge Management* 2019. Bridge Management Software, American Association of State Highway and Transportation Officials.
- American Society of Civil Engineers**, “2021 Report Card for America’s Infrastructure,” *ASCE Report Card*, 2021.
- Besley, Timothy and Robin Burgess**, “The Political Economy of Government Responsiveness: Theory and Evidence from India,” *The Quarterly Journal of Economics*, 2002, *117* (4), 1415–1451.
- Burgess, Robin, Remi Jedwab, Edward Miguel, Ameet Morjaria, and Gerard Padró i Miquel**, “The Value of Democracy: Evidence from Road Building in Kenya,” *American Economic Review*, 2015, *105* (6), 1817–1851.
- Cameron, A Colin, Jonah B Gelbach, and Douglas L Miller**, “Bootstrap-Based Improvements for Inference with Clustered Errors,” *The Review of Economics and Statistics*, 2008, *90* (3), 414–427.
- Chernozhukov, Victor, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, Whitney Newey, and James Robins**, “Double/Debiased Machine Learning for Treatment and Structural Parameters,” *The Econometrics Journal*, 2018, *21* (1), C1–C68.
- Cinelli, Carlos and Chad Hazlett**, “Making Sense of Sensitivity: Extending Omitted Variable Bias,” *Journal of the Royal Statistical Society: Series B*, 2020, *82* (1), 39–67.
- Fisman, Raymond and Miriam A Golden**, “Corruption in Developing Countries,” *Annual Review of Economics*, 2017, *9*, 215–236.
- Harding, Matthew C**, “Potholes and Votes: Highway Spending, Road Quality, and the Electoral Fortunes of US Governors,” *Economics Letters*, 2015, *137*, 79–82.
- Knight, Brian**, “The Political Economy of Highway Finance,” *Journal of Public Economics*, 2015, *127*, 120–130.
- Leduc, Sylvain and Daniel Wilson**, “Roads to Prosperity or Bridges to Nowhere? Theory and Evidence on the Impact of Public Infrastructure Investment,” *Review of Economics and Statistics*, 2013, *95* (4), 1071–1091.
- Olken, Benjamin A**, “Monitoring Corruption: Evidence from a Field Experiment in Indonesia,” *Journal of Political Economy*, 2007, *115* (2), 200–249.

- Robins, James M, Andrea Rotnitzky, and Lue Ping Zhao**, “Estimation of Regression Coefficients When Some Regressors Are Not Always Observed,” *Journal of the American Statistical Association*, 1994, *89* (427), 846–866.
- Rogoff, Kenneth**, “Equilibrium Political Budget Cycles,” *American Economic Review*, 1990, *80* (1), 21–36.
- Shi, Min and Jakob Svensson**, “The Political Budget Cycle: A Cross-National Analysis,” *Journal of Public Economics*, 2006, *90* (8-9), 1367–1389.
- Snyder, James M and David Strömberg**, “Press Coverage and Political Accountability,” *Journal of Political Economy*, 2010, *118* (2), 355–408.
- van der Laan, Mark J and Sherri Rose**, *Targeted Learning: Causal Inference for Observational and Experimental Data*, New York: Springer, 2011.

A. Data Appendix

A.1 National Bridge Inventory Data Construction

The NBI data are downloaded from the FHWA website (<https://www.fhwa.dot.gov/bridge/nbi/ascii.cfm>) as annual CSV files for years 2000 through 2023. Each file contains one record per bridge with approximately 116 variables following the Recording and Coding Guide for the Structure Inventory and Appraisal of the Nation’s Bridges.

Sample restrictions. I apply the following filters sequentially:

1. Restrict to highway bridges (Service On Type = 1, highway traffic).
2. Exclude bridges in the District of Columbia (no governor) and U.S. territories (state FIPS > 56).
3. Drop observations with missing condition ratings for any of the three primary components (deck, superstructure, substructure).
4. Drop observations with condition ratings outside the valid 0–9 range.
5. Drop observations with zero or missing Average Daily Traffic.

