

NBER WORKING PAPER SERIES

WHERE THE RUBBER MEETS THE ROAD:
PAVEMENT DAMAGE REDUCES TRAFFIC SAFETY AND SPEED

Margaret Bock
Alexander Cardazzi
Brad R. Humphreys

Working Paper 29176
<http://www.nber.org/papers/w29176>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 2021

We thank participants at the 2020 NBER Economics of Transportation in the 21st Century Conference, Ed Glaeser, and Matt Turner for valuable comments and the NBER for financial support. This research was funded through an interagency agreement between the US Department of Transportation and the National Science Foundation, award no. 1559013. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2021 by Margaret Bock, Alexander Cardazzi, and Brad R. Humphreys. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Where the Rubber Meets the Road: Pavement Damage Reduces Traffic Safety and Speed
Margaret Bock, Alexander Cardazzi, and Brad R. Humphreys
NBER Working Paper No. 29176
August 2021
JEL No. R41

ABSTRACT

Road maintenance constitutes a significant component of public transportation spending at all levels of government. Formulation of efficient transportation infrastructure policy requires information about factors affecting road and traffic conditions. We generate the first causal evidence that decreasing pavement quality impacts vehicle crash rates and decreases average speed. Results from Instrumental Variable models using spatially and temporally disaggregated data from Federal-Aid Highway System (FAHS) roads in California show statistically and economically significant increases in vehicle crash rates and decreases in average vehicle speed caused by road damage. These impacts imply significant increases in social costs attributable to road damage.

Margaret Bock
Department of Economics
Goucher College
1021 Dulaney Valley Road
Baltimore, MD 21204
margaret.bock@goucher.edu

Brad R. Humphreys
Department of Economics
West Virginia University
1601 University Avenue, P.O. Box 6025
Morgantown, WV 26506-6025
brhumphreys@mail.wvu.edu

Alexander Cardazzi
Department of Economics
West Virginia University
1601 University Avenue
Morgantown, WV 26506-6025
alex.cardazzi@gmail.com

1 Introduction

Highway safety represents a major concern to motorists and all levels of government. Substantial resources go toward ensuring its provision. For example, highway maintenance, a form of transportation infrastructure investment which includes highway resurfacing, pothole filling, guardrail replacement, and other actions that improve damaged pavement, accounted for \$51.4 billion in spending by all levels of government in 2017. This included \$5.5 billion in maintenance spending on the Interstate Highway System, and \$16.3 billion in maintenance spending on the larger Federal-Aid Highway System that includes US Highways and Minor Arterial, Urban Collector, and Rural Major Collector Highway systems. 2017 expenditure on new lanes and roads totaled \$104 billion. States spend heavily on road maintenance. For example, in 2018-19, the California Department of Transportation (“Caltrans”) increased annual spending on road maintenance from \$421 million to \$576 million and spending on road rehabilitation from \$424 million to \$994 million.

This level of government spending raises important questions about the relationship between road quality and travel. Travel speed represents a key component of overall travel costs (Couture et al., 2018), and rougher roads might reduce travel speeds, leading to longer trips for personal and commercial vehicles and diminished response times for emergency and commercial vehicles, generating deadweight loss. Maintaining smooth roads may lead to increased traffic speeds and decreased crash rates.

Road quality represents a choice variable for government transportation agencies (Small and Winston, 1988) and road damage can be repaired by maintenance. A few papers empirically analyze specific road maintenance projects in terms of cost functions (Link, 2014) and from a cost-benefit analysis perspective (Alqadhi et al., 2018). This literature generally ignores any external social costs of road damage in terms of traffic safety and travel speed.

A large body of economic research analyzes highway safety but none focused on the impact of road damage on safety and speed. Newbery (1988) investigated the economics of road damage from a theoretical perspective, highlighting the fundamental importance of road damage, but little subsequent empirical research addressed it. Duranton et al. (2020) recently surveyed the state of US transportation infrastructure and economic research on this topic. They identified research on the impact of road quality in terms of roughness as an important area needing more attention. The lack of economic evidence on the impact of road damage on traffic outcomes limits our understanding of the importance of road quality in the context of transportation infrastructure investment policy (Winston, 2020).

We address this lack of evidence by empirically analyzing the relationship between highway roughness and traffic outcomes using spatially and temporally disaggregated data from Federal-Aid Highway System (FAHS) roads in California over the period 2011-2019 to estimate reduced form empirical models. Geospatial matching of pavement quality data from the Caltrans Automated Pavement Condition Survey (APCS), traffic conditions and crash data from the Performance Measurement System (PeMS) generated a comprehensive data set for about 16,000 miles of FAHS highways in California at the one mile/one month level.

Results from an instrumental variables approach using proxy variables for distance to groundwater under road segments reveal a substantial impact of road damage on both vehicle crash rates and average speed. A one standard deviation increase in road roughness, as measured by the International Roughness Index (IRI), causes an 11% decrease in average

speed, holding other factors affecting speed like precipitation and temperature constant. A one standard deviation increase in IRI causes an increase in vehicle crashes large enough to move a safe road segment with no crashes in an average month to about 0.5 standard deviations above the average crash rate over the sample period, a substantial decrease in traffic safety. These results imply that road damage generates substantial social costs in terms of higher vehicle operating costs, longer travel times, and less traffic safety.

2 Empirical Analysis

Three seminal papers on the economics of highway infrastructure investment motivate our empirical analysis of the effect of road quality on vehicle crashes and speed. Road roughness represents a measurable factor that can be improved by maintenance spending. Vickrey (1969) identified traffic congestion, an externality, as an important driver of road investment spending and also identified vehicle crashes as an important component of congestion costs. Vickrey (1969) pointed out the lack of evidence that road quality affected traffic outcomes. Developing evidence that road quality affects vehicle crashes and speed addresses this dearth of evidence.

Vickrey (1969) motivated economic research on transportation infrastructure by identifying vehicle crashes as a key source of congestion costs and analyzing the impact of transportation infrastructure investment projects on congestion. Vickrey (1969) primarily focused on congestion due to bottlenecks generated by limited capacity on specific road segments and a comparison of toll-based solutions to construction of additional lanes on existing roads to address congestion. Subsequent research on transportation infrastructure also focused on the impact of adding additional lanes to existing highways to address congestion (Duranton and Turner, 2011).

Newbery (1988) identified another important externality, damage to pavement quality caused by vehicles using roads, that affects road investment spending. This paper explicitly modeled the process through which roads become rougher with use and directly linked road roughness to vehicle operating costs. Speed effects the rate at which vehicles burn fuel, an important component of vehicle operating costs. Repair or vehicle replacement due to crashes represents another important vehicle operating cost component.

Small and Winston (1988) pointed out that the rate at which pavement becomes rougher over time reflects design decisions with important economic components and applied economic methods to determine optimal pavement durability. This paper also identified key differences between rigid and flexible pavement in terms of road deterioration. Small and Winston (1988) paper estimated the costs of road deterioration and concluded that widespread underinvestment in pavement durability exists. Evidence on the impact of rougher roads on vehicle crashes and speed can help to understand the broader implications of underinvestment in pavement durability.

These papers generally emphasized the importance of road maintenance expenditure aimed at improving pavement quality, how to price the impact of road use to finance maintenance, and how to inform optimal road construction and maintenance policies. Our analysis extends this literature by focusing on the underlying relationship between road quality, which reflects both construction/maintenance spending and past vehicle use, and traffic outcomes

that reflect both social congestion costs and vehicle operation costs. This relationship is fundamental to economic models that inform optimal highway construction and maintenance policies.

Durantón et al. (2020) pointed out the lack of empirical research on pavement roughness and traffic outcomes in economics. The existing literature generally focuses on highway spending and traffic volume (vehicle miles traveled), but not on crashes or speed. Two papers investigated the relationship between road infrastructure spending and congestion cost estimates based on traffic volume. Winston and Langer (2006) analyzed the relationship between annual state level highway spending and congestion costs, proxied by a formula accounting for delay time, fuel costs, and the value of passengers' time. Lewis and Bajari (2011) analyzed the relationship between individual highway maintenance projects and social costs of delay proxied by a function of traffic volume, ad hoc estimates of delay times, and delay costs per hour using data from California highways over the period 2003-2008. Koh and Lee (2020) analyzed the impact of bridge closures in Pennsylvania on a similar measure of social costs of delay. Kim (2020) found that speed increased on freeways in California following maintenance-related lane closures, but the impact lasted for only about one year.

None of these papers analyze data on pavement quality, and none focus on crash rates or speed. While traffic volume clearly affects congestion costs, a complex relationship exists between between traffic volume and average speed, speed directly impacts vehicle operating costs, and recent research points out the importance of vehicle speed independent of traffic volume (Couture et al., 2018; Bento et al., 2020).

A few papers in the traffic engineering literature examine the relationship between measured International Roughness Index (IRI) on road segments and both vehicle speed and crash rates. Wang et al. (2014) investigated the relationship between free-flow vehicle speed and IRI in a linear regression model using data from California freeways over the period 2000 to 2011. The paper found a weak negative association between IRI and free-flow speed; an increase of 63 inches/mile in IRI only decreased free flow speed by 0.3 mph across all California freeways. Results from subsamples for northern, central and southern California revealed decidedly mixed results, including a positive and statistically significant relationship between IRI and free-flow speed in northern California, a negative and statistically significant relationship in southern California, and no relationship in central California.

Yu and Lu (2014) analyzed IRI and average vehicle speed data for 32 separate, randomly selected one mile segments of interstate highway in California for the first week of July and the first week in September in various years over the period 2000 to 2006 using a fixed effects regression model. The analysis found a small, negative, and statistically significant association between IRI and speed. The results indicated that a IRI increase of 1 inch/mile reduced average speed by 0.008 MPH. Based on the 63 inch/mile increase reported by Wang et al. (2014), that parameter estimate implies a 0.5 mph decrease in speed. Both papers reported very weak associations between IRI and speed.

The evidence on association between IRI and crash rates in the traffic engineering literature is even thinner. Li and Huang (2014) and Mamlouk et al. (2018) undertook simple analyses of the IRI/crash relationship. Li and Huang (2014) collected IRI and crash data from 14,000 highway segments in Texas in 2008 and 2009. Segments were grouped into three IRI quality bins (good, fair, poor). Simple difference in means tests showed higher crash rates per 100 million VMT on road segments in the poor quality group compared to the other

two. Mamlouk et al. (2018) collected IRI and crash rate data for one mile/one year highway segments in Arizona, North Carolina and Maryland for various years from 2013 to 2015. Fitted value plots from a sigmoidal function showed an increasing, nonlinear relationship between IRI and crash rates that only increased at IRI values higher than 150 in/mi.

This limited evidence provides no clear consensus on the association between IRI and speed, and the unconditional nature of the evidence on the IRI-crash rate relationship reduces its usefulness. This literature makes no use of causal inference methods, a limitation we address in this paper.

This paper focuses on Interstate Highways and US Routes in California in the Federal-Aid Highway System. The California Department of Transportation maintains 60,756 lane miles of Interstate Highway and 158,000 lane miles of Federal-Aid Highways, which includes Interstate Highway and US Route lane miles.

California represents an interesting setting in terms of empirical research on pavement quality because of the state’s expansive highway system and statewide variation in climate and terrain. According to the Federal Highway Association, California has almost 400,000 road miles in total, behind only Texas (680,000 miles) and ahead of Illinois (300,000 miles). Further, California is the most populous state by about 10 million people. California, unlike Texas and Illinois, contains a variety of climate types which might age pavement differently. Climate zones have this effect by way of humidity, precipitation, and temperature. California has an environment conducive to studying highway conditions on a large scale, as well as exploring the impacts of environmental factors on highway quality.

A 2018 report by the nonprofit organization TRIP, which researches surface transportation issues, assessed the condition of roads in California. The TRIP report reviewed the state’s roads and bridges and evaluated congestion, safety, and pavement condition (TRIP, 2018). The TRIP report noted that 68% of major roads and highways are of poor or mediocre condition. TRIP estimated that California’s road quality cost drivers over \$22 billion dollars in total, or \$843 per driver, each year in additional vehicle operating costs using the Highway Development and Management Model. This places a dollar value on one component of road damage externalities.

The TRIP report also estimated the annual economic costs of traffic fatalities where roadway features were a contributing factor, which represented about 1/3 of all fatalities, at almost \$10 billion. Clearly, road conditions have large effects on California drivers. These costs also justify our focus on California highways.

2.1 Data Description

We assemble an analysis data set from multiple sources in Caltrans reflecting conditions on FAHS road segments in the state over the period 2011-2019. The final data set spatially and temporally merges data on road quality, traffic conditions, crashes, weather, and geographical features of the terrain below and around road segments. In general, the merging occurs at low levels of temporal and spatial aggregation using both latitude and longitude and highway post mile to match data from different sources.

Spatial matching uses a master geo-location file from Caltrans, the SHN Postmile 0.1 mile interval file, referred to as the Postmile file. This file identifies each segment of road in California by latitude and longitude, generally for 0.1 mile segments. This comprehensive

road catalog serves as a crosswalk between traffic outcomes like vehicle speed, miles traveled, crashes, and pavement quality data.

2.1.1 Traffic Outcomes

Data on traffic outcomes, including crashes, come from the Caltrans Performance Measurement System (PeMS). PeMS collects and reports traffic conditions based on input from over 39,000 individual induction loop sensors on highways in California. Induction loop sensors, oval or rectangular shaped wire arrays embedded in roadways with current running through them, detect the presence of individual vehicles and the time each vehicle spends over the sensor because the metal in vehicles interrupts the current flowing through the sensors. Caltrans converts traffic flow data from induction loop sensors into several traffic flow variables at different levels of spatial and temporal aggregation for roads in the state through an imputation process.

PeMS provides integrated real-time traffic conditions and access to archived traffic conditions data through an Archived Data User Service (ADUS) containing more than ten years of historical traffic conditions data. The PeMS ADUS provides detailed information on several traffic flow measures, incidents on roads, and lane closures at user-selected temporal and spatial levels.

The PeMS ADUS provides access to historical data on Vehicle Miles Traveled (VMT) and Vehicle Hours Traveled (VHT) for all vehicles. We extract these variables for one mile road segments averaged over one month time periods, our primary unit of analysis in the empirical work. Extracting data at this spatial and temporal frequency provides us with relatively disaggregated data that minimizes missing observations and facilitates matching to data from other sources.

VMT reflects the total number of miles driven on a road segment by all types of vehicles. For example, a VMT value of 100 over a one mile road segment over a month can reflect 100 vehicles traveling on this one mile road segment each day of the month or 200 vehicles traveling on 15 days out of the (30 day) month and no vehicles traveling through on the other 15 days of the month. Since we collect VMT data for one-mile road segments, VMT also reflects the total number of vehicles that drove over the one mile segment on average on each day in a month.

VHT reflects the number of hours spent on the road segment by all vehicles. For a given VMT on a road segment, a higher VHT implies slower speeds on average, since more time is spent on the segment. A high VMT and low VHT on a segment suggests a lot of cars traveled over the segment and spent little time on it, suggesting faster speeds. Therefore, these two measures can be used to generate average speed estimates over road segments weighted by vehicle miles. VMT and VHT are tracked for total vehicles. PeMS also imputes truck VMT and VHT for road segments.

The PeMS database also contains records for all traffic incidents that occurred on highways in the state. The underlying data come from the California Highway Patrol and PeMS integrates incident data into the ADUS. The PeMS traffic incident data contain information on the date and time of each incident, a basic description code for the incident, and spatial identification information for each incident such as district, county, freeway, direction of travel, post mile marker, and latitude/longitude.

PeMS traffic incidents include a wide variety of events including lane closures, weather and congestion advisories, reports of debris in roadways, and crashes, among many others. We focus only on incidents linked to vehicle crashes by limiting the incidents analyzed to those with descriptions including words like “collision,” “hit-and-run,” and other similar words and phrases.

2.1.2 Pavement Quality

An analysis of pavement quality requires information about the actual condition of road surfaces. Our data on road surface conditions come from the Caltrans Automated Pavement Condition Survey (APCS). The APCS employs customized vehicles equipped with sensing equipment on roads at highway speeds. These vehicles collect continuous data on pavement quality in the form of longitudinal and transverse profiles. We obtained detailed data from the APCS through a direct request to Caltrans.¹ The APCS data spans California’s entire highway system and contains observations collected in three separate scan waves: Wave 1 (2011-12), Wave 2 (2015-16), and Wave 3 (2018-19).

Caltrans transforms continuous APCS road quality data to one-tenth of a mile (0.1 mile or 528 feet) road segments, called data collection segments (DCS). APCS road quality data at the DCS level include International Roughness Index (IRI), other pavement quality measurements like Alligator cracking, mean profile depth, rutting, information about the surface type, and detailed location of each DCS segment’s start and end points.


IRI represents a commonly used measure of pavement quality. IRI reflects the number of inches of vertical suspension travel experienced by a typical car traveling over one mile of road (Duranton et al., 2020, pg. 6). A large IRI value for a road segment signals poor overall pavement quality. Increasing IRI reflects rougher roads. Specific pavement defects like dips, ruts, and potholes all increase a segment’s IRI.

