

Does the US have an Infrastructure Cost Problem? Evidence from the Interstate Highway System*

Neil Mehrotra,[†] Matthew A. Turner,[‡] Juan Pablo Uribe[§]

June 2024

Abstract: Between 1984 and 2008, expenditure per Interstate vehicle mile travelled fell 10%, while the price of new lane miles and pavement quality more than doubled. To reconcile these trends, we describe an Interstate cost function for a planner who minimizes the cost required to deliver a given level of highway services. Using administrative data, we estimate prices for lane miles and pavement quality and evaluate the user cost of the Interstate. User cost fell by half between 1994-2008, largely due to falling interest rates. In this sense, there is no problem with the cost of the Interstate.

Key words: *Interstate highways, Infrastructure costs, construction productivity.*

JEL classification: H54, R42, R53, E22

*We are grateful to Julia Lynn and Margaux Kelley for excellent research assistance and to Patrick McCarthy for helpful comments. We are also grateful for financial support from the National Science Foundation program on ‘Economics of Transportation in the 21st Century’, and Turner gratefully acknowledges the support of a Kenen Fellowship at Princeton University during part of the time this research was conducted.

[†]Federal Reserve Bank of Minneapolis, 90 Hennepin Avenue, MN 55401. email: neil.mehrotra@mpls.frb.org.

[‡]Department of Economics, Box B, Brown University, Providence, RI 02912. email: matthew_turner@brown.edu. Also affiliated with PERC, IGC, NBER, PSTC, and S4.

[§]Cornerstone Research, 699 Boylston Street, 5th Floor, Boston, MA 02116 email: jp.uribe86@gmail.com.

1 Introduction

Between 1984 and 2008 total public expenditure per Interstate vehicle mile travelled (VMT) fell from 2.9 to 2.6 cents while the network expanded and pavement condition improved. During the same period, the price of new lane miles and the price of maintaining pavement quality both at least doubled. These opposing trends highlight conceptual problems with measuring the cost of the Interstate highway system. First, the flow of services that the Interstate provides is a public good. We do not observe the price of ‘Interstate services’, with its implied information about costs, and we must rely instead on indirect measures of this price. Second, because the Interstate is an asset, the timing of investment expenditure need not match the timing of the realized services. This also complicates an evaluation of costs.

Generalizing an ordinary cost function resolves these problems. The resulting Interstate cost function is the solution to the problem of a highway manager who minimizes the discounted present value of the investment required to deliver a specified level of vehicle miles travelled in each period. This allows us to calculate the marginal cost of VMT in each period - the user cost of the interstate. We use administrative data describing the Interstate network and Interstate expenditure to estimate prices for building lane miles and improving pavement quality, and then use these prices to evaluate the user cost of the Interstate between 1992 and 2008.

The user cost of the Interstate has four main components; the opportunity cost of capital invested in lane miles; the opportunity cost of capital invested in pavement quality; depreciation of pavement; and finally, day-to-day routine operating expenditure and maintenance, e.g., traffic management and snow removal. Although pavement quality is the largest share of contemporary Interstate expenditure, user cost is dominated by the opportunity cost of accumulated Interstate capital. Over our study period, the value of Interstate capital increased rapidly, primarily because the price of lane miles rose, but also because the extent of the network increased. In spite of this, user cost fell by nearly half. The increase in the value of Interstate capital was more than offset by a decline in the rate of return to capital and an increase in the number of Interstate users among whom the opportunity cost of capital was shared. In this sense, there is no problem with the cost of Interstate. To the contrary, its cost fell rapidly from 1994 to 2008. This outcome largely reflects changes in the macroeconomy. If interest rates had not fallen, user costs would have risen dramatically. Alternatively, had the price of lane miles stayed at its initial level, user cost would have fallen even further.

Our estimation of the price of lane miles and pavement quality indicates that both at least doubled between 1994 and 2008. Almost the entire increase in the price of pavement quality appears to reflect an increase in the price of paving materials. The rapid increase in the price of new lane miles remains unexplained, although the data do not provide support for three common hypotheses: 1) that the price increase is a pure composition effect resulting from a shift to more urban construction; 2) that it is a consequence of changing exposure to union labor; 3) that it is a consequence of exposure to more intensively regulated, environmentally sensitive areas. On the other hand, the data suggest some hard to observe change in the nature of construction, such as excess scope (i.e. unnecessary or ancillary construction expenditures) may be to blame. It is natural to speculate that this increase in the price of new lane miles reflects a change in the political process responsible for investment planning, as Brooks & Liscow (2023) suggest.

Measuring the cost of the Interstate is important for at least three reasons. First, measuring the cost of the Interstate is a prerequisite to an assessment of the efficiency with which we are producing transportation infrastructure. Extrapolating from high-profile and over-budget projects, there has been speculation that US productivity in infrastructure construction has been stagnant or declining over the past generation. Recent examples of exorbitantly expensive and delayed transportation infrastructure abound: the Boston Central Artery/Tunnel Project (“Big Dig”), New York City’s Second Avenue Subway, San Francisco-Oakland Bay Bridge replacement, Seattle’s Alaskan Way viaduct replacement, and Maryland/Washington DC Purple Line project. Our analysis offers a logically coherent framework in which to evaluate cost changes for these sorts of programs, and, in the case of the Interstate, suggests decreasing user costs in spite of increasing prices for new lane miles and pavement quality.

Second, although an extensive literature in urban and trade economics (e.g., Allen & Arkolakis (2014); Duranton & Turner (2012)) investigates the benefits of highway networks, there is less systematic evidence on the cost of maintaining and building them. Evaluating increases in transportation investment, assessing the productivity of US infrastructure construction, and performing cost-benefit analyses of transportation infrastructure all rely on an understanding of its costs. We improve our understanding of the cost of the Interstate in four ways. First, we develop a theoretical framework for assessing infrastructure cost for a long-lived asset like the Interstate highway system. Second, our analysis of resurfacing and pavement quality is nearly unique. Third, we provide more timely estimates of the cost of lane miles than the previous literature. Finally, we provide annual estimates of the user cost

of the Interstate system and of the prices of its different components.

Third, despite falling over our study period, our theoretically founded measure of user costs is an order of magnitude higher than the user cost implied by the federal gas tax. The existing gas tax is set so that the resulting revenue is the same order of magnitude as annual expenditures on the network, and these are an order of magnitude smaller than the opportunity cost of the accumulated investment in lane miles. At the 2008 price of new lane miles and interest rate, the opportunity cost of Interstate lane miles was about 123 billion 2010 USD per year. Under current policy, almost this entire amount provides an implicit subsidy for users of the Interstate highway system. While our analysis does not extend to welfare analysis nor to public finance, these calculations suggest that our results have important implications for both.

2 Literature

Lewis (1982) and Bennett et al. (2019) calculate running totals of the expenditure on the Interstate highway system. Brooks & Liscow (2023) estimates the cost per mile of new Interstate and find that it increased by about a factor of four between 1970 and 1993. Further, they find that neither proximity to wetlands nor population density explains this trend, but that the trend in the price of Interstate construction follows the price of nearby housing. Based on these findings and some supplementary evidence, they argue that the increase in highway construction costs reflects increased citizen participation in the planning process, a hypothesis they call ‘citizen’s voice’. Our econometric approach is similar to Brooks & Liscow (2023), but we extend their analyses in three ways. First, we evaluate a more recent time period. Second, in addition to reporting on new construction, we also report on pavement quality and resurfacing costs. Third, we combine our price estimates and our analytical framework to calculate user cost.¹

Our attention to pavement quality and resurfacing is nearly unique. Small & Winston (1988) develop and calibrate a model of optimal pavement thickness for roads subject to periodic resurfacing. They provide the only other evidence on the cost of resurfacing that we have seen, at 200,000 USD2010 per lane-mile for *urban* Interstates.² This is considerably

¹In a recent related paper, Goolsbee & Syverson (2023) evaluates the productivity of the US construction sector in general and the housing sector in particular from 1950-2019. A direct comparison between Goolsbee & Syverson (2023), Brooks & Liscow (2023) and our results is difficult, but the Goolsbee & Syverson (2023) conclusion that construction productivity is flat or declining is broadly consistent with Brooks & Liscow (2023), and with our findings that the prices of lane miles and pavement quality are rising.

²Small & Winston (1988) find that the cost to resurface a lane-mile is 113,000 USD1984 per lane-mile for

higher than our estimates, which range from about 40,000 USD₂₀₁₀ per lane-mile in 1992 to about 75,000 USD₂₀₁₀ in 2008 for an average (not urban) lane-mile.

Finally, Smith et al. (1997) and Smith et al. (1999) investigate factors that affect highway construction costs. Smith et al. (1997) is based on a 1996 survey of state transportation departments. This survey asked respondents to evaluate the effect that various types of federal regulation had on the costs of highway construction. These surveys suggest that wetlands, historic sites, endangered species and hazardous waste sites were all associated with higher costs. Smith et al. (1999) assembles Highway Statistics data from 1990-4 describing construction expenditure and lane miles of all public roads (not just Interstates). Using a research design similar to our construction regression and Brooks & Liscow (2023), they investigate how expenditure responds to the count of endangered species in the state-year, to the number of environmental impact statements performed in the state-year, to the number of superfund candidate sites in the state and to the count of national historic register places in the state. They find suggestive evidence that environmental regulation drove up construction costs.³ Our results are based on much more extensive data and do not support this conclusion.

Our analysis of the user cost of the Interstate is organized around an optimal capital stock problem. This exercise seems to have few precedents, although Keeler & Small (1977) resembles it conceptually. Keeler & Small (1977) calibrate a theoretical model developed by Mohring (1970) to estimate the optimal level of highway provision in a fully dynamic model. Keeler & Small (1977) is more general than our model in that it also specifies the value of the highway network. On the other hand, it provides estimates of construction and maintenance costs on the basis of nine California counties between 1947 and 1972, whereas we use more complete national data from 1984 until 2008 and distinguish between new construction, resurfacing, and other maintenance.

3 Interstate cost function

To develop a cost function for the Interstate, we generalize a conventional cost function in three ways. First, instead of producing a vector of goods at a particular time, our ‘firm’, a

urban roads, converting to 2010USD using the Producer Price Index for All Commodities (PPIACO), gives 200,000.

³For clarity, we note that Smith et al. (1999) rely on Highway Statistics Table SF12 to measure construction costs. As we discuss below, this table aggregates the SF12a data that we rely on to measure construction and resurfacing expenditures separately.

highway planner, produces Interstate services, VMT, in each period. Second, in addition to a description of how Interstate services are produced from lane miles and pavement quality, our planner’s technological constraint includes the law of motion for lane miles and pavement quality. Third, rather than making a single decision to minimize one-time costs, the highway planner makes a sequence of decisions, investments in lane miles and capital, to minimize the discount present value of these investments.

These generalizations yield a cost function describing the minimum discounted present value of investment required to supply a given path of Interstate VMT. From this cost function, we can derive the marginal cost of Interstate services at any time for a given path of VMT provision. In particular, we can evaluate the marginal cost of Interstate services along the observed VMT path. Our empirical work revolves around estimating realized changes in the prices of lane miles and pavement quality between 1994 and 2008 and then using these prices to calculate the implied series of user costs. The resulting path of user costs for VMT provides a basis for answering the question posed in the title. That is, does the trajectory of user costs suggest that the US is becoming less efficient at providing transportation infrastructure?

We begin by stating the technological constraint on the highway planner’s choices. This constraint has two parts. The first describes the evolution of the stock of lane miles and pavement quality. The second describes the transformation of lane miles and pavement quality into VMT.

Let t index years and L_t the lane miles of Interstate in year t . Investment in lane miles, denominated in real units, is I_t^L and has price p_t^L . Thus, expenditure on lane miles in year t is $p_t^L I_t^L$. In practice, once built, a lane-mile of Interstate does not ever leave the system, so the equation of motion for lane miles is:

$$L_{t+1} = L_t + I_t^L.$$

To discuss the evolution of pavement quality, we must first define it. The International Roughness Index (IRI) is the Federal Highway Administration’s (FHWA) primary measure of pavement quality. It is defined as the number of inches of suspension travel a typical vehicle would experience while traveling one mile (Federal Highway Administration, 2016). A newly resurfaced Interstate segment rarely has an IRI below 50, and the FHWA considers roads to be in *good*, *acceptable*, or *poor* condition as their IRI is below 95, between 95 and 170, or above 170 inches respectively (US Department of Transportation, 2013). For our purpose,

define the ‘quality of the Interstate’ as the lane mile weighted average of IRI over the whole network, and denote average IRI in year t as q_t . Note that pavement quality is decreasing in q , so we will sometimes work with inverse IRI, q^{-1} .

Pavement degrades with use. The conventional measure of the intensity with which a road is used is Average Annual Daily Traffic (AADT), the number of vehicles passing over a given segment on an average day during a year. Although AADT is what is reported in our data, for our purpose, it is more convenient to rely on a measure of annual usage. In particular, if we let v_t denote total Interstate VMT in year t , then average annual use is simply total VMT divided by total lane miles, $= v_t/L_t$. This is 365 times the system average AADT. Roughness increases approximately in proportion to average annual daily traffic at rate κ . So annual increase in roughness is $q_{t+1} - q_t = \kappa v_t/L_t$.⁴

Denote investment in IRI, i_t^q , in real terms. Because IRI is measured in inches, the units of i_t^q are also inches. Investment in IRI involves periodic resurfacing of highway segments. It causes a reduction in IRI and comes at price p_t^q . Because IRI is a system average, total expenditure on IRI in year t is $p_t^q i_t^q L_t$. Summing up, the equation of motion for pavement quality is

$$q_{t+1} = q_t + \kappa v(q_t^{-1}, L_t) L_t^{-1} - i_t^q.$$

In words, IRI at year $t + 1$ is IRI at year t , plus increased roughness resulting from use, minus decreases from expenditure on resurfacing.

