

State Capacity and Infrastructure Costs

Zachary Liscow, Cailin Slattery and Will Nober*

August 21, 2025

Abstract

Why is it so expensive to build infrastructure in the United States? We collect new project-level data on infrastructure costs and conduct a survey on how states plan, procure, and deliver these projects. While there are many determinants of project costs, the survey results suggest that low state capacity at the agency delivering the projects is a primary cost driver. We investigate this with administrative data that links individual personnel to infrastructure projects. We find that higher-quality government engineers deliver observationally similar projects at significantly lower cost; going from the 25th to 75th percentile of engineer quality is associated with a 14% reduction in project-level costs, amounting to more than three times the average engineer salary. Further, losing expertise to retirement has substantial consequences: the cost increase arising from engineer departures is six times the size of their wages. Our results highlight the value of experience and human capital in public organizations.

*Liscow: Yale University, zachary.liscow@yale.edu; Slattery: UC Berkeley, cailin.slattery@berkeley.edu; Nober: Columbia University, whn2105@columbia.edu. We thank the Department of Transportation and the National Bureau of Economic Research for generous funding, and we thank Michael Best, Ernesto Dal Bó, Erica Bosio, Rob Clark, Ed Glaeser, Joe Gyourko, Karam Kang, Jeff Liebman, Karthik Muralidharan, Jim Poterba, David Schleicher, Steve Schooner, Steve Tadelis, Francesco Trebbi, Shoshana Vasserman, Reed Walker, Guo Xu, Chris Yukins, and anonymous highway contractors and DOT employees for helpful comments. Audrey Cheng, Hannah Lu, Nicholas Whitaker, and Connor Wilke provided excellent research assistance. A previous draft of this paper circulated under the title “Procurement and Infrastructure Costs.” All errors are our own.

1 Introduction

There is widespread concern over the cost and quality of infrastructure provision in the United States. Highway construction costs have increased substantially over time (Brooks and Liscow, 2023; Turner et al., 2023), and the cost of urban rail is about three times more expensive in the U.S. than in other rich countries (Goldwyn et al., 2023). Despite the magnitude of the investment and the importance of infrastructure as a driver of economic growth, experts have emphasized the poor condition of U.S. transit, roads, and bridges (American Society of Civil Engineers, 2021).¹ Leading theories for high infrastructure costs include overregulation and litigation (Klein and Thompson, 2025; Liscow, 2025), citizen opposition (Brooks and Liscow, 2023), and political distortions (Fajgelbaum et al., 2023; Rosenthal, 2017).

We provide an additional explanation for the high cost of U.S. infrastructure: state capacity. We bring three new pieces of evidence to bear. The first is qualitative: a survey of experts shows that state Departments of Transportation (DOTs) across the country are concerned about a decline in the size and experience of their workforce. We also observe these trends in Census data, and find that states with fewer transportation employees per capita have higher costs. The second is an analysis of the individual engineers who manage projects, wherein we create a measure of engineer quality that captures experience and skill. When a higher quality engineer is assigned to a project, costs are reduced substantially. The third piece of evidence comes from variation in district-level retirements. An increase in retirements is followed by an increase in average project costs, cost overruns, and duration. Retaining high-quality employees and preventing early retirements is therefore one possible policy to manage costs.

We begin by studying highway resurfacing projects, which are the bread and butter of American infrastructure spending.² We start with the basics. How much does it actually cost to resurface a road in the United States? We collect data on *project-level* costs in each state, via public records requests, and then extract details from project plans to calculate the cost per mile of resurfacing: a homogeneous measure that we can compare across projects and space. This dataset is the first of its kind; previous work on infrastructure investment across states relies on aggregate spending data. We document significant variation in costs across states. These cost differences are not explained away

¹The fact sheet for the 2021 Infrastructure Investment and Jobs Act stated that “1 in 5 miles of highways and major roads, and 45,000 bridges, are in poor condition.” See <https://www.presidency.ucsb.edu/documents/fact-sheet-the-bipartisan-infrastructure-deal>.