Figure 1 shows the timing of road segment scans across the three APCS waves in the context of the traffic conditions data described above. Each bar represents one month in our sample period. Orange bars show highway segments with complete vehicle mile traveled and vehicle hour traveled for that month from PeMS. Green bars show highway segments with both IRI scans and highway conditions data. The overall sample contains monthly PeMS data on traffic conditions for almost two million one mile highway segments. About 131,000 of these segments contain APCS scans in the corresponding month. The three separate waves of APCS scans can be seen clearly on Figure 1. Wave 3 of APCS scans did not during the winter of 2018-19.

Figure 2 shows the temporal and spatial distribution of APCS scans in the sample. No systematic patterns appear in terms of specific areas always being scanned in the same month of the year. In general, APCS scans of highway segments in the southern part of California occurred in earlier months of the year (late winter and spring) and more northern highway segments in the summer and fall months. This will be discussed in more detail in a later section.

¹These data are only available by request.

Figure 1: Number of Observations with VMT/VHT and IRI Data

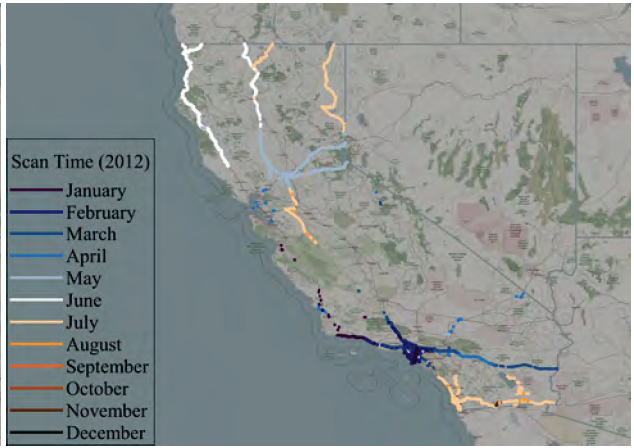


Note: Orange bars reflect the number of 0.1 mile highway segments with VMT/VHT observations in each month of the sample. The green bars reflect the number of 0.1 mile highway segments with both IRI and VMT/VHT observations.

Figure 2: Automated Pavement Condition Survey Scan Timing: 2011-2019

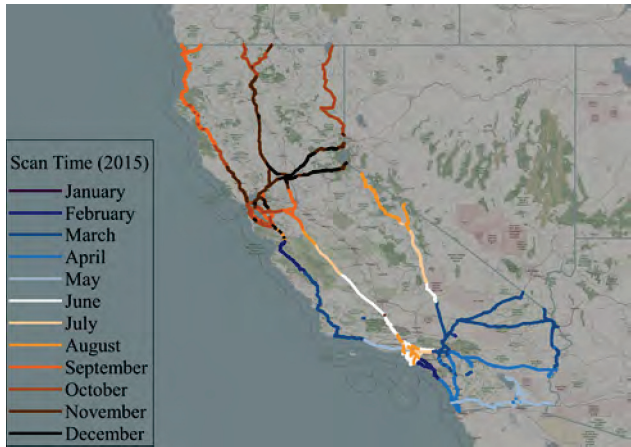
Panel A: 2011 Scan Wave

Panel B: 2012 Scan Wave



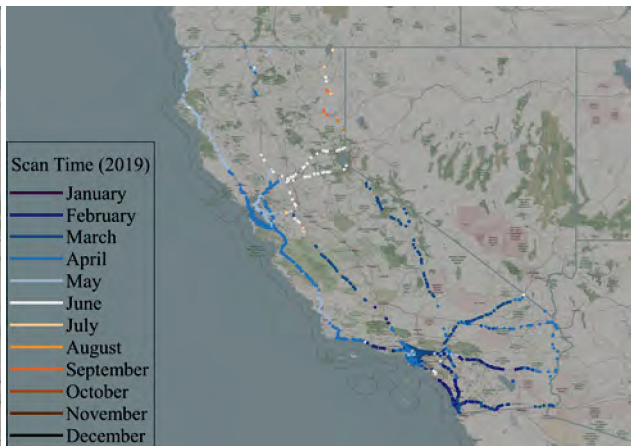
Panel C: 2015 Scan Wave

Panel D: 2016 Scan Wave



Panel E: 2018 Scan Wave

Panel F: 2019 Scan Wave



Note: Figures show Automated Pavement Condition Survey (APCS) scan timing and location by wave.

The APCS data also contain information about the surface material of scanned highway segments. In general, highway pavement can be classified as either flexible or rigid. Flexible and rigid pavements differ substantially on many margins, including construction and main-

tenance costs, durability, and smoothness. In general, rigid pavement costs more per mile and lasts longer, and has lower annual maintenance costs. The APCS generically identifies segments surface type as either asphalt (flexible), concrete (rigid), or bridge, although these terms refer to many different surface types. Caltrans defines asphalt pavement as

Pavement that has a surface consisting of an asphalt-bound layer, including but not limited to hot-mix asphalt (HMA), rubberized HMA (RHMA), open-graded mixes, chip seals, and other asphalt-bound surface treatments. Asphalt pavement can be flexible, semi-rigid and composite pavement structure.

and concrete pavement as

A concrete pavement type that is made of concrete with continuous longitudinal steel reinforcement and no intermediate transverse expansion or contraction joints. There is one type of concrete surface but three types of concrete pavements: jointed plain concrete pavement (JPCP), precast pavement, and continuously reinforced concrete pavement (CRCP). Precast pavements are rated the same way as JPCP.

We will refer to all flexible pavement as “asphalt” and all rigid pavement as “concrete” for the sake of clarity, recognizing the heterogeneity of actual pavement types represented by these terms.

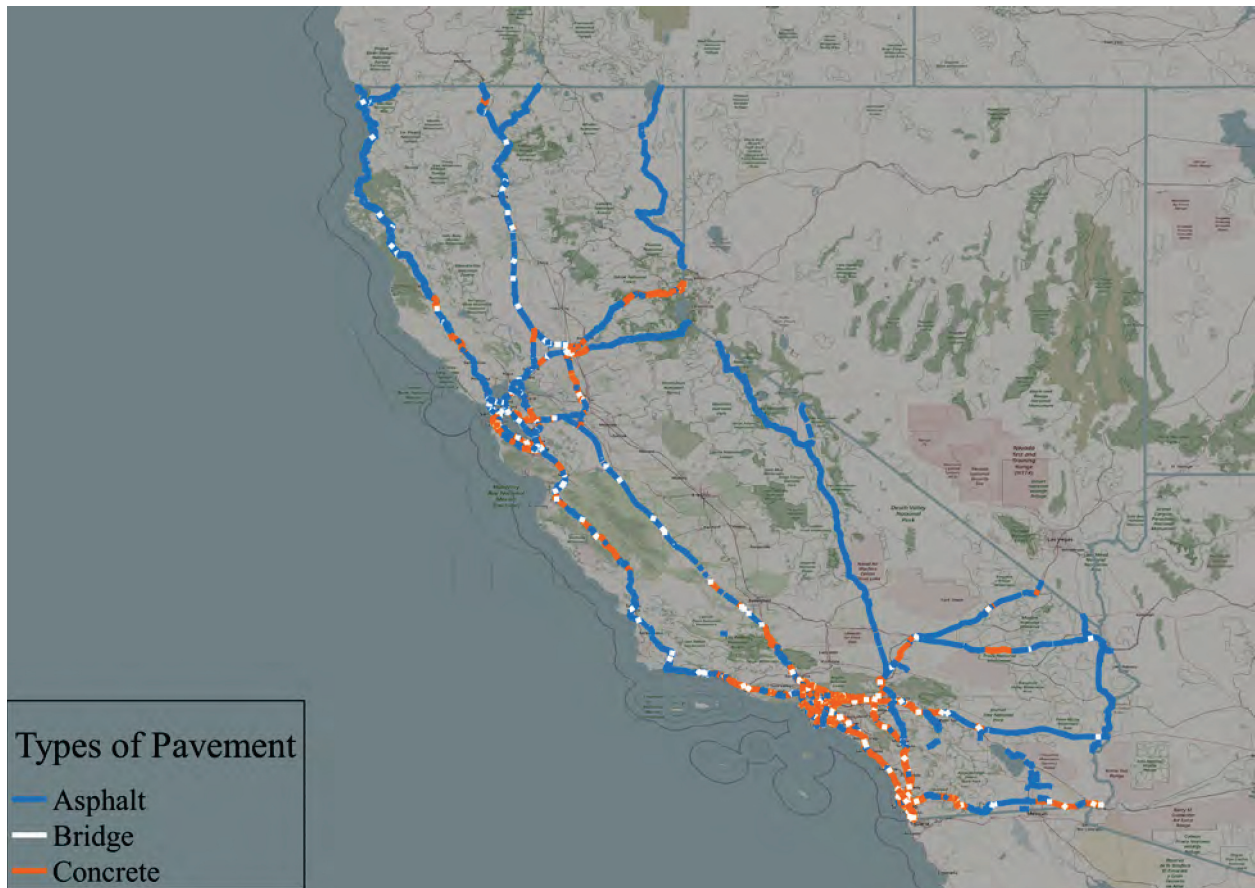
Pavement type represents a design choice made by highway departments. Small and Winston (1988) investigated the impact of different pavement types on pavement durability and found substantial differences. Figure 3 shows pavement types for the analysis sample of FAHS highways. Southern California, and to a lesser extent the Bay and Capital areas, contains mostly concrete pavement in the urban, coastal areas. Segments in less populated areas tend to be asphalt.

We drop all bridge segments from the data used in the empirical analysis. Given the many differences in the characteristics of asphalt and concrete pavement, we estimate regression models using pooled asphalt and concrete segments with an indicator variable for pavement type, and separate regression models for each pavement type.

The APCS manual mentions the presence of *mixed lanes* (“It is possible that two surface material types coexist in the same data collection segment, or that one lane includes a longitudinal joint between JPCP slabs, or strips of CRCP. This often reflects road widening when new lanes use a different surface material than existing lanes. APCS identifies these segments as “mixed lanes.”) We also drop all mixed lane segments from the analysis sample, since no clear method to empirically address this surface heterogeneity exist. Mixed lanes represent 274 one mile segments, or about 1.7% of the sample.

The sample data contain substantially more highway segments with traffic conditions data than segments with traffic conditions and matched contemporaneous IRI data. IRI scan data can be matched to about 6% of the segments with PeMS traffic conditions data. However, expanding the sample to include segments with no contemporaneous IRI scan data would require interpolation of the IRI data to periods between scans. While road roughness generally deteriorates with truck VMT (Small and Winston, 1986; Newbery, 1988), we lack information on truck VMT and resurfacing projects. This complicates interpolation of IRI

Figure 3: Pavement Types



Note: Map shows the spatial distribution of pavement types for California FAHS segments.

for highway segments between IRI scans. The next section addresses the major issues of pavement deterioration over time and interpolation.

Changes in Pavement Quality over Time Standard economic models of pavement construction and maintenance assume that road segments deteriorate over time as vehicles pass over them (Newbery, 1988; Small and Winston, 1988). This assumption rests on accelerated pavement wear experiments conducted by the American Association of State Highway Officials (AASHO) in the early 1960s. These experiments repeatedly drove trucks with variable load weights over test track road segments to simulate pavement wear. The experiments concluded that only truck traffic generated significant declines in the smoothness of road segments and that the deterioration increased with the fourth power of the truck axle loadings, expressed in terms equivalent single axle loads (ESALS), defined in units of 18,000 pounds per axle. Small and Winston (1988) discussed this experiment and described it as “badly flawed.” This remains the primary source of information about the impact of traffic on pavement quality.

Small and Winston (1988) and Newbery (1988) developed models of optimal road durability and pricing to generate funds to repair roads under general conditions. While both models provide considerable insight into economic decisions faced by road policy makers, these models address economic issues beyond the basic idea of the depreciation of roads due to use and weathering. A simple model of pavement deterioration based on standard economic models of capital investment and depreciation can help to understand observed changes in pavement quality over time in our data.

Let $IRI_{s,t}$ represent the quality of pavement segment s in period t as captured by an IRI scan in period t and $IRI_{s,t-k}$ represent the quality of a pavement segment scanned k periods in the past. In this notation $\Delta IRI_{s,t,t-k} = IRI_{s,t} - IRI_{s,t-k}$ represents the observed change in pavement quality for pavement segment s from period $t-k$ to period t . Pavement quality can deteriorate over time ($\Delta IRI_{s,t,t-k} > 0$) because of the number of ESALs, primarily from trucks, that pass over the segment between period $t-k$ and period t and because of weathering occurring during this period. Pavement quality can increase over time ($\Delta IRI_{s,t,t-k} < 0$) because some pavement maintenance or repavement occurred between period $t-k$ and period t .


A simple expression for the observed quality of pavement segment s in period t that reflects these factors can be written

$$IRI_{s,t} = \delta VL_{s,t,t-k} \times IRI_{s,t-k} + \rho R_{s,t-k} \quad (1)$$

where $VL_{s,t,t-k}$ reflects total vehicle loadings traveled over segment s from $t-k$ to t . $R_{s,t-k}$ reflects any maintenance performed on segment s from $t-k$ to t . δ (> 0) captures the depreciation rate of segment s due to both use and weathering. ρ (< 0) captures the effectiveness of maintenance or repavement projects in improving pavement quality.

The parameter δ provides important information about pavement deterioration over time. Small and Winston (1988) emphasized that δ depends on design features of road segments and that choices about these design features affect the size of δ . In principal, Equation (1) could be substituted into the expression for $\Delta IRI_{s,t,t-k}$ to generate a regression model that could be used to estimate δ . However, data limitations make that approach problematic

Figure 4: Scatter Plots of $IRI_{s,t-k}$ versus $IRI_{s,t}$ for Individual Road Segments



Note: These figures show IRI measurements for segment in consecutive scans depending on the length of time between scans. According to Figure A2, the most common lengths of elapsed time are 1-2 months, 9-12 months, and 22-26 months.

in our case. We lack detailed information about repavement and maintenance captured by $R_{s,t-k}$.² We also lack accurate information on truck vehicle miles traveled for road segments. PeMS reports only estimates of total vehicle miles traveled for highway segments. Again, the AASHO experiments conducted in the early 1960s found that only truck ESALS damaged pavement.

We observe both $\Delta IRI_{s,t,t-k}$ and k for a large number of 0.1 mile road segments. We develop stylized facts about $\Delta IRI_{s,t,t-k}$ from this novel data source, abstracting from the impact of vehicle loadings on segment IRI (in other words we assume $VL_{s,t,t-k}$ is constant across time periods and normalize this to 1). From Equation (1), the smaller k , the less likely a pavement maintenance project occurred on a segment between period $t - k$ and t . Our assumption about the impact of $VL_{s,t,t-k}$ on δ also becomes more realistic for smaller values of k .

We first examine actual changes in IRI between scans for different ranges of values for

²PeMS contains information about the timing and location of lane closures for specific road segments. However, these data contain limited information about what activities occurred during these lane closures. Lane closures could reflect mowing, guardrail repair, shoulder work, lighting repair, or other activities not related to pavement maintenance. Kim (2020) analyzes the impact of lane closures identified in PeMS on subsequent traffic volume and speed under the assumption that lane closures represent road work.

k . Appendix Figure A2 shows the distribution of number of months that elapsed between APCS scans (k) for each pavement segment in the sample. A number of notable features appear on this figure. First, from the left end of the figure, the data contain a relatively large number of segments scanned one to two months apart. The next most common time period between IRI scans is roughly 9 to 12 months. This will be useful for comparing APCS data to Highway Performance Management System (HPMS) pavement quality data, which we undertake below. Another smaller peak occurs at roughly 24 months.

Figure 4 shows a scatter plot of IRI measurements for individual 0.1 mile pavement segments at time $t - k$ and t where k is 1 to 2 months on the left panel, 9 to 12 months (roughly 1 year) on the center panel, and 22 to 26 months (roughly 2 years) on the right panel. These groups of time intervals roughly match the distribution of time between scans on Appendix Figure A2 .

The black dashed 45° line on Figure 4 identifies segments with no change in IRI from period $t - k$ to period t (segments where $\Delta IRI_{s,t,t-k} = 0$). Points above this line identify pavement segments with higher IRI (rougher pavement segments) in the current scan relative to the previous scan. Points below this line identify pavement segments with lower IRI (smoother pavement segments) in the current scan relative to the previous scan. The red line represents the fitted values from an OLS regression model where $IRI_{s,t}$ is the dependent variable and $IRI_{s,t-k}$ is the explanatory variable. The slope parameter is a rough, unconditional estimate of δ .

The left panel of Figure 4 reveals substantial noise in segment IRI scan values with very short periods (1 to 2 months) between scans. Conversations with Caltrans traffic engineers confirmed that IRI scans contain substantial noise. These observed IRI differences reflect relatively limited effects of ESALS and weathering, since only 30 to 60 days of loadings occurred between the two scans. Pavement surface maintenance projects likely affect few segments over this short a time period. Rougher segments, located farther to the right on the left panel of Figure 4, would be more likely to undergo some sort of maintenance than smoother segments located further to the left if Caltrans targets extremely rough segments for maintenance of some form.

The slope of the OLS regression line for the left panel of Figure 4 is 0.63, implying that relatively smooth segments tend to get rougher and relatively rough road segments tend to get smoother. The fitted values from the OLS regression model cuts the 45° line at IRI 98. Segments with IRI greater than 98 were, on average, smoother in the next scan. This roughly corresponds with an IRI of 95, the APCS Manual cutoff level between low and medium roughness segments and probably reflects pavement maintenance projects undertaken because of this cutoff.