The depreciation rate, κ , requires further discussion. Let q_0 denote IRI immediately following a resurfacing event and let q_f denote a terminal IRI immediately prior to resurfacing. A section of highway is engineered to withstand K standardized loadings. Following the engineering literature, denominate these loadings as ‘equivalent standard axle loads’ (ESALs), each of which reflects the passage of a typical tractor trailer rig or 2000 passenger cars. Thus, $\kappa \equiv \gamma \frac{q_f}{K} \frac{q_0}{K}$ is a scalar that describes the relationship between average annual traffic and inches of roughness in two steps: γ relates average annual traffic to ESALs and $\frac{q_f}{K} \frac{q_0}{K}$ relates ESALs to changes in IRI. We postpone the calculation of κ to Section 7.

VMT depends on pavement quality and system length according to $v_t = v(q_t^{-1}, L_t)$. We assume that v is constant returns to scale in lane miles and pavement quality (inverse IRI)

⁴This description of the depreciation process is broadly consistent with the more detailed depreciation functions reported in Small & Winston (1988) and Mannering et al. (2007), with two caveats. First, the engineering literature relies on more complicated functions in order to allow the marginal damage of a loading to vary with current road condition and pavement attributes. Second, because damage is sensitive to axle weight, the engineering literature typically considers several classes of users (e.g., combination trucks, single-axle trucks), while we aggregate to a single class.

and increasing in both its arguments. In general, we require no further assumptions on v . For some of our calibration exercises, however, we assume that v takes the form:

$$v(q_t^{-1}, L_t) = A_t (q_t^{-1})^{-\alpha} L_t^\alpha \quad . \quad (3.1)$$

Here, the parameter A_t is a scaling parameter to map IRI and lane miles into VMT. To facilitate calibration, A_t is time varying. The parameter $\alpha < 1$ determines the relative importance of pavement condition and lane miles for the level VMT. We discuss these assumptions below. We are agnostic about the interpretation of $v(q_t^{-1}, L_t)$. Our analysis is consistent with regarding it as either a production function or a demand function.

The highway manager chooses each period's investment in lane miles and quality to produce a given VMT path $(v_t)_{t=0}^T$ in the cost minimizing way. Letting r denote the real interest rate, we can state the highway manager's cost minimization problem as,

$$\begin{aligned} C((v_t)_{t=0}^T; L_0, q_0) &= \min_{I_t^L, I_t^Q} \sum_{t=0}^T \frac{1}{(1+r)^t} (p_t^L I_t^L + p_t^Q I_t^Q L_t) & (3.2) \\ \text{subject to } v_t &\leq v(q_t^{-1}, L_t) \\ L_{t+1} &= L_t + I_t^L \\ q_{t+1} &= q_t + \kappa v(q_t^{-1}, L_t) L_t^{-1} - i_t^Q . \end{aligned}$$

In this problem, the Lagrange multiplier (or shadow price), τ_t , for the constraint involving v_t , is the marginal reduction in cost, in period t dollars, from reducing the amount of VMT provided at period t . It is the marginal cost of providing VMT in period t , conditional on optimizing behavior by the planner and satisfaction of the planner's other constraints. The vector $(\tau_t)_{t=0}^T$ describes the trajectory of the marginal cost of VMT and is the object of interest.

Minimizing the present value of cost subject to a sequence VMT levels, gives rise to the following Euler equations:

$$p_t^L = \frac{1}{1+r} (\tau_{t+1} v_L + p_{t+1}^L - p_{t+1}^Q (q_{t+1} - q_{t+2}) - p_{t+1}^Q \kappa v_L) \quad (3.3)$$

$$p_t^Q L_t = \frac{1}{1+r} (\tau_{t+1} v_Q q_{t+1}^2 + p_{t+1}^Q L_{t+1} - p_{t+1}^Q \kappa v_Q q_{t+1}^2) , \quad (3.4)$$

where $v_L = v_L(q_{t+1}^{-1}, L_{t+1})$ and $v_Q = v_Q(q_{t+1}^{-1}, L_{t+1})$, are the marginal VMT produced from additional lane miles and pavement quality respectively.

The cost-minimizing allocation of lane miles L and pavement quality q satisfy the following steady state first-order conditions:

$$rp^L = \tau v_L - p^q \kappa v_L \quad (3.5)$$

$$rp^q q L = \tau v_q q^{-1} - p^q \kappa v_q q^{-1}. \quad (3.6)$$

Under constant returns to scale, summing these expressions gives:

$$\tau = \frac{rp^L L + rp^q q L + \kappa p^q v}{v} \quad (3.7)$$

which implies that the shadow price of providing steady state Interstate services $\bar{v} = v(q^{-1}, L)$ is the sum of the capital cost of system lane miles, the capital cost of pavement quality, and expenditures to offset depreciation.

Equation (3.7) requires two comments. First, this equation demonstrates the importance of our assumption that v is constant returns to scale. Without this assumption, equation (3.7) would involve derivatives of v , about which little is known, instead of readily observable v . Second, equation (3.7) describes a relationship between equilibrium quantities, and so is not a natural starting point for calculating comparative statics.

Note the relationship between our model and a more basic description of increasing returns to scale. In the basic increasing returns to scale problem, a lumpy fixed investment produces output at low or zero marginal costs. This is exactly the structure of our problem. A lumpy fixed investment in highways provides a flow of services. Our problem differs from the more elementary problem in that we provide an explicit model of how the fixed investment depreciates, and we distinguish between services provided at different times.

In the interests of clarity, the discussion above omits flow maintenance costs, like traffic management and snow removal. Let p^m denote maintenance costs per VMT. To incorporate maintenance costs into our analysis, we add the term $p^m v_t$ to the objective in (3.2). Let $\hat{\tau}$ denote the resulting steady state user cost. Noting the p and $\hat{\tau}$ both appear multiplying the v_t in the associated Lagrangian, it is clear that $\tau = \hat{\tau} + p^m$. Thus, we can accommodate flow maintenance costs in (3.7) by subtracting p^m from the right hand side. Given this, we will see below that p^m is small enough that it is reasonable to ignore it altogether.

Discussion There is little empirical basis for thinking about how pavement quality contributes to travel costs. We justify our assumption of constant returns to scale on two bases. First, it allows us to describe the interstate cost function in a simple, transparent way and

relieves us of the need to estimate partial derivatives of v (about which essentially nothing is known). As the empirical foundations for v improve, our method generalizes in a straightforward, but probably complicated way. Second, Couture et al. (2018) find that the speed of travel in a city is close to constant returns to scale in lane miles and travel time.⁵ The finding in Duranton & Turner (2011) is also relevant. Duranton & Turner (2011) find that total VMT in a city is approximately proportional to lane miles; together with our assumption that v is constant returns to scale, this requires that travel not be very responsive to pavement quality. In the context of (3.1), this requires that α is small and so that $1 - \alpha$ is close to one.

We consider only the costs of the Interstate that are born directly by the government. Two practical considerations motivate this restriction. First, the policy debate centers on the costs born by the public purse, and so we are focusing on issues immediately relevant to this debate. Second, considering the various private and external costs associated with Interstate VMT would greatly increase the complexity and data requirements of our analysis.

Our framework can be modified to provide a foundation for observed growth in VMT and for congestion. To address both issues, we consider the possibility that v is constant returns to scale in pavement quality and lane miles, but decreasing returns to scale in population. A particular formulation of v with these properties is,

$$v = \left(\frac{1}{q_t} \right) L_t^\alpha N_t^{1-\alpha} \quad (3.8)$$

This production function requires that when population grows at a constant rate, then either lane miles grows or VMT per person declines. Using this production function in our earlier statement of the planner's problem, we can show the existence of a balanced growth path where lane miles grow at a constant rate that is a function of population growth. In particular, lane mile growth g satisfies:

$$1 + g = (1 + n)^\alpha, \quad (3.9)$$

where n is the growth rate of the population. In this case, we obtain an analogous steady state condition where τv equals the sum of: opportunity cost of lane miles $rp^L L_t$, the opportunity cost of pavement quality $\frac{r-g}{1+r} p^q q L_t$, and the cost of depreciation $\kappa p^q v$. This model is a generalization of the baseline model; setting $n = 0$ returns the original steady state condition.

⁵Precisely, Couture et al. (2018) estimate that the average speed of travel declines by 15% when the resources devoted to travel double.

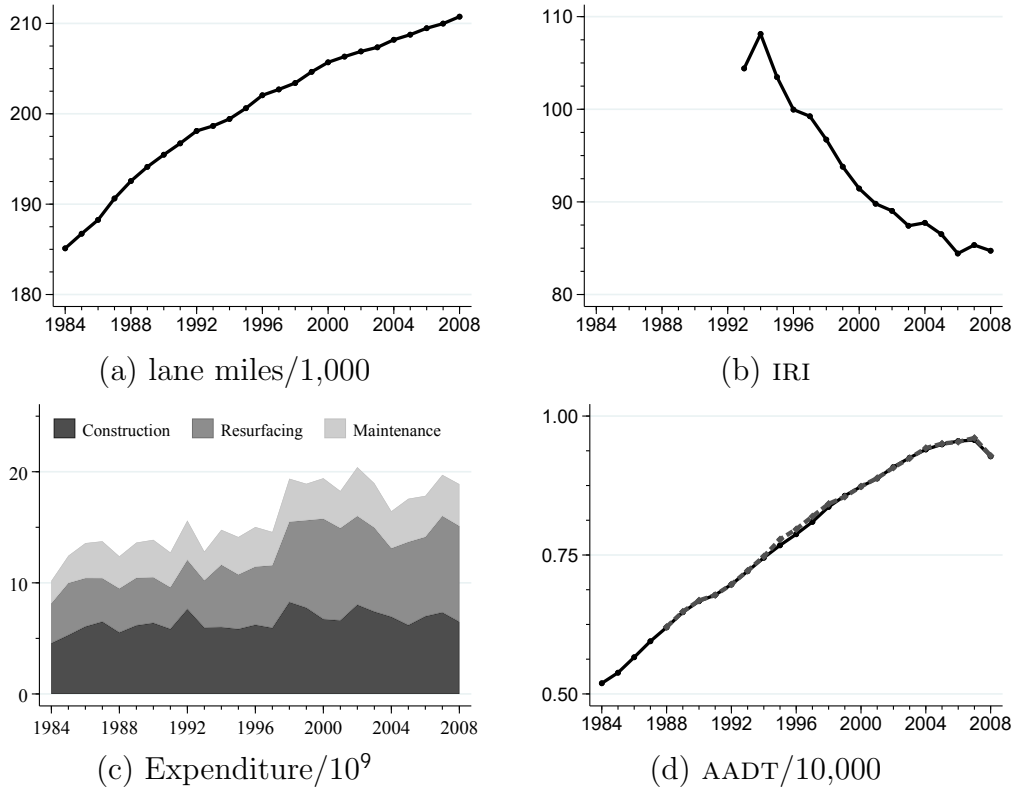
We draw attention to this generalization of our model to demonstrate the possibility, and subtlety, of tailoring our framework more closely to the economic fundamentals of the underlying asset, in our case, the Interstate highway system. We do not consider this model as a basis for a calibration exercise for two reasons. First, to do so would require that we speculate about the population elasticity of VMT, β . Second, Figure 1 presented below demonstrates that VMT and highway miles are not growing at the same rate as population, so it is difficult to defend the claim that current data has converged to a balanced growth path.

A foundational assumption of our approach is that the cost of ‘Interstate services’, here VMT, should be the object of analysis. Other measures of output are also of interest. For example, the vector consisting passenger VMT and ton-miles of freight, or ‘lane-mile-days’, or a vector of VMTs differentiated by region or origin and destination. Our focus on VMT has two advantages; it is a conventional measure of the amount of service provided by the Interstate and data is easily available. Generalizing to other measures of output is easy conceptually, though probably difficult in practice. With this said, note that given the user cost of VMT, it is natural to evaluate the cost of a highway segment on the basis of the marginal VMT that it provides. In this way, our analysis can be applied immediately to project level cost-benefit analyses.

4 Data

We would like to evaluate the user cost of the Interstate as given in equation 3.7 or the non-steady state analog (3.3, 3.4). Most of the variables in these expressions, like lane miles, can be easily observed. Measuring the prices of lane miles and pavement quality, however, is more difficult. To estimate the price of new lane miles and IRI, we require data describing the extent and condition of the Interstate network, the quantity and timing of expenditure, and road characteristics that may affect construction and resurfacing costs. We construct two data sets. The first is organized by segment-year, and we use it to estimate the price of pavement quality. The second is organized by state-year, and we use it to estimate the price of new lane miles. This section describes how we construct these data sets and how the Interstate system evolved over our study period.

Figure 1: Trends in the state of the Interstate highway system



Note: (a) Total Interstate lane miles by year ('000 miles). (b) Lane mile weighted IRI for the whole Interstate by year, (inches per mile). (c) Total Interstate expenditure by year and category. Dark gray is construction, medium gray is resurfacing, light gray is maintenance (10^9 2010USD). (d) Lane mile weighted AADT for the whole Interstate by year (vehicles/day). (a), (b) and the solid line in (d) are based on the HPMS Universe Data. The dashed line in figure (d) is based on the HPMS Sample Data. (c) is based on Highway Statistics.

4.1 Lane miles and IRI

The federal government requires state highway authorities to keep segment-level annual inventories of the system and report them to the Federal Highway Administration. The resulting data are the Highway Performance and Monitoring System maintained by the US Office of Highway Policy Information.⁶

The HPMS consists of two annual data sets, the ‘Universe’ and ‘Sample’ data sets. Both

⁶See, for example, <https://www.fhwa.dot.gov/policyinformation/hpms/shapefiles.cfm>. Our HPMS data came to us on a CD produced by personnel from the Office of Highway Policy Information. In fact, the HPMS tracks all roads for which the federal government has financial responsibility, but the HPMS maintains greater detail about the Interstate network than other federally funded roads.

are available in a consistent format from 1980 until 2008, with 2009 available for a subset of states.⁷ Both are organized by segment-year and have the same basic structure. The Universe Data provides a basic description of every Interstate segment in every year. The Sample Data provide a more detailed description for a random sample of Interstate segments.