²Total capital transit spending across the U.S. is an order of magnitude less than highway spending (70% of which is rehabilitation and maintenance). In 2022, the Federal Transit Administration reported an expenditure of \$27B across all transit authorities, compared to \$266B on highways in the same year (Federal Transit Administration, 2022).

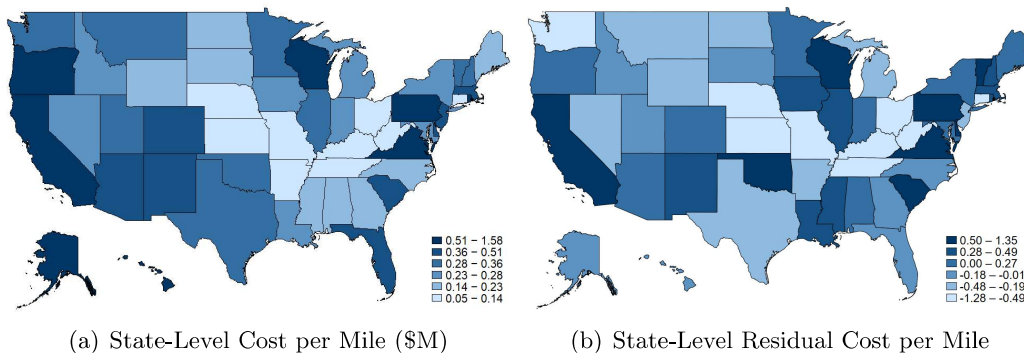
D Cost Data

In this Appendix we review some additional descriptives on costs and then compare our data with state-level aggregates provided by the Federal Highway Administration (FHWA).

We request data on project-level costs directly from each state DOT. The states provide data in a variety of formats, and include different subsets of the variables that we request. For example, in some states we receive data on three cost variables: winning bid, realized cost, and amended contract cost. However, in some states we only receive one or two of the three. In the states that we only observe winning bid, for example, we want to be convinced that it is a good enough proxy for realized costs for the projects in our subsample.


Figure D.1 shows the distribution of cost per mile across states, while Figure D.2 shows the distribution of costs per mile, total cost, and mileage, at the project level. Figure D.3 shows the spending and length of the projects across space. Figure D.4 compares winning bids with realized costs, in the states where we have both (22 states). The panel on the left plots the two objects in the project-level data, and the panel on the right shows the differences, in a histogram. The two objects are highly correlated, with a correlation coefficient of 0.98. Figure D.5 compares winning bids with “current contract” costs, in the states where we have both (17 states). Current contract costs are the winning bids plus or minus any contract amendments. Current contract costs are essentially realized costs, unless the project has not been completed. All of the projects in our data have been completed, therefore we can treat current construction costs as realized costs. The two objects have a correlation coefficient of 0.99. Our final sample uses the realized cost or current contract cost when available (35 states), and the winning bid for the remainder (15 states). As a final exhibit, Figure D.6 shows the correlation between the costs and the duration of the project.

Figure D.1: Resurfacing Costs per Mile (\$M) across States

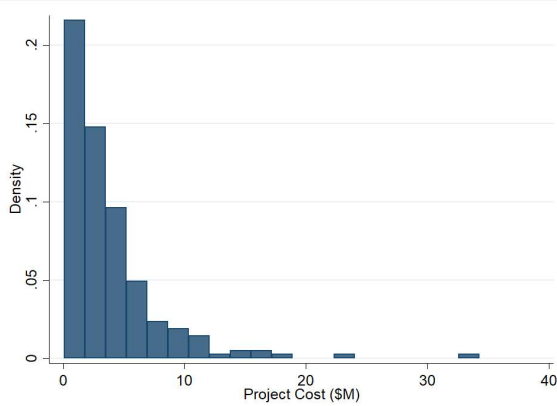


Notes: The figure on the left takes the average of project cost per mile at the state level. The figure on the right displays the state-level average of residualized costs. Residual costs are determined using a project-level regression of project costs per mile on the observables included in Column (2) of Table 1. There are 250 projects in the sample: 5 projects per state. Data collected by the authors.