The regression parameter estimates for $k = 1, 2$ suggest that a smooth pavement segment with IRI of 50 in $t - k$ would have, on average, an IRI of 68 for the next scan, an increase of about 36% over 1-2 months. This increase in pavement roughness seems quite large. However, as we discuss below, this may reflect important differences in traffic loadings on segments scanned every 1-2 months relative to other segments.

The middle panel of Figure 4 uses IRI data from pavement segments scanned roughly 12 months apart. This represents an interesting subsample since the time elapsed between these scans roughly matches the reported values in the HPMS, a key source of pavement quality data. The difference in IRI across scans exhibits less variability than in the subsample using

scans one to two months apart shown on the left panel. Again, segments in this subsample are more likely to have undergone some pavement maintenance between scans relative to the segments in the left panel.

The OLS model parameter estimates predict that a segment with IRI of 50 in the first scan would have, on average, an IRI of 57 in the second scan, an increase of 14%. This reflects a net impact of some segments in this subsample becoming rougher due to traffic loadings and weathering and other segments becoming smoother because of maintenance between the two scans.

The right panel of Figure 4 uses IRI data from pavement segments scanned roughly 24 months apart. Pavement segments in this subsample experience more deterioration from traffic loadings and weathering between scans as well as an increased likelihood of pavement maintenance occurring between scans. The OLS model parameter estimates predict that a segment with IRI of 50 in the first scan would have, on average, an IRI of 58 in the second scan, an increase of 16%, similar to the increase in the subsample of scans occurring roughly 24 months apart.


Figure 5 provides a more complete picture of average $\Delta IRI_{s,t,t-k}$ for all values of k along with the 95% confidence interval for each average shown as whiskers. From Appendix Figure A2, the sample contains many observations for some values of k and relatively few observations for others. For bins with no visible whiskers, the width of the dot is the 95% confidence interval. These bins have many observations. The black dots on Figure 5 identify the values of k that make up the ranges of k depicted on Figure 4: 1-2 months between scans, 9-12 months, and 22-26 months.

Figure 5 highlights the variability of the IRI scan data in the sample. At the left end of the figure, segments scanned one month apart were, on average, significantly rougher for the second scan than for the initial scan but road segments scanned two months apart were, on average, significantly smoother for the second scan than for the first scan. At first glance, this pattern appears to be at odds with the simple model of changes in pavement quality embodied by Equation 1. Road segment quality deteriorates because of axle loadings and weathering. Road segments scanned two months apart experience twice the axle loadings as road segments scanned one month apart, based on our assumption of identical VMT across months for each pavement segment, so they would be expected to be rougher on the second scan compared to the initial scan, absent any road maintenance performed in the period between scans. Road segments scanned between 3 and 5 months apart are also, on average, significantly smoother on the second scan than on the initial scan.

Similar variability can be seen for the road segments scanned between 9 and 12 months apart. Segments scanned nine months apart are significantly smoother on the second scan, segments scanned 10 and 12 months apart are equally as smooth as for the initial scan, and segments scanned 11 months apart are significantly rougher than for the initial scan. Only segments scanned between 22 and 26 months apart, the right most group of solid circles, show unambiguous improvements in pavement quality between scans. One explanation for these patterns is that the inherent noise in APCS scans is large enough to obscure the actual impact of ESALS and weathering on pavement smoothness when the time between scans is one year or less. It is also possible that there are significant differences in ESALS passing over pavement segments with different elapsed times between scans.

The key point revealed by these graphs is that, absent detailed information about the

Figure 5: Average $\Delta IRI_{s,t,t-k}$ By Months Between Scans



Note: This Figure depicts the average $\Delta IRI_{s,t,t-k}$ across road segments for each integer value of months elapsed between scans.

timing and location of pavement maintenance projects and repaving, imputing the month-to-month change in IRI for individual road segments between scans would be a difficult task that could systematically distort the actual change in pavement quality. Consider the differences between segments scanned one month apart and segments scanned two months apart. Segments scanned one month apart got rougher but segments scanned two months apart got smoother. From Equation 1, the increased smoothness should be attributable to pavement maintenance or repaving. But the short period of time that elapsed between these scans reduces the likelihood that some sort of maintenance occurred, unless the results of the initial scan triggered the pavement maintenance. Since we lack information on maintenance, we restrict the empirical analysis to only include outcomes for road segments in the month they were scanned.

Figure 6 displays the distributions of $\Delta IRI_{s,t,t-k}$, the change in IRI, for the three ranges of k values. Table 1 contains summary statistics for these three subsamples. The orange line on Figure 6 reflects road segments with 1-2 months between scans, the purple line 9-12 months between scans, and the blue line 22-26 months between scans. Again, the distribution of $\Delta IRI_{s,t,t-k}$ for road segments scanned between one and two months apart appears to differ substantially from the distribution of the other two.

The variance is largest for the segments with $k \in \{1, 2\}$ which is consistent with the scatter plots on Figure 4. From Table 1, the standard deviation for this subsample exceeds the other two. This is somewhat surprising, since little time passed between these scans. Again, this likely reflects, in part, high variability in IRI values for the same segment due to the way the data are collected. In addition, the distribution of $\Delta IRI_{s,t,t-k}$ for $k \in \{1, 2\}$ has a larger probability mass for positive values than the other two subsamples, indicating right skew in the distribution. The road segments scanned 1 to 2 months apart indicate deterioration in pavement quality over short time periods.


The distribution of $\Delta IRI_{s,t,t-k}$ for the subsample of road segments scanned 9-12 months apart exhibits less variation than the other two subsamples. The mean and median of this distribution are quite close to zero; on average, the quality of these road segments does not change between scans. Again, this reflects the net effect of axle loadings and weathering increasing segment roughness and pavement maintenance reducing roughness. For the subsample of road segments scanned between 22 and 26 months apart, identified by the purple line on Figure 6, the distribution lies between the other two. The mean of this distribution is -8.4 suggesting that these segments, on average, got smoother in the approximately 2 year period between scans. This likely reflects a relatively large number of repaving or maintenance projects occurring in the period between scans.

Table 1: Summary Statistics, Distributions of $\Delta IRI_{s,t,t-k}$ for Ranges of k

Statistic	N	Mean	St. Dev.	Pctl(25)	Median	Pctl(75)
Months Apart ($k \in \{1, 2\}$)	48,144	1.289	39.197	-13.019	1.849	17.102
Months Apart ($k \in \{9, 10, 11, 12\}$)	48,372	0.018	30.809	-8.325	0.101	6.969
Months Apart ($k \in \{22, 23, 24, 25, 26\}$)	25,435	-8.461	35.563	-15.614	-1.734	6.041

Note: This table displays key summary statistics of the distributions of changes in IRI from one scan to the next for scans occurring within 1-2 months, 9-12 months, and 22-26 months.

Figure 6: Distributions of $\Delta IRI_{s,t,t-k}$ for Ranges of k




Note: This Figure shows the distributions of changes in IRI between consecutive scans of 0.1 mile road segments for scans occurring 1-2 months apart (orange line), 9-12 months apart (purple line), and 22-26 months apart (blue line).

While IRI increases with time as pavement deteriorates because of increasing axle loadings, so does the probability of maintenance or resurfacing, which could dramatically decrease IRI. Equation (1) shows the general relationship between pavement quality at time t , wear, and pavement maintenance, absent specific information about the exact time maintenance occurs, we cannot easily distinguish how much wear from traffic and weathering and pavement maintenance individually contribute to observed $\Delta IRI_{s,t,t-k}$.

Figure 7 visually displays information about $\Delta IRI_{s,t,t-k}$ in a different way. Figure 7 averages all values of $\Delta IRI_{s,t,t-k}$ for each initial segment IRI value across the three subsamples identified by time between scans. Again, orange circles identify road segments with 1-2 months between scans, purple circles segments with 9-12 months between scans, and the blue circles with 22-26 months between scans. Each circle is the average change in IRI for all segments with the the same IRI_{t-k} .

The rougher the pavement on the initial scan, circles farther to the right on Figure 7, the more likely it becomes that some pavement maintenance occurred between scans, no matter how much time elapsed between scans. All of these very rough road segments saw substantial increases in quality on the second scan. Patterns in $\Delta IRI_{i,t,t-k}$ on Figure 7 appear generally

Figure 7: Average $\Delta IRI_{s,t,t-k}$ By Initial IRI (IRI_{t-k}) and Months Between Scans



Note: This Figure depicts the average $\Delta IRI_{s,t,t-k}$ for each integer initial IRI value broken down by three subsamples of months elapsed between scans.

consistent with this idea.

The patterns of average values for the orange, purple and blue circles for initial IRI below about $IRI_{t-k} = 100$ run counter to expectations. If these road segments received no maintenance, and experienced the same ESAL volume over time, the purple and blue lines should lie above the orange line, since more time elapsed, and more axle loadings occurred, between scans for the purple and blue circles. Instead, the orange circles generally lie above the purple and blue circles. One possible explanation is that segments scanned 1-2 months apart are high traffic volume/ESAL segments, whereas segments scanned 1-2 years apart are subject to low traffic volume/ESAL segments.

Table 2 contains averages of certain variables reflecting spatial and temporal characteristics of road segments. The first variable, percent of urban cover, is the average urban cover in the county where each road segment is located. Urban cover should be associated with higher axle loadings, since traffic is heavier in more urbanized areas. Segments scanned 1-2 months apart are more likely to be located in counties with a larger percent of urban cover relative to road segments scanned 9-12 and 22-26 months apart. This could be a reason for the substantial increase in roughness for segments scanned 1-2 months apart, especially for segments with low initial IRI values. In addition, the segments scanned 1-2 months apart also have a higher probability of being made of concrete, which ages more quickly than asphalt.

Table 2: Heterogeneity in Scan Group Characteristics

Characteristic	$k \in \{1, 2\}$	$k \in \{9 - 12\}$	$k \in \{22 - 26\}$
% Urban Cover	23.4	17.8	16.5
% Asphalt	54.1	61.2	65.0
Months from June, month $t - k$	3.1	3.0	2.4
Months from June, month t	2.4	2.3	2.7

Note: This table displays spatial and temporal heterogeneity of segment characteristics for road segments scanned 1-2 months, 9-12 months, and 22-26 months apart.

Second, it's possible that the timing of scans in terms of month of the year when the scan occurred could play a role, since more weathering would take place during cold and/or rainy months than during dry and/or warm months. We calculated the average absolute number of months elapsed between the month each segment was scanned and June, a generally temperate and dry month in California. The absolute time from the scan and June decreases on average from 3.1 months for the $t - k$ scan to 2.4 months for the t scan for the group of road segments scanned 1-2 months apart. This indicates that the first scan occurred in the winter or spring. Conversely, even though distance to June for the 9-12 group also decreases from $t - k$ to t , this has a different interpretation. This suggests that the first scan in $t - k$ occurred in autumn, since 9-12 months later the second scan will be closer to June.

Finally, we examine changes in pavement quality between IRI scans by pavement type. In general, wearing differs by pavement type. As discussed above, APCS identifies the surface of all road segments as either flexible (generally asphalt) or rigid (generally concrete). Figure 8 shows the average change in IRI from $t - k$ to t for flexible and rigid surfaces at each initial IRI value. Hollow circles identify rigid segments and dark circles identify flexible segments. Panel A uses the subsample of road segments scanned between 1 and 2 months apart, Panel


B those scanned between 9 and 12 months apart, and Panel C those scanned between 22 and 26 months apart.

As expected, some systematic differences in changes in pavement quality over time between pavement types appear. For relatively smooth pavement segments, those with $IRI_{s,t-k} < 100$ at the left end of the figures, the next scan generally finds rigid segments to be rougher than flexible segments for all three subsamples. The difference is especially pronounced for segments scanned between 1 and 2 months apart, shown in Panel A. This likely reflects the fact that most rigid road segments are in urban areas that experience a large number of axle loadings (see Figure 3) so the pavement quality deteriorated more on those segments.


Figure 8: $\Delta IRI_{s,t,t-k}$ by Initial IRI for Flexible and Rigid Surfaces

Panel A: $\Delta IRI_{s,t,t-k}$ for $k \in \{1, 2\}$

Panel B: $\Delta IRI_{s,t,t-k}$ for $k \in \{9, \dots, 12\}$



Panel C: $\Delta IRI_{s,t,t-k}$ for $k \in \{22, \dots, 26\}$



Note: These figures show the average change in IRI between scans ($\Delta IRI_{s,t,t-k}$) for each initial $IRI_{s,t-k}$ by surface type: flexible (asphalt) or rigid (concrete) for the three subsamples for road segments scanned 1-2 months, 9-12 months, and 22-26 months apart.

For relatively rough pavement segments, those with $IRI_{s,t-k} > 100$ found at the left end of the panels, the next scan generally finds similar improvements in pavement quality for rigid and flexible segments. Smooth flexible segments on Panels B and C show little change in pavement quality between scans.

These differences in pavement quality across pavement types, coupled with the spatial distribution of pavement types, indicates that there may be systematic differences in other traffic outcomes by pavement type. We estimate separate regression models in the empirical analysis described below for each pavement type.

Comparing APCS and HPMS Pavement Quality Data

Our pavement quality data differs from the standard source of pavement quality data, the Highway Performance Monitoring System (HPMS) operated by the U. S. Department of Transportation. The detailed pavement quality obtained from APCS presents a novel and useful context for assessing the nature and quality of pavement smoothness data available in the HPMS. In order to undertake this comparison, we collected HPMS annual data files for California beginning in 1980 from the NBER transportation data archive and the HPMS web site.

While we focus on HPMS road segment quality data from 2011 onward, which matches our APCS data, we also examined all available annual HPMS data from California over the period 1980-2015. In general, the HPMS files contain data on 0.1 mile road segments for all Interstate Highway System roads in the state. State DOTs are only required to submit pavement condition data for Interstate Highway System road segments. Future researchers may be interested in assembling a panel data set from HPMS over a long period of time to analyze log-run changes in pavement quality and the effectiveness of pavement maintenance. Our examination of the annual HPMS data files reveals some challenges to such an undertaking, as the quality and availability of HPMS pavement condition data varies substantially over time in terms of data availability and quality.

The annual HPMS files from 1980 to 1988 for California contain no IRI data. The pavement quality variable in these files reflects the current Present Serviceability Rating (PSR) for road segments. PSR is a subjective 0.0 to 5.0 scale developed by AASHTO where 0.0 represents unpaved or failed road segments and 4.0 to 5.0 new or recently repaved road segments in excellent condition. HPMS annual data files over the 1980 to 1988 period contain no information about pavement resurfacing or maintenance projects carried out in each year.


The annual HPMS files from 1988 to 1998 contain variables capturing IRI, PSR, the year of the last pavement improvement project, and the type of improvement project undertaken for road segments. The existence of information about the most recent pavement improvement project makes this period a good candidate for empirical research on pavement maintenance and road segment surface quality.

The annual HPMS data files from 1998 to 2008 contain variables reflecting IRI and PSR, but no information on road segment improvement projects. These older data can be found at the NBER data repository maintained by the NBER.³ These data may not be comparable to later data, since the FHA switched to a GIS based segment identification system in 2011 from a relatively simple system using only mile posts to identify segments.

Little or no HPMS road segment data are available for 2009 or 2010. The annual HPMS data files from 2011 on contain a variable for road segment IRI only. The year the currently reported IRI scan occurred for each road segment is available only in the 2011 data file.

³<https://www.nber.org/research/data/transportation-economics-21st-century-highway-data>

Figure 9: Year IRI Scan Occurred, 2011 HPMS Annual Data File



Note: Figure shows the frequency of reported IRI scan years in the 2011 California HPMS data. Only 30.4% of the road segments in the 2011 California data file were scanned in 2011.

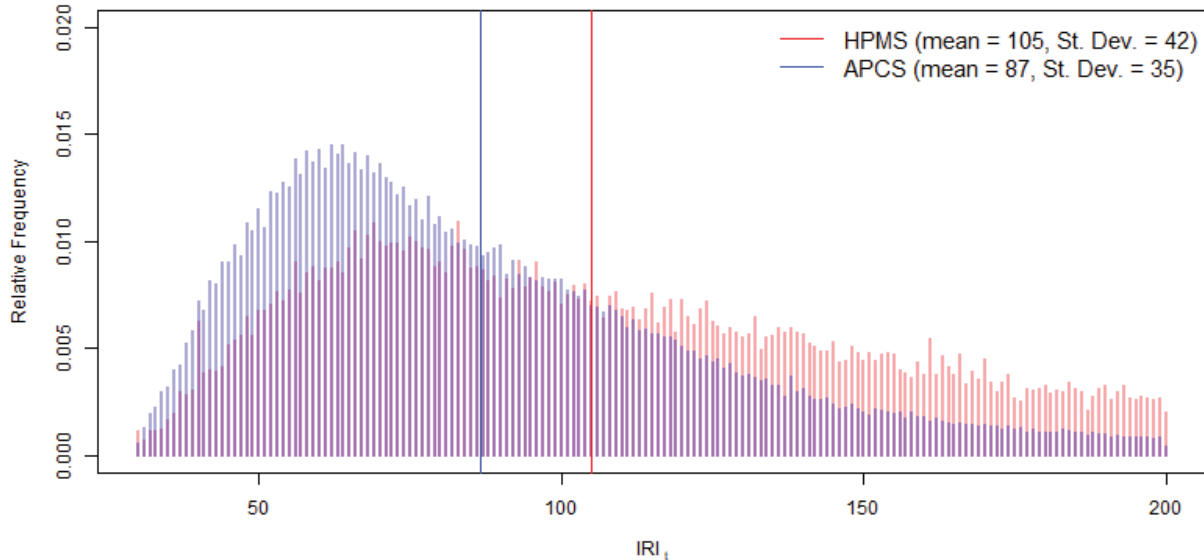
Figure 9 summarizes the year in which each scan reported in the 2011 HPMS data file took place. From Figure 9, only 30% of the highway segment IRI data in the HPMS data file actually came from scans conducted in 2011. 70% of the observations came from other years, including a substantial number of scans from 2012. About 12% of the IRI values in the file were from scans that occurred 10 years earlier. More than 30% of the observations came from scans occurring between 2007 and 2010. Given this temporal disconnect, the reported pavement quality in the annual HPMS data files may not reflect the actual quality of the pavement in that year. Unfortunately, this is the only HPMS annual data file that identifies the year in which the reported scan occurred, so the extent to which these temporal disconnects between annual HPMS pavement quality variable values and the actual condition of those pavement segments cannot be determined.

