Infrastructure Inequality: Who Pays the Cost of Road Roughness?*

Lindsey Currier[†] Edward L. Glaeser[‡] Gabriel E. Kreindler[§]

November 26, 2024

Abstract

Which Americans experience the worst infrastructure? What are the costs of living with that infrastructure? We measure road roughness throughout America using vertical acceleration data from Uber rides across millions of American roads. Our measure correlates strongly and positively with other measures of road roughness where they are available, negatively with driver speed. We find that road repair events decrease roughness and increase speeds. We measure drivers' willingness-to-pay to avoid roughness by measuring how speeds change with salient changes in road roughness, such as those associated with town borders and road repaving events in Chicago. These estimates suggest that one standard deviation of road roughness in the US generates losses to drivers of 33 cents per driver-mile. Roads are worse near coasts, and in poorer towns and in poorer neighborhoods, even within towns. We find that a household that drives 3,000 miles annually on predominantly local roads will suffer $450 per year more in driving pain if they live in a predominantly Black neighborhood than in a predominantly White neighborhood. The relationship between road roughness and both race and income is substantially stronger in less populous and rich places. Road roughness has little ability to explain subsequent road resurfacing in eleven cities, which suggests American rides could be much smoother if the bumpiest roads were fixed first.

*We thank Uber for sharing the data used in this project, especially Elizabeth Mishkin, Jonathan V. Hall, and Mariya Shappo. We also thank Jack Botein, Eliza Glaeser, Rushil Mallarapu, Rohan Nambiar, Alex Min, Taryn O'Connor, Michael Pak, Sam Patterson, Leo Saenger, Rafael Tiara, and Vivian Zhang for excellent research assistance. We thank Clifford Winston for discussing this paper, and to seminar participants for comments. We thank MassDOT for facilitating access to data on vehicle inspection failures in Massachusetts. We thank the Star-Friedman Challenge for Promising Scientific Research for financial support. Under the agreement with the authors, Uber has the right to review the paper for confidential information but not to dispute or influence the findings or conclusions of the paper. Currier worked previously at Uber. JEL codes: R41, R42, L92, O18.

[†]Harvard University. Email: lcurrier@g.harvard.edu

[‡]Harvard University. Email: eglaeser@harvard.edu

[§]Harvard University. Email: gkreindler@fas.harvard.edu

1 Introduction

How bad are America's roads? While the American Society of Civil Engineers (ASCE) lowered their grade given to America's highways from C+ to D- over the period from 1988 to 2009 (ASCE, 2017),[1] Duranton et al. (2020) use International Roughness Index (IRI) data provided by the Department of Transportation to document that our highways have become far smoother since the 1980s. Yet neither that data, nor most of the existing empirical work on US roads (e.g. Small and Winston, 1988) covers the local roads that make up much of the driving experience and seem likely to contain far more of America's potholes.[2]

In this paper, we measure the roughness of America's local roads and highways using vertical accelerometer smartphone data provided by Uber.[3] This data is similar to the vertical acceleration measures of roughness used by the Department of Transportation, but the D.O.T. drivers travel on off-peak hours and maintain a fixed speed. Constant speed improves accuracy because vertical acceleration is a function of both road roughness and speed. The richness of the Uber data means that we have enough variation in speed over the same road segment that we can estimate the entire roughness-speed relationship at the road segment level. We can compute predicted roughness at any given speed, which is our primary measure of road quality for over five million road segments in America, covering half a million kilometers and over one thousand towns in urban areas.

Section 3 presents attempts to validate this data. The correlation between our measure and the D.O.T. measure of road roughness is approximately 0.7 for both highways and arterial roads. Our measures show the same geographic patterns found in the D.O.T. data. Road roughness is higher in populous MSAs, in coastal areas, and in places where winters are cold. Local road segments in Chicago that have an "at grade" railway crossing are 1.3 standard deviations rougher than similar nearby control road segments. Also in Chicago, we document that our roughness measure shows an average improvement of 0.2 log points after a road is resurfaced, based on 611 repaving events in 2021.

Section 4 of this paper presents a model in which drivers value their time and dislike bumpy vertical motion. When road roughness is anticipated, drivers will slow down to reduce their vertical motion. The cost of a large change in roughness has a direct effect because of a bumpier ride, and an indirect effect due to slower speed. We do not observe the direct utility cost in the data, but the model enables us to estimate that direct cost by measuring the speed response to known bumpiness. By slowing down, drivers reduce bumpiness, revealing their willingness to pay in lost time for a smoother ride. To turn a time cost into a dollar cost, we assume that the value of travel time is

[1] That Society also estimates that America must spend over one trillion dollars over the next decade to rescue our "failing transportation infrastructure" (ASCE, 2020).

[2] The Interstate Highway System accounts for 2.5% of U.S. lane miles. In contrast, the DOT estimates 77% of the U.S. roadway is owned by local governments.

[3] Winston and Karpilow (2020) envision the even larger flow of vertical acceleration data that could be generated by autonomous vehicles.

15 dollars per hour. Our key empirical goal is to estimate this impact of salient road roughness on driving speed.

We use median road segment speed from the Uber data and focus on local roads, where roughness is more extreme and there are salient jumps in road roughness. This work builds on the analysis by Bock et al. (2021), who use altitude and proximity to aquifers as instrumental variables and show that roughness leads to lower driver speeds on highways in California. We show results using four sources of variation in road roughness: average differences across towns, differences across towns within a narrow bandwidth of the border, drops in roughness over time due to resurfacing events in Chicago, and roughness spikes induced by at-grade railway crossings in Chicago. In our main exercise, we focus on large roughness changes around 1,285 town borders.

We estimate an elasticity of speed with respect to roughness of around -0.3. We obtain a lower elasticity of -0.14 when we analyze speed changes after repaving events in Chicago, which may indicate less awareness of variation in road roughness due to resurfacing.[4] We estimate our driver model of driver preferences for speed and smoothness with the observed speed elasticity, average speeds, and the elasticity of driver speed with respect to the speed limit, which we use as a proxy for other costs speed.

The estimated model implies that if a road segment becomes one standard deviation rougher than the median roughness level, the total costs of driving on a road segment goes up by 33 cents per mile, with a 95% confidence interval of $[28, 40]$ cents per mile. This cost is directly proportional to the assumed value of travel time and also depends on the estimated elasticities of speed with respect to road roughness and the speed limit. We use our estimated costs of roughness to scale our roughness measure for local roads in the rest of the paper.

We supplement this exercise with more direct cost measures. We correlate our town-level roughness measures for Massachusetts with town-level vehicle inspection results over the course of a year. Towns with rougher roads have a higher index of inspection failures. However, since we also find higher failure rates for damage to windshield wipers or mirrors, this could merely reflect correlation between car quality and road quality. We then turn to safety costs, and find that towns with rougher roads have fewer road crash fatalities, even after controlling for population, income, and measures of employment and commute trips to work. This suggests that road roughness may generate social benefits that are beyond the scope of our paper. The negative relationship between roughness and observed road fatalities could be explained by lower speeds or a reduction in the amount of driving. If we assume that the number of miles-traveled is independent of roughness conditional on our other controls, then the mortality-roughness and the speed-roughness relationships together yield an implied elasticity of fatal crashes with respect to speed of 1, which is lower than many of the speed-mortality estimates in the literature.[5]

[4] Alternatively, drivers who are particularly averse to bumpiness may be able to avoid resurfaced roads more readily than town crossings.

[5] Bock et al. (2021) find that roughness increases collisions on highways in California.

Section 5 of the paper then documents basic patterns of road roughness across towns and neighborhoods within the 100 largest metropolitan areas. Cities and towns with more employment levels that are closer to the center of the metropolitan area typically have rougher roads. Places with lower median incomes have worse roads, and areas with large shares of minorities have significantly rougher roads. As the share of the Black residents in a tract increases from zero to 100% within an MSA, the average cost of roughness on local roads increases by 15 cents per road mile, or 450 dollars per year for a household that drives 3,000 miles per year on local roads.

We next look at the relationship between neighborhood characteristics and road roughness within cities and towns, where the roads are the responsibility of the same authority. We find that within town, the relationship between tract income and road roughness is three times larger in less populous towns than it is in bigger places. The relationship between race and road roughness is ten times larger in towns that are richer than the median in our sample, compared to within poorer towns.

These results contribute to efforts to document and understand spatial disparities by income and minority status (Brueckner et al., 1999; Cutler et al., 1999; Gobillon et al., 2007), especially with respect to commuting and road quality (Fu et al., 2023; Government Accountability Office, 2022).

In Section 6, we look at how cities and towns decide which roads to fix. First, we analyze road repaving decisions in 11 large cities, where we have detailed data on road repaving events that occur after our road roughness measure. Our baseline measure of local road roughness weakly predicts where road segments get resurfaced in New York City, Denver, and Philadelphia, but in the other cities there is no relationship between roughness and repaving. We compare observed repaving with the implications of an algorithm which allocates road repaving on the basis road quality and traffic levels (as measured by Uber use). The observed repaving patterns in New York City are compatible with 10-16% of road repairs being based on these criteria and the remainder being based on orthogonal considerations. We cannot rule out that these repaving patterns are the solution to an optimal repaving algorithm with additional considerations.

We ran a survey of repaving practices in 25 towns in Massachusetts and 96 towns nationwide to learn how they assess road quality and how they target road repairs. Typically, towns have a road survey measure of road roughness, and they take several factors into account when deciding which roads to resurface, including road conditions, traffic intensity, upcoming utility work, and other factors. In practice, many towns and cities report repaving a small fraction of roads that need repaving, and spending significantly more on resurfacing compared to maintenance work. These reported repaving strategies are compatible with our data, and they may be inefficient relative to targeting the worst roads first, taking into account traffic patterns.

Section 7 concludes by discussing how the rise of big data has made it increasingly possible to measure and improve aspects of urban life (Glaeser et al., 2018). The availability of road roughness data from sources such as Uber may reduce the need to pay outside providers for road assessments. While there are other interpretations of our repaving data, one possible interpretation is that it could be possible to radically improve the quality of America's roads, especially in poor and minority neighborhoods, simply by resurfacing the rougher roads first. Better targeting may be particularly valuable given the large and growing cost of transportation infrastructure (Brooks and Liscow, 2023).

2 Data

2.1 Vertical Acceleration Smartphone Data

Our main data source for measuring road roughness is smartphone location and acceleration data collected by Uber from Uber driver smartphones. GPS receivers and accelerometer sensors are ubiquitous and precise in modern smartphones (Grouios et al., 2022). Uber records and stores compressed accelerometer data from the duration of Uber trips. The sensors collect acceleration data measured in meters per second squared, with a frequency of approximately 5Hz (5 observations per second).

We are mainly interested in *vertical* acceleration as a raw measure of experienced road "bumpiness."[6] Smartphones record acceleration along three axes aligned with the physical device.[7] In order to identify sharp horizontal acceleration events such as harsh braking, Uber created a proprietary algorithm to re-align these axes to the vehicle. This algorithm is aided significantly by the fact that drivers typically keep their phones fixed on a dashboard mount. In the aligned data, the X axis points toward the right of the car, the Y axis points forward in the direction of travel, and the Z axis is vertical pointing upward. We use measures of acceleration along the re-aligned Z axis (or vertical acceleration) as our primitive measure to construct road roughness.

The vertical acceleration data contains rich information on road irregularities. Figure 1 shows vertical acceleration during a single Uber trip that uses local roads in Chicago for the start and end of the trip and uses the highway for the middle portion. The middle portion has noticeably lower dispersion in vertical acceleration.

[6] Our approach builds on proof-of-concept work by Aleadelat et al. (2018), who study 35 road segments and show that accelerometer data from smartphones can be a reliable proxy for IRI measures of road roughness. Our contribution is to use the data collected by Uber to perform this type of exercise at scale for the entire country, and to extensively validate this type of measurement. More generally, this method is related to research using distributed sensors to measure transportation-related outcomes with high resolution. Akbar et al. (2023) use Google Maps queries for hypothetical trips to measure speed across millions of road segments across the world and compile them to city and country speed indices. Apte et al. (2017) measures air pollution on streets around Oakland, CA, using air quality sensors mounted inside Google Street View vehicles.

[7] When holding a smartphone, the X axis points rightward, the Y axis upward, and the Z axis protrudes from the screen.