We restrict attention to the 48 states of the continental US. Because only about half of the states report any HPMS data in 2009, we end our study period in 2008. Beginning in 1988, the Sample Data required states to report IRI, for every segment. States were slow to comply with the new reporting requirement, and IRI reporting is substantially incomplete until 1992, when we begin our analysis of the price of IRI. Our expenditure data, described below, does not begin until 1984, and so our analysis of the price of lane miles begins in this year, four years after the start of the HPMS Universe Data.

The Universe Data form the basis for our estimates of the price of new lane miles. These data report the length and number of lanes for every segment of the Interstate in every year, allowing us to calculate lane miles of Interstate by state-year. Road segments are rarely promoted to (or demoted from Interstate status) and the HPMS tracks such status changes at the segment level. This allows us to avoid confusing changes in the administrative status of roads with the construction of new lane miles. The HPMS does not record segments leaving the Interstate system for any other reason. This means that we can measure new construction of lane miles as year-over-year increases in state lane miles.

The Sample Data form the basis for our estimates of the price of pavement quality. As their name suggests, the Sample Data provide a detailed description of a carefully constructed random sample of Interstate segments (Office of Highway Policy Information, 2016), along with sampling weights that permit the construction of state-year means. Each segment is identified by a unique segment-ID, and we are able to track these IDs over time. Figure 1(d) reports lane-mile weighted AADT by year calculated from the Universe Data and estimated from the Sample Data. The close agreement between the two series validates the HPMS sampling methodology.

Over time, the accuracy with which a sample of segments represents the Interstate-network deteriorates as the characteristics of the sample and the population diverge. In addition, changes in the network need not reflect segment definitions. For example, adding a lane to half the length of a sample segment requires the creation of two ‘sub-segments’ to keep

⁷The HPMS went through three revisions between 1980 and 2009. These revisions preserved the basic structure of both data sets. During 2009-10, the Federal Highway Administration converted the HPMS from its original tabular form to a GIS based data model. As a consequence of this conversion, data for 2010 is not available, and post-2010 HPMS data is not directly comparable to the older data.

track of the change, and so the complexity of any given sampling frame increases over time. To address these problems the HPMS periodically updates the population of segments and segment-IDs from which the Sample Data are drawn. This sometimes interrupts our ability to track particular segments. Because new segment-ID's can reflect either new construction or a revision of the sampling frame, we cannot use the Sample Data to track new construction at the segment level.

The HPMS Sample Data does not report expenditures on highways. However, for each segment-ID and year, they report a categorical variable indicating whether a segment experienced one of 14 different types of improvement, and this classification system is stable across years.⁸ Of these 14 categories, 10 refer explicitly to reconstruction, restoration, rehabilitation or resurfacing. We code segments that experience these improvements as being 'resurfaced' during the relevant year. We classify three categories of improvements as construction; 'major widening, 'new route' and 'relocation'.⁹ Only one of the 14 improvement categories remains, 'minor widening'. Because HPMS codebooks typically list 'minor widening' under the sub-heading of reconstruction, itself a resurfacing event, we also count 'minor widening' as a resurfacing event.

4.2 Investment

The Federal Highway Statistics series contains annual reports of expenditure on the national highway system.¹⁰ Highway Statistics Table SF12 reports total state and federal Interstate expenditure by year.¹¹ Table SF12a allows us to decompose total expenditure. This table reports expenditure on 'Right of Way and Engineering', 'New Construction' and 'Major Widening'. We sum these three categories for our measure of construction expenditure. Table SF12a also reports expenditure on 'Reconstruction' and 'Rehabilitation, Restoration and Resurfacing'(3R). We sum these two categories for our measure of resurfacing expenditure. Note the close, and probably purposeful, correspondence between categories of expenditure

⁸See, for example, *Archive Highway Performance Monitoring System (HPMS) Data Item Descriptions: 1993 - 1998*, item 50.

⁹We do not attempt to use segment level data on 'major widening, 'new route' or 'relocation' to estimate the amount of new construction. This is possible in theory, but these events are so rare that resulting estimates of state-year totals are too noisy to form a basis for analysis.

¹⁰These reports are available from the Federal Highway administration almost continuously from 1946 until the present at <https://www.fhwa.dot.gov/policyinformation/statistics.cfm> and <https://www.fhwa.dot.gov/policyinformation/hsspubsarc.cfm>.

¹¹In fact, Table SF12 reports total expenditure under two main headings, 'capital outlay' and 'maintenance'. Despite their names, the capital outlay and maintenance expenditure reported in Table SF12 does not correspond neatly to new construction and resurfacing.

in Highway Statistics Table SF12a and the categories of improvement in the HPMS. HPMS improvement categories map transparently into categories of expenditure reported in Table SF12a.

Table SF12a separately reports expenditure on ‘Bridge Work’ and the HPMS does not report on bridges at all. Our measure of maintenance expenditure is the difference between the sum of resurfacing and construction expenditure, and total expenditure net of expenditure on bridges. In this way, we use the categories reported in Tables SF12 and SF12a to classify expenditure to correspond with the new construction and resurfacing that we observe in the HPMS.¹²

Figure 1 (c) shows annual expenditure on the Interstate system by year, across the three classes of expenditure: new construction, resurfacing and maintenance. The height of the bottom region indicates expenditure on construction in billions of 2010USD. The intermediate region indicates expenditure on resurfacing and the upper region indicates expenditure on maintenance. The upper envelope of the figure indicates total expenditure.¹³ Between 1984 and 2008, total annual expenditure increases from 10.2 billion to 18.9 billion, expenditure on construction increases more slowly than does expenditure on resurfacing, and maintenance expenditure is about constant. In 2008, expenditures on maintenance, resurfacing and new construction were 3.8, 8.6, and 6.5 billion. This is 20%, 46%, and 34%, respectively.

Highway Statistics begins reporting the detailed expenditure breakdown of Table SF12a in 1984. This is well before the 1992 start of complete IRI reporting in the Sample Data, but four years after the 1980 start of the Universe Data.

We merge the HPMS data sets and Highway Statistics by state-year. One can imagine that Highway Statistics might record expenditures in a different year than the HPMS records the associated road work. An indicator of this problem would be ‘impossible state-years’ in which either expenditure occurs in Highways Statistics but there is no new construction or resurfacing in the relevant HPMS data, or no expenditure occurs in Highways Statistics but we observe new construction or resurfacing in the relevant HPMS data. Appendix A discusses this issue in detail. Briefly, this problem is rare in the data that matches Highway Statistics and the HPMS Universe Data. However, it affects about 30% of state-years in the data that matches Highway Statistics and the HPMS Sample Data. This appears to primarily reflect the fact that the HPMS Sample Data describes a *sample* of segments, while Highway

¹²For more detail on bridge expenditure and maintenance, see Duranton et al. (2020).

¹³Between 1998 and 2008, our data (not shown) report expenditure on right of way separately. During this period, right of way expenditures are only 10-15% of construction expenditures and do not show an obvious trend.

Table 1: Lane mile weighted means of network characteristics

	1984	1992	2008
% Urban (HPMS)	30.2	33.5	42.7
% Urban (NLCD)	13.1	13.2	13.3
Grade (HPMS)	.	1.2	1.1
Water (NLCD)	7.3	7.4	7.5
Elevation	456.2	449.0	440.1
% New miles	0.7	0.5	0.2
Structural Number	.	6.6	6.9
% Flexible	.	21.8	24.3
% Rigid	.	40.9	26.8
% Composite	.	37.3	48.9
Unionization	24.7	20.1	15.9

Note: Variables from the HPMS Universe or GIS data are reported for 1984, 1992, and 2008. Those based on HPMS Sample Data are reported for 1992 and 2008.

Statistics describes all expenditure.

4.3 Segment and network characteristics

Much of our data on system attributes derives directly from the HPMS. The Universe Data reports segment length and number of lanes by state-year. Aggregating, we obtain the estimates of system length reported in figure 1(a). The Interstate consisted of about 185,000 lane miles in 1984, increasing to about 210,000 by 2008 primarily by the addition of expansion lanes to the existing network. The Universe Data reports AADT. To calculate lane-mile weighted average AADT, we multiply segment level AADT by the number of lanes and length, sum over segments, and divide by system lane miles. This yields the solid line in figure 1(d). Lane mile weighted mean AADT increased from about 5200 vehicles per day in 1984, to nearly 9000 vehicles per day in 2008. Panel (b) reports the lane-mile weighted average of the International Roughness Index (IRI) for the whole Interstate highway network. Mean IRI declines from about 110 inches per mile to about 85 inches per mile between 1992 and 2008. The surface quality of the Interstate system has improved over time, from just above the good-acceptable threshold to just below.

Using the Universe Data, we estimate the share of all new construction that is new route miles as opposed to expansion lanes on existing routes.¹⁴ Table 1 reports this share in 1984,

¹⁴We observe the change in route mileage and the change in lane miles for each state-year. If we assume that all new mileage in a state-year has the same number of lanes as an average segment in the preceding

1992 and 2008. It is less than 1 percent throughout our study period and trending down. Most Interstate construction during our study period involves the expansion of existing routes. The Universe Data also report whether each segment is urban or rural according to whether it lies in an urbanized area or not.¹⁵ Table 1 shows that the share of urban lane miles increased from 0.30 to 0.43 between 1984 and 2008. This change partly reflects the construction of lane miles in urbanized areas and partly reflects the expansion of urbanized area boundaries.

The Sample Data provide a detailed description of each segment, e.g., shoulder width, subsurface drainage. We focus attention on a handful of variables that are likely to affect construction or resurfacing costs. These are, grade, construction type, and structural number (defined below).

Table 1 reports the lane-mile weighted mean grade calculated in 1992 and 2008. New construction shifts toward flatter areas over time.

The Sample Data reports a categorical variable describing construction type as either ‘rigid’, ‘composite’, or ‘flexible’. A ‘rigid’ segment is one that consists primarily of steel reinforced concrete slabs. A flexible segment is one that consists primarily of asphaltic concrete, i.e., blacktop. A composite road consists of a combination of such layers, for example, a layer of asphaltic concrete over a concrete base. The share of flexible and composite lane miles increases at the expense of rigid roads. Much of this change probably reflects the conversion of rigid roads to composite by the addition of an asphaltic concrete layer during a resurfacing event.

Closely related, the Sample Data reports the ‘structural number’ of each segment. Structural number is an engineering index used to measure the durability of a road (Mannering et al., 2007, Ch. 4). It is a weighted sum of the thicknesses of the various layers of gravel, concrete and asphaltic concrete that make up each segment.¹⁶ For example, each inch of asphaltic concrete contributes about 0.41 to a segment’s structural number, depending slightly on the quality of the material. Over the period during which we observe these data, 1992 to 2008, lane mile average structural number increases from about 6.6 to about 6.9. This increase is consistent with the construction of progressively more durable roads or the accumulation of paving material as a consequence of ongoing resurfacing. Because a one inch layer of asphaltic concrete will contribute about 0.4 to the structural number of a road, the

year, we can use this value to estimate the share of all new lane miles that are part of new segments.

¹⁵The Federal Highway Administration maps of urbanized area are based on the corresponding census maps, but are slightly adjusted (Federal Highway Administration, 2013).

¹⁶Structural number is simply the thickness of concrete in inches for rigid roads.


trends in structural number are consistent with an average Interstate lane-mile consisting of about an extra 0.75 inches of asphaltic concrete in 2008 than in 1992.

To investigate the role of exposure to unionized labor markets, we rely on the Current Population Survey's report of the share of the labor force that is in a union by state and year.¹⁷ Table 1 reports a lane-mile weighted national mean of these state level unionization shares. The dramatic decline in this mean reflects both changes in the national unionization rate and changes in the distribution of lane miles across states.

We also calculate network attributes from GIS data. Starting from the 2005 NHPN planning map of the Interstate (Federal Highway Administration, 2005), we create a buffer extending 2.5 miles on either side of the Interstate. We use this buffer to calculate the attributes of land within the buffer from three GIS based data layers. First, from the 2001 NLCD (United States Geological Survey, 2011), we calculate the share of land within the buffer that is classified as urban. Second, also from the 2001 NLCD, we calculate the share of land within the buffer that is water or wetlands. Third, using a digital elevation map (United States Geological Survey, 2010), we calculate the mean elevation of the Interstate within a state, as of 2005. Note that these measures vary only at the state level and time series variation in national means entirely reflects changes in how the lane miles are distributed across states. Table 1 shows that over time progressively larger fractions of Interstate lane miles lay in states that had more urban cover in a buffer near the 2005 Interstate, that had more water or wetlands in this buffer, and where the route of the 2005 Interstate was at a lower elevation. In particular, in 1980, the 2.5 mile buffer strip on either side of an average lane-mile of Interstate was about 13.1% in urban cover and 7.1% water and wetlands cover in 1980, and these shares increased slightly but steadily to 13.3% and 7.5% by 2008. The elevation of a similarly average strip fell by about 40 feet over this time.

Summing up, over the course of our study period, the network shifted towards areas that were flatter, lower, wetter and more urban. The Interstate's exposure to unionized workers decreased dramatically. New construction became even more focused on expansion lanes rather than new mileage, the structural number of an average lane-mile increased and the type of surface shifted from rigid toward flexible pavement.

Figure 2: Effect of resurfacing expenditure on IRI by year



Note: x-axis is years, y-axis is inches per resurfaced mile from one million dollars per resurfaced mile of resurfacing expenditure. The solid line plots the trend in change in IRI and millions of dollars of expenditure per lane mile between 1992 and 2008 estimated in column 6 of table 2. This plot is the basis for the time series of p^q that we use in our calibration exercise in section 7 and report in table B4.