The 2011 onward data files contain no variable for PSR and no information on surface improvement projects. 2016 HPMS data is missing for California. In the comparison below, we use HPMS data files for FAHS highway segments in California from 2011 to 2015. These are generally available for 0.1 mile road segments.

First, consider the distributions of road segment IRI values from the HPMS and APCS. Figure 10 contains histograms for each integer IRI value in the two samples. Orange bars

represent observations from HPMS and blue bars observations from the APCS. The tall vertical lines represent the sample means. On average, HPMS road segments are rougher than APCS segments.

Figure 10: HPMS and APCS Road Segment IRI Histograms



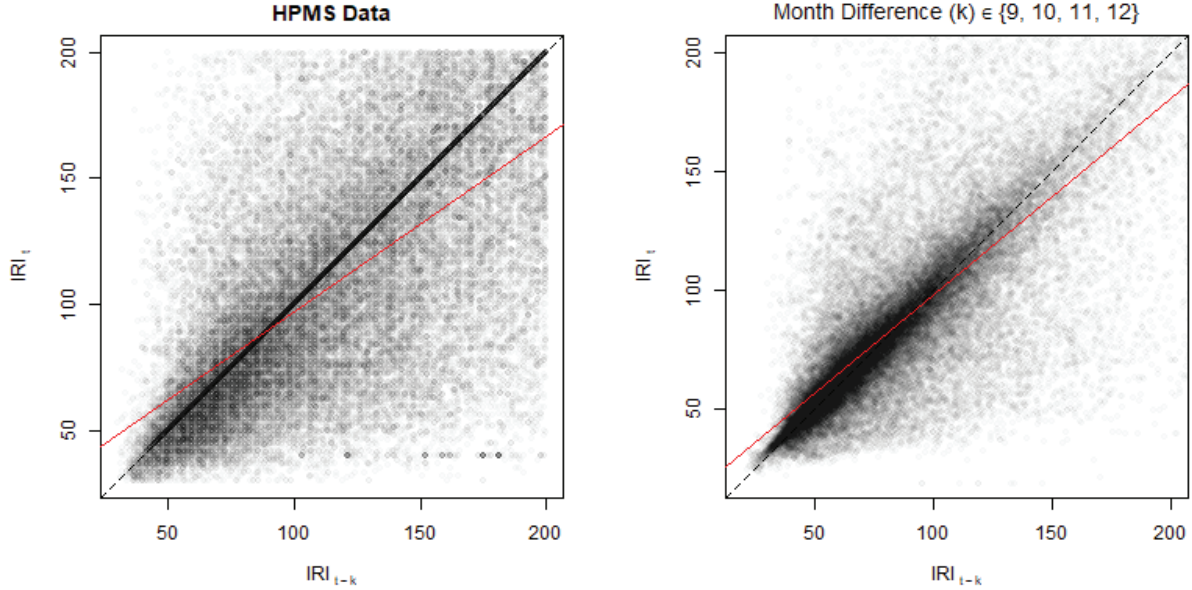
Note: This Figure displays the relative frequency of IRI values for FAHS road segments from the HPMS and APCS datasets.

Figure 10 reveals several other interesting features. First, the shapes of the distributions differ substantially. The distribution for HPMS segments is flatter, and has a fatter right tail than the APCS segments. This can also be seen by the difference in the sample variances for the two data sources. Second, the HPMS data contain more relatively rough road segments, and the APCS data contain relatively more smooth road segments. Both contain relatively long right tails.

We next compare scatter plots of individual road segment IRI values in period $t - k$ and t from the two data sources. The right panel of Figure 11 duplicates the middle panel of Figure 4 above, using APSC data for road segments scanned between 9 and 12 months apart. The right panel shows the scatter plot for HPMS road segments. Again, the red lines represent the fitted values for an OLS regression with dependent variable IRI_t and explanatory variable IRI_{t-k} .

Two notable features stand out on Figure 11. First, the dark thick 45° line for the HPMS data reflects the fact that many road segments in the HPMS data have identical IRI values in consecutive years. More than one in four road segments (27.2%) in the HPMS data have identical year to year IRI values. This is not the case in the APCS data, where only 4.7% of the road segments have the same IRI value from one scan to the next when rounding to the nearest integer. The HPMS Manual instructs transportation agencies to fill in missing IRI scan data with the value from the previous scan when no current IRI data exist. While this may reflect the underlying pavement quality for road segments with low axle loadings and

Figure 11: Scatter Plot of IRI_{t-k} and IRI_t , Annual HPMS and APCS Data



Note: The figure contains scatter plots of IRI_{t-k} and IRI_t for the HPMS data 2011-2017 (left panel) and APCS data for scans occurring between 9 and 12 months apart (right panel). Red lines represent fitted values from OLS regression models: $IRI_{s,t} = b_0 + b_1 IRI_{s,t-k} + e_{st}$.

little weathering, this practice also reduces the variation in IRI values, which will impact any empirical analysis using these data in two ways. First, reducing the variability of a key variable will reduce the estimated standard errors in any regression model. Second, for road segments with large numbers of axle loadings and weathering, replacing missing IRI data with the IRI from the previous year introduces systematic, non-classical measurement error into this variable since these road segments are not randomly distributed across the sample.

The parameter estimates of the regression lines estimated using data from HPMS and APCS also differ. The intercept and slope of the regression line from the HPMS data are 27.456 and 0.696. For the APCS data, the parameter estimates are 15.25 and 0.829.⁴ Both the intercepts and, more importantly, slopes are significantly different from one another. The difference in estimated slope parameters implies that pavement quality declines more rapidly over time in the APCS data relative to the estimates from the HPMS data.


In addition, variation in $IRI_{s,t-k}$ explains more of the variation in observed $IRI_{s,t-k}$ in the APCS data. R^2 for the model estimated with HMPMS data is about 0.455 compared to 0.581 for the model estimated with APCS data.

We finally compare average change in segment pavement quality ($\Delta IRI_{s,t,t-k}$) at each initial IRI value in the HPMS data to the averages in the APCS data, which appears in Figure 7 above. Figure 12 shows average $\Delta IRI_{s,t,t-k}$ for each $IRI_{s,t-k}$ value. In addition to

⁴These estimates are relatively stable across the three subsamples identified by months between scans, although the slope estimates are slightly larger for the 10 and 11 months between scan subsample by about .02 and .04, respectively.

plotting the average values from the APCS, we also plot average values from the HPMS data from California. The most analogous group in the APCS data to the HPMS data would be the subset of APCS data with 9 to 12 months between scans, identified by purple circles on Figure 12.

Figure 12: Average $\Delta IRI_{s,t,t-k}$ for APCS and HPMS Segments by $IRI_{s,t-k}$



Note: This figure adds average $\Delta IRI_{s,t,t-k}$ for HPMS Road Segments to Figure 7. HPMS segment averages appear as an “x.” HPMS reports IRI annually but contains no information about the actual timing of scans.

The average change in IRI across years in the HPMS data generally resembles the changes in the APCS data, especially for relatively smooth road segments between about $IRI=75$ and $IRI=125$ which constitute a substantial part of total observations in both samples (40% of HPMS observations and 36% of APCS observations). Average $\Delta IRI_{s,t,t-k}$ for relatively smooth HPMS road segments, at the left end of Figure 12, tend to be larger than the averages in the APCS for the subsample of scans occurring 9-12 and 22-26 months apart and smaller than the averages for the APCS subsample scanned 1-2 months apart. In general, this comparison indicates that HPMS pavement quality data generally resemble the detailed APCS pavement quality data analyzed in this paper.

However, several caveats apply. The large number of observations in the HPMS data lying on the 45° line on Figure 11 raises some concerns in terms of empirical analysis of pavement quality. Nearly a third of the IRI observations in the HPMS are identical across years. The middle panel of Figure 11 indicates that is not the case in the APCS data. This

substantially reduces the variability of observed IRI in the HPMS. As discussed above, this likely reflects non-classical measurement error in the HPMS generated by the practice of replacing missing IRI data with the IRI value for that segment in the HPMS in the previous year. The HPMS contains less variation in pavement segment quality.

Another limitation of the HPMS pavement quality data is a lack of any information about the exact timing of the pavement segment scans. USDOT instructions for state agencies to collect data over the course of a calendar year and submit data from that year in April of the following year. HPMS contains no information on exactly when pavement segment scans occurred. So, for example, IRI for a road segment scanned in July of 2012 and 2013 would represent the impact of 12 months of axle loadings and weathering. IRI for a segment scanned in January of 2012 and in December of 2013 would actually reflect 23 months of axle loadings and weathering despite appearing in two consecutive annual HPMS data files. IRI for a segment scanned in December 2012 and again in February of 2013 would reflect only two months of axle loadings and weathering. The heterogeneity of changes in IRI by number of months elapsed between scans shown on Figure 5 suggests that interpreting HPMS IRI scans across years as homogeneous may not match the actual pavement conditions.

2.1.3 Weather and Geographic Features

Both weather and the geographic characteristics of road segments can affect pavement quality and traffic outcomes. Weather data come from the National Oceanic and Atmospheric Administration (NOAA), and include average monthly temperature and total monthly precipitation by county. Weather can influence traffic outcomes in a myriad of ways, so controlling for these channels is important to our analysis. APCS contains the elevation of each road segment in feet above sea level.

Groundwater can affect pavement quality by damaging the road base and sub-base. We collected data on groundwater from the United States Geological Survey (USGS) National Water Information System (NWIS). The groundwater database consists of more than 850,000 records of wells, springs, test holes, tunnels, drains, and excavations in the United States and indicates the distance from the surface to groundwater in feet. Appendix Figure A3 shows the location of all groundwater measurement wells in California. Nearby NWIS data exist for only about 5% of the road segments in the sample. Most of these wells exist far from road segments in the sample, so road segments are matched to the closest observation well based on straight-line distance.

As an alternative to groundwater data, we obtained geo-referenced shape files defining the boundaries of all aquifers in California from the USGS. Aquifers contain layers of rock that easily transmit water. The increased ability of aquifers to transmit water could impact pavement quality through impacts on the road's base and sub-base layers. Burchfield et al. (2006) pointed out the influence of aquifers on urban form. Duranton and Turner (2018) use the presence of aquifers as an instrument to predict urban density in an empirical analysis of driving behavior.

Figure 16 shows the locations of all six California aquifers based on 2003 USGS maps and highways in the sample. Figure 16 reveals substantial variation in the presence and absence of aquifers under highways in the sample. While continuous stretches of highway run over the Central Valley aquifer, highways in the northeast and southeast parts of the state, and along

the central coast, run over aquifers with many non-contiguous section, providing substantial variation across highway segments. About 80% of the analysis sample ($N = 12,745$ one mile segments) are located above an aquifer.

2.2 Analysis Sample Construction

We empirically analyze the relationship between pavement quality and both vehicle crashes and vehicle speed. APCS contains extensive information on pavement conditions on roads in California. PeMS contains extensive information on traffic conditions and crashes. We merge data from these two sources, augmented with data from other sources like weather conditions and geological characteristics, to construct an analysis data set.

Again, APCS contains road quality measures for 0.1 mile pavement segments at a point in time while the PeMS ADUS interface provides traffic outcome measures at user determined levels of spatial and temporal aggregation through imputation. We initially extract PeMS data for 1 mile segments over a one month period to facilitate matching with APCS pavement quality data and the Postmile file. PeMS identifies these road segments by route and post miles identified at 1 mile intervals. In cases where highway segments identified in APCS and PeMS do not align exactly with 0.1 mile highway segments in the Postmile file, we use weighed averaging to spatially merge the data. The weights reflect the distance in each APCS 0.1 mile segment and PeMS 1 mile segment that overlaps with 0.1 mile segments in the Postmile file. Appendix section A1.6 contains details on the matching process.

While 0.1 mile highway segments facilitate spatial matching, this represents a less than ideal distance in terms of a unit of analysis in regression models. A vehicle traveling 60 mph takes just 6 seconds to travel 0.1 miles. The impact of pavement roughness on traffic outcomes may not be reflected in road segments this short. In general, we could construct a data set at virtually any level of spatial and temporal aggregation. We use data for one mile road segments over one month as the basic unit of analysis. We also limit the sample to one mile road segments with an APCS scan in that month. This reduces the sample to those segments in green on Figure 1.

This choice reflects several important factors. The shorter the temporal dimension, the more zeros appear in the vehicle crash rate variable. For a one mile, one day sample, more than 70% of the vehicle crash/rate cells contain zeros. This level of temporal aggregation would require count data models to estimate, and preliminary calculations show that the Poisson model requirement of equal mean and variance would not hold. Also, monthly precipitation and average temperature data by county are readily available but daily precipitation and temperature data must be obtained from individual weather monitoring stations. These stations would need to be matched to highway segments and many of these segments have no nearby monitoring stations. Controlling for these factors would be difficult at a daily frequency.

Restricting the sample to months when pavement scans occurred reduces measurement error in the pavement quality estimates. Over time, weathering and road use increase measured IRI (Newbery, 1988). Expanding the sample to include months when no pavement scans occurred means that the actual road quality differs from the measured road quality. While IRI in non-scanned months could be imputed, this approach risks adding systematic measurement error to the IRI variable. Restricting the sample to only months when

pavement scans occurred provides the cleanest possible measure of pavement quality.

We focus on highway segments from either the Interstate Highway System (IHS) or the United States Numbered Highway System (USNHS), which primarily consists of surface-level trunk roads. The FAHS includes both types of highway. APCS and PeMS also contain data on state routes, but we do not include these highways in the analysis because of differences in maintenance responsibility and spotty pavement quality data. The state maintains IHS road segments; both state and local governments maintain USNHS road segments.

These two different types of highways use different design standards and may need to be treated differently in the empirical analysis. IHS highways in the dataset include: I10, I110, I15, I205, I210, I215, I280, I380, I405, I5, I580, I605, I680, I710, I780, I8, I80, I805, I880, I905, and I980. USNHS highways include: US101 and US50.

Figure 13 provides a basic overview of the key variables in the analysis data set. The unit of observation is a one mile segment of road over a one month period. To construct the final analysis data set we dropped all one mile segments with IRI values less than 20 and greater than 300, and all segments with estimated average monthly speed greater than 90 mph. These represented a very small number of one mile road segments. We also drop road segments located above 2000 feet in elevation from the sample for reasons discussed in the identification section below. This represents less than 5% of the sample segments.

Figure 13 shows average IRI, vehicle crash rates, and speed over all highway segments and months in each calendar year in the sample. VMT is aggregated over all segments and months. The IRI values reported on Figure 13 are interpolated in 2013-2014 and 2017 because no APCS scans occurred in those years. The annual values for crash rates in 2013 were interpolated because of missing crash data in that year.

Several clear patterns emerge from Figure 13 in terms of road quality and traffic outcomes in California over the sample period. From Panel A, road quality on FAHS highways generally improved over the period, for both concrete and asphalt segments. Duranton et al. (2020) reported a similar pattern in national data over the period 1980-2006 and note that this runs against the widely held belief of “crumbling” US infrastructure. Concrete segments generally were rougher than asphalt segments. From Panel B, vehicle crash rates per 100,000 VMT generally decreased over the sample period, suggesting an overall improvement in traffic safety.

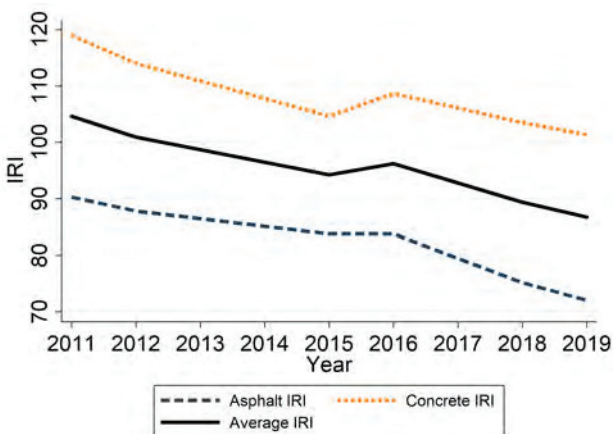
From Panel C, highway use increased over the sample period overall, and on asphalt and concrete road segments, although vehicle miles traveled on asphalt segments fell slightly after 2015. Duranton et al. (2020) documented a similar upward trend in national data. From Figure 3, asphalt highway segments primarily exist in rural parts of the state.