Figure D.2: Project-Level Costs(\$M) and Mileage



(a) Project-Level Cost per Mile



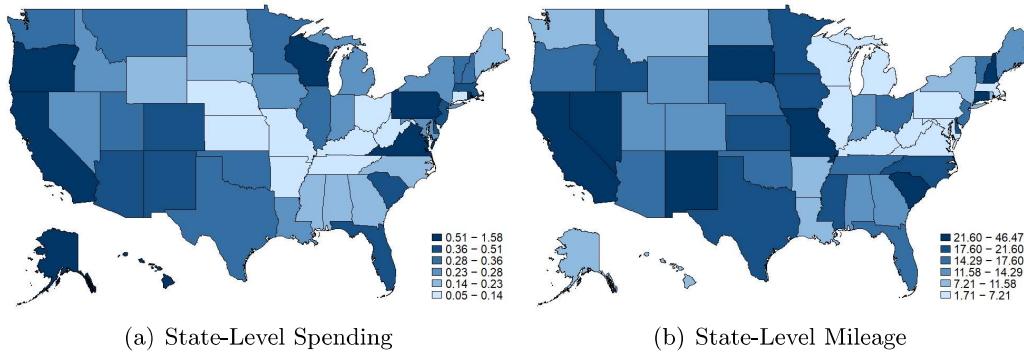
(b) Project-Level Total



(c) Total vs Mileage

Notes: The top figure (a) shows the histogram of project-level resurfacing costs per mile. There are 250 projects in the sample: 5 projects per state. The histogram on the left (b) shows the distribution of project-level costs in the data collected by the authors. The scatter plot on the right (c) takes each project cost and plots it against total mileage of the project (in lane miles).