Table 2: Resurfacing expenditure and IRI

	(1)	(2)	(3)	(4)	(5)	(6)
$1_{ist(q)}i_{st}^q$	-619.29 (38.80)	-607.60 (38.07)	-646.12 (41.90)	-921.00 (93.70)	-922.91 (92.92)	-992.86 (104.83)
t				-0.02 (0.07)	-0.02 (0.07)	0.00 (0.08)
$1_{ist(q)}i_{st}^q \times t$				27.46 (7.01)	27.29 (6.89)	29.96 (7.64)
State FE	No	No	No	No	Yes	No
State-Year FE	No	Yes	Yes	No	No	No
Segment id FE	No	No	Yes	No	No	Yes
N	186,055	186,054	181,235	186,055	186,055	181,236

Note: Estimations of variants of equation (5.2). We drop segments that occur in just one year in specifications that include segment fixed effects, columns 3 and 5. Standard errors in parentheses, clustered at the state-year Level. $^+ p < 0.10$, $p < 0.05$, $p < 0.01$, $p < 0.001$.

5 The price of pavement quality and lane miles

We now turn to estimating the annual price for the reduction of Interstate roughness, p_t^q , and the price of a new lane-mile of Interstate, p_t^l . These prices are of independent interest

¹⁷Data constructed by Hirsch & MacPherson (2003) updated annually at unionstats.com.

and are inputs into the calculation of the user cost of the interstate.

5.1 The price of pavement quality, p^q

Our estimates of the price of IRI are based on data organized by road segment, state and year; $j \in J$, $s \in \{1, \dots, 48\}$ and t . Let L_{jst} indicate total lane miles of Interstate highway for segment j in state s and year t , let L_{st} indicate lane miles of Interstate highway in state s and year t , and let Δ indicate first differences. Thus, $\Delta L_{st} = L_{st} - L_{st-1}$ is change in state lane miles. We rely on segment-year level measurements of IRI, q_{jst} , to measure pavement quality. Let x_{jst} denote a vector of other attributes of a given segment highway, and x_{st} the corresponding state-year aggregate. Let $\mathbf{1}_{jst}(q)$ be an indicator for whether a segment was resurfaced in each year.

Let I_{st} indicate total expenditure for a state-year, and I_{st}^L , I_{st}^q , and I_{st}^m , be expenditure on new lane miles, expenditure on resurfacing, and expenditure on maintenance respectively.

Our data on resurfacing and IRI is at the level of a segment-year, while our expenditure data is at the state-year level. Our challenge is to devise a regression framework that allows us to use these data to estimate a yearly national average price of IRI.

We can easily estimate how the effect of resurfacing on IRI changes over time,

$$\Delta q_{jst} = C_0 + C_1 \mathbf{1}_{jst}(q) + C_2 [\mathbf{1}_{jst}(q)t] + C_3 x_{jst} + \epsilon_{jst}. \quad (5.1)$$

C_1 is the conditional mean difference in IRI between resurfaced and unresurfaced segments when $t = 0$ (1992) and C_2 is the rate at which this difference changes over time. The x_{jst} represents a subset of the controls; state indicators, year indicators, state-year indicators and segment indicators.

In equation (5.1), C_1 and C_2 describe the time path of the effects of *resurfacing* on pavement quality. To estimate a time path of the effect of *expenditure on resurfacing* on pavement quality, let L_{st}^q denote lane miles of Interstate resurfaced in a state-year and calculate millions of dollars of resurfacing expenditure per resurfaced mile as,

$$v_{st}^q \equiv \frac{I_{st}^q}{L_{st}^q}.$$

We regress change in IRI on the interaction of resurfacing expenditure per mile and the

resurfacing indicator,

$$\Delta q_{jst} = A_0 + A_1 [1_{jst}(q) \iota_{st}^q] + A_2 [1_{jst}(q) \iota_{st}^q t] + A_3 x_{jst} + \epsilon_{jst}. \quad (5.2)$$

Because the left hand side is denominated in inches per mile and the units of $1_{jst}(q) \iota_{st}^q$ are millions of dollars per resurfaced lane-mile, the units of A_1 are inches per million dollars. A_2 is the same as A_1 , but it measures the rate at which A_1 changes, i.e., inches per million dollars per year. Thus, $A_1 = 1/p_{1992}^O$ and $A_1 + A_2 t = 1/p_{1992+t}^O$, so that a simple transformation lets us calculate the time path of the price of IRI, p_t^q , from an estimate of equation 5.2.

We experiment with other parameterizations of the trend in the price of IRI. The data do not allow us to determine whether the rate of change is different in different parts of our study period. Given this, we restrict attention to the linear specification.

Figure 2 illustrates the evolution of the effect of resurfacing expenditure on pavement quality. It reports the coefficients of the following regression,

$$\Delta q_{jst} = \sum_{=1992}^{2008} A [1_{jst}(\tau = t) 1_{jst}(q) \iota_{st}^q] + \epsilon_{jst}. \quad (5.3)$$

The units of IRI and ι_{st}^q are inches per mile and millions of dollars per resurfaced mile. It follows that the units for the A are inches per million dollars. As in regression (5.2), these regression coefficients are inverse prices. Confidence intervals are based on standard errors clustered at the state-year level.

Figure 2 shows a clear upward trend. In the early 1990s, one million dollars of expenditure reduced IRI by about 900 inches. By the end of our sample, the same million dollar expenditure reduced IRI by about 450 inches. That is, the raw data suggest that the price of reducing IRI about doubles between 1992 and 2008.¹⁸

Table 2 estimates the effect of one million dollars per resurfaced mile of resurfacing expenditure on IRI, that is, equation (5.2). Interpreting these results requires careful attention. Decreases in IRI are good, so if the price of resurfacing goes up, the coefficient A_1 of $1_{jst}(q) \iota_{st}^q$ will increase to become a negative number with a smaller magnitude. Second, the units for

¹⁸For completeness, Figure B1 illustrates the evolution of the effect of resurfacing on IRI. For almost all years a resurfacing event reduces the IRI of a segment by between 20 and 40 inches per mile. Confidence intervals for the different years usually overlap and this figure shows at most a small positive trend. Resurfacing events in 1992 were not much different than in 2008. Table B1 estimates variants of equation (5.1). On average, resurfacing reduces IRI by about 33 inches. There is a small positive change in the effect of resurfacing expenditure that is barely distinguishable from zero. That is, resurfacing results in a slightly smaller reduction in IRI in 2008 than in 1992.

A_1 are inches per mile, per million dollars of expenditure per lane-mile. This is an inverse price, so as A_1 increases in magnitude the price of IRI falls. Similar comments apply to interpreting the coefficient of $1_{ist}(q)_{st}^q$.

In column 1, we estimate that one million dollars per lane-mile of resurfacing expenditure reduces IRI by about 637 inches. This magnitude does not vary as we add state-year indicators in column 2, or segment and state-year indicators in column 3. In column 4, we allow for a trend and an interaction between the trend and expenditure. That the coefficient on the interaction is 27 means that one million dollars of expenditure eliminates 27 fewer inches in each successive year. Thus, in 1992 one million dollars eliminates about 900 inches of IRI. By 2008, this falls to about 400 inches. These estimates are almost unchanged in columns 5 and 6 where we add state and segment indicators.

For later reference, table B1 presents estimates of the relationship between resurfacing events and IRI given in (5.1). Broadly, this table shows that an average resurfacing event reduces segment level roughness by about 35 inches per mile, and that a resurfacing event results in slightly smaller reductions in IRI over time.

5.2 Price of lane miles, p^L

We would also like to estimate the price of new lane miles, p^L . We proceed much as we did for the price of IRI, adjusting for the fact that our data on lane miles is at the state-year level. In particular, we estimate,

$$\Delta L_{st} = A_0 + A_1 I_{st}^L + A_2 [I_{st}^L t] + A_3 t + \epsilon_{st}. \quad (5.4)$$

This equation relates state-year change in lane miles to state-year construction expenditure. We denominate expenditure on lane miles in millions of 2010 dollars per year. Because the dependent variable is measured in lane miles, A_1 gives lane miles per million dollars of expenditure when $t = 0$ (1984). A_2 gives the rate at which this inverse price changes over time. As for our resurfacing regression, this is an inverse price, with $A_1 = 1/p_{1984}^L$ and $A_1 + A_2 t = 1/p_{1984+t}^L$. Increases in A_1 indicate that a million dollars of construction expenditure buys more, so the price is lower.

The data show that the price of new lane miles increases between 1984 and 2008. They do not allow conclusions about whether this rate of increase is faster or slower in different parts of our study period. Given this, as with our IRI regressions, we present only the linear specification. Because these data are relatively coarse, our ability to include control variables

is limited, however, in some specifications, we include state indicator variables.

To describe the increase in construction costs, define $1_{st}(\tau)$ to be one in year τ and zero otherwise. Next conduct the following regression,

$$\Delta L_{st} = \sum_{=1984}^{2008} A [1_{st}(\tau)I_{st}^L] + \epsilon_{st} \quad (5.5)$$

In this regression, the A are the mean number of lane miles constructed per million of expenditure by year. Figure 3 plots these inverse prices by year. This figure shows a decline in the number of lane miles purchased by one million dollars of expenditure.


Table 3 presents regressions based on equation (5.4). Column 1 presents a regression of ΔL_{st} on I_{st}^L . Column 2 adds state fixed effects. The dramatic change in the coefficient of expenditure confirms the importance of state level variation in construction costs documented in Brooks & Liscow (2023). In column 3, we add a trend (year-1984) and an interaction of the trend with expenditure. As suggested by Figure 3, one million dollars buys fewer lane miles in each successive year. Column 4 repeats column 3 but restricts the sample to 1992-2008 in order to match the sample we use to investigate resurfacing. Consistent with what we observe in Figure 3, lane miles per million dollars declines more slowly during the later part of the study period. We use this estimation to calculate the time series of p^q that we use in our calibration.¹⁹ This has a negligible effect on our estimate of the trend downward in lane miles per million dollars of expenditure.

In Appendix B.2 we also experiment with an instrumental variables estimation strategy based on Leduc & Wilson (2013). Table B2 reports on a specification like column 4 of Table 3, but instruments terms involving expenditure with corresponding terms involving the four year lag of total Interstate appropriations. This change in estimating technique has little impact on our estimates of the trend in prices, and reassures us that mis-measurement of expenditure is not causing economically important changes in our results.

It remains to document the level and changes in expenditure on maintenance. Figure B3 shows the results of a regression that is similar to equation (5.5), but which predicts annual maintenance expenditure as a function of year indicators. From the figure, maintenance costs are about 0.01×10^6 or about 10,000\$ per lane-mile per year. These costs have been steady or declining over time.

¹⁹We cannot use the estimation in column 3 for this purpose because the implied value of p_t^q becomes negative at the end of the sample. This reflects the pattern we observe in Figure 3. If we instead allow for different functional forms in Table 3 we arrive at similar estimates for p_t^q during 1992-2008. We revisit this issue in Section 7

Figure 3: Effect of construction expenditure on lane miles by year



Note: x-axis is years, y-axis is lane kilometers per million dollars. 95%CI's based on robust standard errors. The solid line plots the trend in lane miles per million dollars between 1992 and 2008 estimated in column 4 of table 3. This plot is the basis for the time series of p^L that we use in our calibration exercise in section 7 and report in table B4.

6 Explaining the price increases

6.1 Explaining the increase in the price of pavement quality

Our data indicate an increase in the price of pavement quality. To attribute this increase to possible causes, we allow the trend in the inverse price of IRI to vary with segment or state characteristics. For a given segment or state attribute x_{jst}^0 , this leads to a generalization of our earlier estimating equation,

$$\begin{aligned} \Delta q_{jst} = & A_0 + A_1 [1_{jst}(q)\iota_{st}^q] + A_2 [1_{jst}(q)\iota_{st}^q t] + A_3 t \\ & + B_1 [1_{jst}(q)\iota_{st}^q x_{jst}^0] + B_2 [1_{jst}(q)\iota_{st}^q x_{jst}^0 t] + B_3 x_{jst} + \epsilon_{jst}. \end{aligned} \quad (6.1)$$

If $x^0 = 0$ this equation collapses to our earlier regression equation (5.2), and the interpretation of regression coefficients is also similar to equation (5.2). As x^0 varies, B_1 measures the mean change in base year, price and B_2 measures the rate at which the trend in price changes with changes in x^0 . For example, we generally find that if x^0 is a measure of how urban is the state or segment, then $B_1 > 0$ and $B_2 < 0$. This means that, all else equal, one million dollars reduces IRI by a smaller amount on more urban roads in 1992, but that this urban penalty decreases over time.

We now investigate explanations for the upward trend in p_t^q . It is well known that road construction is more expensive in urban areas (Ng & Small, 2012). In Table 1 and Figure 1

Table 3: Construction expenditure and new lane miles

	(1)	(2)	(3)	(4)
I_{st}^L	0.0472 (0.0230)	-0.0008 (0.0134)	0.1135 (0.0328)	0.0512 (0.0363)
t			-0.6487 (0.2910)	-0.4112 (0.3119)
$I_{st}^L t$			-0.0045 (0.0012)	-0.0018 (0.0022)
State FE	No	Yes	No	No
N	1,171	1,171	1,171	808

Note: OLS Estimations of variants of equation (5.4) Standard errors in parentheses, clustered at the state-year Level. All regression are based on the 1984-2008 period when we observe state-year lane miles, except column 4. In column 4 we restrict attention to 1992-2008 to match the study period we use to analyze IRI. $^+ p < 0.10$, $p < 0.05$, $p < 0.01$, $p < 0.001$.

we see that over time the average lane-mile of Interstate is more heavily used, more likely to be designated urban, and is in a state where the area near the 2005 Interstate had a higher fraction of urban cover in 2001. By all three measures, the network becomes ‘more urban’. This suggests that the price of IRI is rising because resurfacing is occurring on more expensive urban roads.