Figure 1: Uber Data Has Signal: Different Types of Roads

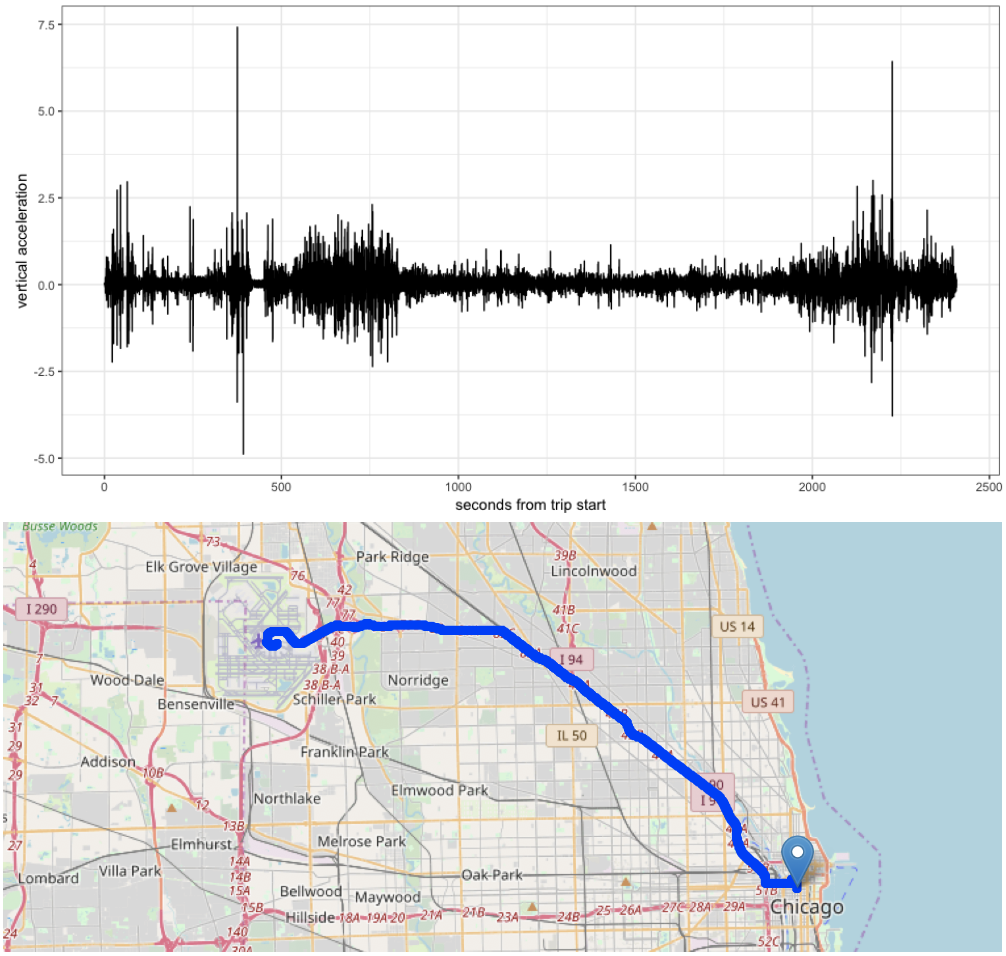

Note: These graphs show raw measures of vertical acceleration during a single Uber trip. The Z-acceleration measure is less variable during the middle portion of the trip, along highways, consistent with smoother pavement.

We begin with anonymous data for all Uber trips across all Uber US markets from August, 2021. After cleaning the raw data with Uber's alignment algorithm, we produce a dataset organized by a driver identifier, timestamp, and vertical acceleration in meters per second squared. We also have data in Cook county, Illinois for March - August 2021, and April 2018.

The GPS location data contains a driver identifier, timestamp and geographic coordinates. We match GPS coordinates, which in general contain noise, to a map of road segments derived from Open Street Map (OSM), an open source mapping service. To do this, we apply a proprietary algorithm created by Uber based on Newson and Krumm (2009). Since OSM continually updates, we match all data to a "stable" basemap from April 15, 2021. Our road segments correspond roughly to a block and are a partition of the larger OSM segments. The median segment in our data for a local road is 49 meters long. We then match acceleration measurements to a given road segment based on timestamps.

The full data set at the level of acceleration observations matched to road segments is so large as to pose significant computational challenges. To reduce the data to a size we can feasibly analyze (2 Terabytes), we aggregate to the level of an Uber trip by road segment by taking the standard deviation of acceleration.

2.2 Estimating Road Roughness at the Road Segment Level

We use the standard deviation of the vertical acceleration to measure road roughness. For each trip i and road segment r, we compute $Z^{\sigma}_{ri} = SD(Z_{rit})$, where Z_{rit} is the vertical (aligned) acceleration at time t. This measures uses our multiple observations during the same Uber trip within a single block-length road segment. We use the terms "vertical acceleration," "bumpiness," and "experienced roughness" interchangeably to refer to this measure.

Vertical acceleration reflects the bumpiness of the road, but it also reflects characteristics of the car and the speed of the car.

The richness of our data enables us to control flexibly for driver identity and speed, relying on significant variation in the driving speed on the same road segment across journeys. We use the following specification:

$$Z^{\sigma}_{ri} = \mu_r + \gamma_r \cdot \text{speed}_{ri} + \psi_{d(i)} + \phi_{t(i)} + \theta_{h(i)} + \epsilon_{ri} \tag{2.1}$$

where Z^{σ}_{ri} is the standard deviation of vertical acceleration of trip i on road segment r, μ_r are road segment fixed effects, speed_{ri} is the vehicle's median speed on the road segment on trip i and γ_r is a road segment specific slope of experienced roughness on speed, and $\psi_{d(i)}$, $\phi_{t(i)}$, and $\theta_{h(i)}$ are (mean-zero) fixed effects for driver $d(i)$, date $t(i)$, and hour of the day $h(i)$. Driver fixed effects both control for particular driving styles and for attributes of the driver's vehicle. Equation (2.1) allows us to estimate the "speed-roughness" relationship separately for each road segment in the data. We will use the road segment fixed effects μ_r and the speed slopes γ_r to compare predicted roughness between road segments at a fixed speed.

This specification constrains our estimation by imposing a linear relationship between speed and vertical acceleration that varies by road segment, but not by driver. In principle, the linearity and constant driver effect assumptions could be relaxed, given sufficient data for estimation. However, there is little evidence that suggests important non-linearities in the relationship between speed and acceleration (Figure A.1), or that there are driver-specific segment slopes (Appendix Figure A.2).

We estimate equation (2.1) separately for each Uber city.[8] For computational tractability we

[8]There are 238 Uber cities in the US in our data. An "Uber city" is usually centered around a large city and includes surrounding towns.

use a maximum of 20 randomly-chosen observations per day per road segment. This procedure implies a maximum of 620 observations per road segment in our one month of data, and we also restrict the sample to all road segments with at least 50 trip observations. We obtain a sample with 2.2 billion trip-by-segment observations in the entire US, and over 78 million observations in New York City alone. We estimate equation (2.1) using a linear least squares estimator accounting for the high-dimensional fixed effects and high-dimensional linear interactions, using the method of alternating projections (Correia, 2014).[9] Including driver fixed effects in equation (2.1) means that we are essentially comparing road segments holding the quality of the vehicle fixed.

We obtain estimates for 8,199,741 directed road segments in the US, covering 5,705,777 undirected road segments. Of all undirected segments, 9.3% are highways, 67.8% are arterial roads, and 22.4% are local roads.[10] We discuss the implications in terms of coverage and representativeness in the next subsection and in Figure A.4 and Table A.1.

Our measure of road roughness is *predicted* roughness at a fixed speed. To keep our predictions in sample, we use average speeds by road type: 20 mph for local roads, 32 mph for arterial roads, and 48 mph for highways. For example, the measure for a local road segment r is $\widehat{Z}_r^{20} = \widehat{\mu}_r + \widehat{\gamma}_r \cdot 20$ mph. These measures are purged of the *actual* speed used by drivers and only rely on the road segment fixed effects and speed slopes estimated in specification (2.1).

The map in Figure 2 plots predicted roughness for local and arterial road segments in Cook County, Illinois. Each road segment is drawn with a color that corresponds to the decile of the distribution of predicted roughness ($\widehat{Z}_r^{20}$ for local roads, and $\widehat{Z}_r^{32}$ for arterial roads). Warmer colors indicating rougher roads. Roads inside the city of Chicago are significantly rougher than those outside, and there is a pronounced discontinuity at the city border, a pattern that we will return to later in the paper in Section 4.2.1. There is heterogeneity in road roughness within Chicago, and we will explore patterns by income and race in Section 5.

Our measure of predicted roughness will reflect any factor that affects the speed-roughness relationship, including speed bumps and brick roads.[11] While it seems appropriate to include these forms of "chosen bumpiness" when we estimate the link between roughness and travel time, we anticipate that chosen bumpiness will bias relationships between roughness and either income or race, since we anticipate that chosen roughness will be more common in wealthier areas.

Much of the analysis in the rest of the paper will focus on differences in road roughness between nearby areas, such as towns or tracts. For a subset of analysis, we will be comparing road roughness across larger areas, such as metropolitan areas across the entire US. In those cases, we need to make

[9] Driver, date and hour fixed effects are normalized to be mean-zero. We do not include a constant so the road segment fixed effects μ_r are not constrained to be mean-zero. The segment-specific speed slopes do not require a normalization. Overall, in each Uber city, average predicted roughness from equation (2.1) using actual median speeds is approximately equal to the mean of the outcome variable on the estimation sample.

[10] Among directed road segments, we have 1,833,151 local, 5,776,089 arterial, and 551,703 highway segments. For undirected road segments, we have 1,279,204 local, 3,870,754 arterial, and 528,296 highway segments.

[11] For example, see Appendix Figure A.7.

an additional assumption that the mean of driver fixed effects in these different regions do not vary systematically with our explanatory variables of interest. For example, if we are interested in the correlation between distance to coast and road roughness, our estimates would be biased if Uber vehicles were systematically bumpier in MSAs near the coast.

2.3 Robustness Checks.

We perform several robustness checks on our estimation strategy. Later, we perform a series of external validation tests (section 3).

How persistent are our measures of road roughness? In Chicago, where we have multiple months of data, we estimate predicted roughness using data from March 2021 and April 2021, using completely distinct data sets. Table A.3 reports that when we regress the April measure on the March measure we see an R^2 of 0.82 for highway segments, 0.9 for arterial roads, and 0.84 for local roads. The R^2s using August instead of April data are 2-4 percentage points lower, suggesting a role (though a limited one) for changing roads over time. The binned scatter plots in Figure A.3 show a linear relationship that is very close to the identity line. Our measure of road roughness appears to capture a persistent characteristic of the road segment, not estimation noise.

Does the fact that we estimate road roughness controlling for driver speed bias our results when we regress driver speed on road roughness? We perform a split-sample exercise to check these issues. In Chicago, we use 75% of our data (the "estimation" sample) to estimate equation (2.1), and the remaining 25% of the data (the "test" sample) to estimate median driving speed. We then further restrict the test sample only to off-peak hours, when congestion is significantly less likely (outside 5am to 9am and 4pm to 8pm). We then regress median driving speed on predicted roughness, using all three definitions of speed, separately for each type of road.

We find very similar slopes regardless of which sample we use to compute speed (Table A.4). For example, on local roads, the coefficient on predicted roughness is -20.5 when we compute speed on the estimation sample, -21.2 when we use the test sample, and -21.1 when we use the off-peak test sample. While some of these estimates are distinct, the magnitude of the difference is small, at only 3% of the size of the coefficient. While we do not think that using speeds computed using the same sample significantly biases our results, this may constitute an issue for inference. For this reason, in most of the analysis using speeds as an outcome, we use a bootstrap procedure to account for the fact that the speed data is used to estimate the roughness measures.