Figure D.3: Spending and Project Characteristics

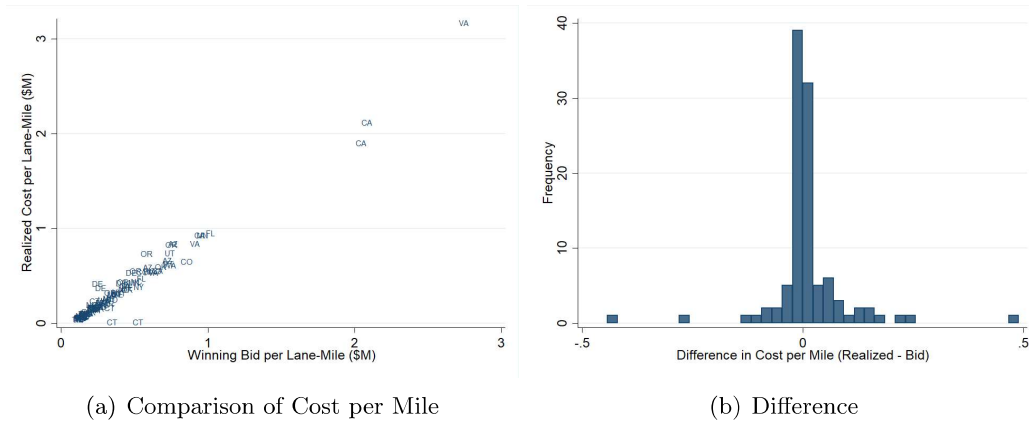


(a) State-Level Spending

(b) State-Level Mileage

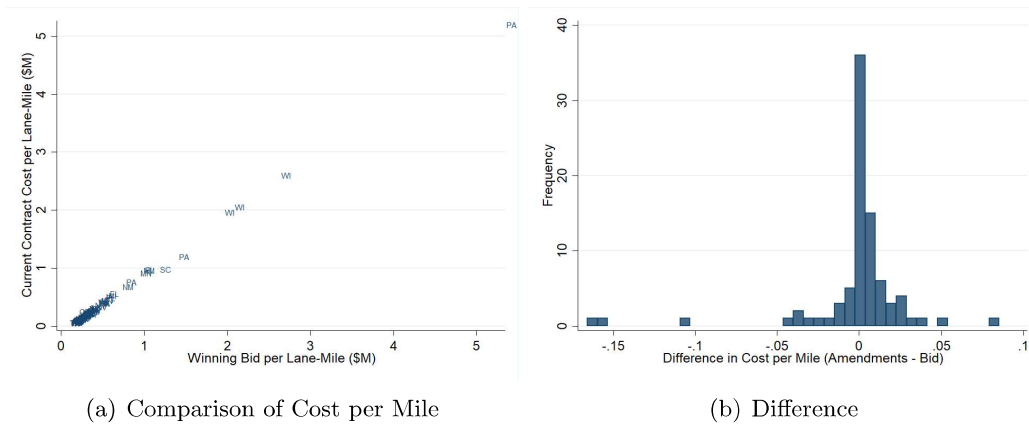
Notes: The map on the left shows project costs per mile, and the right shows the mileage for these projects (reproduced from Figure D.2 in the main text). The cost data collected by the authors is for resurfacing projects started in 2018 or 2019, with a length between 1 and 20 miles, on a non-interstate highway. There are 5 such projects per state, and the state average is used for this map. Interestingly, project length does seem to vary by state, with some states undertaking much longer projects, on average, than others.

Figure D.4: Bids vs Realized Costs




Notes: This figure compares winning bids with realized costs. The panel on the left plots the two objects in the project-level data, and the panel on the right shows the differences, in a histogram.

Figure D.5: Bids vs Amended Contracts



Notes: This figure compares winning bids with “current contract” costs. Current contract costs are the winning bids plus or minus any contract amendments. The panel on the left plots the two objects in the project-level data, and the panel on the right shows the differences, in a histogram.

Figure D.6: Costs and Project Duration

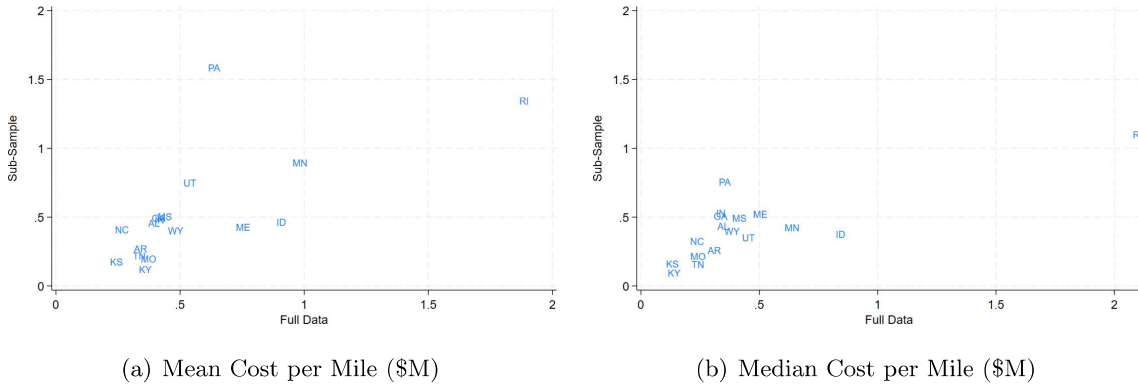


Notes: This is a binned scatter plot of project-level costs and duration. The y-axis shows the cost per mile, in \$M. The x-axis shows the duration per mile, in days. The data was collected by the authors, as part of the project-level cost data. The project duration variable is only available for 33 states.