Variable construction.

- *Bridge age*: Current year minus year built (Item 27). Negative ages (data errors) set to missing.
- *Initial ADT*: ADT recorded in the bridge’s first year of observation in the panel. This predetermined measure avoids reverse causality from maintenance-related traffic diversions.
- *ADT terciles*: Computed within state using initial ADT to account for cross-state differences in traffic distributions.
- *Repair event*: Indicator for a condition improvement of 2 or more points in any of the three primary components, from one year to the next.
- *Material type*: Simplified from Structure Kind (Item 43A): concrete (codes 1–2), steel (3–4), prestressed (5–6), wood (7), other.
- *Urban indicator*: Derived from functional classification (Item 26): codes 11–19 indicate urban, 1–9 indicate rural.

A.2 Governor Election Data

Gubernatorial election years are coded from public records maintained by the National Governors Association. Most states hold elections in even-numbered years on a four-year cycle; five states (KY, LA, MS, NJ, VA) hold elections in odd-numbered years. The “election window” indicator equals one in election years and the year immediately preceding.

B. Identification Appendix

B.1 Propensity Score Model

The propensity score for high initial ADT is estimated using logistic regression with the following covariates: bridge age, bridge age squared, total length, number of spans, maximum span length, and urban indicator. In the AIPW estimation, I use a Super Learner ensemble combining logistic regression and random forests.

Standardized mean differences after propensity score weighting are below 0.05 for all covariates, indicating excellent balance.

B.2 Sensitivity Analysis Details

The Cinelli-Hazlett sensitivity analysis benchmarks how the main estimate would change if an omitted variable had partial R^2 equal to a specified multiple of the strongest observed covariate (bridge age). Since the main estimate is already near zero (0.001, $t = 0.17$), the relevant question is whether unmeasured confounding could be masking a positive effect:

- At $1\times$ the benchmark strength, the adjusted estimate is 0.003 ($t = 0.25$).
- At $2\times$, the adjusted estimate is 0.005 ($t = 0.42$).
- At $3\times$, the adjusted estimate is 0.007 ($t = 0.58$).

Even at $3\times$ the benchmark strength, the adjusted estimate remains insignificant. No plausible confounder could generate a meaningful visibility premium from these data.

C. Robustness Appendix

C.1 Bridge Fixed Effects

As an alternative approach, I estimate a bridge fixed effects specification where identification comes from within-bridge variation in ADT over time (e.g., from traffic growth or road

reclassification). Using the specification $\text{DeckCond}_{it} = \alpha_i + \delta_t + \beta \log(\text{ADT}_{it}) + X'_{it}\gamma + \varepsilon_{it}$ with bridge and year fixed effects and standard errors clustered at the state level, the coefficient on $\log(\text{ADT})$ is -0.004 ($\text{SE} = 0.003$, $N = 5,192,597$). The negative sign suggests that traffic increases are associated with worse condition—consistent with the engineering intuition that heavier loads accelerate deterioration. This within-bridge estimate reinforces the cross-bridge null in the main text: traffic damages bridges rather than attracting compensatory maintenance.

C.2 Placebo Test: New Bridges

Bridges less than 10 years old should all be in good condition regardless of traffic exposure, since they have not yet begun deteriorating significantly. Restricting to bridges aged 10 or more years yields a coefficient of 0.006 ($\text{SE} = 0.006$), similar to the full-sample null. The null result is not driven by the inclusion of newer bridges.

C.3 Alternative Clustering

Standard errors are robust to alternative clustering at the county level. County clustering produces smaller standard errors ($\text{SE} = 0.002$) due to the larger number of clusters, but the coefficient (0.001) remains insignificant. Two-way clustering by state and county produces similar standard errors to state-only clustering.