Figure 2: Predicted Roughness in Cook County, IL

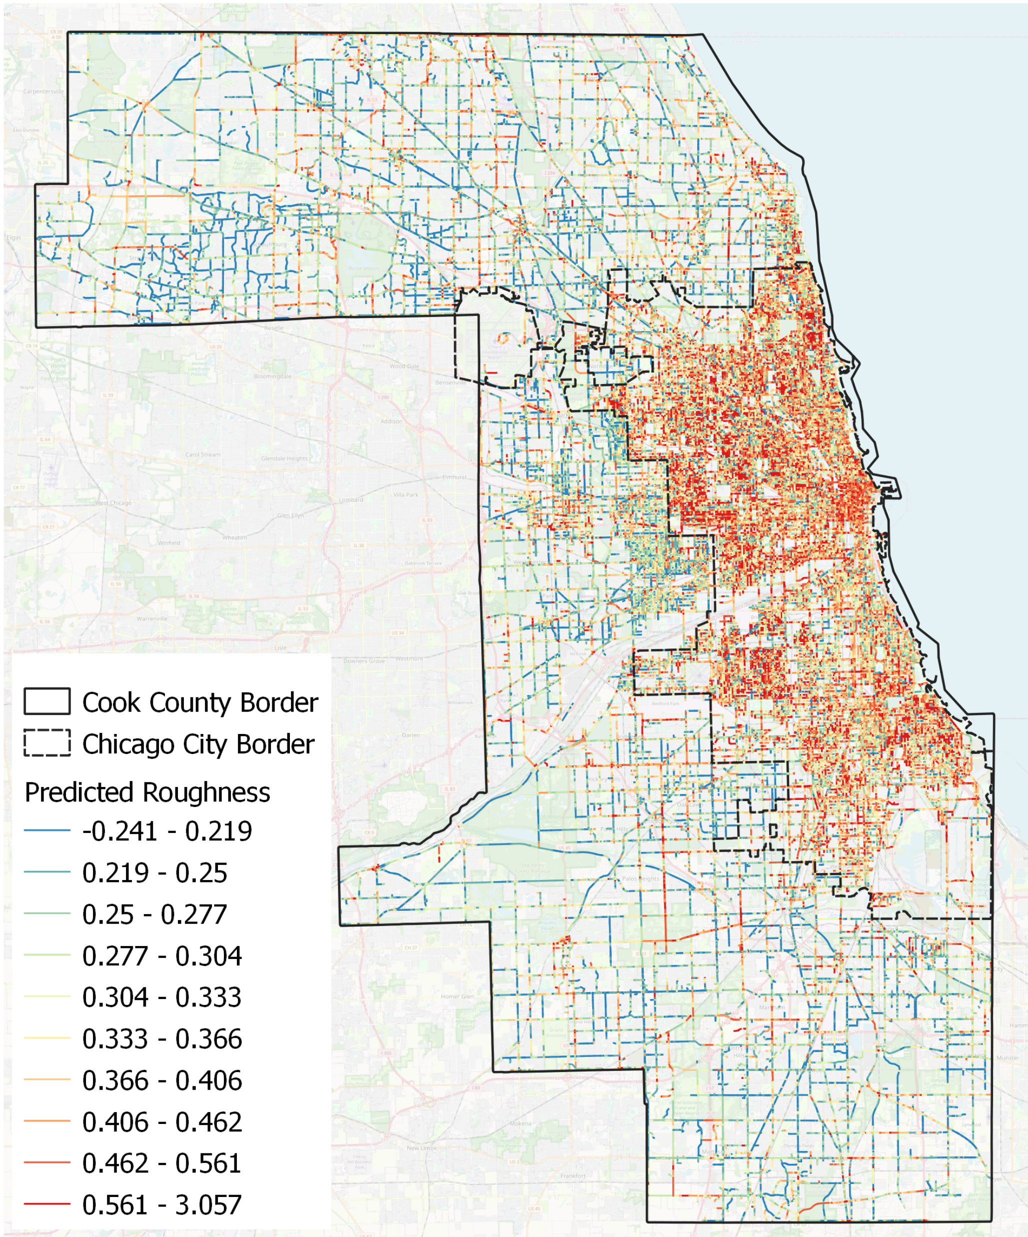

Note: This map plots predicted road roughness for all local and arterial road segments in Cook County. Colors correspond to deciles of the predicted roughness distribution. The city boundary of Chicago is shown in the dashed black line. Figure A.6 shows an analogous map for highway roughness.

Our Uber data only covers a subset of all road segments in the US, and the selection of which roads get into our data may impact our results. Using the Uber data, we obtain fixed effect estimates for 5.7 million undirected road segments, which represents 14.8% of all the 38.1 million undirected road segments in our entire base map. Table A.1 shows that when we weight by town population, segment coverage is 12-14% for local roads, and 63-75% for arterials and highways. Figure A.4 shows that our measure's coverage increases rapidly in the population level.

To test for selection into the sample, we check whether more Uber coverage correlates with rougher roads in our sample. We do not find this to effect to be quantitatively important. Figure A.5 and Table A.2 focus on our full sample of over 1.8 million directed road segments on local roads, and correlate the log number of Uber trips we observe on the road segment, and the road segment's predicted roughness. We observe a slight gradient over the full sample, with one log point of trips per segment associated with 0.073 SD of predicted roughness. However, within Uber cities this correlation falls to 0.016 SD per log point of trips per segment. Importantly, the correlation vanishes completely when we look at road segments with below median number of trips. This lack of selection among road segments with low coverage suggests that road segment selection into the Uber sample is not biasing our results.

2.4 Other Data

Throughout the paper we use several other data sets, including alternate measures of road roughness for validation, road resurfacing data for validation and to characterize targeting patterns, vehicle inspections in Massachusetts, fatalities from vehicle crashes, and census data. We introduce each data set in the section where we use it. Appendix C includes detailed information on all the data sets used in the paper.

3 Validation

In this section, we provide four tests of whether our measure captures road roughness. First, we correlate our measure with the US Department of Transportation (DOT) International Roughness Index measures of highway road roughness in Cook County, IL. Second, we compare the geographic patterns for the Uber measure and for IRI data nationwide. Third, in the city of Chicago we test whether our measure of road roughness is significantly higher on road segments that intersect railroad tracks. Fourth, again in Chicago, we look at whether our measure of road roughness improves after the city reports repaving the road. These tests collectively make us comfortable interpreting our road roughness measures as a solid, if imperfect, measure of the state of city streets.[12]

[12]We report two additional validation exercises in the appendix. In Appendix D.1 we compare our measure to the IRI nationwide, and in Appendix D.2 we compare our measure with New York City's Pavement Condition Index (PCI), a visually-based measure of road roughness, and we correlate both of our measures with driver speeds.

3.1 Correlation with Arterial and Highway IRI in Cook County

We first compare our estimates to the International Roughness Index (IRI), a measure of roughness collected by the DOT's Federal Highway Administration (FHWA). The federal government distributes around $50 billion each year to state transportation departments, and it requires states to biennially measure the IRI, a widely used measure of road quality (Sayers et al., 1986), for the subset of roads that constitute the National Highway System.[13] This data is generally produced by state employees who drive vans with laser sensors and onboard computers at fixed speeds.

We report results from Cook County, Illinois. We use 2018 IRI data directly published by the Illinois Department of Transportation, which allows us to also look at arterial roads in addition to highways. We match road segments from the two sources based on the distance separating them, a fuzzy name match of the road names, and the angle between them. (See Appendix C.5 for further details.) For consistency, we use a sample of Uber data in Cook County from March, 2018.

Figure 3 plots the data. We are comforted by the relatively high level of segment level correlation, despite the imperfect nature of our fuzzy match and the different time periods. Table A.19, Panel (b) shows that the correlation coefficients between IRI roughness and our measures are close to 0.7 for both highways and arterial roads.

Figure 3: Uber Roughness and IRI Segments

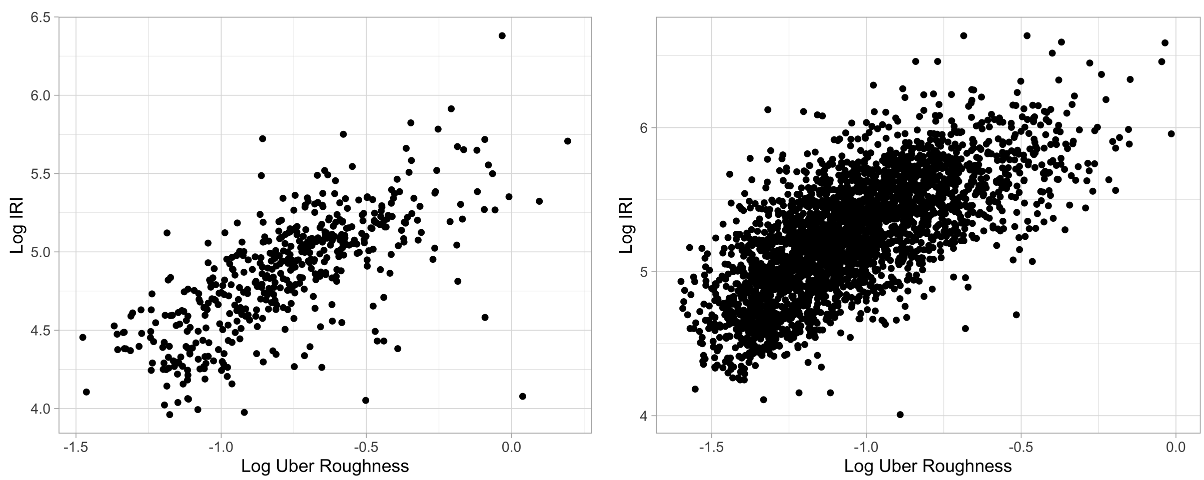

(a) Highway Roads (b) Arterial Roads

Note: These graphs show the correlation of IRI and Uber roughness for road segments in Cook County, IL in 2018, based on the matching procedure described in Appendix C.5. Highways include OSM classifications "motorway" and "trunk," and arterials include "primary," "secondary," and "tertiary." We drop OSM segments for which there are no overlapping IDOT basemap segments of similar length.

[13]The National Highway System is a network of highways across the U.S. defined by historical and strategic demands. It covers around 225,000 miles of the 4 million mile road network.

3.2 Geographic Patterns

Our next validation exercise is to show that our data displays similar regional patterns as the D.O.T. roughness data for highways and arterial roads. Our method to estimate road roughness controls for driver effects within Uber regions, where we have overlap between segments driven by the same driver, but not across Uber regions. Thus, one limitation of this exercise is that differences across regions in estimated roughness using our measure may, in principle, reflect differences in vehicle quality or driving style. Moreover, there is far more variation in road quality within metropolitan areas than across metropolitan areas, especially for local roads (Table A.5).

To compare the regional patterns between the two data sources, we take average road roughness for each MSA for both Uber and D.O.T. data, weighting by road segment length, We then regress this index on climate variables in Table 1. Appendix Figures A.8, A.9 and A.10 plot the average road roughness by MSA for local roads, arterial roads, and highways, respectively.

Table 1 looks at patterns across the 100 largest metropolitan areas in our sample. Our only non-climactic control is metropolitan area population: we focus on race and income later. We run separate regressions for local roads, arterials and highways. Certain variables are put in the form of z-scores.

The first regression shows that a one standard deviation increase in January temperatures is associated with a decrease in road roughness of .32, .29 and .23 standard deviations for highways, arterial roads and local roads in the Uber data. In the D.O.T. data, a one standard deviation increase in January temperature is associated with a .41 and .58 standard deviation decrease in road roughness for highways and arterial roads, respectively. All of these coefficients are strongly negative, but cold appears to be far more harmful for arterial roads in the D.O.T. data than in our data. For decades scholars have noted that "highway deterioration, particularly the formation of potholes in winter in the presence of water, is quite often associated with freeze-thaw cycles" (Hershfield, 1979). As groundwater freezes, the ground expands and this causes asphalt to shift.

In the second row of the table, we find a slight positive coefficient on January precipitation which is compatible with the view that water is also a determinant of road roughness. The highway coefficients are almost identical in the two data sources (.068 and .087), but the coefficient in the regression using D.O.T. measures of arterial roughness is substantially higher than in the regression using Uber data (.137 vs. .014).

The third row shows coefficients on July precipitation which are all negative. We have no good causal interpretation for why wet Julys would lead to smoother roads, and we suspect that this is capturing omitted factors that lead some drier western metropolitan areas to have rougher roads. The coefficients on highways and arterials are $-.24$ and $-.36$ respectively in the Uber data. With D.O.T. data, the coefficients are $-.16$ and $-.10$. Again, the two data sets show the same broad patterns but the magnitudes of coefficients do vary significantly.

The fifth variable shows that metropolitan areas that are within 50 miles of a coast have highways that are .62 standard deviations rougher than inland roads in the Uber data and .5 standard deviations rougher than inland roads in the D.O.T. data. For arterial roads, the Uber coefficient is .64 and the D.O.T. coefficient is .87. This effect could capture groundwater and the nature of soil that is closer to the coast. Cities like New Orleans and Houston, which are both close to the coast, suffer from extremely rough roads, at least partially because of the prevalence of clay soils that expand when wet. Coastal cities may also get more truck traffic because they have ports.

Table 1: Road Roughness and Climate

	Road Roughness (z-score)				
	highway Uber	arterial Uber	local road Uber	highway IRI	arterial IRI
	(1)	(2)	(3)	(4)	(5)
January temperature (z-score)	-0.322^{***}	-0.285^{***}	-0.227^{**}	-0.409^{***}	-0.579^{***}
	(0.099)	(0.090)	(0.097)	(0.098)	(0.085)
January precipitation (z-score)	0.068	0.007	0.014	0.087	0.137^{*}
	(0.093)	(0.085)	(0.091)	(0.092)	(0.080)
July precipitation (z-score)	-0.240^{**}	-0.357^{***}	-0.057	-0.156^{*}	-0.103
	(0.094)	(0.086)	(0.093)	(0.094)	(0.081)
Close to coastline	0.623^{***}	0.639^{***}	0.469^{**}	0.502^{**}	0.873^{***}
	(0.219)	(0.200)	(0.215)	(0.219)	(0.188)
Log population	0.358^{***}	0.478^{***}	0.576^{***}	0.377^{***}	0.404^{***}
	(0.111)	(0.102)	(0.110)	(0.111)	(0.096)
Observations	100	100	100	100	100
Adjusted R^2	0.239	0.365	0.265	0.242	0.438

Note: This table shows coefficients from regressing the z-score of average predicted road roughness in an MSA on characteristics of its climate and log population. The sample is the 100 largest MSAs in mainland US with Uber data. (We exclude Alaska, Hawaii, and Puerto Rico.) In columns 1-3, the outcome is computed using Uber data. We compute predicted road roughness at 20mph for local roads, 32mph for arterials, and 48mph for highways, and then compute the z-score over all MSAs. In columns 4-5, we use DOT IRI data and also compute the z-score. We include 5-year averages of January average temperature, January average precipitation and July average precipitation, from NOAA. (Temperature in January and July are highly correlated, whereas precipitation is not, so we do not include July temperature.) The data is provided at the county level, so we average across counties in each CBSA. We control for an indicator for closer than 50 miles to the coastline, and log population from 2019. Robust standard errors in parentheses, *p<0.1; **p<0.05; ***p<0.01.