D.1 Comparing Our Subsample to All Available Data

We use a subsample of 5 projects per state to facilitate the collection of detailed information for each project. However, we have more than 5 projects in 22 states. Figure D.7 plots the mean and median cost per mile of our subsample against the mean and median cost per mile observed in the full data. For this comparison, we restrict the full data to match our selection criteria for the subsample (non-interstate resurfacing projects between 1 and 20 mile that began in 2018 or 2019). The correlation coefficient for mean cost per mile is 0.67 and the correlation coefficient for median cost per mile is 0.78. In this figure we restrict California, which is an outlier (the subsample data exhibits higher costs than the full sample). This allows us to visualize the remainder of the states. However, when we include all of the data from California in Table 1 we do not find different results.

Figure D.7: Subsample vs Full Data by State



Notes: The figure on the left displays the mean cost per mile of our subsample versus the mean cost per mile of all available data, by state. The figure on the right performs the same comparison with median cost per mile. Both figures restrict to states with at least 5 projects observed in the full data. This leaves us with 18 states. California is omitted from both figures to better visualize the other states. Projects included in the subsample are excluded from the full data sample. Data collected by the authors.

D.2 Comparison with the FHWA Data

We can compare our data with state-level cost data from the Federal Highway Administration (FHWA). The FHWA requires states to report spending on a variety of categories and types of roads via an annual survey. As such, these data are self-reported by the states. A potential concern is that states do not all report resurfacing spending, for example, in the same narrow category. A second concern is that we do not have a denominator of miles resurfaced; we only have total miles in the state. Therefore, it is hard to know if costs are higher in one state because they completed more projects than other states or because the projects are more expensive. However, this is the only comprehensive dataset on state-level costs and is often used in reports about spending across states

(e.g., Reason Foundation, 2018).

The FHWA data includes spending by category, as we mentioned earlier. Therefore we need to select certain categories to compare with our sample of resurfacing projects. Due to potential reporting issues, we are fairly broad in the categories we include in the series we create. We define our series of interest as “Resurfacing + Maintenance cost per vehicle-mile.” This includes all maintenance cost variables and certain capital outlay variables.¹⁰⁹ Vehicle-miles are the estimated number of miles traveled by vehicles on roads in that state. We use all road types in all definitions, but the series looks similar when we exclude local roads. These data are very similar to the series used by Turner et al. (2023). For the FHWA data, we think spending per vehicle-mile is a better measure of cost efficiency, as it implicitly controls for the expected degradation due to heavy road use.

Figure D.8 shows the variation in spending across states for the two sources. There is a striking amount of variation in the FHWA data with Maine, Delaware, and Nebraska spending 4 times per vehicle mile than Mississippi, Georgia, and Alabama. Our series, although still exhibiting variation across states, does not exhibit as much heterogeneity. Importantly, there is very little correlation between the two series. Since we have directly collected resurfacing cost data from the states, this suggests the FHWA data is not a good proxy.


Moreover, it does not seem that costs are higher in areas where roads are rougher in the FHWA data, while our data show a stronger correlation between costs and road quality (Figure D.9). We would expect maintenance costs per lane mile to be higher on rougher, poorer-quality roads. We also compare the FHWA total spending measure to cost totals in our data for California, a state where our data covers the universe of highway construction projects.¹¹⁰ Here, the total spending in the FHWA data and our own are very similar.¹¹¹ These comparisons bolster our confidence that the project-level cost data that we collected consists of the individual projects underlying the aggregate FHWA data series.

¹⁰⁹The capital outlay variables are “Relocation,” “Reconstruction: Added Capacity,” “Reconstruction: No Added Capacity,” “Major Widening,” “Minor Widening,” “Restoration & Rehabilitation,” and “Resurfacing.”

¹¹⁰For most other states, our raw data is not as comprehensive and we are not confident that they cover the universe of DOT projects. One exception is Texas, but we are missing a measure of project completion year for over half of the projects, limiting our ability to compare it to the FHWA totals.