From Panel D, average speeds generally declined over the sample period on both types of pavement. This is consistent with the increasing traffic volume shown on Panel C. Duranton et al. (2020) also found similar trends in average speeds and volume nationwide over a longer period of time.


Figure 14 shows density plots for vehicle crash rates, speed, IRI, and elevation. The crash rate variable shows censoring at zero. About 4,000 of the one mile/one month segments experienced zero accidents on average. Average speed on these one mile/one month road FAHS highway segments is just under 60 mph. The speed histogram reveals some left skew in the distribution because of some urban one mile/one month road segments with very low average speeds. Average road segment IRI is 94 with a long right tail. Again, this

Figure 13: California Pavement and Traffic Outcomes

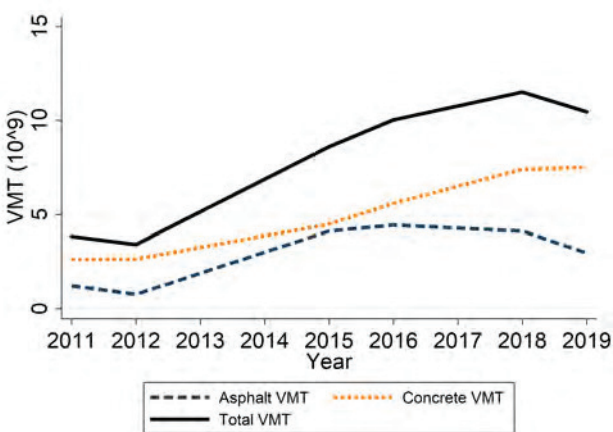
Panel A: IRI



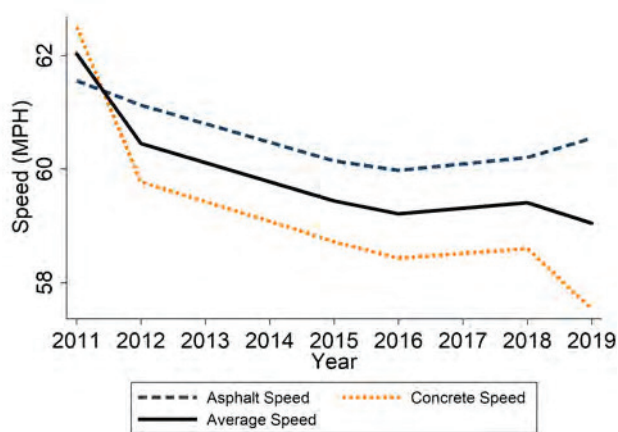
Panel B: Crash Rate per 100,000 VMT



Panel C: VMT



Panel D: Speed (MPH)

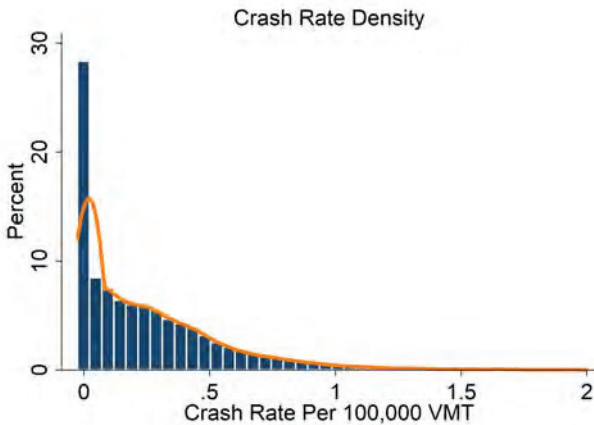


Note: Panels A-D are based on APCS and PeMs traffic data. Panel A: annual average IRI. Values for 2013-2014 and 2017 represent linear interpolation due to no APCS scans in those years. Panel B: crash rate = $\frac{Crash\ Count_y}{VMT_y/100,000}$. 2013 dropped due to some missing data; line is interpolated between 2012 and 2014. Panel C: sum of VMT. Panel D: average speed (MPH). In Panels A-B and D, solid black lines show the state average; in Panel C, solid black lines show the state total. Dashed blue lines show statistics for asphalt segments and orange dashed lines concrete segments.

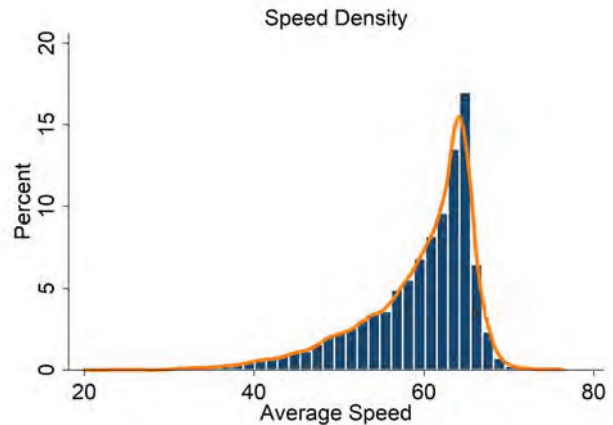
value lies at approximately the standard threshold separating good quality road segments from acceptable. The elevation distribution contains a long right tail reflecting a few road segments at very high altitude in the western part of the state.

Figure 14: Pavement and Traffic Density Plots and Histograms

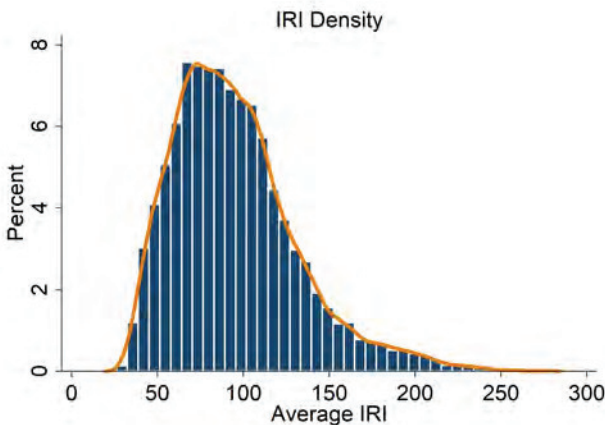
Panel A: Crash Rate



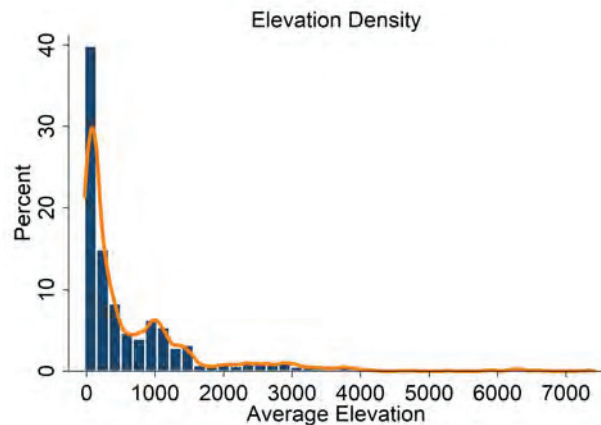
Panel B: Traffic Speed



Panel C: IRI



Panel D: Road Segment Elevation



Note: Panels A-D are density plots and histograms using APCS and PeMs road segment and traffic data for one mile/one month road segments.

Table 3 shows summary statistics for the full analysis sample as well as for asphalt and concrete road segments separately. The full sample IRI average, 94, falls just below the standard threshold separating good ($IRI < 95$) from acceptable ($IRI \geq 95$) pavement quality. Rigid (concrete) segments tend to be rougher than flexible (asphalt) segments and traffic tends to be slightly slower on rigid segments. Roughly 80% of the road segments in the sample lie over aquifers.

2.3 Empirical Strategy

We first undertake a reduced form analysis of the effect of pavement quality, measured by IRI, on vehicle crash rates and speed. The basic reduced form regression model takes the

Table 3: Summary Statistics - Pavement Quality and Traffic Outcomes

	Mean	S.D.	Min.	Max.
Panel A: Full Sample, N=16,002				
IRI	94.18	36.64	27.2	284.6
Crash Rate	0.255	0.321	0.0	9.421
VMT (100,000s)	29.4	14.8	7695	193.0
VHT	51197	28371	130.1	310000
Speed (mph)	59.4	6.753	21.98	76.25
Elevation (ft)	661.9	972.3	-26.0	7406.2
Over Aquifer	0.796	0.403	0	1
Panel B: Asphalt Segments, N=7,581				
IRI	80.4	30.3	27.2	265.8
Crash Rate	0.201	0.283	0.000	5.289
VMT (100,000s)	27.9	15.5	7695	193.0
VHT	47732	28197	130.1	310000
Speed (mph)	60.0	6.27	27.24	76.25
Elevation (ft)	655.7	1109.8	-26.0	7406.2
Over Aquifer	0.770	0.421	0	1
Panel C: Concrete Segments, N=8,421				
IRI	106.6	37.4	31.1	284.6
Crash Rate	0.304	0.344	0.0	9.421
VMT (100,000s)	30.8	14.0	21018	162.0
VHT	54315	28166	350.3	249000
Speed (mph)	58.89	7.12	21.99	75.16
Elevation (ft)	667.4	829.2	-24.2	6640
Over Aquifer	0.821	0.384	0	1

Note: Summary statistics for 1 mile/1 month FAHS segments over the 2011-2019 sample period. Panel A contains all California road segments; Panels B and C split the sample into asphalt and concrete segments respectively. IRI = international roughness index; VMT = vehicle miles traveled; VHT = vehicle hours traveled. Crash rate is crashes per 100,000 VMT.

form

$$TO_{smy}^i = \alpha_f + \alpha_1 IRI_{smy} + \alpha_2 precip_{smy} + \alpha_3 temp_{smy} + \epsilon_{smy}, \quad (2)$$

where TO_{smy}^i is one of two traffic outcomes, either vehicle crash rate ($i = 1$) or vehicle speed ($i = 2$) on California FAHS road segment s in month m in year y . IRI_{smy} reflects pavement quality of road segment s based on a scan that occurred in month m of year y . ϵ_{smy} , a mean zero, possibly heteroscedastic random variable, captures all other time varying factors that affect traffic outcomes on road segment s in month m in year y .

α_f represents a vector of fixed effects indicator variables to control for unobserved heterogeneity in factors affecting traffic outcomes on each highway segment. This vector contains indicator variables for month of year, year, route, county, and pavement type (asphalt of concrete). α_1 represents the parameter of interest capturing the impact of pavement quality of traffic outcomes. We control for the effects of weather on the traffic outcomes by including the average monthly temperature and average monthly inches of precipitation for each one mile highway segment.

We also estimate causal effects of road quality on traffic outcomes. Naive OLS estimates of Equation (2) cannot credibly deliver plausibly causal estimates of the relationship because of the existence of unobserved, time varying factors that affect both pavement quality and traffic outcomes. These factors will be captured by the error term, ϵ_{smy} , and their presence means that the OLS estimate of α_1 cannot be interpreted as causal.

2.3.1 Identification

We use an instrumental variables (IV) approach to estimate the causal impact of pavement roughness on crash rates and speed. Implementing IV requires one or more variables correlated with observable pavement roughness and uncorrelated with unobservable, time varying factors that affect vehicle crash rates and speed.

Pavement quality could be influenced by the presence of groundwater (Christopher et al., 2006; Titus-Glover et al., 2019). If a road segment is relatively close to ground water, the segment has a higher risk of water seeping into the pavement’s base or sub-base. In cold weather environments, these segments could also be susceptible to frost heave which creates potholes. Water penetrating the roadbed can lead to internal damage which reduces the structural integrity of the road segment. Both Christopher et al. (2006) and Titus-Glover et al. (2019) contain substantial evidence of reduced pavement quality associated with road segments built over shallow ground water.

Distance to groundwater is excludable to traffic outcomes such as crashes and traffic speed. The depth of groundwater under a road segment should not be correlated with unobservable factors affecting the number of vehicle crashes on that segment since ground water levels cannot be observed by drivers. It is plausible that the only reason that distance to groundwater affects crashes or travel speeds is through its strong relationship to road quality.

Unfortunately, of the over 5.5 million road segments in APCS, only about 5 percent ($N = 284,096$) are within 0.5 miles of a groundwater measurement well. Again, Appendix Figure A3 shows the road segments in our sample and the location of all groundwater measurement

wells in California. No nearby groundwater measurement wells exist for a substantial number of sample road segments. Matching segments to the nearest groundwater measurement well would lead to either a significantly restricted sample or a weak instrument because of this lack of proximity.

While distance to ground water appears to be an ideal instrument, we lack access to adequate data on this outcome. Instead, we use an instruments known to be related to the water table, the road segment's elevation and an indicator variable for the presence of an aquifer under road segments.

Elevation APCS contains information on the elevation of each scanned road segment. The geology literature contains evidence on the association between elevation and the depth of the water table (Chen et al., 2010; Condon and Maxwell, 2015). Given this strong association, elevation represents a possible instrument for pavement quality. Road segment elevation should not be correlated with traffic outcomes. However, elevation could be correlated with other factors that affect traffic outcomes, most notably weather.

We include variables capturing the average monthly precipitation and temperature for the county where each segment exists in all regression models in order to control for this potential limitation in our instrument.

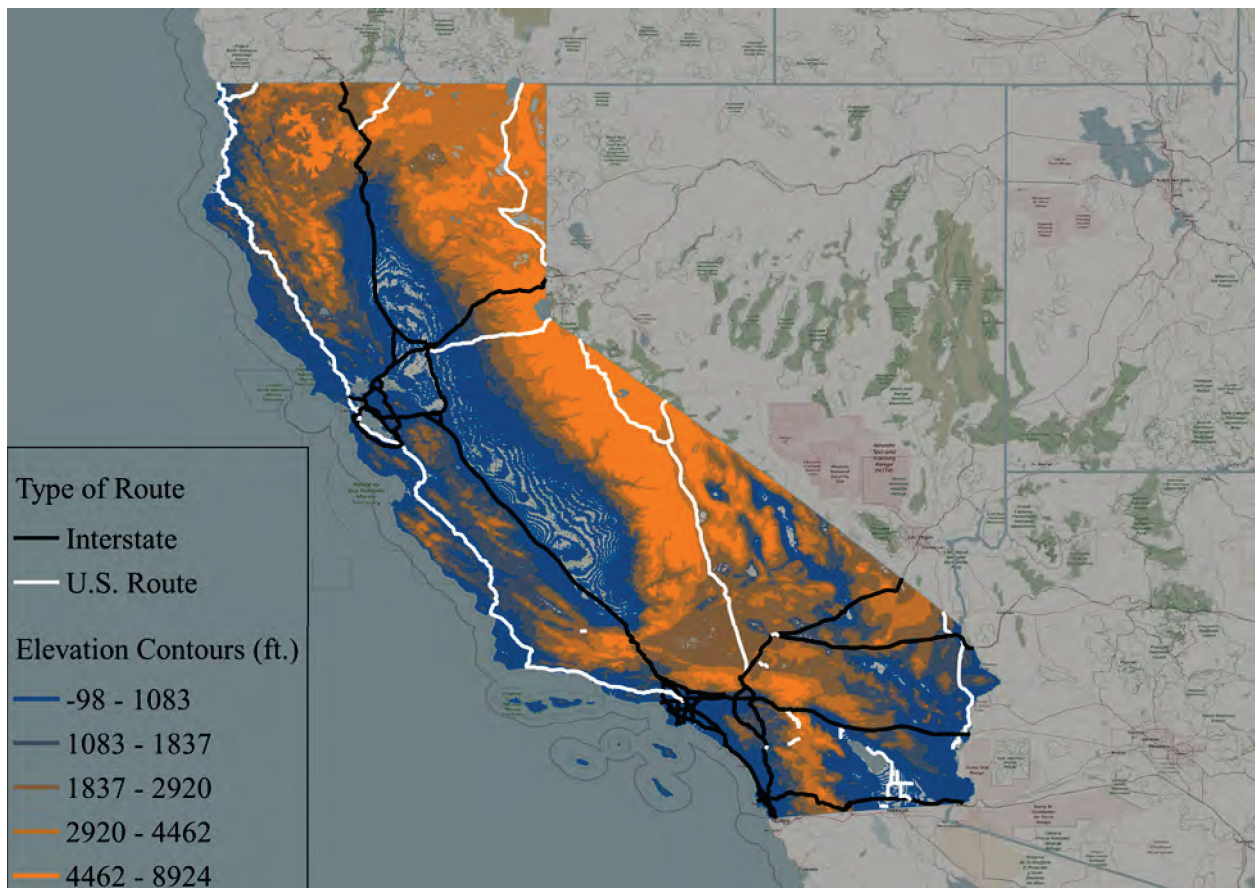
Figure 15 shows sample road segments and elevation contours for feet above sea level in California. The state contains substantial variation in elevation. Coastal areas, and the Central Valley, are low elevation areas. The eastern part of the state contains many highway segments at high elevation, mostly in the Sierra Nevada mountain range. This variability should provide a strong instrument for predicting pavement roughness, holding precipitation and temperature constant.

Aquifers In some cases, elevation may not fully reflect distance to groundwater. Very high elevation represents one possible case where this holds. Table 3 and Figure 15 show some sample road segments at very high elevations, in excess of 6,000 feet. Most of these segments are in the Sierra Nevada mountains in eastern California. Because of this, we drop road segments above 3000 feet in elevation from the sample. This constitutes less than 5% of the sample segments. The results presented below are robust to this exclusion, but including these segments affects instrument strength.