The final row looks at the coefficient on the logarithm of metropolitan area population. A one log point increase in population is associated with an increase in roughness ranging from .38 to

.48 standard deviations for the arterial roads and highways. The coefficients for the comparable D.O.T. and Uber regressions are remarkably close. For highways, the two coefficients are .36 and .38. For arterials, the two coefficients are .48 and .40. The coefficient in the local roads regression is substantially higher.

Table 1 shows similarities and differences between the two data sets. Every coefficient has the same sign with the two data sets, and in many cases, the magnitudes are also similar.

3.3 Road Roughness and Railroad Crossings in Chicago

We next turn to analyzing experienced road roughness on road segments that include an at-grade railway crossing, focusing on Cook County.

Figure A.11 pools all the data on a road segment that has a railway crossing. Vertical acceleration is significantly more dispersed exactly around the railway crossing (Panel A). Taking the standard deviation of vertical acceleration for each trip and each section of the road segment shows that the entire distribution is significantly higher around the railway crossing (Panel B). Average speed also declines at the railway crossing.

We have assembled a data set of all the local road segments with railway crossings in Cook County, based on Illinois Department of Transportation data, and nearby and comparable control segments. For the control group, we consider all road segments within 100 and 200 meters away from the nearest treated road segment. Our final sample contains 121 local road segments with crossings, and 828 control local road segments.

Predicted road roughness at 20 mph is 63% higher on road segments with a railway crossing relative to nearby control segments (Figure 4). The entire distribution of predicted roughness is shifted to the right, with only about 5% of road segments with railways having predicted roughness below the median value for control segments. Driving speeds are an average 0.97 mph or 0.21 standard deviations lower on road segments with a crossing than on nearby control segments. These results average over the entire road segment, so the differences at the crossing itself are likely to be even starker. Again, these results confirm that the Uber measure is capturing road roughness.

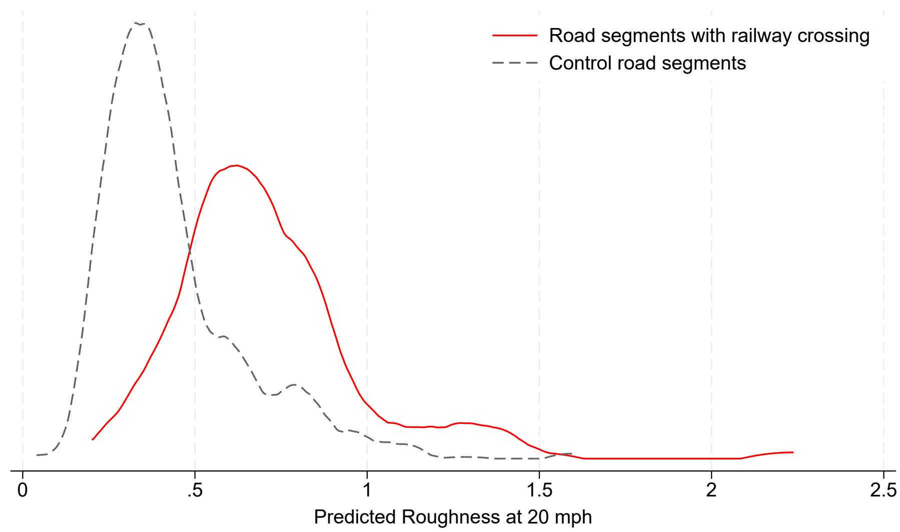

Figure 4: Roughness on Roads with Railway Crossings

Note: This figure shows the distribution of predicted road roughness at 20mph on local roads with at grade railway crossings (red, solid) and control segments between 100 and 200 meters away (gray, dashed).

3.4 Impact of Resurfacing in Chicago

Our final validation exercise looks at the impact of road resurfacing on measured road roughness. We use data on road resurfacing in 2021 in Chicago based on its publicly available street work moratorium data.[14] We match moratorium addresses to our road segments, and we infer the construction end date from the moratorium dates. Appendix C.8 has details.

We analyze all local road resurfacing events in the city between May and mid-July 2021. Our control group consists of the set of all road segments that are between 100 meters and 200 meters away from the closest repaved road segment. Our final analysis sample consists of 611 repair events, covering 1,716 repaved local road segments and 3,218 unique control local road segments. We also use Uber data in Chicago that spans the period April-August 2021.

Figure 5 plots experienced roughness and driving speed over time for a set of road segments with inferred resurfacing date in May 2021, relative to control road segments. We plot the average over calendar time, rather than relative to each road segment's precise inferred repaving date, because the moratorium dates are an imprecise proxy of the exact dates of construction. We want high-frequency data, so we use the standard deviation of the vertical acceleration. This measure

[14] We have confirmed in discussions with the Chicago Department of Transportation that the moratorium data covers both *road reconstruction*, which means work that replaces the asphalt base layer and the surface layer, and *road resurfacing*, which covers only replacing the surface layer, and excludes other maintenance work such as pothole filling, slurry seal, etc.

will tend to attenuate differences in roughness because drivers slow down to compensate for high roughness.

Figure 5: Experienced Roughness and Speed on Repaved Road Segments

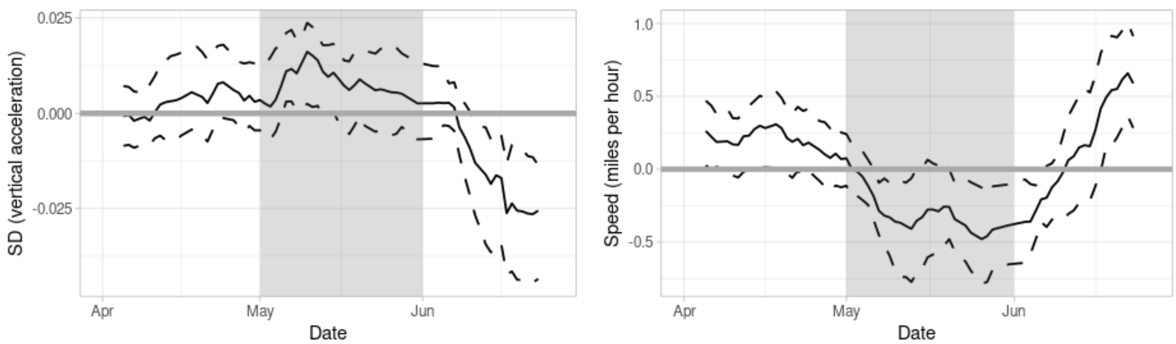

(a) Experienced Roughness (b) Speed

Note: These panels plot seven day moving averages of daily means of the treatment group minus the control group over time. Control road segments are from a "donut" area between 100 and 200 meters from the nearest treated segments. We plot point-wise 95% confidence bands based on bootstrapping the entire procedure at the level of clusters of road repair events.

Experienced roughness first increases in early May then gradually decreases to a level below its baseline level (panel (a)). The initial increase likely captures a period of road construction, when car rides may be more bumpy. Drivers slow down considerably at the onset of road repair construction in early May, but speeds trend upward, eventually reaching a level higher than in the baseline period (panel (b)). We will return to this setting when we estimate the model and analyze this data using long differences in time, using road repaving events interacted with the post period as instruments for road roughness.

4 The Speed Costs of Rough Roads

4.1 A Model of Driver Optimal Road Segment Speed

In this section, we present a model in which drivers reveal their willingness to pay for smoother rides by slowing down. Well-paved roads deliver a direct benefit of less discomfort and damage, which we do not empirically observe, and an indirect benefit of enabling faster trips. We observe speed and roughness in our data, and can combine this with external estimates of the value of time, to estimate the indirect benefit of smoother roads.

In our model, the utility cost of driving over a road segment is additive in time cost, which depends on inverse speed, in speed costs due to bumpiness, and in other speed costs that do not depend on bumpiness.

The marginal cost of increasing road roughness has a simple and intuitive form. Given that drivers are optimizing speed, the envelope theorem states that we are interested in the direct effect of roughness holding speed constant. While we do not empirically observe this direct utility cost of bumpiness, we can use agents' optimizing behavior to back it out. In a simple version of the model where the cost of speed only depends on bumpiness, the semi-elasticity of total user costs with respect to roughness is equal to the observed time cost. In the full model, we also need to subtract the marginal non-bumpiness cost of speed. For larger (non-marginal) changes in road roughness, we can integrate this expression to obtain the change in total driver cost.

The key empirical object is how drivers adjust their speed in response to exogenous and salient changes in road roughness. We estimate this relationship using four sources of variation in road roughness: variation between towns, variation around town borders where roughness changes discontinuously, the effect of road resurfacing in Chicago, and variation induced by at-grade railway crossings in Chicago. We estimate the two versions of our model by matching the estimated elasticity of speed to roughness, and the elasticity of speed to the speed limit, which we estimate using granular data on speed limits that we collect.

We use our estimates to quantify the dollar value that drivers place on road quality, and we use that estimate for the rest of the paper. Our speed elasticity estimates reflect Uber driver behavior, so we also report sensitivity analysis for our headline cost numbers. Readers who are skeptical about the moments we use can use these numbers to adjust our cost estimates accordingly.

4.1.1 Model Setup and Optimal Speed

We consider the optimal speed of a driver on a 1 mile-long road segment r. The utility cost of driving on segment r for driver i is the sum of the cost of driving time, the cost of bumpiness, and other speed costs:

$$cost_{ri}(s;\gamma_r, s_r^{\text{LIM}}) = \underbrace{\frac{v_i}{s}}_{\text{time cost}} + \underbrace{B\left(\gamma_r s\right)}_{\text{cost of bumpiness}} + \underbrace{K\left(\frac{s}{s_r^{\text{LIM}}}\right)}_{\text{other speed costs}}, \tag{4.1}$$

where v_i is the value of (driving) time in dollars per hour, and s is the chosen speed in miles per hour. Bumpiness costs B depend on the experienced roughness term, $\gamma_r s$, which is linear in speed s with a slope γ_r.[15] We let other speed costs depend on the driver's speed relative to the road segment's speed limit s_r^{LIM}. This captures the facts that driving over the speed limit increases the probability of receiving a fine and that speed limits could be a a proxy for other costs of driving fast, such as the risk of accidents. The function $K(\cdot)$ determines the importance of other speed

[15] This relationship has the intuitive property that that experienced roughness is zero at zero speed.

costs scale relative to time costs.[16]

The driver knows v_i, observes the road roughness slope parameter γ_r and speed limit s_r^{LIM}, and chooses the speed that minimizes cost, as shown in Figure 6a. The first-order condition implies that optimal speed is positive, finite, and given by

$$\frac{v_i}{(s_r^*)^2} = \gamma_r B'(\gamma_r s_r^*) + \frac{1}{s_r^{\mathrm{LIM}}} K'\left(\frac{s_r^*}{s_r^{\mathrm{LIM}}}\right). \tag{4.2}$$

Drivers choose to go slower on roads where experienced roughness increases faster with speed (as captured by γ_r), when the speed limit is lower, if they are more sensitive to bumpiness (high B') or to other speed costs (high K'), and when the value of time v_i is lower.[17]

For a small change in road roughness $d\gamma_r$, the envelope theorem implies that the change in the cost to the driver is:

$$\frac{dcost_{ri}^*}{d\gamma_r} = s_{ri}^* B'(\gamma_r s_{ri}^*) = \frac{v_i}{\gamma_r s_r^*} - \frac{s_r^*}{\gamma_r s_r^{\mathrm{LIM}}} K'\left(\frac{s_r^*}{s_r^{\mathrm{LIM}}}\right), \tag{4.3}$$

where the last step follows from substituting B' from the first order condition (4.2).

Intuitively, the marginal cost of roughness depends on B' holding speed constant, as show in Figure 6b. While we do not observe that direct cost empirically, the first order condition reveals it in terms the value of time, the observed speed, and the marginal other cost of speed (which needs to be estimated).

Consider next a larger change in road roughness from γ_0 to γ_1. We can integrate equation (4.3) to obtain the full change in cost:

$$cost_{ri}^*(\gamma_1) - cost_{ri}^*(\gamma_0) = \int_{\gamma_0}^{\gamma_1} \left[\frac{v_i}{\gamma s_{ri}^*(\gamma)} - \frac{s_r^*}{\gamma_r s_r^{\mathrm{LIM}}} K'\left(\frac{s_r^*}{s_r^{\mathrm{LIM}}}\right)\right] d\gamma \tag{4.4}$$

Figure 6c depicts the intuition for this integral where we apply the envelope theorem to small changes in roughness over and over again between γ_0 and γ_1. The integral is the limit when those changes become arbitrarily small.