¹¹¹Between 2015 and 2019, the total spending observed in California in our data is \$28.1 billion and the total spending observed in California in the FHWA data is \$24.9 billion. We believe the small difference in the total is partially due to some imprecision in how we assign project costs across years. With our data, all costs for a project will be attributed to a single year, even if some spending occurred outside of that year. With the FHWA data, the state reports totals by spending year, not project completion year.

Figure D.8: Comparison of Cost Data



Notes: This is a scatter plot with cost data collected by the authors on the y-axis and cost data collected by the FHWA on the x-axis. The cost data collected by the authors is for resurfacing projects started in 2018 or 2019, with a length between 1 and 20 miles, on a non-interstate highway. There are 5 such projects per state, and the state average is used for this plot. The FHWA cost data is spending on resurfacing and maintenance per vehicle mile.

Figure D.9: Costs and Road Roughness



(a) Cost per Mile and Roughness

(b) FHWA \$ per Vehicle-Mile and Roughness

Notes: The panel on the left shows a scatter plot of state-level resurfacing cost per mile and road roughness (IRI). The cost data collected by the authors is for resurfacing projects started in 2018 or 2019, with a length between 1 and 20 miles, on a non-interstate highway. There are 5 such projects per state, and the state average is used for this plot. The panel on the right shows the same using cost from Highway Statistics (FHWA). To create the cost series we include a subset of capital outlay costs (relocation, reconstruction, major widening, minor widening, restoration, rehabilitation, resurfacing) plus maintenance costs. We use the International Roughness Index (IRI) as the quality measure.

D.2.1 Comparison of FHWA and BidX Data

To further probe the reliability of the Highway Statistics spending data, we compare it to spending totals that we aggregated from BidX, a private-sector service that many state DOTs contract for its construction bidding software. Data from BidX has also recently been used for research purposes (e.g., Bolotnyy and Vasserman, 2023; Kroft et al., 2022). BidX posts winning bids for state DOT construction projects for 38 states. They post descriptions for each project, allowing us to select

projects of a similar scope to those we look at in FHWA Highway Statistics.¹¹²

As a validation exercise, we compare cross-state spending differences in the BidX data with those from Highway Statistics. We first examine Georgia and South Carolina, a pair of states that are similar on observables but have divergent levels of spending per vehicle mile (VM) levels in the Highway Statistics data. In 2018 and 2019, spending per VM is higher in South Carolina in both data sources, but the difference is much larger in Highway Statistics than in BidX.

We repeat the exercise with Georgia and Alabama, two states with very similar spending per VM in Highway Statistics. In the averaged totals across 2014, 2018, and 2019, we find a 15 percent difference between the two states in BidX, whereas the averages in Highway Statistics are virtually identical.

Lastly, we compare New York and North Carolina, which both have high spending per VM in Highway Statistics, to Georgia and Alabama, which both have low spending levels. In 2019, the two datasets disagree: Highway Statistics has the spending per VM significantly higher in NY and NC, whereas BidX has spending per VM significantly higher in Alabama and Georgia. The last comparison also raised a red flag about the internal consistency of the BidX data, as NY has just over 1/3 of the spending per VM as Alabama, which is unlikely to reflect the full universe of spending.

We compare magnitudes of spending in addition to cross-state relative differences. To improve precision, we focus on two categories of spending: resurfacing, rehabilitation, and restoration (3R) and widening/reconstruction. These activities are accounted separately in Highway Statistics and are often grouped together in BidX projects. We choose Georgia for this exercise due to its detailed project descriptions on BidX. In this exercise, the BidX magnitudes are well below those we observe in Highway Statistics.

Overall, it appears that BidX does not typically include the universe of spending laid out in Highway Statistics, and that the degree of the coverage gap in BidX data varies widely by state.

¹¹²The following keywords are indicative of the projects we are looking for: “resurface,” “rehabilitation,” “widening,” in addition to “mill,” “surface,” “CMRB.” We focused on states that had the most detailed descriptions.