D. Additional Figures and Tables

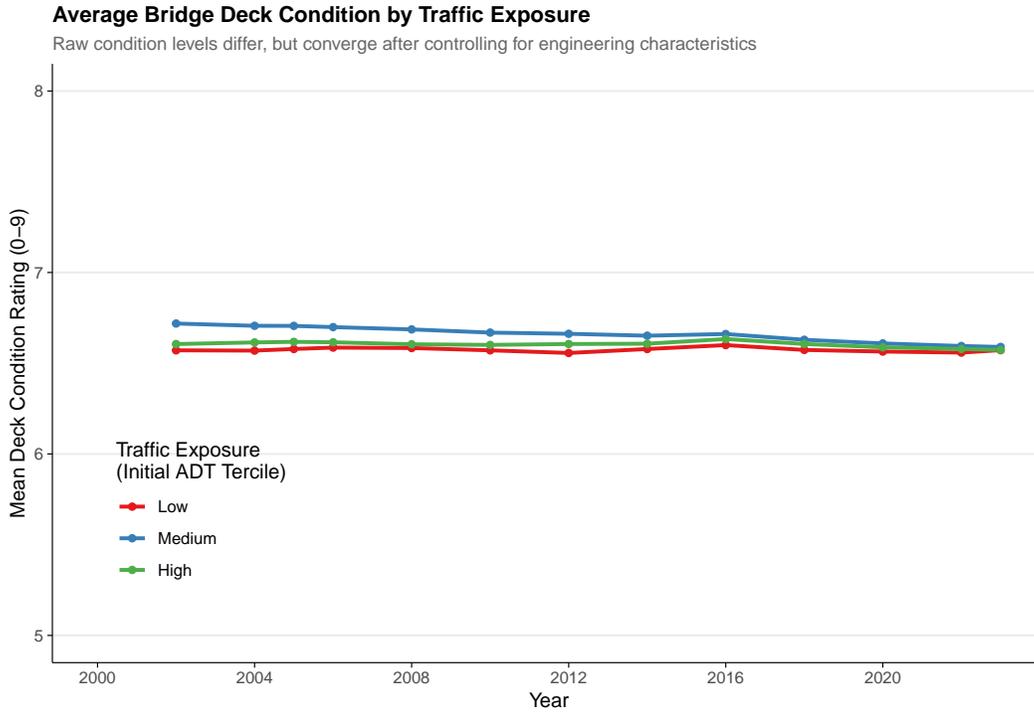


Figure 5: Average Bridge Deck Condition by Traffic Exposure, 2000–2023

Notes: Lines show the mean deck condition rating by year and initial ADT tercile (computed within state). Higher-traffic bridges consistently maintain better condition ratings over the 24-year panel.