We will estimate two versions of this model. First, we estimate a simple model where we assume $K' = 0$, that is, the only cost of speed is due to bumpiness. More ambitiously, we estimate the full integral in (4.4) using two key empirical objects. First, we need to know how optimal driver speed $s_{ri}^*(\gamma)$ depends on roughness γ. Second, we need to estimate how driver speed varies when the speed limit varies, in order to learn about the cost term K.

[16] We assume B is increasing and differentiable, and as s increases, $s^2 B'(s)$ is bounded away from zero. (This rules out an infinite optimal speed.) We assume that K is non-decreasing and differentiable.

[17] Variation in driver speed on a given segment may arise because of heterogeneity in v_i which measures how rushed a driver is, from driver-specific or idiosyncratic cost shocks to the B function, or from random shocks to realized speed.

Figure 6: Optimal Speed and the Costs of Road Roughness

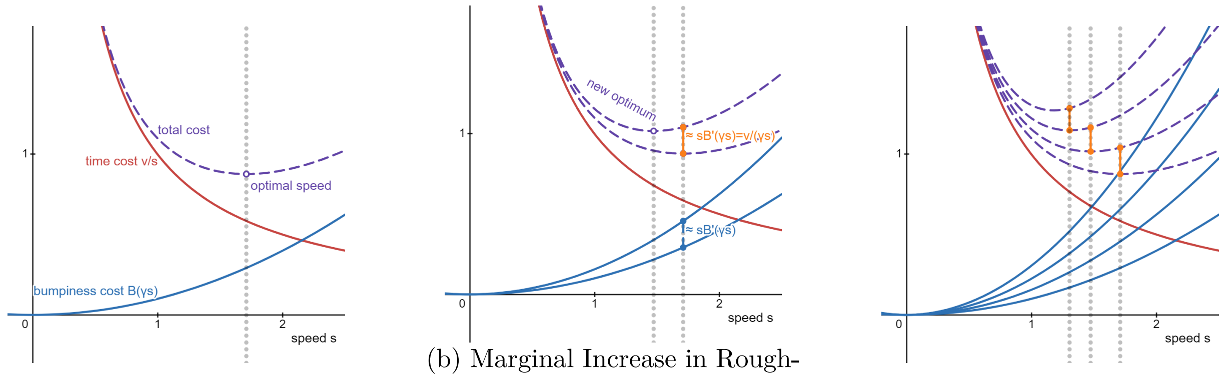

(a) Time and Bumpiness Costs

(b) Marginal Increase in Roughness

(c) Larger Increase in Roughness

Note: Graphic representation of the model and cost of roughness in the simple model with $K = 0$.

4.2 Empirical Results

Our main empirical exercise is to estimate how driver speed varies with salient changes in road roughness, holding other road characteristics fixed.[18]

We use four sources of variation in road roughness: (1) average differences in roughness between towns, (2) roughness close to town borders, (3) road repaving events in Chicago, and (4) roads with railroad crossings in Chicago.

We present this analysis in the next three subsections. Whenever possible, we also estimate how variation in speed limits predicts change in driver speed. In section 4.4 we estimate the two versions of the model by matching these reduced form empirical results in the model, and we use the estimated model to quantify the total private costs of roughness.

4.2.1 US-wide Town Border Discontinuities

To calculate the impact of roughness on driver speed, we use sharp changes in road roughness at town borders. Different cities and towns may have systematically different levels of road quality, and those differences should be salient to Uber drivers, as we saw earlier on the map of Chicago (Figure 2). These changes are particularly useful because in our model, drivers *anticipate* the level of roughness and choose optimal speed based on roughness.

To illustrate our approach and provide a first round of evidence of this channel, we first show

[18]We use the speed data as a control when estimating the relationship between speed and bumpiness at the road segment level, and we are now using the median driver speed of a segment as an outcome. In principle, using the speed data this way could bias our estimates, yet this is unlikely to be a problem in practice. In section 2.2, we showed that the relationship between predicted roughness and speed is very similar when we measure speed in a hold-out sample not used to estimate roughness.

results from the border of Chicago. Figure 7 shows binned scatter plots of predicted roughness and speed up to 4 kilometers away from the boundary of the city of Chicago. There is a sharp increase in predicted road roughness of approximately half a standard deviation just inside the city of Chicago, while speeds on local roads are lower by about 3 miles per hour. The changes are sharp around the boundary and remain relatively stable away from the boundary.

Figure 7: Predicted Roughness and Speed around the Chicago border

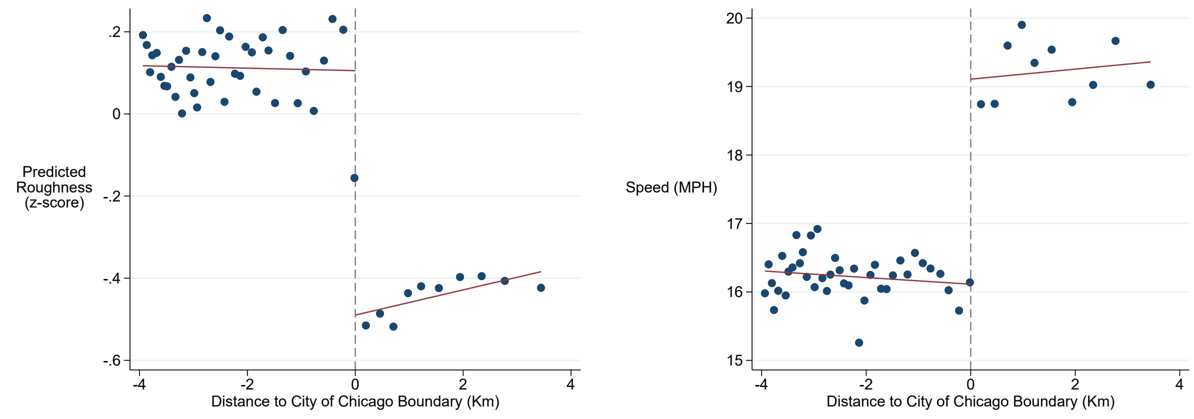

Panel (A) Predicted Roughness

Panel (B) Speed

Note: Panel A plots a binned scatterplot of predicted road roughness as a function of the distance to the boundary of the city of Chicago, for local road segments. Negative distances correspond to roads inside the city of Chicago. Panel B uses speed in miles per hour as the outcome. Appendix Figure A.12 repeats the exercise for arterial roads..

4.3 Town Borders Across the US

We next investigate the relationship between road roughness and speed across the US. We use the variation in γ_r induced by sharp changes at town borders. We then estimate how much slower a driver travels when they enter a town with roads where driving faster leads to a larger increase in experienced roughness. We also estimate the effect of varying speed limits.

We restrict to pairs of towns that share a border and that have local road segments within 250 meters of the border on both sides of the border. We focus on a sample of all local roads in these towns, and a sample of local road segments within 500 meters to the border.[19] We have 1,516 towns in the first sample. In our second “town border” sample, we have 1,209 towns and 1,285 border pairs, covering 72 Uber cities.

[19]We exclude town pairs that lie in different Uber cities. In order to reduce estimation noise, we consider only road segments that have above median number of observations (at least 97 observations per segment), and only border pairs that have above median number of segments on the side of the border with the smallest number of observations (at least 10 road segments on each side of the border).

We estimate a log linear equation at the town- or town times border pair level. For road segment r in town c, located in Uber region Uber(c) and with town border pair b, we estimate

$$\log(s_{rc}) = \alpha \log(\gamma_r) + \beta \log(s_c^{\mathrm{LIM}}) + \mu_{\mathrm{uber}(c)} + \epsilon_{rc} \tag{4.5}$$

$$\log(s_{rbc}) = \alpha \log(\gamma_r) + \beta \log(s_c^{\mathrm{LIM}}) + \mu_b + \epsilon_{rbc}, \tag{4.6}$$

where s_{rc} and s_{rbc} denote the median speed on road segment r measured in the Uber data. In the first specification we include Uber city fixed effects. In the second specification, we use (finer) border pair fixed effects, which means we only use variation within town border areas. We compute median speed limit s_c^{LIM} at the town level, based on the speed limit data from HERE.com, a mapping data platform (Appendix C).

We construct the slope of the speed-roughness relationship, γ_r, at the road segment level using all our Uber data. When we estimated the segment-specific slopes in section 2.2, we modeled the relationship between speed and experienced roughness including an intercept term $\hat{\mu}_r$ in addition to the slope term $\hat{\gamma}_r$. To construct the measure γ_r that we use here, we regress the estimated slope $\widehat{\gamma}_r$ on predicted roughness $\widehat{Z}_r^{20} = \widehat{\mu}_r + 20 \cdot \widehat{\gamma}_r$ at 20 miles per hour, and define γ_r as the fitted values.[20] We construct γ_r this way to reduce estimation noise, taking advantage of the fact that Z_r contains information on the roughness relationship from both $\widehat{\mu}_r$ and $\widehat{\gamma}_r$.

We estimate 4.5 and 4.6 using an instrumental variable approach, respectively using town and town-by-border-pair indicators as instruments for the road roughness term $\log(\gamma_r)$. For the first equation, the exclusion restriction that we impose is that within Uber regions and controlling for the official speed limit, average roughness at the town level is uncorrelated with omitted town-level factors that correlated with average log speed. For the second equation, we assume that within border pair and controlling for the speed limit, average roughness is uncorrelated with omitted factors that correlated with average log speed.

4.3.1 The Impact of Road Resurfacing in Chicago

We introduced the resurfacing events in Section 3.4. We focus on a balanced panel of road segments for road repaving events p between May and mid-July 2021. We measure road roughness slopes γ_{rpt} for all road segments r that are repaved during event p, as well as control road segments in a donut around the repaved segment. We include two time periods $t = 0, 1$ covering the periods before repaving (April 2021) and after repaving (August 2021). We estimate the following relationship:

$$\log(s_{rpt}) = \alpha \log(\gamma_{rpt}) + \varphi_r + \phi_p \times Post_t + \epsilon_{rpt}, \tag{4.7}$$

where φ_r are road segment fixed effects, ϕ_p are fixed effects for all observations associated with event p, both the repaved and the control segments, and $Post_t$ is an indicator for $t = 1$. We

[20] 20mph is approximately the median speed for local roads in the sample.

estimate this equation using an interaction variable $Repaved_r \times Post_t$ as instrument for $\log(\gamma_{rpt})$ where $Repaved_r$ measured whether the segment r was repaved. The exclusion restriction requires that speed cost factors do not change differentially over time for repaved relative to the control road segments.

While we control for the speed limit for each road segment in our sample, speed limits do not predict actual speeds, possibly because of noise created by the match between speed limit and road segment and possibly because variation in speed limits on local roads within Chicago is modest (Appendix Table A.7).

4.3.2 Empirical Results: The Impact of Roughness on Speed

Table 2 reports the results from our analysis of the impact of road roughness and speed limits on driver speed. The first column reports IV estimates that use across-town variation based on equations (4.5). Confidence intervals are constructed using town-level Bayesian bootstrap (Rubin, 1981).

We find an elasticity of speed with respect to roughness of -0.46, which means that driver speed fall by 4.6% on a road segment with 10% larger roughness slope. We also find a positive elasticity of 0.24 with respect to the speed limit. However, the exclusion restriction at the town level is demanding, as other town-level factors may correlate with both road speed and roughness.

In the second column, we focus on variation in road roughness, within 500 meters of town borders. We estimate equation (4.6), and we instrument for log roughness using town by border pair indicators. We perform inference using a border-pair level Bayesian bootstrap. The elasticity of speed with respect to roughness drops by around a third to -0.31 and continues to be precisely estimated. The effect of speed limit drops slightly to 0.19. We report more related specifications in Appendix Table A.6. These results show that speed limits appear to be orthogonal to roughness at town borders.

In the third column, we report results using variation from road resurfacing in Chicago, based on equation (4.7). We estimate a lower elasticity of -0.14.[21] This estimated elasticity may be smaller either because resurfacing is less salient or because the spatial scale here is much smaller than in our town border analysis.

In the last column we use the large variation in road roughness induced by at-grade railway crossings. We discussed the sample construction in section 3.3 and showed the results on roughness visually in Figure 4. Uber driver speeds are significantly lower on road segments with railway crossings, and we estimate an implied elasticity of -0.23.

Overall, these results show that people drive significantly slower on rough road segments, and speed limits are also a determinant of driver speeds. These estimated parameters allow us to compute the total user costs of roughness when drivers choose speed optimally.

[21] Appendix Table A.7 reportes the first stage and alternate specifications.