To investigate this possibility, Table 4 presents three estimates of equation (6.1) in which the extra segment attribute is, from column 1 to 3, segment-year level AADT, the HPMS segment-year urban indicator, and the NLCD state level impermeable cover measure. We include segment indicators as controls in all of the results we present in Table 4. In unreported results, we replicate each of the specifications in Table 4 for the combinations of fixed effects that we use in Table 2. Parameter estimates are stable across specifications.

Beginning with column 1, we see that AADT has two effects. First, as expected, the level effect of AADT on inches per million dollars is positive. Increasing AADT by 1, here 10,000 vehicles per average day, decreases the amount of roughness repaired by one million dollars by about 43 inches. To the extent that busier roads are more urban, this confirms our prior that urban construction is more expensive. Second, the mean annual change in this AADT premium is small, -1.54 inches, and not distinguishable from zero. The signs on the two terms involving AADT are opposite so that the premium for smoothing high AADT segments is weakly decreasing over time.

The next two columns of Table 4 consider the HPMS and NLCD based urban measures. By

Table 4: Composition effects resurfacing

	(1)	(2)	(3)	(4)	(5)	(6)
	AAADT	HPMS Urban	NLCD Urban	Rigid surface	Structural Number	Unionization
$1_{ist(q)}v_{st}^q$	-1076.50 (127.36)	-1005.20 (126.08)	-1662.48 (220.56)	-928.00 (99.31)	-635.53 (195.73)	-1849.20 (252.40)
$1_{ist(q)}v_{st}^q \times t$	28.56 (8.58)	25.19 (8.22)	55.60 (15.46)	26.27 (7.33)	14.34 (13.81)	92.81 (16.83)
t	0.04 (0.08)	0.00 (0.08)	-0.00 (0.08)	0.05 (0.08)	0.00 (0.08)	-0.03 (0.13)
x^0	-0.38 ⁺ (0.20)			3.54 (1.17)	0.37 (0.32)	-0.09 (0.42)
$x^0 \times 1_{ist(q)}v_{st}^q$	24.58 ⁺ (14.62)	26.98 (136.38)	4,504.98 (980.24)	-346.82 (168.65)	-53.77 ⁺ (31.92)	50.12 (13.32)
$x^0 \times 1_{ist(q)}v_{st}^q \times t$	-0.29 (0.92)	10.68 (9.45)	-166.79 (75.38)	19.88 (12.37)	2.39 (2.12)	-3.79 (0.93)
Segment id FE	Yes	Yes	Yes	Yes	Yes	Yes
N	181,236	181,236	181,236	181,068	181,236	181,236

Note: Estimations of variants of equation (6.2). Column headings indicate the interaction variable x^0 . The NLCD based urban measure varies at the state level but not at the state-year level and the HPMS sampling frame requires that segment-id change when urban status changes so HPMS urban status also does not vary within segment. We omit the levels of these variables in columns 2 and 3 because they are colinear with segment fixed-effects. Standard errors in parentheses, clustered at the state-year Level. ⁺ $p < 0.10$, $p < 0.05$, $p < 0.01$, $p < 0.001$.

either measure, one million dollars repairs fewer inches of IRI as segments are more urban. The premium for urban segments is decreasing over time and is distinguishable from zero for the NLCD based measure of urbanization.²⁰

The first three columns of Table 4 confirm that lowering IRI is more expensive in urban areas. They also show that the urban premium decreases over time. In Table 1 we see that, however measured, the Interstate is becoming more urban over time. Thus the trend toward

²⁰The NLCD based urban measure varies at the state level but not at the state-year level. In addition, the HPMS sampling rule requires that segment-id change when urban status changes, so HPMS urban status also does not vary within segment. We omit the levels of these variables in columns 2 and 3 of table 4 because they are collinear with segment fixed-effects.

a more urban Interstate and the decrease in the urban price premium work against each other. Indeed, the fact that $1_{jst}(q) \nu_{st}^q$ remains significant and of almost the same magnitude as we see in columns 4, 5 or 6 of Table 2 suggests that the two trends approximately cancel each other out. While urban status is important for determining the price of IRI, it does not explain the trend in this price.


A second candidate explanation for the increase in the price of IRI involves increased exposure to union labor. We see in Table 1 that the average lane-mile is in a state where union share of employment is lower at the end of our study period than at the beginning. If union exposure is to explain the increase in the price of pavement quality, the union premium must increase over time. To investigate this possibility, the last column of Table 4 considers the effect of state-year union share of all employment. The pattern of coefficient estimates is the same as we saw for AADT and the two urban measures. The price of IRI is higher in state-years with higher union shares and this premium is declining over time. Changes in union exposure also work against the increase in the price of pavement quality.

Columns 4 and 5 consider the physical characteristics of segments. Column 4 considers an indicator that is one if the segment is rigid, i.e., a concrete slab. We see that it is less expensive to make such segments smooth, and this discount decreases over time. Column 5 considers the role of structural number. Increasing structural number by one means that one million dollars reduces IRI by an extra 53.77 inches. This discount decreases over time by 2.39 inches per year per unit of structural number. Alone among the composition variables, for the structural number specification there is no unconditional trend in the price of pavement quality. Columns 4 and 5 together suggest that something about the physical characteristics of the segment may be behind the increase in the price of pavement quality. Structural number seems particularly deserving of further investigation, and we will return to it below.

Unreported results like those in Table 4 investigate the role of proximity to water, average grade and elevation. Neither average grade nor elevation is important for the level or trend in the price of pavement quality. Given the uniformity of the Interstate, that resurfacing costs are not sensitive to the range of grade and elevation that exists within the system seems intuitive.

Proximity to water is more interesting because it helps to shed light on the role of environmental regulation on costs. Enacted in 1972, the Clean Water Act is one of the nation's more important pieces of environmental regulation. Intended to protect the quality of surface water, it requires permits for storm water discharges from construction activities

Figure 4: Inverse price of asphaltic concrete over time



Note: Minus one times lane miles worth of asphaltic concrete per million USD2010 of expenditure, assuming a one inch thick resurfacing layer.

and management of non-point source run-off from roads.²¹ If the Clean Water Act were responsible for the increase in highway construction costs, we would expect the price of IRI to rise faster for roads in wetter areas. Our results do not support this hypothesis. While the price of IRI is higher in wetter areas, proximity to water or wetlands does not explain the trend in this price.²² This does not obviously support the hypothesis that the trend in the price of IRI is due to environmental regulation.

Comparison to engineering based estimates

Our regressions show increases in the price of IRI and suggest that these increases reflect changes in structural number. Flexible roads consist entirely of asphaltic concrete. Because the relationship between the structural number of asphaltic concrete and pavement thickness is known, for flexible roads, we can do two back-of-envelope calculations to see how changes in the price of asphaltic concrete should affect the cost of resurfacing. The first of these calculations allows us to plot the number of miles that can be resurfaced with one millions dollars worth of asphaltic concrete as prices change. We can then compare this time series to our estimated time series for the price of IRI. The second calculation uses our regression results to estimate how the cost to resurface one mile changes over time, and to compare this change with an estimate of the change in price of materials required to resurface a lane mile. Both calculations validate our regression results.

To measure the national average price of asphaltic concrete we combine two price series

²¹<https://www.epa.gov/npdes/stormwater-discharges-construction-activities>, May 15, 2020.

²²We note that positive effect of proximity to water on price is similar to the finding in Smith et al. (1999).

(Federal Highway Administration (1987) (1975-2006), U.S. Bureau of Labor Statistics (2020) (2006-8)), and use the Producer Price Index for All Commodities (PPIACO) series to convert all prices to 2010 dollars. The price of asphaltic concrete stays close to 50 dollars per ton from 1980 until the early 1990s and then increases rapidly to about 125 dollars per ton.

For the sake of illustration, suppose a resurfacing event involves the application of exactly one inch of material. Since an average lane of Interstate is 12 feet wide,²³ resurfacing one lane-mile requires about 196 cubic yards of asphaltic concrete. At about two tons per cubic yard, this is 392 tons of paving material. The price of asphaltic paving material was 44.63 per ton in 1992 and 116.74 per ton in 2008. Multiplying the difference, 70.01, by 392 tons per lane-mile, we have an increase of 28,228 dollars per lane-mile because of increases in the price of asphaltic concrete. Note that Figure 2 shows regression based annual estimates of the number of inches of IRI repaired by one million dollars of resurfacing expenditure. Using our asphaltic concrete price series, and assuming 392 tons per lane-mile of resurfacing, we can calculate the number of lane miles of paving material per million dollars of expenditure on the basis of each years' price for asphaltic concrete. To compare this price series with our regression-based estimate (in Figure 2) we multiply by minus 1 and plot in Figure 4. Comparing the two figures, we see that the inverse price of IRI tracks the inverse price of paving material closely. Table 4 establishes that, at least in a purely statistical sense, structural number alone can explain the change in the price of pavement quality. Figure 4 confirms this conclusion by a different argument.

We can also directly compare an engineering-based cost estimate of the price of IRI to our regression-based estimates.

There are 2552 segments for which we observe a resurfacing event and also observe the segment for at least two years before and after resurfacing. Of these, 926 have flexible pavement. Figure B2 presents the results from the event study showing how structural number changes around resurfacing events for these 926 segments. For reference, the figure also shows the corresponding event study for IRI.

The figure shows a sharp increase in structural number around resurfacing events. This increase is between about 0.2 and 0.4, depending on whether we look at the change over the preceding one or two years. Taking the larger of these two values, and recalling that one inch of asphaltic concrete contributes about 0.4 to structural number, this means that the calculation performed above applies. Thus, observed changes in structural number around resurfacing events (for flexible segments) together with realized changes in the price of as-

²³See, for example, Highway Statistics 2006, table HM-53.

phaltic concrete imply an increase in the cost to resurface a lane-mile of Interstate of 28,228 dollars between 1992 and 2008.

Our regressions also imply a per lane-mile increase in the price of resurfacing. The regressions of equation (5.1) reported in table B1 describe the change in IRI that results from resurfacing. From column 3, resurfacing reduced IRI by 34.18 inches per mile. Similarly, from table 2 column 6, one million dollars of expenditure reduced IRI by 922.86 inches in 1992 and 443.50 inches in 2008.²⁴ Taking the ratios of each year’s values, we conclude that on average one million dollars of expenditure resurfaced $922.86/34.18 = 27.00$ lane miles in 1992 and $443.50/34.18 = 12.97$ in 2008. Inverting, this is 37,037 dollars per lane-mile in 1992 and 77,101 in 2008. Taking the difference, the increase in per lane-mile resurfacing costs implied by our regressions is 40,064 dollars per lane-mile. The engineering-based estimate, 28,228, is about 70% as large as regression-based estimate. This seems quite close, particularly when we consider that paving material is not the only input into resurfacing.

Thus we have three pieces of evidence in support of the hypothesis that increases in the price of IRI largely reflect increases in materials costs. First, in Table 4 changes in structural number completely explain the trend in the price of IRI in a statistical sense. Second, we see in Figure 4 that the inverse price of IRI closely tracks an appropriately transformed national price index for asphaltic concrete. Third, a comparison of changes in the price of resurfacing implied by an engineering estimate and derived from our regressions correspond closely.

6.2 Explaining the increase in the price of new lane miles

As we did for IRI, we would like to understand the increase in the price of new construction. To accomplish this, we include an interaction term, much as we did in equation (6.1). Letting x_{st}^0 denote the state level attribute of interest, we estimate,

$$\begin{aligned} \Delta L_{st} = & A_0 + A_1 I_{st}^L + A_2 [I_{st}^L t] + A_3 t \\ & + B_1 [I_{st}^L x_{st}^0] + B_2 [I_{st}^L x_{st}^0 t] + B_3 x_{st} + \epsilon_{st}. \end{aligned} \tag{6.2}$$

In this regression, the interpretation of A_1 and A_2 are about the same as in (5.4). B_1 and B_2 measure how the price level varies with x , and B_2 is a ‘cross-partial’ term that measures how the difference in price between ‘high x^0 ’ and ‘low x^0 ’ roads evolves over time.

Table B3 reports estimates of equation (6.2) and parallels Table 4 by examining the role of composition in the increasing price of new lane miles. We estimate the effect of changes

²⁴That is, $922.86 - 16 \times 29.96$.

in the following variables on the change in construction costs: grade, elevation, proximity to water, proximity to urban land cover, urban classification, unionization, AADT, share of new mileage in construction, mean structural number and, finally, the share of rigid pavement.²⁵

Only a handful of the estimated interaction effects are different from zero. Construction is more costly as the share of lane miles classified as urban increases. As for IRI, construction costs are higher in more urban places. Also, like IRI, the trends in the urban premium decrease over time, so the shift toward more urban construction does not explain the trend up in the price of new construction. States with greater exposure to unions or higher mean structural number do not have measurably different costs. For each of HPMS grade, elevation, water proximity, and share of new-miles, we see that the coefficient on the interaction term $I^L x^0$ is indistinguishable from zero. These variables do not explain the trend up in the price of new construction.

The last two columns of Table B3 investigate the role of structural number and share rigid. These two construction variables are the only ones for which the interaction term, $I^L x^0 t$, is distinguishable from zero and $I^L t$ is not. That is, in a purely statistical sense, trends in these variables explain the trends in the price of new construction. In addition, structural number is the only variable for which the sign on $I^L t$ is positive. Over time a million dollars buys progressively more miles of low structural number highway and less of high structural number highway. The precision of this term is such that it not distinguishable from zero at conventional levels, but is distinguishable from the corresponding trend for an average segment, -0.0044 , that we estimate in Table 3.