A lack of any groundwater represents a second reason that elevation may not capture distance to groundwater. To address this limitation, we obtained a map of California aquifers from the USGS. Using this map, we identified all road segments existing over an aquifer (coded as one) and all road segments not over an aquifer (coded as zero). Over 70% of the road segments in the analysis sample exist over an aquifer. We posit that these segments could be rougher than segments not over aquifers because of the increased likelihood that ground water seeped into the road bed under these segments.

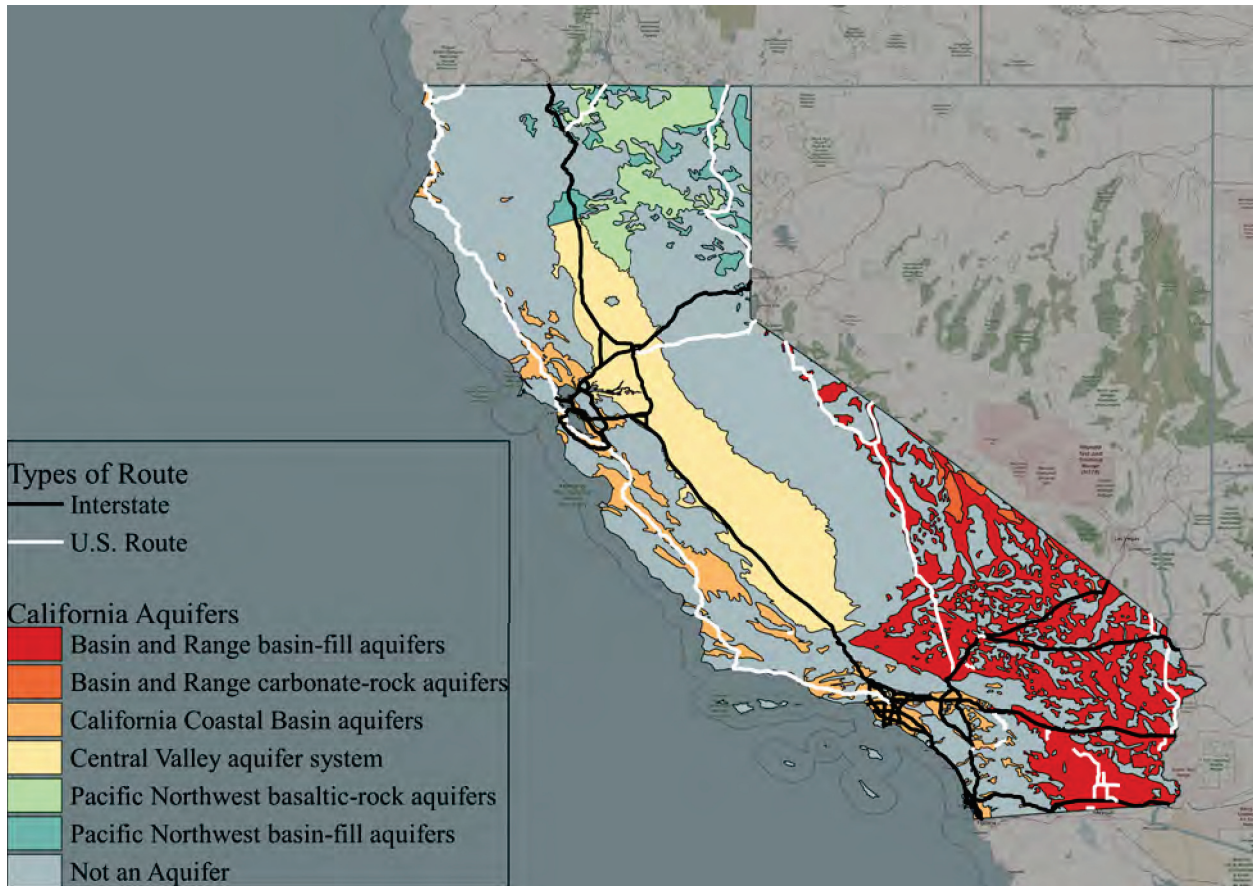
Figure 16 shows the sample California road segments and the aquifers in the state. The sample contains substantial spatial variation in the presence of aquifers under road segments, including in the south west part of the state over the Basin and Range aquifers (in red), in the north west part of the state, over the Pacific Northwest aquifers (in green and blue), and along the coast, under the California Coastal Basin aquifers (in orange). This variation

Figure 15: California Elevation Contours



Note: Map shows the elevation levels in California (feet above sea level).

Figure 16: California Aquifers



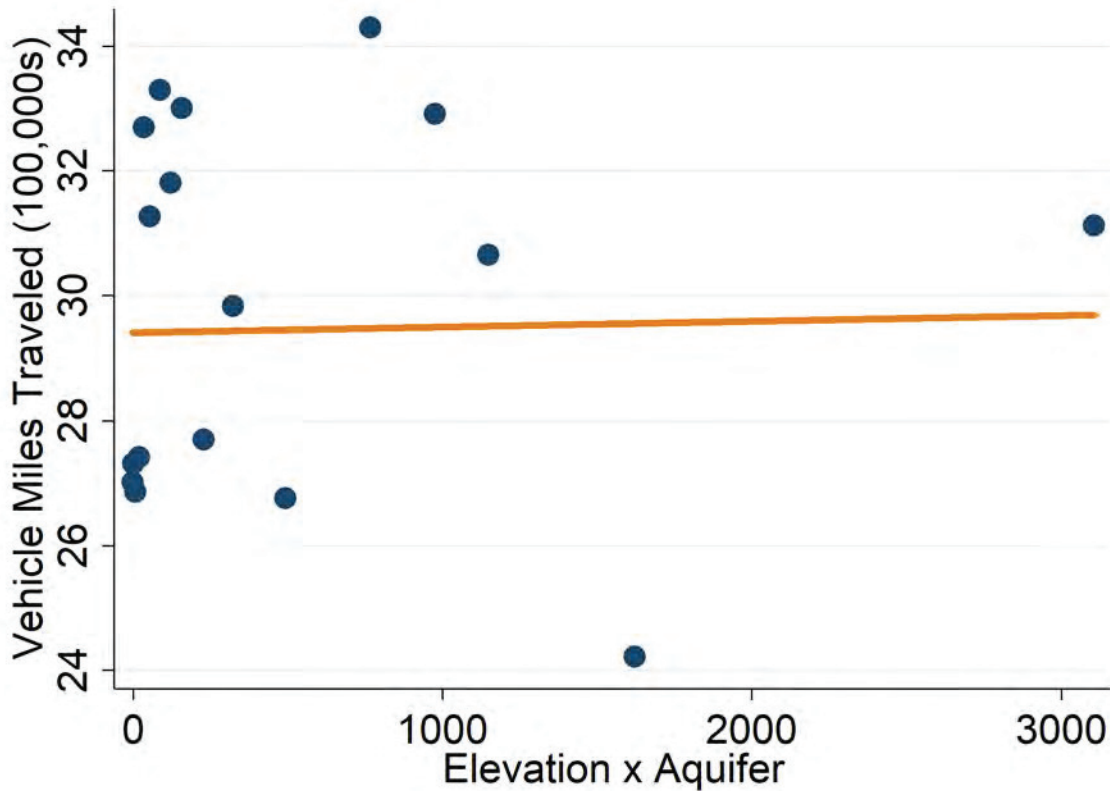
Note: Map shows the location of all California aquifers and FAHS highways in the sample. Aquifer data come from the U.S. Geological Survey in 2003.

occurs at relatively small spatial levels and many contiguous road segments will not be above an aquifer. While some long stretches of road exist over aquifers, most notably for the Central Valley aquifer, this should also provide enough variation to make this a strong instrument for predicting road roughness.

The use of these geographic variables as instruments follows a growing literature making use of similar or identical instruments. Bleakley and Lin (2012) used distance to the fall line on rivers as instrument for portage sites in an analysis of factors causing increases in population density in cities. Feyrer and Sacerdote (2009) used wind speed as instrument for colonial rule in an analysis of how colonial rule caused changes in current income in countries. Dinkelman (2011) used elevation as instrument for electrification in South Africa in an analysis of how electrification caused employment growth in Africa.

Burchfield et al. (2006) found that the presence of aquifers affected urban form, in terms of population density. Other research used aquifers as instruments or explanatory variables. Duranton and Turner (2018) used an aquifer map as an instrument predicting urban population density. The basic idea behind these papers is that the presence of an aquifer facilitates residential construction by limiting the need for construction of water systems.

Figure 17: Relationship between VMT and Road Segment Elevation x Aquifer Presence



Note: Figure shows the binned scatter plot for the instrument of choice, elevation x aquifer, and VMT for 1 mile/1 month FAHS highway segments in the sample.

One potential violation of the IV exclusion restriction would occur if our instrument for pavement roughness, elevation of road segments located over an aquifer, is correlated with unobservable factors that affect traffic outcomes. Traffic outcomes, such as the number of miles driven (VMT), may be associated with underlying geographic characteristics if these characteristics are related to the “urban form” of an area. Also, urban location and sprawl may be directly impacted by elevation or aquifer presence, which, in turn, could influence how much individuals drive and, therefore, any resulting traffic outcomes.

Figure 17 plots a binned scatter plot relationship between our preferred instrument (elevation x aquifer location) and VMT on one mile/one month highway segments. Figure 17 shows no relationship between the two variables. No significant correlation ($\rho = 0.0048$) exists between our instrument and VMT on California FAHS road segments. This mitigates any concerns that the exclusion restriction is violated through this channel.

2.3.2 Instrumental Variables Model

We use the instrumental variables (IV) approach to estimate the causal impact of pavement quality on highway outcomes. The first stage relationship in our IV approach explains

observed pavement roughness. This regression model takes the form

$$IRI_{smy} = \gamma_f + \gamma_1 Z_{smy} + \gamma_2 precip_{smy} + \gamma_3 temp_{smy} + \varepsilon_{smy}, \quad (3)$$

where the dependent variable IRI_{smy} is the International Roughness Index for road segment s scanned in month m in year y . Like in Equation (2) above, γ_f is a vector of indicator variables to control for unobserved heterogeneity in factors affecting pavement roughness for each highway segment. This vector contains indicator variables for month of year, year, route, county, and pavement type (asphalt of concrete). $precip_{smy}$ is total precipitation and $temp_{smy}$ is the average temperature in month m and year y in the county where road segment s is located. The error term, ε_{smy} , a mean zero, possibly heteroscedastic random variable, captures all other time varying factors that affect pavement roughness on road segment s in month m in year y .

Z_{smy} contains one or more instruments that identify exogenous variation in pavement roughness. We have two possible instruments: the elevation of road segment s and an indicator variable equal to 1 if segment s is over an aquifer. Both these variables have been linked to distance to ground water. We estimate an overidentified model containing both variables separately and a just identified model using these two variables interacted. This variable is simply segment elevation for all segments over aquifers and zero for other segments.

The second stage in our IV model uses the fitted values from Equation (3) in place of actual IRI for each segment, the standard 2SLS estimator. The second stage regression model takes the form

$$TO_{smy}^i = \beta_f + \beta_1 \widehat{IRI}_{smy} + \beta_2 precip_{smy} + \beta_3 temp_{smy} + u_{smy}, \quad (4)$$

where TO_{smy}^i is again one of two traffic outcomes, either vehicle crash rate ($i = 1$) or vehicle speed ($i = 2$) on California FAHS road segment s in month m in year y . $precip_{smy}$ is total precipitation and $temp_{smy}$ is the average temperature in month m and year y in the county where road segment s is located. u_{smy} represents a mean zero, possibly heteroscedastic random variable capturing all other time varying factors that affect traffic outcomes on road segment s in month m in year y .

β_f again represents a vector of fixed effects indicator variables to control for unobserved heterogeneity in factors affecting traffic outcomes on each highway segment. This vector contains indicator variables for month of year, year, route, county, and pavement type (asphalt of concrete). β_1 represents the parameter of interest capturing the impact of pavement quality of traffic outcomes.

β_1 , the parameter of interest, reflects the causal impact of pavement roughness on crash rates and speed on California roads. $\widehat{\beta}_1$ reflects a causal effect if the usual IV exclusion restrictions hold. These require the instrument to be correlated with observed pavement roughness and uncorrelated with unobservable factors affecting traffic outcomes.

2.4 Results

2.4.1 Naive OLS and Nonparametric Relationships

We first establish basic patterns in our data using a naive OLS model, Equation (2), and basic graphical depictions using average values for the key variables sorted into bins (binscatter).

Table 4 contains estimates from a naive OLS model using the full 1 mile/month analysis sample. The top panel defines the crash rate, in terms of average monthly crashes per 100,000 VMT on each segment, as the dependent variable. The bottom panel defines average speed in MPH as the dependent variable. Both models contain indicator variables for month of year, year, route, county, and pavement type (asphalt or concrete).

Note that these results could be interpreted as causal if the location and timing of scans of road segments were randomly assigned. APCS scans cover relatively long continuous sections of highway each day, and the vehicles performing the scans generally stay in the same area for several months, based on patterns apparent in Figure 2. But the APCS scans also do not follow obvious patterns in terms of when areas are scanned in terms of months of the year.

The results on Table 4 indicate a positive relationship between road roughness and crash rates and a negative relationship between road roughness and speed. Rougher road segments experience more vehicle crashes, and vehicle speeds decrease on rougher road segments. Both results indicate that rougher road segments impose costs on motorists, in terms of reduced traffic safety (more crashes) and increased delay.

The parameter estimates appear small because they represent the effect of a one inch change in IRI over a one mile road segment. From Table 3, the average IRI in the sample is about 95 inches. Based on the parameter estimates on Table 3, a one standard deviation change in IRI, about 36.5 inches rougher over a one mile segment, would increase the crash rate on that segment by 0.02, or about 8.5%, and decrease average speed by about 0.85 mph, or about 1%. These effect sizes are similar to those reported in the traffic engineering literature discussed above.

From Figure 14, the crash rate variable is left censored at zero. OLS may not be an appropriate estimator with censored dependent variables. We re-estimated Equation (2) using the Tobit estimator to correct for this right censoring. The parameter estimate on the IRI variable was 0.0008, which translates to an 11% increase in crash rates. It appears that the right censoring of the crash rate variable has little impact on the OLS estimates.

We next show that our identification approach approximately holds, in unconditional terms, in the data. We partition the IRI, crash rate, and speed data into bins and take the average of values in each bin. We then generate scatter plots and simple regression lines, for these bins. This approach is a flexible, nonparametric method for examining relationships between pairs of related variables.

We focus on three specific relationships that form the basis for our IV estimates: the reduced form, which shows the relationship between the instrument and the outcome variable, the first stage which shows the relationship between the instrument and the endogenous explanatory variable, and the second stage, which shows the relationship between the endogenous explanatory variable and the outcome variable.

Figures 18 and 19 visually depict the reduced form, first stage, and second stage rela-

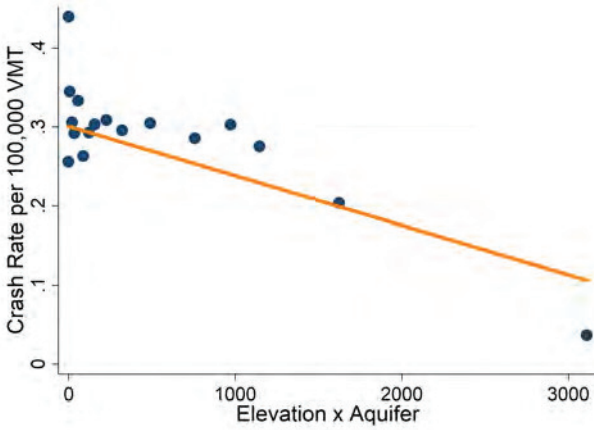
Table 4: Naïve OLS Estimates of IRI Association
with Traffic Outcomes

Panel A: Crash Rate	
IRI	0.00064*** (0.00010)
Precipitation	0.00279*** (0.00026)
Temperature	0.00381*** (0.00092)
Observations	16,002
R^2	0.1498
Panel B: Speed	
IRI	-0.02327*** (0.00165)
Precipitation	-0.00909* (0.00504)
Temperature	-0.04863*** (0.01532)
Observations	16,002
R^2	0.3451

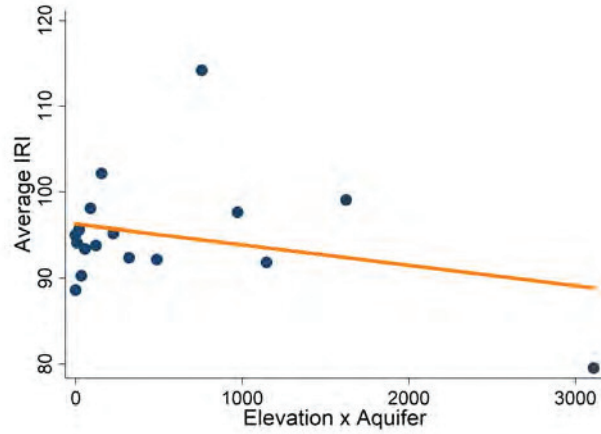
Notes: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Dependent variable is aggregated to the 1 mile/monthly level. Segments with average IRI values < 20 or > 300 are dropped from the sample. Segments with average speeds greater than 90 MPH are also dropped. Panel A dependent variable: vehicle crash rate per month per 100,000 VMT. Panel B dependent variable: average speed in MPH. Separate month, year, route, county, and pavement type fixed effects are included. Controls for average precipitation and temperature are also included. Heteroskedastic-robust standard errors are used in every panel.