Table 2: Empirical Moments: Uber Speed, Roughness and Speed Limits

	Log speed (mph)			
	(1)	(2)	(3)	(4)
Log roughness slope	-0.46***	-0.31***	-0.14***	-0.23***
	(0.02)	(0.02)	(0.03)	(0.05)
Log speed limit	0.24***	0.19**		
	(0.04)	(0.06)		
Sample	Town	Borders	Chicago	Railways
Sample restriction:		< 500m		
Uber City FE	Yes			
Border pair FE		Yes		
Road segment FE			Yes	
Repaving event $\times$ Post FE			Yes	
Segment group FE				Yes
Estimator	IV	IV	DDIV	OLS
Uber cities	90	72		
Towns	1,516	1,209		
Border pairs		1,285		
Repaving events			611	
Observations	657,734	68,593	19,878	946

Note: This table reports estimates of the elasticity of speed with respect to road roughness and the speed limit using four sources of variation on local roads. In column 1, the sample is all road segments in our sample of towns. In column 2, the sample is restricted to road segments within 500 meters from a town border. The sample in column 3 is road segments that were resurfaced between May and mid-July 2021, and nearby control segments from a donut area 100-200 meters away from resurfaced segments. The sample in column 4 is local road segments with at-grade railway crossings in Chicago, and nearby control segments in a donut 100-200 meters away from railways. The specifications for columns 1, 2, and 3 are (4.5), (4.6), and (4.7), respectively. For inference, standard errors are constructed using town-level and town-by-border-pair-level Bayesian bootstrap procedure in columns 1 and 2 (Rubin, 1981), and clustered at the level of repaving events in column 3. Column 4 has bootstrapped standard errors at the level of segment group.

4.4 Model Estimation and The Total Costs of Roughness

To connect the data with the models discussed in Section 4.1, we make the following parametric assumptions.[22] We assume the bumpiness cost function is a power function:

$$B(\gamma s) = \beta_1 (\gamma s)^{\beta_2},$$

with $\beta_1, \beta_2 > 0$, and the other cost of speed is linear $K(s/s^{\text{LIM}}) = \kappa \cdot s/s^{\text{LIM}}$.

We use classic minimum distance (CMD) to estimate the three parameters $(\beta_1, \beta_2, \kappa)$ by matching three moments from our reduced form analysis: the elasticity of speed with respect to roughness, the elasticity of speed with respect to the speed limit, and the average speed in our sample. We also estimate the simple model without other speed costs ($\kappa = 0$) using only the first and third moments.

Our benchmark model uses the sample and specification from column 3 from Table 2. To run CMD, for each candidate parameter values, for each road segment in the sample, we compute model-predicted log speed given the segment's road roughness and the cost parameters, and we use this as the outcome to run the same specification as we do with the true speed data. We then find the parameters for which we obtain the same reduced form moments.

For inference, we propagate the Bayesian bootstrap uncertainty from our reduced form estimates through the model. Specifically, we repeat the entire procedure 200 times, each time estimating parameters to match a different set of reduced form moments.

Table A.8 reports the results. In the full model, we find $\widehat{\beta}_2 = 1.06$ and we cannot reject a linear cost of bumpiness, while with the simple model we find a precise concave cost of bumpiness based on $\widehat{\beta}_2 = 0.47$.

We next use the estimated model to quantify the total user costs of roughness expressed in dollars per mile driven, when drivers choose speed optimally, as a function of a road segment's roughness slope γ_r. We use this estimate for the rest of this paper to put a dollar value on differences in road quality across different areas.

To standardize analysis, we will evaluate the cost of roughness for a road segment of 1 mile and we use a measure of the value of driving time of $v_i = 15\$/\text{hour}$, which is roughly one half of the hourly wage in the US in the first quarter of 2023.[23] We will evaluate the cost at the median speed limit ($\overline{s}^{\text{LIM}} = 25\text{mph}$).

We use the estimated model parameters from column 1 of Table A.8. We also use the model to construct optimal speed, which allows us to focus on changes in driving speed that happen exclusively due to changes in road roughness.

Figure 8 shows the implied relationship between the roughness slope γ_r and cost. We include

[22] Appendix E.1 discusses non-parametric identification.

[23] Average weekly earnings are 1,068 USD and average hours are 34.7 in August 2021. We are following Ken Small on willingness to pay equal to half the hourly wage.

the histogram of γ_r across road segments in the border sample for reference. The solid red line plots total user costs relative to costs at median roughness. These costs include the cost of time $\frac{v_i}{s^*_{ri}}$ as well as the speed costs due to bumpiness and due to other speed costs, and are computed at the optimal speed. Gray dashed lines indicate bootstrapped confidence bands.

Figure 8: Total User Cost of Road Roughness (Local Roads)

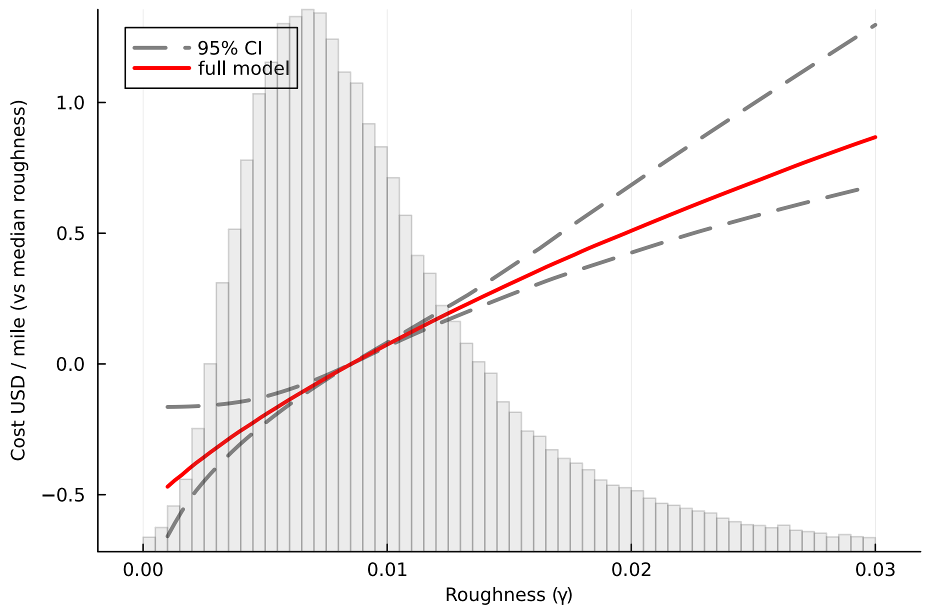

Note: This graph plots the total user costs of roughness per mile of road from equation based on the full model. We compute optimal speed at different road roughness slopes γ_r in the model using the estimates from column (1) in Table A.8, at the median speed limit (25 mph). The histogram of the road segment slope γ_r in the town border pair sample is displayed in the background. The dashed gray lines indicate point-wise 95% confidence intervals from bootstrapping the entire estimation procedure at the level of town border pairs.

How costly is road roughness? Table 3 summarizes the impact of increasing roughness by one standard deviation, starting at the median value. Through the lens of the full model, starting from median roughness, a one standard deviation of road roughness increases user costs by 0.33 USD/mile, with bootstrapped 95% confidence interval of [0.28, 0.40] USD/mile. The results from the simple model are 0.31 USD/mile, slightly smaller but not statistically distinguishable.[24]

[24]This may seem surprising because equation (4.1) implies that setting $K = 0$, as in the simple model, always increases the cost of roughness. Figure A.13 confirms this when we use the full model parameters and set $K = 0$. However, the *estimated* simple model yields different parameters and predictions for optimal speed than the full model.

Table 3: Total User Costs of Road Roughness (Local Roads)

	Full Model	*Simple Model*
Cost of +1SD	0.33	0.31
(USD/mile)	[0.28, 0.40]	[0.27, 0.34]

Note: This table reports the total cost increase in road roughness due to a one standard deviation increase in roughness starting from median road roughness. 95% confidence intervals are based on model parameter uncertainty, in turn based on Bayesian bootstrap at the level of town border pairs.

In the rest of the paper, we will use the estimates from Figure 8 for local roads. These are based on our town border analysis and as elasticity of speed with respect to roughness of -0.31. We also show that the cost of an additional SD of road roughness is increasing in the absolute value of this moment (Figure A.14 and Table A.9). Readers may use these results to adjust our estimates of cost of roughness in the rest of the paper accordingly.

4.5 Other Costs of Rough Roads

The logic of the model implies that our estimated costs will capture both time costs and all other costs associated with road roughness that are borne by the driver, including vehicular damage. To obtain a more direct measure of costs, we look at damage found in Massachusetts' vehicle inspections.

Road roughness could also generate positive or negative externalities, which would not be captured by our model. During congested periods, one car slowing down can delay other cars and cost them time. Slower driving can also mean fewer accidents, which is why some towns build speed bumps, which appear as road roughness in our data. We will look at the link between road roughness and fatal vehicle crashes.

4.5.1 Other Costs: Vehicle Maintenance

We obtained non-commercial vehicle inspection failure rates for all inspection stations in Massachusetts for ten indicators for one year (May 2021 to April 2022). We pre-selected seven "main" forms of inspection failure that we hypothesized could be linked to vehicle damage due to rough roads: brakes, front end, steering and suspension frame, muffler and exhaust system, bumpers/fenders/exterior sheet metal, and tires. We also pre-selected three "placebo" types of failure that we thought unlikely to be related to rough roads: windshield wipers and cleaner, windshield, and rear view mirror. We focus on the average failure rate over the main indicators, and for the average over the placebo indicators, but we also show disaggregated results in Table A.10. For each station, we also obtained the average vehicle age and the total number of inspections.

The original data covers 1,753 inspection stations and over 4.6 million inspections. After geocoding the station locations based on their names and address, we are left with 1,263 stations, covering 3,665,733 inspections. For each station, we also obtained the average vehicle age and the total number of inspections. We then merge this data with our Uber data on road roughness. Our final analysis samples cover 926 stations in 112 towns where we have Uber data at the local road level. We then aggregate and run the analysis at the town level.

Figure 9 shows the correlation between road roughness on local roads for a given town, and that town's log average failure rate, where the average is taken over all inspections in the town in our data, and over the seven main indicators. The correlation is positive (0.21) and on the edge of statistical significance (p-value 0.105).

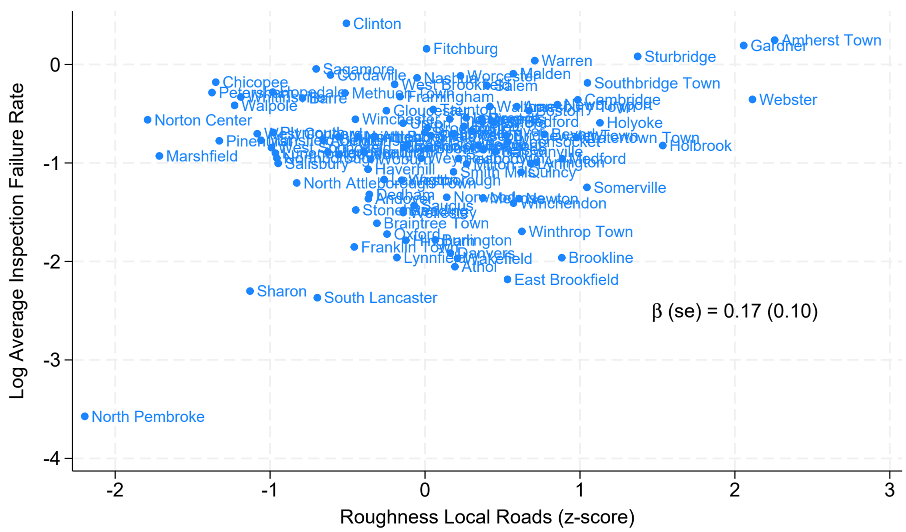

Figure 9: Inspection Failure Rate vs Local Road Roughness (Main Indicators)

Note: This graph plots town-level inspection failure rates versus local road roughness. The failure rate is the average over the seven main inspection indicators.

To use our placebo design, we calculated the average failure rate for the "main" measures and the "placebo" measures for each town and estimated the following regression using Poisson Pseudo Maximum Likelihood (PPML):

$$\log(\text{Failure Rate}) = -\underset{(0.06)}{1.72} + \underset{(0.04)}{1.02} \cdot \text{Main} + \underset{(0.09)}{0.06} \cdot \text{Local Roughness} + \underset{(0.06)}{0.07} \cdot \text{Local Roughness} \times \text{Main}.$$

There were 224 observations (112 towns and two failure rates per town). Local roughness does correlate more strongly with the main effects than with the placebo effects, but the interaction is not statistically significant. Moreover, there is also a positive and insignificant correlation between roughness and the placebo failure rate. These results suggest that rough roads are associated with slightly higher failure rates, but there is little reason to have confidence that the relationship is causal.[25]

4.5.2 Other Costs: Vehicle Crashes

To test whether slower speeds reduce the number of collisions, or make collisions less deadly, we use the universe of fatal injuries suffered in motor vehicle traffic crashes in the US in 2021 from the Fatality Analysis Reporting System (FARS) maintained by the The National Highway Traffic Safety Administration (NHTSA). We know for each crash the coordinates and the type of road where it happened. We assign crashes to towns and type of road (local, arterial or highway), and then estimate the regression:[26]

$$\log(\mathbb{E}Fatalities_i^r) = \alpha^r + \beta^r Z_i^r + \gamma^r X_i + \epsilon_i^r$$

where $Fatalities_i^r$ is the number of fatalities in 2021 in town i on road type r, Z_i^r denotes the z-score of predicted road roughness for road type r, and X_i is a vector of controls for town i. We run this analysis separately by road type r (local, arterial, highway). Due to the presence of zeros in the outcome variable, we estimate this equation using PPML.