Summing up, the results in Table B3 are largely negative. Terrain, urban share, union exposure and the share of new miles do not seem to explain the increase in the price of new construction. Our estimates for the effect of share rigid and structural number are imprecise, but suggest that these variables, structural number, in particular, may be related to the increase in the price of new construction.

²⁵The estimations in Table B3 are qualitatively similar to those in Brooks & Liscow (2023), but differ in a number of particular ways. First, we study a more recent time period, 1984 to 2008 versus 1960 to 1993. Second, they analyze highway miles, while our more detailed data allow us to analyze lane miles. Third, our construction expenditure data exploits the extra detail that is available in the more recent Highway Statistics volume to exclude expenditure on the Interstate that is not explicitly related to new construction. Fourth, the details of our specification and the source data for our variables differ in a number of ways that seem minor. Finally, our data does not include a measure of housing prices, while Brooks & Liscow (2023) do not observe construction materials or quantities.

Comparison to engineering based estimates

We now turn our attention to the contribution of materials costs to the increase in the price of new lane miles. Our two measures of the physical attributes of the state highway network, ‘structural number’ and ‘share rigid’, were the only variables in Table B3 for which the interacted trend term was measurably negative and the un-interacted trend term ceased to be distinguishable from zero. Thus, from a purely statistical point of view, a change in the physical characteristics of new lane miles is our best guess to explain the trend in construction price.

As above, we focus attention on flexible roads because they are simple. A typical flexible segment of the Interstate consists of 12 inches of asphaltic concrete.²⁶ Using the same conversion as above, this means that each lane-mile of flexible Interstate construction requires 4692 tons of material. The price of a ton of asphalt in 1984 was 48.00 (slightly higher than 1990) so the change in price per ton from 1984 to 2008 was 68.74. Multiplying tons by the change in the price per ton, we have that the price of asphalt required to build a lane-mile of flexible Interstate increased by about 323 thousand dollars between 1984 and 2008. We can read our regression-based estimate from Figure 3. In 1984, one million dollars bought about 0.2 lane miles, and by 2008 this had fallen by about a factor of five to 0.04 lane miles. Inverting, the price of a lane-mile increased from about 5 to about 20 million dollars. This increase is orders of magnitude larger than 323 thousand dollars per lane-mile that we can ascribe to the price of paving materials.

Summing up, of all of the variables we consider in Fable B3, only the two describing the physical attributes of the roadway appear to be related to the trend in prices, although this relationship is not particularly strong. From an engineering standpoint, the measured increase in materials costs does not explain the magnitude of the change in price of new lane miles. It appears more likely that pavement type and thickness are correlated with other changes in highway construction methods or materials that explain the price of new lane miles. For instance, more stringent noise mitigation may add to the cost of lane miles and be correlated with pavement type/thickness choices.

Figure 5: Steady state user cost per vehicle mile travelled over time



Note: User cost of Interstate capital per vehicle mile implied by steady state condition (3.7). Figure based on data in table B4.

7 Calculating the user cost per Interstate vehicle mile travelled

We are now able to evaluate the steady state user cost of the interstate given by equation (3.7),

$$\tau = \frac{rp^L L + rp^q qL + \kappa p^q v}{v}$$

The right hand side of this expression consists of observed quantities. Total lane miles, L , is described in Figure 1. IRI is our measure of quality. We calculate total VMT from data on AADT and segment lengths.

The federal government funds much of highway construction and borrows at the risk free rate. To estimate this rate, r , we use the best linear fit to the January average of the 10 year Treasury rate, net of the annual inflation rate calculated from the CPI.²⁷ We note that the choice of discount rate is the subject of longstanding debate (e.g. Baumol (1968)) the details of which are beyond the scope of this paper. While the level of costs is sensitive to the level of interest rate, our interest in changes in costs depends primarily on changes in the interest rate. For this purpose, choosing exactly the right discount rate is less important than the fact that most defensible choices for the planner's rate of time preference will track the risk free rate. With this said, we consider a range of values below.

Much of our econometric effort has been directed to the estimation of p^L and p^q . For our

²⁶Table IV-3 of HPMS item descriptions for 1993-8.

²⁷The raw data and linear fit are reported in Figure B4.

Table 5: Sensitivity and Counterfactuals

		2007	1992	2007/1992
A.	Baseline	0.19	0.33	0.56
B. Counterfactuals	VMT ₉₂	0.27	0.33	0.81
	p_{92}^L	0.09	0.33	0.27
	p_{92}^g	0.18	0.33	0.56
	r_{92}	0.53	0.33	1.60
C. Sensitivity	IV 92-08	0.06	0.15	0.40
	IV All	0.08	0.15	0.50
	Non parametric (Smooth)	0.07	0.14	0.50

Note: Values of τ in 1992, 2007, and percentage change between the two years. Panel A gives baseline values based on the same data and calculation as presented in figure 5. Panel B considers four counterfactual cases identical to the baseline, except with a single variable held fixed. Panel C considers three cases identical to the baseline except for the technique used to estimate p^L .

baseline calibration, we rely on annual values of p^g and p^L calculated from column 4 of Table 3 and column 6 of Table 2. We calculate maintenance costs per vehicle mile travelled, p^m , by dividing mean annual maintenance expenditure per lane-mile from figure B3 by lane-mile weighted mean annual AADT. We report these data in table B4.

It remains to evaluate κ . As a first step, we evaluate γ , the number of ESALs per vehicle. An ESAL is caused by the passage of a typical tractor trailer rig or about 2000 passenger cars. Assume a truck share of AADT of 12%, consistent with national averages towards the end of our study period. In this case, a segment experiences $.12 + .88/2000 \approx 0.12$ ESALs per average vehicle. A typical design for an Interstate segment will withstand 9m ESALs (Mannering et al., 2007). During its lifetime, we expect a road to increase from an initial IRI around 50 to the acceptable/poor threshold of 170. These are q_0 and q_f . Thus we have $\kappa = 0.12 \times (170 - 50)/9,000,000 \approx 0.0000016$ inches of IRI per average vehicle. Given this value of κ , a new segment experiencing an about average AADT of 8000 depreciates in about 26 years.

We use the data in Table B4 to evaluate the right hand side of (3.7) in each year from 1992 until 2008 and plot the results in figure 5. The units on the y axis of this figure are dollars per vehicle mile travelled. Steady state user cost per mile falls from about 33 to about 19 cents over our 1992 to 2008 study period.

To develop some intuition about this conclusion, Table 5 describes a number of counterfactual results. For reference, the top line of the table describes the baseline case reported

in Figure 5. In this case, the user cost in 2008 is 59% of its 1992 value. Panel B reports initial and terminal steady state user costs when we fix a single quantity at its 1992 level but otherwise replicate the baseline calculation of user cost. If we fix VMT at its 1992 level, user cost declines more slowly than in the baseline case, but is still just 82% of its 1992 value in 2008. If we fix p^L , the price of new lane miles, at its 1992 level, then user cost declines even more rapidly than the baseline case and user cost is 29% of its initial value in 1992. The next two results are more surprising. Fixing the price of IRI at its initial level has only a tiny effect on the 2008 user cost, while fixing interest rates at their higher 1992 level not only undoes the baseline decrease in user costs, but leads to a 2008 user cost that is 160% of the initial value.

The intuition underlying these results is transparent if we consider the relative magnitudes of the different terms that make up the right hand side of equation 3.7. If we let $o(k)$ denote a term of order 10^k , then by using Table B4 we can evaluate the approximate order of magnitude of the three terms in the numerator of equation 3.7,

$$\begin{aligned} rp^L L &\sim o(-2) \times o(7) \times o(5) = o(10) \\ rp^q q L &\sim o(-2) \times o(3) \times o(2) \times o(5) = o(8) \\ \kappa p^q \text{VMT} &\sim o(-6) \times o(3) \times o(11) = o(8). \end{aligned}$$

The first term of equation 3.7 is about two orders of magnitude larger than the second and third terms, so we can ignore the second and third terms when thinking about user costs: only the first term matters. The first term reflects the opportunity cost of lane miles, so it is the components of this term, p^L and r that have the largest impact on user cost. Conversely, the opportunity cost of pavement quality and depreciation are not important determinants of user costs in a neighborhood of observed values. User costs reflect the cost of capital embedded in lane miles. Because lane miles don't depreciate, and because the price of lane miles has increased dramatically, the cost of lane mile capital is the most important part of steady state user costs. In the context of this analysis, this conclusion seems obvious. When we consider that the flow of expenditure on resurfacing is now larger than that on construction, it is more surprising.

Panel C of Table 5 presents robustness tests. These tests focus exclusively on different estimates of p^L for two reasons. First, we have seen that user costs are sensitive to this variable. Second, the two other important quantities, VMT and r , are observed directly but p^L is estimated and so is more uncertain.

The three rows of panel C in Table 5 each replicate the baseline evaluation of steady state user costs using a different method to estimate p_t^L . In our baseline evaluation of equation (3.7) we rely on prices calculated from column 4 of Table 3 and plotted as the solid line in Figure 2. In the top row of panel C, we rely on estimates (not shown) of p^L based on the same specification and sample, but where we instrument for expenditure using lagged appropriations, as in Table B2. In the second row of panel C, we rely on estimates of p^L based on Table B2. In the final row of panel C, we use a locally weighted linear regression to smooth the annual coefficients presented in Figure 2 and estimate p^L from the resulting regression line. Figure B5 shows both the underlying estimations and the derived price series for each case. Although the level of the user costs varies with our estimate of p^L , the basic conclusion that we draw from the baseline case does not: the steady state user cost of the Interstate fell by about half between 1992 and 2008.

For completeness, Figure B6 presents an evaluation of the dynamic equations (3.3) and (3.4) using the same data as we use to evaluate the baseline steady state case and the particular functional form for v given in (3.1).²⁸ This figure suggests the following conclusions. First, like the steady state baseline, calibrations of both (3.3) and (3.4) indicate decreasing user costs over time. Second, in the baseline case, both (3.3) and (3.4) are negative by the end of the study period, so the optimal user cost is a subsidy. This is an implication of intertemporal arbitrage - if price increases in lane miles are expected to persist, a subsidy is justified to build more lane miles.

There are three natural benchmarks against which to compare our estimates of user cost. Both Allen & Arkolakis (2014) and Duranton & Turner (2012) are primarily interested in the benefits of the Interstate system, but also estimate its costs. On the basis of a 1982 estimate of 590 billion USD2010 of total construction cost in Lewis (1982), and 69 billion of annual maintenance, Allen & Arkolakis (2014) estimate that the total annual cost of the Interstate system is about 106 billion USD2010 per year. From Figure 1(a), the extent of the network in 1982 was about 170,000 lane miles. Dividing, we have a total annual cost per lane-mile of about 0.62 million. On the basis of 2006 construction cost estimates reported in Ng &

²⁸Evaluating these equations requires that we pick a value for α , the pavement quality elasticity of VMT. We have no empirical foundation for this choice. One simple calculation suggests that α is likely to be small. Consider a segment with an IRI value of 100, just above the good/acceptable threshold. For this segment, a 1% decrease is about equal to a one inch change. Such a change is probably almost imperceptible, and it is natural to suspect that it will elicit a change in travel volume of much less than one percent. This suggests values of α on the order of 0.1 or 0.01. On the other hand, this tiny elasticity seems inconsistent with the fact that resurfacing is the largest component of Interstate expenditure in 2008, so it also seems reasonable to think that α is close to the about 0.4 resurfacing share of highway expenditure. Figure B6 shows plots of equations (3.3) and (3.4) for $\alpha = 0.1$ and 0.4.

Small (2012), Duranton & Turner (2012) conclude that construction costs are between 27 and 89 million USD2007 per mile. Using an estimate of maintenance costs similar to that of Allen & Arkolakis (2014) and annualizing construction costs (also with a 5% interest rate), Duranton & Turner (2012) estimate total annual costs per lane-mile of between 2.1 and 5.5 million USD2010 per lane-mile.

In Table 5 we estimate a user cost of about 0.19 2010USD per vehicle mile travelled (vmt) in 2007. From Table B4, we calculate that an average lane mile of Interstate served about 3.37 million vehicles per year, and hence 3.37 million vmt. Multiplying, our estimates suggest a cost per lane mile of $0.19 \times 3.37 = 0.64\text{m}$ 2010USD for an average lane mile. This is much smaller than Duranton and Turner and almost perfectly coincides with Allen and Arkolakis. This conclusion requires two caveats. First, Duranton and Turner restrict attention to the urban Interstates, while we are reporting on an average Interstate lane mile. Second, Allen and Arkolakis use an interest rate of 5% in their estimate, while our estimate is based on the about 1.4% risk free rate that prevailed in 2007. Re-evaluating our cost estimate with a 5% interest rate would increase it by about a factor of three, much larger than Allen and Arkolakis and close to the Duranton and Turner estimate for urban highways.


Finally, we can compare our estimates of user cost to the gas tax, its (approximate) real life analog. We start with the annual total of all user fees and taxes (primarily gas tax revenue) from Highway Statistics Table FE9. We discount by the fraction of all VMT carried by the Interstate.²⁹ Finally, we divide by total annual Interstate VMT to arrive at the (federal) user fee per Interstate vehicle mile travelled. The y axis is dollars per mile, so this actual user fee ranges between 1 and 1.5 cents per mile. Comparing Figure 6 to 5, we see that the federal user fee per mile is about one order of magnitude below the level of the steady state user fee required to rationalize the network in every year of our study period.

8 Conclusion

While the benefits of the Interstate network are the subject of a now large literature, its costs have been less well studied. To the extent that an understanding of the costs of the Interstate are important for cost benefit analysis and public finance, this is an important gap in our knowledge. To fill this gap, we investigate the cost of the Interstate highway system with conventional tools from producer theory and asset valuation. This allows us to

²⁹We estimate this share annually by using VMT calculations from the HPMS Sample Data. It varies between 25 and 29 percent.

Figure 6: Observed user cost per mile based on federal gas tax.