Figure 18: Relationships between Key Variables: Elevation, IRI, and Crash Rate

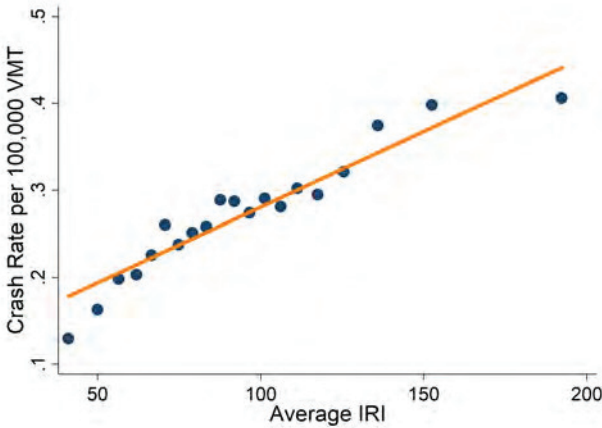
Panel A: Reduced Form: Elevation x Aquifer and Crash Rate



Panel B: First Stage: Elevation x Aquifer and IRI



Panel C: Second Stage: IRI and Crash Rate




tionships detailed in Equations 3 and 4 for the traffic outcomes of interest: crash rates and average speed.

Figure 18 summarizes the relationship between our preferred instrument, road segment elevation interacted with an indicator variable for segments located over aquifers, road roughness, and average crash rates. Panel A shows a negative relationship between the instrument and crash rates. Panel B shows a negative first stage relationship between the instrument and the endogenous explanatory variable, road segment roughness. Panel C shows a positive relationship between the endogenous explanatory variable, road roughness, and crash rates, the traffic outcome variable of interest.

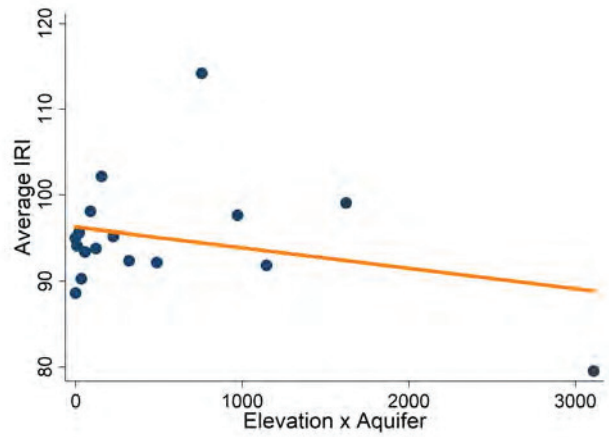
Figure 19 summarizes the relationship between our preferred instrument, road roughness, and average vehicle speed. Panel A shows a positive relationship between the instrument and average speed. Panel B shows a negative first stage relationship between the instrument and the endogenous explanatory variable, road segment roughness. Note that this is identical to Panel B on Figure 18. Panel C shows a negative relationship between the endogenous explanatory variable, road roughness, and average speed, the second traffic outcome variable

Figure 19: Relationships between Key Variables: Elevation, IRI, and Speed


Panel A: Reduced Form: Elevation x Aquifer and Speed



Panel B: First Stage: Elevation x Aquifer and IRI



Panel C: Second Stage: IRI and Speed



of interest.

2.4.2 IV Results

We next estimate the causal impact of road roughness on crash rates and average speed for the entire sample of road segments in California. We use the 2SLS approach with a first stage regression model defined by Equation (3) and a second stage regression model defined by Equation (4). If the exclusion restrictions hold for the instrument Z , then the parameter estimate $\hat{\beta}_1$ from Equation (4) reflects the causal impact of changes in road roughness on traffic outcomes.

Table 5 contains IV results for the full sample using our preferred instrument, road segment elevation interacted with an indicator variable identifying segments over aquifers, a proxy for distance to groundwater for each segment. Both Equation (3) and Equation (4) contain variables reflecting average monthly precipitation and temperature in the county where each segment exists and the full set of fixed effects. The top panel on Table 5 defines crash rates per 100,000 VMT as the dependent variable and the bottom panel defines average speed as the dependent variable.

The key IV results appear in Column (3) of both panels. The first stage regression F statistic (Kleibergen-Paap F) indicates sufficient instrument strength to identify a causal effect. The basic relationship between road roughness and the two traffic outcomes found in the OLS results and bin scatter analysis hold up in the IV analysis. Rougher roads cause more traffic accidents and cause vehicles to drive more slowly, other things equal. Allowing pavement to deteriorate imposes costs on motorists.

The IV parameter estimates are larger in absolute value than the OLS estimates on Table 4. The second stage parameter for the model using average speed as the dependent variable implies that a one standard deviation (36 inch) increase in IRI on a road segment reduces average speed on that segment by 6.4 MPH, about an 11% decrease.

The distribution of the crash rate variable has a fair number of one mile/one month road segments with zero crashes per month on average and a relatively long right tail. This somewhat complicates characterization of the effect size. The second stage parameter estimate for the model using average vehicle crash rate as the dependent variable implies that a one standard deviation increase in IRI would move a road segment with an average of zero crashes per month to an average crash rate about 0.6 standard deviations above the overall sample average monthly crash rate. In other words, this increase in IRI would change a safe stretch of road where no crashes occur to a stretch of road with an above average crash rate.

Alternative IV models that use the segment elevation and aquifer identifier variables separately, in other words an overidentified IV model, generate similar results, in terms of parameter signs and statistical significance, but the Kleibergen-Paap F statistics are smaller. Alternative IV models using the Tobit estimator for the second stage generate similar results for the crash rate model. The parameter estimate on the IRI variable in the second stage IV Tobit model was 0.02 and statistically significant at the 1% level. The results for the causal impact of road damage on crash rates remain robust to correcting for left censoring of the crash rate variable.

These results have important implications for traffic safety, vehicle operating costs, and

Table 5: IV Estimates of Causal Effect of IRI on Traffic Outcomes

	(1) Reduced Form	(2) 1st Stage	(3) 2nd Stage
Panel A: Crash Rate			
Instrument: Elevation x Aquifer	-0.00004*** (0.00001)	-0.00331*** (0.00058)	
\widehat{IRI}			0.01287*** (0.00217)
Observations	16,002	16,002	16,002
Kleibergen-Paap F			33.089
Panel B: Speed			
Instrument: Elevation x Aquifer	0.00058*** (0.00010)	-0.00331*** (0.00058)	
\widehat{IRI}			-0.17646*** (0.03526)
Observations	16,002	16,002	16,002
Kleibergen-Paap F			33.089

Notes: * p<0.1, ** p<0.05, *** p<0.01. Dependent variables aggregated to the 1 mile/1 month level. Segments with average IRI values < 20 or > 300 are dropped from the sample. Segments with average speeds greater than 90 MPH are also dropped. Column 1 shows reduced form results (elevation x aquifer on the traffic outcome), Column 2 shows first stage results (elevation x aquifer on IRI), and Column 3 shows second stage results (\widehat{IRI} on each traffic outcome). Panel A dependent variable is the crash rate per 100,000 total VMT. Panel B dependent variable is average speed in miles per hour. Separate month, year, route, and county fixed effects and controls for average precipitation and temperature are also included. Heteroskedastic-robust standard errors are reported in every column and panel.

consumer welfare. IRI increases because of vehicle use and weathering. Transportation departments can reduce IRI through maintenance projects. Delaying maintenance aimed at reducing IRI means rough road segments will experience continued IRI increases. This will cause increases in crash rates on those segments, reducing traffic safety and increasing vehicle operation costs for all motorists using those road segments. The increased roughness will also reduce vehicle speed, generating welfare loss from increased travel times (Couture et al., 2018) and increased vehicle operating costs because vehicles traveling at lower speeds burn more fuel.

2.5 IV Results by Pavement Type

Small and Winston (1988), and the traffic engineering literature, point out differences in roughness and rate of deterioration for flexible (asphalt) and rigid (concrete) pavements. Figure 13 shows some systematic differences for the two pavement types in terms of roughness, crash rates, and speed. This suggests the possibility of different relationships between road segment roughness and traffic outcomes for these two pavement types. The APCS data identify surface type for all road segments.

Figure 3 shows the spatial distribution of pavement types in the sample. From Figure 3, long contiguous stretches of asphalt road segments exist outside metropolitan areas in California, particularly in the northern and eastern part of the state. In urban areas, particularly in southern California, the Bay Area, and around Sacramento, pavement types vary at small spatial scales. Splitting the sample into subsamples by pavement type will only generate long contiguous heterogeneous stretches of road for those asphalt road segments outside the urban centers in the state. This could affect the estimates.

In addition, splitting the sample by pavement type could affect instrument strength. Taken together, Figure 3, Figure 16, and Figure 15 show clear spatial patterns in the distribution of pavement types and aquifers by elevation. We find this to be the case.

Appendix Table A1 shows the results from separate IV models estimated for the asphalt and concrete road segment subsamples. Columns (1) - (3) contain results for asphalt road segments and Columns (4) - (6) results for concrete segments. Kleibergen-Paap F statistics decline substantially in both subsamples, and we lose instrument strength in the concrete subsample. The first stage parameter estimates across both models changes completely relative to those for the full sample on Table 5.

Further reducing the sample to include only road segments at relatively lower altitude clearly shows the role played by altitude in terms of instrument strength. Table 6 shows IV separate results for asphalt and concrete when the sample is restricted to include only road segments at 2000 feet or lower. Splitting the sample at 2000 feet roughly divides the state in half east-west. Note that this restriction does not reduce the sample sizes by much. The number of asphalt road segments declines to 6,733 from the full sample asphalt segment count of 7,581. The number of concrete road segments falls to 7,820 from 8,421 full sample concrete road segments.

The results on Table 6 strongly resemble the full sample results on Table 5 in terms of parameter estimate sign and significance. The Kleibergen-Paap F statistics increase relative to those on Table 5. The size of the second stage parameters generally match those for the full sample. Rougher pavement causes a relatively larger decline in average speed on asphalt

Table 6: IV Estimates of Causal Effect of IRI on Traffic Outcomes at Lower Elevations

	(1) Reduced Form	(2) 1st Stage	(3) 2nd Stage
Panel A: Crash Rate			
Instrument: Elevation x Aquifer	-0.00010*** (0.00001)	-0.00868*** (0.00131)	
\widehat{IRI}			0.01031*** (0.00178)
Observations	14,553	14,553	14,553
Kleibergen-Paap F			44.000
Panel B: Speed			
Instrument: Elevation x Aquifer	0.00320*** (0.00021)	-0.00868*** (0.00131)	
\widehat{IRI}			-0.36820*** (0.05676)
Observations	14,553	14,553	14,553
Kleibergen-Paap F			44.000

Notes: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Dependent variable is aggregated to the 1 mile/monthly level. Segments with average IRI values < 20 or > 300 are dropped from the sample. Segments with average speeds greater than 90 MPH are also dropped. Column 1 shows reduced form results (elevation x aquifer on the traffic outcome), Column 2 shows first stage results (elevation x aquifer on IRI), and Column 3 shows second stage results (\widehat{IRI} on the traffic outcome). Panel A examines the crash rate per 100,000 total VMT. Panel B examines average speed. Separate month, year, route, and county fixed effects are used in every panel. Controls for average precipitation and temperature are also included. Segments with an average elevation above 2000 ft. are dropped. Heteroskedastic-robust standard errors are used in every column and panel.

segments. Rougher pavement causes a relatively larger increase in crash rates on concrete segments.

In general, the full sample results are robust to pavement type, as long as the analysis focuses on road segments at an elevation of 2000 feet or less, which only removes a small number of road segments from the sample. For completeness, Appendix Table A2 shows IV results for the pooled asphalt and concrete segments for the subsample of road segments at 2000 feet or less. These results also closely resemble the IV results on Table 5.

Note the larger Kleibergen-Paap F statistics on Appendix Table A2 compared to those on Appendix Table A1. Our preferred instrument is stronger at lower altitudes and deteriorates at higher altitudes, controlling for precipitation and temperature differences across these areas. This likely reflects a deterioration in the relationship between altitude and distance to ground water at high altitudes.

3 Conclusions

We develop evidence that rougher roads cause increases in vehicle crash rates and decreases in average vehicle speed on Federal Aid Highway System roads in California. Results from Instrumental Variables models using the interaction of road segment elevation with an indicator variable identifying road segments over aquifers indicate that a one standard deviation increase in a road segment's International Roughness Index, an increase of about 36 inches in vertical motion over a one mile road segment, decreases average vehicle speed by 11%, a substantial reduction. A one standard deviation increase in IRI also causes a substantial increase in vehicle crash rates. This increase would move a one mile road segment with an average of zero crashes per month, a very safe road segment, to an average crash rate approximately 0.6 standard deviations above the average crash rate in the sample. The deterioration of road surface quality generates substantial negative impacts on drivers and passengers.

To our knowledge, our results represent the first causal estimates of the impact of road damage on traffic outcomes. This paper begins to fill the sizable gap in the transportation literature identified by Duranton et al. (2020) by addressing the fundamental relationship between deteriorating pavement and both vehicle crashes and speed. The limited evidence in the transportation engineering literature made no use of causal inference methods and found a weak relationship between these variables. The OLS results in this paper also indicate a small impact of road damage on traffic outcomes. The larger IV estimates provide important new evidence on the extent to which road damage affects traffic safety, vehicle operating costs, and travel time.

These results have important implications for transportation infrastructure policies. Winston (1991) pointed out that efficient transportation infrastructure policy requires appropriate pricing of user fees and optimal investment, including new lanes and roads and road maintenance. Efficient user fees reflect the full marginal cost of transportation infrastructure use, including the cost of damage inflicted on roads due to use. Optimal transportation infrastructure investment trades off construction costs with future maintenance costs, so understanding the consequences of failure to maintain roads, resulting in rougher road surfaces, represents an important part of this puzzle. The results developed here suggest a larger full

marginal cost of road use than previously thought, implying a need for more transportation infrastructure investment, especially for road maintenance.

Small and Winston (1988) showed that pavement design practice leads to underinvestment in pavement durability. In practical terms, this means rougher existing roads than would have existed across the entire highway network under different pavement durability decisions. The results in this paper represent an important extension of our understanding of the consequences of this underinvestment. In addition to the increased direct maintenance costs, the results here reveal significant social costs, in the form of reduced traffic safety, increased vehicle operating costs, and longer travel times. These costs impact individuals and firms, and underscore the need for changes in current transportation infrastructure investment policies.

The detailed pavement quality data from APCS collected for this research allow for a novel assessment of the pavement quality data in the U. S. Department of Transportation's widely used Highway Performance Monitoring System. While the IRI data in the HPMS generally resemble the APCS data, our comparison reveals several potential limitations in the HPMS. The time elapsed between pavement scans in the HPMS pavement quality data cannot be determined, which complicates empirical analysis of changes in pavement quality over time. Also, the practice of filling in missing annual IRI values with the values reported in previous years reduces the variation in IRI data in HPMS which also impacts empirical analysis of these data.

Finally, the analysis of changes in road segment IRI over time highlights the importance of information about highway maintenance and repavement projects for understanding changes in road segment quality over time. Pavement deterioration negatively affects traffic safety and speed, highlighting the importance of improving the quality of rough road segments through maintenance or repaving. The observed changes in IRI analyzed here represent the impact of both axle loadings reducing pavement quality and unobserved maintenance that improves pavement quality. Older HPMS data contain information on both pavement quality and the time of the most recent maintenance, but the maintenance timing variable disappeared from the HPMS data files years ago. Development of data that contain both pavement quality and maintenance information should be a priority for transportation agencies.

References

- Alqadhi, S., Ghahari, S. A., Chen, S., Volovski, M., and Woldemariam, W. (2018). Costs and benefits of highway resurfacing: a case study of Interstate 465 in Indiana, USA. *Infrastructure Asset Management*, 5(2):45–55.
- Bento, A., Roth, K., and Waxman, A. R. (2020). Avoiding traffic congestion externalities? the value of urgency. Working Paper w26956, NBER.
- Bleakley, H. and Lin, J. (2012). Portage and path dependence. *The Quarterly Journal of Economics*, 127(2):587–644.
- Burchfield, M., Overman, H. G., Puga, D., and Turner, M. A. (2006). Causes of sprawl: A portrait from space. *The Quarterly Journal of Economics*, 121(2):587–633.