The first six regressions of Table 4 report the semi-elasticity of the number of fatal crashes in a town with respect to the z-score of road roughness. The first three regressions include controls only for total population and median household income, using the ACS. The next three regressions include controls for the two best estimates of road traffic that we have: the number of car commuters, or residents driving to work, measured from the ACS, and total local employment constructed from the Census Bureau's ZIP Codes Business Patterns. Controlling for car commuters has two problems: it fails to capture the non-commute drives which may be more important for accidents, and the variable is endogenous with respect to road quality.

The first three regressions shows that a one standard deviation increase in road roughness is associated with .13 log points fewer deaths on local roads, .19 log points fewer deaths on arterial roads, and .16 log points fewer deaths on highways. Regressions (4)-(6) show that controlling for car commuters and employment reduces the estimated coefficients to -.12, -.12 and -.06 for local, arterial and highway roads respectively. The fact that controlling for these imperfect measures of

[25]Table A.10 reports the correlation between local road roughness and the failure rates for each category separately. None of the estimated coefficients is statistically distinct from zero, and the coefficient on window wipers tests (a placebo) are almost as large as the coefficients on suspension or tire failures.

[26]The crashes data includes the DOT functional class category of the road segment where the crash took place. We assume that arterial roads include the categories: "Major Collector," "Minor Collector," and "Minor Arterial." We assume that highway roads include the categories: "Interstate," "Principal Arterial - Other Freeways and Expressways," and "Principal Arterial - Other." The remaining category is "Local," which we map to local roads in our sample.

traffic volume causes the estimated coefficients to drop by one to two thirds for arterial roads and highways suggests that an important reason why roughness is linked to fewer fatalities may be that people drive less in areas with rough roads. We hope that future work will combine better data on road traffic with roughness measures and crashes.

Table 4: Crash Fatalities and Road Roughness at the Town Level

	Fatal Crashes						Log Speed		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Local Roughness (z-score)	-0.133*			-0.124			-0.137***		
	(0.066)			(0.068)			(0.014)		
Arterial Roughness (z-score)		-0.188***			-0.122***			-0.109***	
		(0.054)			(0.036)			(0.003)	
Highway Roughness (z-score)			-0.162**			-0.065*			-0.070***
			(0.053)			(0.029)			(0.005)
Log Residents	0.988***	1.016***	1.000***	0.784***	0.386	0.143	-0.038*	-0.010	-0.043**
	(0.026)	(0.032)	(0.023)	(0.178)	(0.252)	(0.152)	(0.016)	(0.009)	(0.016)
Log Income	-1.542***	-1.061***	-0.923***	-1.537***	-1.148***	-1.010***	0.016*	0.016***	0.009
	(0.127)	(0.064)	(0.062)	(0.131)	(0.081)	(0.065)	(0.006)	(0.004)	(0.007)
Log Residents Who Drive to Work				0.015	0.503*	0.583***	0.057***	0.012	0.059***
				(0.167)	(0.253)	(0.153)	(0.015)	(0.009)	(0.015)
Log Employment				0.182**	0.131***	0.280***	-0.002	-0.000	0.007*
				(0.064)	(0.035)	(0.033)	(0.003)	(0.003)	(0.003)
Observations	6,165	9,056	5,570	6,165	9,056	5,570	6,165	9,056	5,570

Note: This table reports the semi-elasticity of fatalities and speed with respect to road roughness for local, arterial and highway roads. The sample in each column is all towns in the US with Uber road roughness data for that type of road segment, excluding Puerto Rico. The outcome measures the number of fatalities on the specified type of road (local, arterial, highway). The controls are log town population, log mean income at the town level, and log of the number of town residents who drive to work, from the ACS, and log town employment from the Census Bureau's ZIP Codes Business Patterns. Columns 1-6 are estimated using PPML, and columns 7-9 are estimated using OLS. Robust standard errors in parentheses, *p<0.1; **p<0.05; ***p<0.01.

Are these results compatible with external estimates of the connection between speed and fatalities? Columns (7)-(9) find that the semi-elasticity of speed with respect to roughness is -.14 for local roads, -.11 for arterials and -.07 for highways.

The coefficients of speed on roughness would suggest our coefficients of fatalities on roughness if the elasticity of fatalities with respect to speed was approximately one. Most external estimates estimate a much larger impact of speed on fatalities. Elvik (2005) and Elvik et al. (2019) both provide comprehensive meta-analyses of the copious literature on the link between speed and traffic accidents. Elvik (2005) reports an elasticity of fatalities with respect to traffic speeds average 3.65 across "well-controlled studies." Elvik et al. (2019) update this study and find a somewhat higher elasticity of between 5.5 and 7, based on 18 new studies published after the year 2000.[27]

[27] Ashenfelter and Greenstone (2004) shows that a speed limit increase led to a 4% increase in speeds and 35%

Our fatalities impact is much lower than the impact of roughness on speed might suggest, perhaps because speed reductions from road roughness are not as "effective" at reducing fatal crashes compared to interventions such as changes in speed limits or enforcement, which do not involve physical changes to the road. Our results are therefore consistent with roughness having a *positive* effect on fatalities conditional on speed, which could occur if drivers have less control over their vehicles in rough areas.

5 Infrastructure Inequality in the US

We now turn to our original question: which Americans experience the worst infrastructure? How big are the costs imposed by rough roads on disadvantaged Americans? In this section, we focus on the cost of road roughness on minorities and low-income Americans. We focus on tract level variation, and ask whether road quality is worse in lower-income tracts, and in tracts that are disproportionately inhabited by Black households. We then differentiate between variation across towns and variation within town.

For the most part, we do not try to differentiate effects due to race with effects due to income. Race and income are strongly correlated across space, and we do not have exogenous variation in either variable. This section documents the extent to which poorer Americans and minorities experience worse roads and then whether obvious tract or town level variables seem to explain why the roads in their neighborhoods are worse.

The inequality in local road quality that we document by income and minority status links our paper to several literatures that study spatial disparities (e.g. Brueckner et al., 1999). A congressional report prepared by the GAO in 2022 documented that highway road quality was lower in census tracts with a high share of minorities or a high share of poor households (Government Accountability Office, 2022). Fu et al. (2023) uses Census data to show that Black commuters who commute by car face longer commutes, especially in large, congested, and expensive cities. These patterns are in part explained by residential segregation and spatial mismatch (Cutler et al., 1999; Gobillon et al., 2007). In our analysis, we highlight how lower local road quality in Black neighborhoods imposes direct costs, including through lower speeds. More generally, our work is related to a broad literature studying spatial disparities in neighborhood outcomes and access to amenities.

5.1 Rough Roads, Income and Race

The results in Section 4 allows us to transform the units of our road roughness measure from standard deviations of vertical acceleration to dollars per mile. This measure captures the cost

increase in fatalities, which means an elasticity of around 9. See also DeAngelo and Hansen (2014) and Van Benthem (2015).

to drivers from driving slower to compensate for roughness, as well as the direct cost of driving over rough roads. For each local road segment with Uber data, we convert its estimated roughness measurement to a per mile cost, using the formula described in Section 4.4.

Our main analysis is at the Census tract level, which contain roughly 4,000 people, but we also report results for the localities that actually administer local roads.[28] At each level of geography, we compute the average cost generated by road roughness per mile. We focus on the one hundred most populous Metropolitan Statistical Areas (MSAs), because our sample of local roads is large but not comprehensive. The median MSA in the sample has 136 tracts with roughness data, or 62% of all the tracts in the MSA. We miss tracts on the outskirts of the metropolitan areas, where few Uber drivers travel. The median tract in the sample has roughness data for 13% of its local road segments.

There is certainly incomplete measurement of any aggregate geography in our data, but Table A.2 shows that among road segments with below median Uber coverage, coverage is not correlated with roughness. This fact suggests that our sample is not unusually rough or smooth, although it remains possible that roads never used by Uber drivers are either much rougher or much smoother. Appendix Tables A.11 and A.12 show analogous results to those presented below where we restrict the sample by requiring tracts or towns to meet several coverage thresholds.

Figure 10 visualizes our tract-level results for Chicago and New York City. In Chicago, the core urban area near to Lake Michigan is smoother, as is the near south side, which is disproportionately Black. The north west side, which is more industrial, is rougher and so is the far south side of the city. In New York, the core is rough, especially downtown Manhattan and the periphery is smooth, especially Staten Island.

Table 5 shows the correlation between the average cost per mile of local roads in a Census tract and median household income, and the percent of residents in the tract who are Black, Hispanic, or Asian. The first three columns focus on income. The coefficient of $-.061$ in regression (1) means that as income increases by 1 log point, the cost per mile of local roads drops by 6.1 cents. If a household drivers 3,300 miles per year on local roads, then a one log point increase in income is associated with 201 dollars less harm from bad roads per year.[29] In the second regression, which controls for metropolitan area fixed effects, the coefficient increases in magnitude to $-.074$. Richer metropolitan areas, such as New York and San Francisco, have worse roads and so the income gradient is steeper within area than across the entire US. This coefficient implies that a household that drives 3,300 miles on local roads per year suffers 244 dollars less harm as its income increases

[28] We use Census place data, which includes all legally bounded entities such as cities, towns, and villages (depending on the state), as well as Census Designed Places (CPDs), which are statistical entities such as unincorporated communities.

[29] US households drive an average of 21,553 miles per year (summed up over all vehicles in the household), based on 2017 data from the National Household Travel Survey (NHTS). Estimates from the federal Department of Transportation's Highway Statistics reveal that 15.3% of all vehicle miles traveled in the US occur on local roads, and 44.4% on local and arterial roads combined. We thus estimate that each year, the average household drives 3,297 miles on local roads and 9,569 on local and arterial roads.

by one log point.

Figure 10: Road Roughness by Census Tract

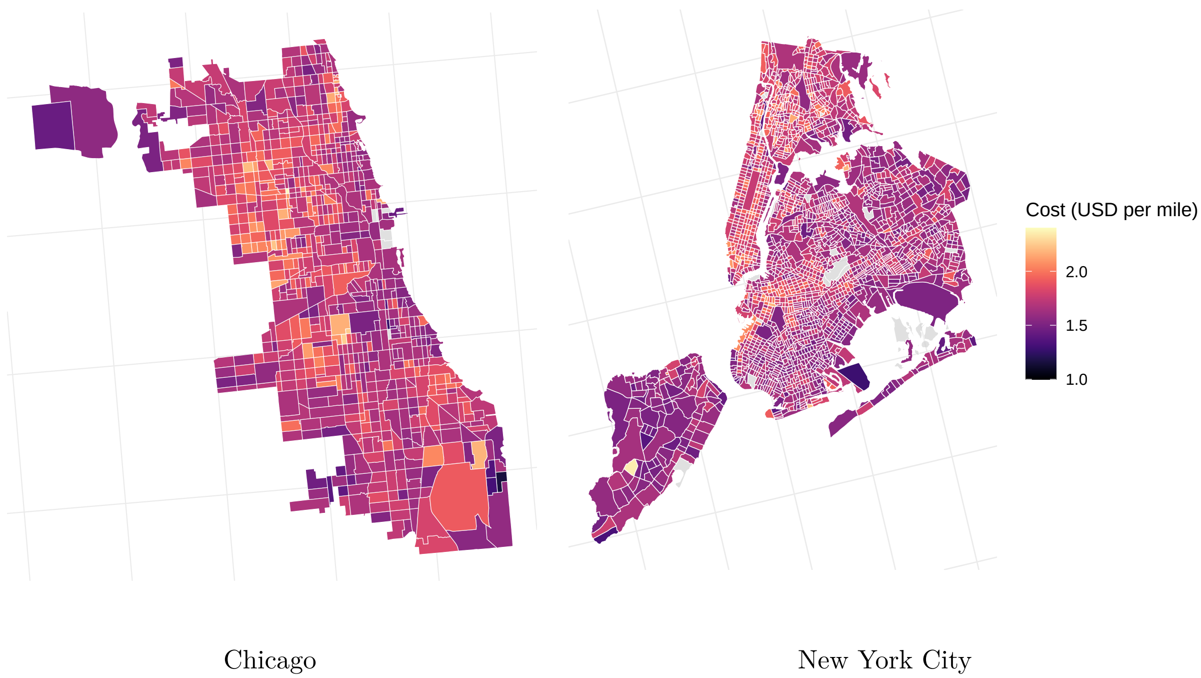

Note: These maps show costs of local road roughness at the Census tract level for Chicago and New York City. To construct them, for each local road segment we compute the cost per mile using the model from Section 4, and take the average at the tract level.