Note: User cost per mile in 2010USD based on federal gas tax. This user cost is calculated from the annual total of all user fees and taxes (mainly gas tax revenue) from Highway Statistics Table FE9. We discount this sum by the fraction of all VMT carried by the Interstate as calculated from the Sample Data. To arrive at a per user mile value, we divide by total annual Interstate VMT.

construct a cost function for the interstate. This cost function, in turn, allows us to evaluate the marginal cost of Interstate services along any trajectory of investment, prices and output. That is, the user cost of the interstate.

This user cost consists of four components; the opportunity cost of lane miles, the opportunity cost of pavement quality, depreciation; and flow maintenance. Using administrative data describing the road network and expenditures, we estimate the prices of new lane miles and pavement quality for each of the years in our main 1992-2008 sample. These estimates allow us to evaluate the user cost of the Interstate.

In spite of the fact that resurfacing is now the largest share of Interstate expenditure, only the price of lane miles is important for determining user cost. Over our study period, this price increased rapidly. In spite of this, user cost fell by nearly half. The increase in the price of lane miles was more than offset by a decline in the market rate of return to capital and an increase in the number of Interstate users. In this sense, there is no problem with the cost of Interstate. To the contrary, its cost fell rapidly. This outcome largely reflects changes in the macroeconomy. If interest rates had not fallen, user costs would have risen dramatically. Alternatively, had the price of lane miles stayed at its initial level, user cost would have fallen even further.

Our estimates also provide a check on other estimates of the cost of the interstate in the literature. Our estimates are close to those of Allen & Arkolakis (2014), but are much

higher if we use the same interest rate as they do. Our estimates are lower than those in Duranton & Turner (2012), although this likely reflects the fact that they restrict attention to urban Interstates. More generally, our results provide a foundation upon which to base cost estimates against which benefits estimates can be compared.

The rapid increase in the price of pavement quality appears to be largely a consequence of increases in materials prices, not a problem with construction productivity. The rapid increase in the price of new lane miles remains unexplained, although the data do not provide support for three hypotheses: 1) that the price increase is a pure composition effect resulting from a shift to more urban construction; 2) that it is a consequence of changing exposure to union labor; 3) that it is a consequence of exposure to more intensively regulated environmentally sensitive areas. On the other hand, the data suggest that something correlated with structural number may be to blame. This, in turn, suggests some hard to observe change in the nature of construction, such as excess scope (i.e. unnecessary or ancillary construction expenditures).

The increase in the price of lane miles suggests that concern about construction productivity is warranted. However, the possibility that the price increase reflects a change in way roads are constructed invites further research on the question, and hopefully, cost-benefit analyses for any changes in Interstate construction that come to light.

By focusing on steady state costs, our analysis largely abstracts from the dynamics of investment in the Interstate. To the extent that we have explored these issues, these dynamics appear to be economically important. The Interstate is a scarce and appreciating asset. This is probably the sort of public investment a country should seek out. These issues would seem to be natural topics for further research.

Finally, we would underscore the order of magnitude divergence between our estimates of the user cost of the Interstate and user cost implied the current level of the federal gas tax. This divergence primarily reflects the opportunity cost of lane miles. The actual policy is intended to, more or less, finance year-to-year expenses. On the other hand, in our calculation, the largest portion of user cost is the opportunity cost of lane miles. Under the current policy, nearly the entire return to the country's generations long investment in highways is an implicit subsidy to current drivers. Although an analysis of these issues is beyond the scope of this paper, they would seem to have important implications for welfare and public finance.

References

- Allen, T., & Arkolakis, C. (2014). Trade and the topography of the spatial economy. *The Quarterly Journal of Economics*, 129(3), 1085–1140.
2, 35, 36, 37
- Baumol, W. J. (1968). On the social rate of discount. *The American Economic Review*, 58(4), 788–802.
32
- Bennett, J., Kornfeld, R., Sichel, D., & Wasshausen, D. (2019). Measuring infrastructure in bea’s national economic accounts. In *Economics of infrastructure investment*. University of Chicago Press.
3
- Brooks, L., & Liscow, Z. (2023). Infrastructure costs. *American Economic Journal: Applied Economics*, 15(2), 1–30.
2, 3, 4, 22, 30, 44, 45
- Couture, V., Duranton, G., & Turner, M. A. (2018). Speed. *Review of Economics and Statistics*, 100(4), 725–739.
9
- Duranton, G., Nagpal, G., & Turner, M. (2020, May). *Transportation infrastructure in the us* (Working Paper No. 27254). National Bureau of Economic Research. Retrieved from <http://www.nber.org/papers/w27254>
14
- Duranton, G., & Turner, M. A. (2011). The fundamental law of road congestion: Evidence from us cities. *American Economic Review*, 101(6), 2616–52.
9
- Duranton, G., & Turner, M. A. (2012). Urban growth and transportation. *Review of Economic Studies*, 79(4), 1407–1440.
2, 35, 36, 38
- Federal Highway Administration. (1987). *Price trends for federal-aid highway construction, publication number fhwa-if-06-048*. US Department of Transportation. (<https://www.fhwa.dot.gov/programadmin/pt2006q4.cfm>)
28
- Federal Highway Administration. (2005). *National highway planning network*. www.fhwa.dot.gov. (Accessed: 2014-05-24)
17

- Federal Highway Administration. (2013). *Highway functional classification concepts, criteria and procedures*. US Department of Transportation Washington, DC.
16
- Federal Highway Administration. (2016). *Measuring and specifying pavement roughness*. US Department of Transportation FHWA-HIF-16-032.
5
- Goolsbee, A., & Syverson, C. (2023). *The strange and awful path of productivity in the us construction sector* (Tech. Rep.). National Bureau of Economic Research.
3
- Hirsch, B. T., & MacPherson, D. A. (2003). Union membership and coverage database from the current population survey: Note. *ILR Review*, 56(2), 349-354. Retrieved from <https://doi.org/10.1177/001979390305600208>
18
- Keeler, T. E., & Small, K. A. (1977). Optimal peak-load pricing, investment, and service levels on urban expressways. *Journal of Political Economy*, 85(1), 1–25.
4
- Leduc, S., & Wilson, D. (2013). Roads to prosperity or bridges to nowhere? theory and evidence on the impact of public infrastructure investment. *NBER Macroeconomics Annual*, 27(1), 89–142.
22, 47
- Lewis, D. L. (1982). The interstate highway system: issues and options..
3, 35
- Mannering, F., Kilareski, W., & Washburn, S. (2007). *Principles of highway engineering and traffic analysis*. John Wiley & Sons.
6, 16, 33
- Mohring, H. (1970). The peak load problem with increasing returns and pricing constraints. *The American Economic Review*, 60(4), 693–705.
4
- Ng, C. F., & Small, K. A. (2012). Tradeoffs among free-flow speed, capacity, cost, and environmental footprint in highway design. *Transportation*, 39(6), 1259–1280.
23, 35
- Office of Highway Policy Information. (2016). *Highway performance monitoring system field manual*. Office of Management & Budget (OMB).
12

- Small, K. A., & Winston, C. (1988). Optimal highway durability. *The American economic review*, 78(3), 560–569.
3, 6
- Smith, V., Von Haefen, R., & Heintzelman, M. (1997). Environmental compliance costs for highways. *Center for Environmental and Resource Economics Research Note*(1-97).
4
- Smith, V., Von Haefen, R., & Zhu, W. (1999). Do environmental regulations increase construction costs for federal-aid highways? a statistical experiment. *Journal of Transportation and Statistics*, 2(1), 45–60.
4, 27
- United States Geological Survey. (2010). *Global multi-resolution terrain elevation data 2010 (gmted)*. topotools.cr.usgs.gov/gmted_viewer/viewer.htm. (Accessed: 2013-06-17)
17
- United States Geological Survey. (2011). *NLCD 2001 land cover version 2.0*. www.mrlc.gov. (Accessed: 2014-01-24)
17
- U.S. Bureau of Labor Statistics. (2020). *Producer price index by industry: Asphalt paving and roofing materials manufacturing [pcu3241232412]*. St. Louis Federal Reserve Bank. (retrieved from FRED, Federal Reserve Bank of St. Louis, <https://fred.stlouisfed.org/series/PCU3241232412>)
28
- US Department of Transportation. (2013). *Status of the nation's highways, bridges, and transit: Conditions and performance, report to congress*. US Department of Transportation.
5

A Data construction

A.1 Merging HPMS and Highway Statistics

We merge HPMS and Highway Statistics data on the basis of the state-year in which expenditure and road work are reported. We here describe the details of this process.

Merge of Universe and Highway Statistics Data We begin by describing the merge of the HPMS Universe Data with Highway Statistics, the data we use to analyze new construction. Table A1 describes the initial samples of HPMS and Highway Statistics, along with the estimating sample that results from merging the two data sets.

Our data covers 25 years and 48 states in the continental US, 1200 state-years in all. In panel A we see that all 1200 state-years are present in both the HPMS Universe and Highway Statistics. Trivially, when we merge on state-year we are left with a sample of 1200 state-years. When we drop ‘impossible’ state-years, those where there is no construction expenditure and an increase in length, or conversely, we see in column 3 that we lose 29 state-years and are left with a sample of 1171.

Table A1 also reports sample totals for lane-mile-years and aggregate expenditure over all years. The HPMS reports about 5m lane-mile-years and 1.6t dollars of construction expenditures; these are the integrals of the curves reported in figure 1 (a) and (c). In panel A we see that the final estimation sample reports all of the expenditure recorded in Highway Statistics, but drops about 2% of lane-mile years; all of the impossible state-years involve increases in lane miles in the absence of expenditure. Together with the similarity of OLS and IV results reported columns 3 of table 3 and table B2, this suggests that mismeasurement of expenditure is not an important problem.

Merge of Sample and Highway Statistics Data We next consider the merge of the HPMS Sample Data with Highway Statistics. Our data covers 17 years and 48 states in the continental US, for 816 state-years in all. In panel B we see that all 816 state-years are present in the Highway Statistics data, but that only 815 state-years are present in the HPMS Sample Data. This is because Virginia did not report HPMS Sample Data in 1998. It follows that when we merge on state-year we are left with a sample of 815 state-years. Of these, we drop 240, and are left with a sample of 575. This is a loss of 29% of state-years, 28% of segment-years, 31% of lane-mile-years, and 23% of total resurfacing expenditure.

Of the 240 state-years that we drop, there are nine where the states did not report IRI

Table A1: Description of Merge of HPMS and Highway Statistics

	All	Merge	Final
A. HPMS Universe 1984-2008			
N	1,200	1,200	1,171
Lane Miles	5,012,646	5,012,646	4,929,418
Highway Statistics 1984-200			
N	1,200	1,200	1,171
Construction	162,790	162,790	162,790
B. HPMS Sample 1992-2008			
N	815	815	575
Segments	257,490	257,490	186,055
Lane Miles	3,462,979	3,462,979	2,389,689
Highway Statistics 1992-2008			
N	816	815	575
Resurfacing	116,158	116,044	89,393

data. There are two state-years which record neither expenditure nor resurfacing events (and do not contribute to the estimation of resurfacing effects). There are 12 state-years that report resurfacing, but no expenditure. These are ‘impossible’ years and reflect a misreporting of timing of expenditure or resurfacing. This leaves 217 state-years where we record resurfacing expenditure but no resurfacing events. Table A2 reports more detail.

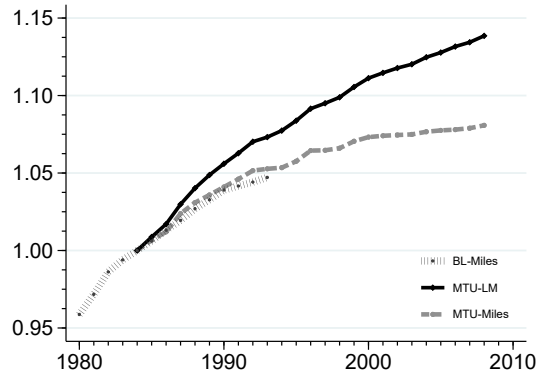
The HPMS Sample Data is a *sample* and reports on a sample of segments, and resurfacing events are rare, they affect about 1% of segment years. On the other hand, Highway Statistics reports all expenditure, not expenditure on sampled segments. This means that we should expect that states where the rate of resurfacing expenditure is low, will sometimes report zero under the sampling rule of the HPMS Sample Data, even if expenditure is positive. The fact that a lower share of resurfacing expenditure than state-years are affected by this problem buttresses this logic. Incompleteness in the way we merge the HPMS Sample and Highway Statistics appears to primarily reflect sampling error in the Sample Data.

Table A2: Accounting for state-years in merge of HPMS Sample and Highway Statistics

	N
No missings	575
No expenditure	12
No resurfacing events	217
No resurfacing and no expenditure	2
No IRI	7
No IRI no resurfacing	2
Total	815

A.2 Correspondence to Brooks & Liscow (2023)

Figure A1: Comparing PR511 and HPMS aggregate mileage



Note: Light gray dashed line is total miles of Interstate by year from the PR511 data on which Brooks & Liscow (2023) is based. Medium gray line is corresponding quality from the HPMS data on which this paper is based. Black line in lane miles of Interstate from the HPMS. All three series are normalized to 1 in 1984, the first year we study. We see that the two mileage estimates track each other closely. Lane miles, however, grow more quickly.

We rely on the Highway Performance and Monitoring System data, while Brooks & Liscow (2023) use the PR511 data. This leads our estimations to differ from Brooks & Liscow (2023) in three important ways. First, their study period ends in 1992, while ours extends to 2008. Second, our data reports a long list of segment characteristics, while the PR511 data reports only length. Third, Brooks & Liscow (2023) do not observe or analyze resurfacing. During most of our study period, resurfacing is a larger fraction of highway expenditure than new construction.