- Chen, C.-H., Wang, C.-H., Hsu, Y.-J., Yu, S.-B., and Kuo, L.-C. (2010). Correlation between groundwater level and altitude variations in land subsidence area of the Choshuichi Alluvial Fan, Taiwan. *Engineering Geology*, 115(1-2):122–131.
- Christopher, B. R., Schwartz, C., and Boudreau, R. (2006). Geotechnical aspects of pavements reference manual. Technical Report FHWA NHI-05-037, Federal Highway Administration.
- Condon, L. E. and Maxwell, R. M. (2015). Evaluating the relationship between topography and groundwater using outputs from a continental-scale integrated hydrology model. *Water Resources Research*, 51(8):6602–6621.
- Couture, V., Duranton, G., and Turner, M. A. (2018). Speed. *Review of Economics and Statistics*, 100(4):725–739.
- Dinkelman, T. (2011). The effects of rural electrification on employment: New evidence from South Africa. *American Economic Review*, 101(7):3078–3108.
- Duranton, G., Nagpal, G., and Turner, M. A. (2020). Transportation infrastructure in the US. In *Economic Analysis and Infrastructure Investment*. University of Chicago Press.
- Duranton, G. and Turner, M. A. (2011). The fundamental law of road congestion: Evidence from US cities. *American Economic Review*, 101(6):2616–52.
- Duranton, G. and Turner, M. A. (2018). Urban form and driving: Evidence from US cities. *Journal of Urban Economics*, 108:170–191.
- Feyrer, J. and Sacerdote, B. (2009). Colonialism and modern income: Islands as natural experiments. *The Review of Economics and Statistics*, 91(2):245–262.
- Kim, J. (2020). Do roadworks benefit road users? evidence from California freeways. Working paper, California State University, Long Beach.
- Koh, Y. and Lee, G. M. (2020). Infrastructure investment and travel time. *Economics Letters*, 187:108901.
- Lewis, G. and Bajari, P. (2011). Procurement contracting with time incentives: Theory and evidence. *The Quarterly Journal of Economics*, 126(3):1173–1211.
- Li, Y. and Huang, J. (2014). Safety impact of pavement conditions. *Transportation Research Record*, 2455(1):77–88.
- Link, H. (2014). A cost function approach for measuring the marginal cost of road maintenance. *Journal of Transport Economics and Policy (JTEP)*, 48(1):15–33.
- Mamlouk, M., Vinayakamurthy, M., Underwood, B. S., and Kaloush, K. E. (2018). Effects of the international roughness index and rut depth on crash rates. *Transportation Research Record*, 2672:418–429.

- Newbery, D. M. (1988). Road damage externalities and road user charges. *Econometrica: Journal of the Econometric Society*, pages 295–316.
- Small, K. A. and Winston, C. (1986). Efficient pricing and investment solutions to highway infrastructure needs. *The American Economic Review*, 76(2):165–169.
- Small, K. A. and Winston, C. (1988). Optimal highway durability. *The American Economic Review*, 78(3):560–569.
- Titus-Glover, L., Darter, M. I., and Von Quintu, H. (2019). Impact of environmental factors on pavement performance in the absence of heavy loads. Technical report, Federal Highway Administration.
- TRIP (2018). California transportation by the numbers: Meeting the state’s need for safe, smooth and efficient mobility. Report, TRIP: A National Transportation Group.
- Vickrey, W. S. (1969). Congestion theory and transport investment. *The American Economic Review*, 59(2):251–260.
- Wang, T., Harvey, J., Lea, J., and Kim, C. (2014). Impact of pavement roughness on vehicle free-flow speed. *Journal of Transportation Engineering*, 140(9):1–11.
- Winston, C. (1991). Efficient transportation infrastructure policy. *Journal of Economic Perspectives*, 5(1):113–127.
- Winston, C. (2020). Efficient transportation infrastructure policy revisited: A thirty-year update motivated by Brooks And Liscow. In *Economic Analysis and Infrastructure Investment*. University of Chicago Press.
- Winston, C. and Langer, A. (2006). The effect of government highway spending on road users’ congestion costs. *Journal of Urban Economics*, 60(3):463–483.
- Yu, B. and Lu, Q. (2014). Empirical model of roughness effect on vehicle speed. *International Journal of Pavement Engineering*, 15(4):345–351.

Appendix

A1.1 Elevation, Pavement Type, and Instrument Strength

Table A1 shows full sample IV results for the asphalt and concrete subsamples.

Table A1: IV Estimates of Causal Effect of IRI on Traffic Outcomes by Pavement Type

	(1) Asphalt RF	(2) Asphalt 1st Stage	(3) Asphalt 2nd Stage	(4) Concrete RF	(5) Concrete 1st Stage	(6) Concrete 2nd Stage
Panel A: Crash Rate						
Elevation x Aquifer	-0.00001*** (0.00000)	0.00230*** (0.00073)		-0.00005*** (0.00000)	0.00051 (0.00069)	
\widehat{IRI}			-0.00590*** (0.00186)			-0.10642 (0.15229)
Observations	7,581	7,581	7,581	8,421	8,421	8,421
Kleibergen-Paap F			15.929			0.495
Panel B: Speed						
Elevation x Aquifer	-0.00017 (0.00013)	0.00230*** (0.00073)		0.00144*** (0.00010)	0.00051 (0.00069)	
\widehat{IRI}			-0.07254 (0.05732)			2.80568 (3.99944)
Observations	7,581	7,581	7,581	8,421	8,421	8,421
Kleibergen-Paap F			15.929			0.495

Notes: * p<0.1, ** p<0.05, *** p<0.01. Dependent variable is aggregated to the 1 mile/monthly level. Segments with average IRI values < 20 or > 300 are dropped from the sample. Segments with average speeds greater than 90 MPH are also dropped. Panel A examines crash rate per 100,000 total VMT. Panel B examines speed. Column 1 shows reduced form (RF) results (elevation x aquifer on the traffic outcome), Column 2 shows first stage results (elevation x aquifer on IRI), and Column 3 shows second stage results (\widehat{IRI} on the traffic outcome) for asphalt road segments. Columns 4 - 6 show the same for concrete. Separate month, year, route, and county fixed effects are used in every panel. Controls for average precipitation and temperature are also included. Heteroskedastic-robust standard errors are used in every panel and column.

Table A2: IV Estimates of Causal Effect of IRI on Traffic Outcomes at Lower Elevations by Pavement Type


	(1) Asphalt RF	(2) Asphalt 1st Stage	(3) Asphalt 2nd Stage	(4) Concrete RF	(5) Concrete 1st Stage	(6) Concrete 2nd Stage
Panel A: Crash Rate						
Elevation x Aquifer	-0.00008*** (0.00001)	-0.00658*** (0.00178)		-0.00007*** (0.00001)	-0.00284 (0.00161)	
\widehat{IRI}			0.01196*** (0.00386)			0.02616 (0.01611)
Observations	6,733	6,733	6,733	7,820	7,820	7,820
Kleibergen-Paap F			14.958			2.602
Panel B: Speed						
Elevation x Aquifer	0.00406*** (0.00027)	-0.00658*** (0.00178)		0.00285*** (0.00026)	-0.00284 (0.00161)	
\widehat{IRI}			-0.61683*** (0.16260)			-1.00290* (0.60555)
Observations	6,733	6,733	6,733	7,820	7,820	7,820
Kleibergen-Paap F			14.958			2.602

Notes: * p<0.1, ** p<0.05, *** p<0.01. Dependent variable is aggregated to the 1 mile/monthly level. Segments with average IRI values < 20 or > 300 are dropped from the sample. Segments with average speeds greater than 90 MPH are also dropped. Panel A examines crash rate per 100,000 total VMT. Panel B examines speed. Column 1 shows reduced form (RF) results (elevation on the traffic outcome), Column 2 shows first stage results (elevation on IRI), and Column 3 shows second stage results (\widehat{IRI} on the traffic outcome) for asphalt road segments. Columns 4 - 6 show the same for concrete. Separate month, year, route, and county fixed effects are used in every panel. Controls for average precipitation and temperature are also included. Segments with an average elevation above 2000 ft. are dropped. Heteroskedastic-robust standard errors are used in every panel and column.

A1.2 Supplemental Information: Scans and Road Quality

Figure A1 shows the average number of times one mile highway segments were scanned over the sample period.

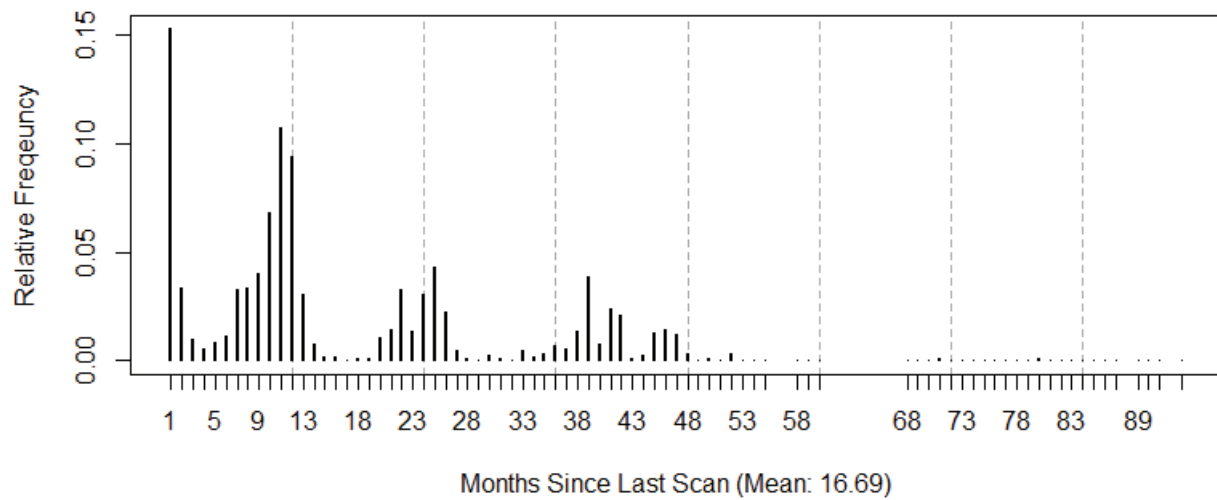
Figure A1: Distribution of Number of Scans for Each .1 Mile Segment



Note: This is a distribution of the number of unique months a segment was scanned. In reality, a segment could be scanned a handful of times within a month, but we average across those scans and only count it as 1. There should be very little, if any, variation within months.

Figure A2 shows the distribution of the observed time between APCS scans for each pavement segment in the sample. The distribution shows three distinct peaks at 1-2 months between scans, 9-12 months between scans, and 22-26 months between scans.

Figure A2: Distribution of Number of Months between IRI Scans



Note: Distribution of the number of months elapsing between APCS scans for individual segments. APCS scans some segments several times over a short period. It is possible that a segment is scanned two times in a week, but those two days are part of different months. Three distinct peaks can be seen at 1-2 months, 9-12 months, and 22-26 month intervals.

A1.3 Supplemental Information: Pavement Types

Table A3 shows distribution of pavement types in the sample.

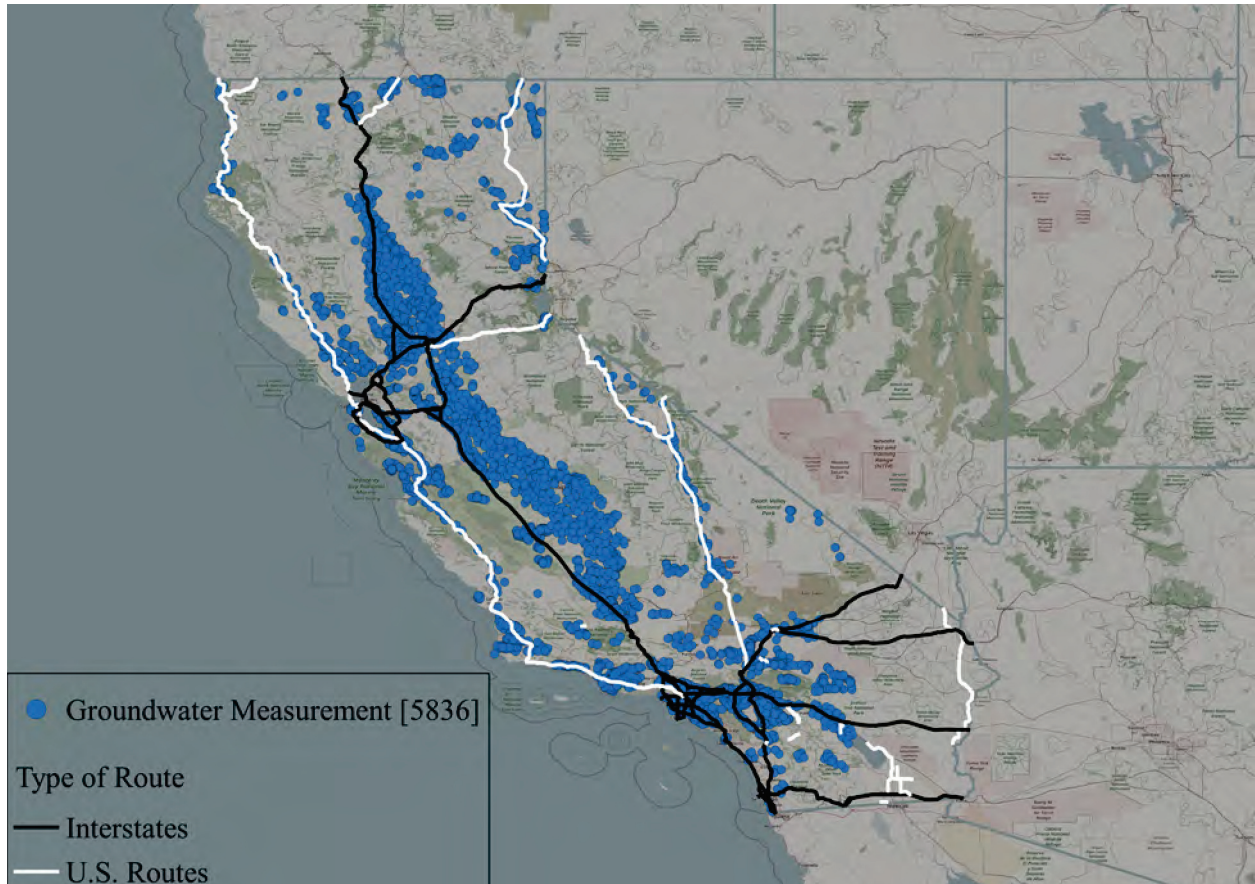
Table A3: Frequency of Route and Pavement Types

	Asphalt	Concrete	Total
Interstate	5210 (32.6%)	7631 (47.7%)	12841 (80.2%)
U.S. Route	2371 (14.8%)	790 (4.9%)	3161 (19.8%)
Total	7581 (47.4%)	8421 (52.6%)	16002 (100.0%)

Note: Frequency of route-pavement type combinations for 1 mile segments shown above.

A1.4 Supplemental Information: Location of Groundwater Measurement Wells in California

Figure A3: California Groundwater Measurement Wells



Note: Map depicts the 5836 groundwater measurement wells in California.

A1.5 Additional Sample Information: T-tests of Differences in Means

Table A4 shows differences in means of all variables across pavement types.

Table A4: Differences in Means

	(1) Mean	(2) Mean	(3) Difference
Panel A	Asphalt	Concrete	
IRI	80.383	106.617	-26.235***
Crash Rate	0.201	0.304	-0.103***
VMT	2.79e+06	3.08e+06	-2.91e+05***
VHT	47732.368	54315.487	-6583.119***
Speed (mph)	60.001	58.896	1.105***
Elevation (ft)	655.753	667.428	-11.675
Over Aquifer	0.770	0.821	-0.051***
Panel B	IRI < 200	IRI ≥ 200	
Crash Rate	0.254	0.309	-0.055***
VMT	2.93e+06	3.60e+06	-6.69e+05***
VHT	50874.881	72243.989	-2.14e+04***
Speed (mph)	59.511	53.493	6.018***
Elevation (ft)	662.743	606.519	56.224
Over Aquifer	0.795	0.863	-0.068**

Note: T-tests for difference in means is displayed above for different groups of the 1 mile - monthly data. Panel A compares asphalt roads to concrete roads. Panel B compares segments with an observed IRI of less than 200 to those with an IRI of more than 200. IRI = international roughness index; VMT = vehicle miles traveled and VHT = vehicle hours traveled. Speed = VMT/VHT.

A1.6 APCS-PeMS Matching Procedure

We matched APCS road segment data to PeMS road segment data using the following method:

- APCS and PeMS data are collected by 0.1 mile road segments and month.
 - Incidents are simply matched to the lowest 0.1 mile segment in the Postmile file. Consider two incidents recorded at post mile (PM) 1.99 and 2.01 in the PeMS data. These two incidents are matched to PM 1.9 and 2.0 respectively in the Postmile file.
 - APCS data are more nuanced since they are point estimates for an entire road segment and while most DMS are 0.1 mile in length, some are longer or shorter. DCS do not always begin exactly at 0.1 PM Postmile file values. Rather, there is a distribution of DCS start positions relative to Postmile file values over the sample. This requires a different spatial matching approach for APCS segment locations to Postmile file values. There are three possible ways an APCS DCS can overlap a 0.1 mile Postmile file segment.
 1. The DCS lies completely within the 0.1 mile segment. For example, the DCS runs from Postmile file values .21 to .29.
 2. The DCS only partially overlaps the 0.1 mile segment. For example, the DCS runs from .17 to .27.
 3. The DCS completely overlaps the .1 mile Postmile file segment. For example, the DCS runs from PM .17 to .37.

A weighted average is constructed for each DCS to estimate IRI for each .1 mile segment in the Postmile file. The weights reflect the fraction of distance in each DMS segment contained in the Postmile file 0.1 mile segment.

- Once the raw data is matched to .1 mile segments, data is then aggregated to 1 mile segments.
- Some sections of highway have realigned postmiles. This is due to road alterations that might change the length of a section of road. We drop road realignments since they might not be consistent over time.