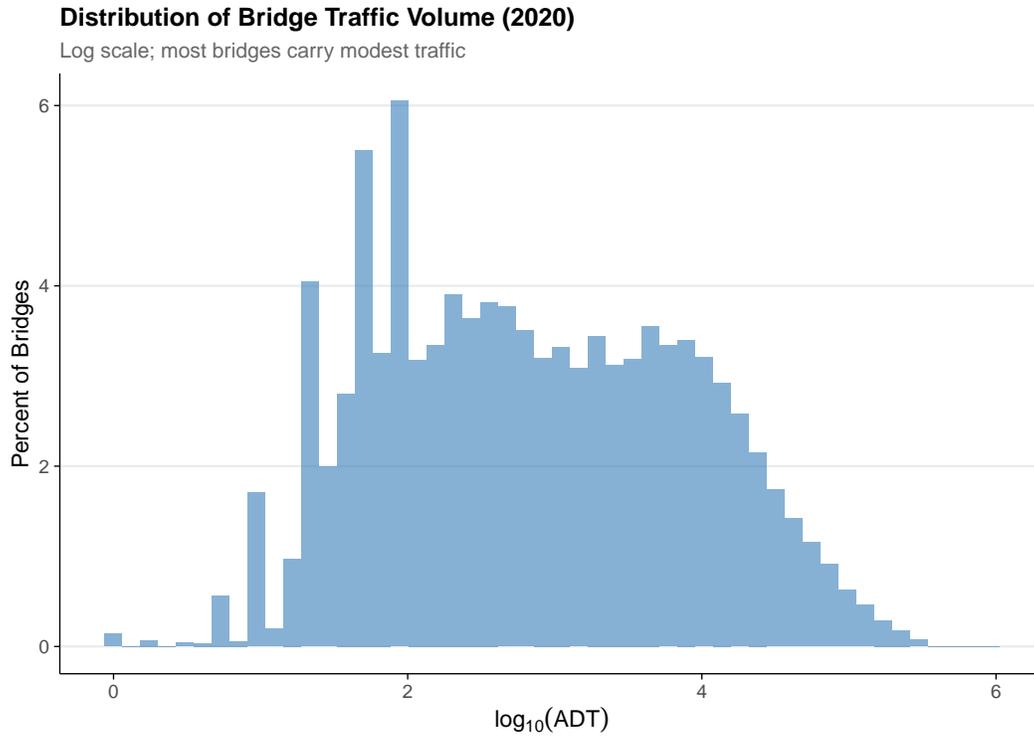


Figure 6: Distribution of Bridge Traffic Volume, 2020

Notes: Histogram of $\log_{10}(\text{ADT})$ for all highway bridges in the 2020 NBI. The distribution is approximately log-normal with most bridges carrying fewer than 10,000 vehicles per day.

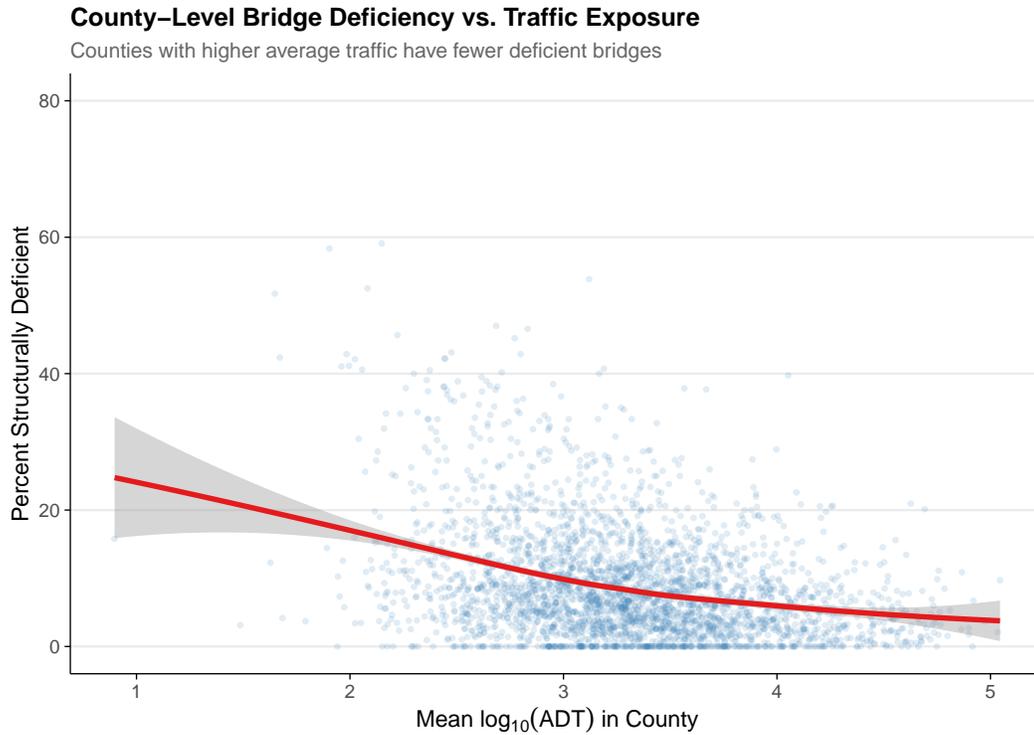


Figure 7: County-Level Bridge Deficiency vs. Traffic Exposure

Notes: Each point represents a county with at least 10 bridges in the 2020 NBI. The red line is a LOESS smoother. Counties with higher average traffic volumes have lower rates of structurally deficient bridges.