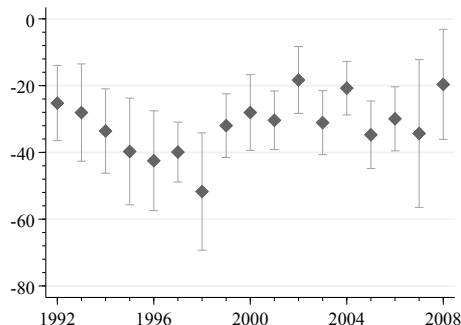
Figure A1 compares mileage in the PR511 and HPMS data. They match closely. Lane

miles, however, grow more quickly than does mileage during our study period. A comparison of the Brooks & Liscow (2023) expenditure data to ours indicates greater divergence. This is in part due the fact that they base their analysis on federal expenditures while we rely on the sum of state and federal expenditure.³⁰

B Supplemental tables and figures

B.1 Resurfacing and IRI

Figure B1: Effect of resurfacing on IRI by year



Note: x-axis is years, y-axis is inches. This figure reports the results of an estimation of equation (B.1) and illustrates the effect of resurfacing on the IRI of resurfaced segments. The figure shows a barely discernible trend upwards, so that a resurfacing event leads to a marginally smaller reduction in IRI at the end of the sample than the beginning.

Figure B1 illustrates the evolution of the effect of resurfacing on IRI. To construct this figure, we estimate the regression

$$\Delta q_{jst} = \sum_{\tau=1992}^{2008} A [\mathbf{1}_{jst}(\tau = t)\mathbf{1}_{jst}(q)] + \epsilon_{jst}. \quad (\text{B.1})$$

Because the indicator variable $\mathbf{1}_{jst}(q)$ is zero for any segment year where the segment is not resurfaced, these coefficients A give the mean change in IRI for resurfaced segments by year. Figure B1 (a) plots these coefficients and 95% CIs based on errors clustered by state-year.

Although we see some variation in point estimates, for the most part, confidence intervals for the different years overlap. For almost all years a resurfacing event reduces the IRI of a

³⁰We are grateful to Leah Brooks and Zachary Liscow for sharing their data for the purpose of this comparison.

segment by between 20 and 40 inches per mile. This figure shows at most a small positive trend so that resurfacing events in 1992 were not much different than in 2008.

Table B1 presents estimates of variants of equation (5.1). Column 1 of Table B1 presents a regression of segment-year change in IRI on an indicator for whether the segment was resurfaced, a simplified version of (5.1) omitting terms involving time. On average, resurfacing reduces IRI by about 33 inches. Column 2 refines Column 1 by including state-year indicator variables. Column 3 repeats column 1, but includes segment and state-year indicators. Although the identifying variation in each of these regressions is quite different, the estimated effect of resurfacing is not.

Column 4 estimates equation (5.1) including the terms involving time. There is a small positive change in the effect of resurfacing expenditure that is barely distinguishable from zero. Column 5 replicates the regression of Column 4 while including state indicators. Column 6 replicates column 4 while including segment indicators. Consistent with the barely visible trend that we see in figure B1 (a), these regressions indicate a barely detectable trend in the effect of resurfacing expenditure. In column 6, given the point estimate of about 0.83 on the interaction of time and the resurfacing indicator, the effect of resurfacing decreases from 43.28 inches in 1992 to 29.17 inches in 2008.

Table B1: Resurfacing and IRI

	(1)	(2)	(3)	(4)	(5)	(6)
$1_{ist}(q)$	-32.66	-31.94	-34.18	-39.33	-39.31	-43.28
	(2.02)	(1.97)	(2.20)	(4.57)	(4.55)	(5.24)
t				-0.05	-0.04	-0.02
				(0.07)	(0.07)	(0.08)
$1_{ist}(q) \times t$				0.69 ⁺	0.67	0.83 ⁺
				(0.41)	(0.41)	(0.47)
State FE	No	No	No	No	Yes	No
State-Year FE	No	Yes	Yes	No	No	No
Segment id FE	No	No	Yes	No	No	Yes
N	186,055	186,054	181,235	186,055	186,055	181,236

Standard Errors in Parentheses Clustered at the State-Year Level.

⁺ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Table B2: Construction expenditure and new lane miles, TSLS estimate

I_{st}^L	0.1584 (0.0458)
t	-0.8271 (0.2815)
I_{stt}^L	-0.0037 (0.0025)
State FE	No
N	1,171
F	20.65

Note: IV Estimation of equation (5.4) Standard errors in parentheses, clustered at the state-year Level. ⁺ $p < 0.10$, $p < 0.05$, $p < 0.01$, $p < 0.001$.


B.2 Instrumental variables estimation of p^L

We are concerned about measurement error that arises from imprecision in our matching of construction expenditure, from Highway Statistics, to new construction, from HPMS Universe data. Such a mismatch can occur either because the HPMS and Highway Statistics record the expenditure for a project in a different year than it is recorded as coming into service in the HPMS, or because of imprecision in the correspondence between Highway statistics expenditure categories and HPMS improvement categories.

To address this, we conduct an instrumental variables estimate where we instrument for current expenditure with the four year lag of total state interstate highway appropriations. The rationale for this instrument is similar to that given in Leduc & Wilson (2013). Instrument validity requires that lagged appropriations predict the expenditure, but not be related to measurement error. In fact, lagged appropriations strongly predict expenditure, and it seems reasonable to suppose that they do not anticipate mismeasurement of expenditure. (We do not conduct these IV regressions for our investigation of IRI because first stage predictive ability is too low.)

Table B2 repeats column 3 of table 3, but instruments terms involving expenditure with corresponding terms involving the four year lag of total Interstate appropriations. This change in estimating technique has little impact on our estimates of the trend in prices, and reassures us that mis-measurement of expenditure is not causing economically important changes in our results.

Figure B2: Event study of structural number and IRI




Note: Changes in IRI and structural number around resurfacing events for all segments with flexible pavement.

B.3 Other supplemental results

There are 2552 segments for which we observe a resurfacing event and also observe the segment for at least two years before and after resurfacing. Of these, 926 have flexible pavement. Figure B2 presents the results from the event study showing how structural number changes around resurfacing events for these 926 segments. For reference, the figure also shows the corresponding event study for IRI. Except for the different sample, the about 25 inch drop in IRI around resurfacing that we see in figure B2 is comparable to the within segment estimate in table B1 column (6). Note that we can use this same research design to check whether the change in structural number from resurfacing is constant throughout our sample. The data do not indicate that the amount of paving material used for resurfacing changes over our study period.


Table B3 parallels table 4 and examines the role of composition in the increasing price of new lane miles. We estimate the effect of changes in the following variables on the change in construction costs; grade, elevation, proximity to water, proximity to urban land cover, urban classification, unionization, AADT, share of new mileage in construction, mean structural number and, finally, the share of rigid pavement.

Figure B3: Millions of dollars per mile of Interstate maintenance expenditure by year.



Note: Plot of average state maintenance expenditure per lane-mile over time.

Figure B4: Risk free interest rate




Note: Dots indicate January average of the 10 year Treasury rate net of the annual inflation rate calculated from the CPI. Solid line is a local linear regression. Dotted line is the best linear approximation that yields the values of r_t that we use in our calibrations exercise and report in table B4.

Table B3: Composition effects in construction costs

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	HPMS-Grade	Elevation	NLCD-Water	NLCD-Urban	HPMS-Urban	Unionization	AADT	New-miles	Structural Number	Share rigid
I^L	0.167 (0.108)	0.049 (0.032)	0.024 (0.046)	0.156 ⁺ (0.092)	0.201 (0.089)	0.121 (0.074)	0.207 ⁺ (0.105)	0.047 (0.022)	-0.120 (0.167)	0.009 (0.028)
I^{Lt}	-0.006 (0.004)	-0.003 (0.002)	-0.003 (0.002)	-0.010 (0.004)	-0.013 (0.004)	-0.009 (0.003)	-0.012 (0.005)	-0.003 (0.001)	0.011 (0.007)	-0.000 (0.002)
t	0.041 (0.289)	-0.602 (0.284)	-0.690 (0.272)	-0.525 ⁺ (0.266)	-0.276 (0.337)	-0.560 (0.334)	0.371 (0.430)	-0.729 (0.313)	-0.367 (0.310)	-0.168 (0.265)
x	2.470 (6.976)				-44.410 (49.636)	0.279 (1.326)	-0.138 ⁺ (0.081)	-0.272 (0.268)	2.177 (3.403)	31.373 (23.323)
I^Lx	-0.060 (0.057)	0.000 (0.000)	0.304 (0.300)	-0.483 (0.371)	-0.289 (0.138)	-0.004 (0.004)	-0.000 ⁺ (0.000)	0.017 (0.017)	0.028 (0.027)	0.165 (0.118)
I^{Ltx}	-0.000 (0.003)	-0.000 (0.000)	-0.010 (0.011)	0.029 ⁺ (0.016)	0.019 (0.007)	0.000 ⁺ (0.000)	0.000 ⁺ (0.000)	-0.000 (0.001)	-0.002 ⁺ (0.001)	-0.011 ⁺ (0.006)
State FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
N	799	1,171	1,171	1,171	1,171	1,171	1,171	1,006	988	988


Note: Standard errors in parentheses, clustered at the state level. ⁺ $p < 0.10$, $p < 0.05$, $p < 0.01$, $p < 0.001$. This table parallels table 4 and examines the role of composition in the increasing price of new lane miles. Each column reports an estimate of equation (6.2). The interaction/composition term used in each column is indicated in the column heading. Construction is more costly as the share of lane miles classified as urban increases and the urban cost premium mean decreases over time. States with greater exposure to unions or higher mean structural number do not have measurably different costs. States where union share declines faster see slightly faster cost increases. For all of HPMS Grade, elevation, water proximity, and share of new-miles, we see that the coefficient on the interaction term I^{Ltx^0} is indistinguishable from zero. Structural number and share rigid are only reported in the Sample Data. These data start in 1992 instead of 1984, and so the sample size for columns (9) and (10) is smaller. Grade is also reported only for a subset of years.

Figure B5: Estimates of lane miles per million dollars of expenditure and of p^L



Note: Panel (a) plots the different regressions that we use to estimate p^L . Dots reproduce mean miles of new construction per million dollars by year from figure 2. The dashed black line is based on column 3 and is used in our baseline calibration exercise. The solid black linear fit is based on column 2. The heavy black non-linear curve is constructed by applying LOWESS to annual means, i.e., the dots in the figure. The green line is based on a replication of column 3 of table 3 (not shown) but instruments for expenditure using lagged appropriations. Panel (b) presents identical information as in panel (a) but transforms regression estimates of lane miles per million dollars of expenditure into prices, millions of dollars per lane-mile by year, using the transformation described in the text. The dashed black line in (b) matches the values of p^L that we report in table B4 and use in our baseline calibration.

Figure B6: User cost per vehicle mile travelled over time



Note: (a) User cost of Interstate capital per vehicle mile implied by Euler condition (3.3) (b) User cost of Interstate capital per vehicle mile implied by Euler condition (3.4). All figures rely on the data in table B4. In both panels the solid black line indicates calibration to actual data with $\alpha = 0.1$; dashed grey line is $\alpha = 0.4$.

Table B4: National variables for the calibration

	$VMT \times 10^9$	L_t	q_t	τ^{gas}	r	$p^L \times 10^6$	p^q	m_h
1984	350.79	185,108.6	.	.	0.054	.	.	0.006
1985	366.78	186,723.3	.	.	0.053	.	.	0.007
1986	388.90	188,257.7	.	.	0.051	.	.	0.008
1987	413.83	190,627.3	.	.	0.049	.	.	0.008
1988	435.57	192,557.6	.	0.011	0.047	.	.	0.007
1989	458.46	194,128.3	.	0.011	0.046	.	.	0.007
1990	475.77	195,470.1	.	0.009	0.044	.	.	0.007
1991	486.87	196,727.8	.	0.010	0.042	.	.	0.006
1992	503.91	198,103.8	108.57	0.011	0.040	20.25	1,038.5	0.007
1993	522.62	198,654.8	109.48	0.011	0.039	21.02	1,071.9	0.005
1994	542.34	199,429.1	112.49	0.009	0.037	21.85	1,107.4	0.006
1995	561.95	200,617.4	104.17	0.011	0.035	22.75	1,145.4	0.006
1996	580.67	202,051.2	104.50	0.013	0.033	23.72	1,186.1	0.006
1997	598.98	202,696.3	103.86	0.011	0.031	24.79	1,229.8	0.005
1998	621.10	203,407.3	96.07	0.016	0.030	25.95	1,276.9	0.006
1999	639.85	204,643.5	97.40	0.016	0.028	27.23	1,327.7	0.005
2000	655.53	205,697.6	95.56	0.016	0.026	28.64	1,382.7	0.006
2001	668.57	206,328.8	94.43	0.013	0.024	30.21	1,442.4	0.005
2002	685.89	206,905.1	95.24	0.015	0.023	31.96	1,507.6	0.006
2003	700.08	207,355.3	92.85	0.014	0.021	33.92	1,578.9	0.006
2004	714.24	208,194.7	94.30	0.014	0.019	36.14	1,657.3	0.005
2005	723.18	208,755.4	91.99	0.014	0.017	38.67	1,743.9	0.005
2006	730.12	209,471.6	89.94	0.012	0.016	41.58	1,840.0	0.005
2007	733.38	209,982.2	90.74	0.013	0.014	44.96	1,947.3	0.005
2008	713.50	210,751.0	89.70	0.011	0.012	48.95	2,068.0	0.005

Note: Annual values of all variable used in calibration exercise of section 7. L is total lane miles. q is system average IRI. τ^{gas} is actual gas tax revenue per vehicle mile and reported in figure 6. r is the real interest rate. p^L is millions of 2010USD per lane-mile. p^q is inches of roughness eliminated per million dollars of 2010USD expenditure. m is non-resurfacing maintenance expenditure per Interstate vehicle mile travelled.