In the third regression, we control for town fixed effects. In this case, the coefficient declines in magnitude to $-.023$. The reduction in magnitude suggests that two-thirds of the relationship between income and road roughness occurs because richer towns have smoother roads and one-third reflects the fact that richer people live in neighborhoods that are particularly smooth within their town. The town level connection between income and road roughness is not surprising, as richer towns should have more resources for repaving. The remaining tract-level connection within cities could represent richer people choosing to live further away from city centers or in lower density areas, or the city directing road repaving efforts to disproportionately wealthy neighborhoods. We will evaluate these hypotheses in later tables.

In regression (4), we turn to race and ethnicity. The coefficient of .151 on the fraction Black means that as the share of the population moves from 100 percent White to 100 percent Black, the cost associated by road roughness increases by 15 cents per mile. A household that drives 3,000 miles annually on local roads would pay an extra roughness cost of 450 dollars per year.[30] The

[30]Black households drive an average of 19,386 miles per year, based on 2017 NHTS, which implies that Black

second regression shows that controlling for metropolitan area fixed effects does not change the coefficient. Black households live disproportionately in some of America's roughest metropolitan areas, such as New Orleans, and in some of America's smoothest metropolitan areas, such as Jackson, Mississippi. Hispanic and Asian households also live in areas that have rougher roads, with slightly smaller coefficients than for Blacks. In the case of Asian-Americans, metropolitan area fixed effects reduce the coefficient by more than half.

Table 5: Tract Roughness, Income, and Race

	Dependent variable:					
	Cost (USD per mile)					
	(1)	(2)	(3)	(4)	(5)	(6)
Ln median income	-0.061^{***}	-0.074^{***}	-0.023^{***}			
	(0.002)	(0.002)	(0.002)			
Fraction Black				0.151^{***}	0.147^{***}	0.021^{***}
				(0.004)	(0.005)	(0.005)
Fraction Hispanic				0.128^{***}	0.108^{***}	0.042^{***}
				(0.005)	(0.005)	(0.006)
Fraction Asian				0.120^{***}	0.050^{***}	-0.050^{***}
				(0.010)	(0.010)	(0.011)
Climate controls	Yes			Yes		
MSA Fixed effects		Yes			Yes	
Town Fixed effects			Yes			Yes
Observations	32,959	33,126	33,126	32,959	33,126	33,126
Adjusted R^2	0.053	0.147	0.352	0.072	0.147	0.352

Note: This table shows coefficients from regressing average tract roughness cost on tract income and the fraction of the tract that is minority (Black, Hispanic, Asian). The sample includes tracts in the top 100 MSAs. In columns 4-6, the omitted category is people who do not identify as any of "Not Hispanic or Latino: Black or African American alone", "Asian alone", or "Hispanic or Latino". On average, this population is 97% "Not Hispanic or Latino: White alone". Columns (1) and (4) control only for baseline climate variables (January temperature, January precipitation, July precipitation, and an indicator for being close to the coast, all at the MSA level). Columns (2) and (5) add MSA fixed effects, which absorb climate controls. Column (3) and (6) add town fixed effects, which absorb MSA fixed effects. Robust standard errors in parentheses, *p<0.1; **p<0.05; ***p<0.01.

Regression (6) includes town fixed effects and the coefficient on share Black drops to .021. The gap between regression (5) and regression (6) suggests that 85 percent of the connection between race and rough roads is explained by difference across local jurisdictions. This fact suggests that the

households on average drive 2,966 miles on local roads and 8,607 on local and arterial roads. See footnote 29 for details.

large gaps in road quality across the US are better understood by sorting into different towns than by either discriminatory behavior by local public works companies or by selection into neighborhoods with worse roads. The coefficient also falls significantly for Hispanics, and it reverses sign for Asian households.

Table 6 re-examines the relationship between racial composition and the cost of road roughness, and looks at whether controlling for measures of income, including median income, poverty and unemployment, eliminates the link between roughness and race. Regressions (1)-(3) control for metropolitan area fixed effects; regressions (4)-(6) control for town fixed effects. Regression (1) reproduces Table 5 Regression (5), and regression (4) reproduces Table 5 Regression (6).

Table 6: Tract Roughness and Race, Controlling for Income

	Dependent variable:					
	Cost (USD per mile)					
	(1)	(2)	(3)	(4)	(5)	(6)
Fraction Black	0.147***	0.092***	0.096***	0.021***	−0.004	−0.005
	(0.005)	(0.006)	(0.006)	(0.005)	(0.006)	(0.006)
Fraction Hispanic	0.108***	0.054***	0.055***	0.042***	0.014**	0.013*
	(0.005)	(0.006)	(0.006)	(0.006)	(0.007)	(0.007)
Fraction Asian	0.050***	0.035***	0.034***	−0.050***	−0.063***	−0.063***
	(0.010)	(0.010)	(0.010)	(0.011)	(0.012)	(0.012)
Ln median income		−0.044***	−0.047***		−0.021***	−0.019***
		(0.003)	(0.003)		(0.003)	(0.003)
Fraction poverty			−0.004***			0.003
			(0.001)			(0.002)
Fraction unemployed			−0.032***			0.004
			(0.008)			(0.008)
MSA Fixed effects	Yes	Yes	Yes			
Town Fixed effects				Yes	Yes	Yes
Observations	33,126	33,126	33,126	33,126	33,126	33,126
Adjusted R^2	0.147	0.154	0.155	0.352	0.353	0.353

Note: This table shows coefficients from regressing average tract roughness cost on the fraction minority and other controls. See notes for Table 5. Robust standard errors in parentheses, *p<0.1; **p<0.05; ***p<0.01.

Regressions (2) and (3) show that controlling for income leads the estimated coefficient on the share of the population that is Black to fall from .147 to .092. Consequently, 37% of the link between race and road roughness across neighborhoods and towns could be accounted for by the fact that neighborhoods with a higher share of Black households are on average lower-income. By comparing the coefficient on income in Regression (2) with the coefficient on income in Table 5, Regression (2), we also see that controlling for race appears to reduce the measured coefficient on income by 40%. Adding in further controls in Regression (3) for poverty and unemployment does little to the estimated coefficient. These variables have little explanatory power, which is compatible with the view that local resources to pave roads are mostly associated with the average income.

Regression (5) in Table 6 shows that controlling for income leads the coefficient on the Black population share to flip sign and become statistically indistinguishable from zero, when we control for town-specific fixed effects. The income coefficient remains essentially unchanged from Table 1. Again, adding further controls for the poverty and unemployment rates does little to the estimated coefficients on race and median income. Taken together these results seem to suggest that race and income are both powerful correlates of road quality across the U.S. as a whole, but the correlation with the Black population share is driven almost entirely by variation across, rather than within towns, especially when we control for income. However, within towns there is still some correlation between income and smooth roads.

Table 7 looks at tract roughness within town, and asks what can explain the relatively modest within-town correlations between road roughness and income, and race. Columns (1) and (4) reproduce columns (3) and (6) in Table 5 for comparison purposes. Columns (2) and (5) include tract level controls, including population and land area (which together combine to form density), employment and proximity to the Central Business District of the MSA, which we approximate by the coordinates of city hall. These variables all have the expected signs. Population and employment both predict worse roads. Land area is correlated with better roads. Roads that are further away from the city center all have better roads, presumably because they get less use. Columns (3) and (6) add controls for the share of the population that drives to work and the number of Uber observations we have in each tract, which is a measure of Uber use at the tract level and presumably also a measure of the demand for mobility across the tract. Roads are much smoother in areas where people drive to work, but we don't know if this reflects lower levels of bus or truck travel in these neighborhoods or if people drive more often when the roads are nice. The coefficient on Uber segments is also negative. Given the evidence in Table A.2 that measured road roughness is not correlated with Uber usage at the road segment level, we interpret the negative coefficients on tract-level Uber usage in Table 7 as indicating that areas with high traffic tend to have smoother roads. Controlling for these variables does cause the coefficient on income to attenuate, but not the coefficient on fraction Black.

Table 7: Tract Roughness and Tract Characteristics

	Dependent variable:					
	Cost (USD per mile)					
	(1)	(2)	(3)	(4)	(5)	(6)
Ln median income	-0.023^{***}	-0.009^{***}	-0.007^{***}			
	(0.002)	(0.002)	(0.002)			
Fraction Black				0.020^{***}	0.024^{***}	0.020^{***}
				(0.005)	(0.005)	(0.005)
Ln miles to CBD		-0.014^{***}	-0.012^{***}		-0.015^{***}	-0.012^{***}
		(0.002)	(0.003)		(0.002)	(0.003)
ln population		0.011^{***}	0.012^{***}		0.012^{***}	0.013^{***}
		(0.002)	(0.002)		(0.002)	(0.002)
Ln employment		0.007^{***}	0.008^{***}		0.008^{***}	0.008^{***}
		(0.001)	(0.001)		(0.001)	(0.001)
Ln area (miles^2)		-0.026^{***}	-0.027^{***}		-0.028^{***}	-0.028^{***}
		(0.002)	(0.002)		(0.001)	(0.002)
Fraction drive to work			-0.084^{***}			-0.084^{***}
			(0.010)			(0.010)
Ln avg Uber usage			-0.031^{***}			-0.031^{***}
			(0.003)			(0.003)
Town Fixed effects	Yes	Yes	Yes	Yes	Yes	
Observations	33,126	33,126	33,126	33,126	33,126	33,126
Adjusted R^2	0.352	0.362	0.366	0.350	0.363	0.366

Note: This table shows coefficients from regressing average local roughness cost on log median income of the Census tract, the percent of the tract population that is Black, Hispanic, or Asian, and a series of tract characteristics. Robust standard errors in parentheses, *p<0.1; **p<0.05; ***p<0.01.

In Table 8, we examine heterogeneity in the within town relationship between roughness and income and race. For each town, we create variables indicating whether that town is above the median population, the median household income, and the median fraction Black, across towns. We then interact these indicators with the tract level income and race covariates. The results suggest the within town relationship between income and better roads is stronger in less populated towns and in higher income towns.

Table 8: Tract Roughness and Tract Characteristics, Town Heterogeneity

	Dependent variable:							
	Cost (USD per mile)							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Ln tract income	-0.067^{***}	-0.022^{***}	-0.021^{***}	-0.052^{***}				
	(0.015)	(0.006)	(0.002)	(0.016)				
Ln tract income × High town population	0.044^{***}			0.042^{***}				
	(0.016)			(0.016)				
Ln tract income × High town fraction Black		-0.002		-0.012^{*}				
		(0.007)		(0.007)				
Ln tract income × High town income			-0.019^{***}	-0.020^{***}				
			(0.006)	(0.006)				
Tract fraction Black					0.195^{***}	0.053	0.014^{***}	0.092
					(0.049)	(0.066)	(0.005)	(0.080)
Tract fraction Black × High town population					-0.176^{***}			-0.145^{***}
					(0.049)			(0.050)
Tract Fraction Black × High town fraction Black						-0.032		0.065
						(0.067)		(0.068)
Tract fraction Black × High town income							0.128^{***}	0.124^{***}
							(0.022)	(0.022)
Town Fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	32,622	32,622	32,608	32,608	32,622	32,622	32,608	32,608
Adjusted R^2	0.350	0.350	0.350	0.350	0.348	0.348	0.349	0.349

Note: This table shows coefficients from regressing average local roughness cost on log median income of the Census tract, with interaction terms for three town level indicator variables, for town population, town median income, and town fraction Black. For each, the variable in the regression is an indicator for being above the median across all towns. 89% of tracts are in above median population towns, 74% are in above median percent Black population towns, and 30% are in above median income towns. Robust standard errors in parentheses, *p<0.1; **p<0.05; ***p<0.01 .

We now turn to the much larger relationships between road roughness and race and income at the town level. There is sizable heterogeneity across towns, and heterogeneity is correlated substantially with both race and income. Figure 11 panels (a) and (b) shows bar plots for towns and cities in the New York-Jersey City metro area of road cost by income and race quantiles. Within the New York metropolitan area, the gap between the richest and poorest quartiles of towns is 10 cents per mile. Panels (c) and (d) show the analogous plot for all US towns, after residualizing by MSA. We then add back the mean to facilitate comparison across plots. The gap between the richest and poorest towns is 3 cents per mile, which is somewhat smaller than in New York. The race related gaps seem somewhat smaller in these pictures, but that partially reflects the fact that the quartile with the highest percent Black contains all of those towns where more than 22 percent of the population is Black.

Figure 11: Road Quality by Town Income and Percent Black

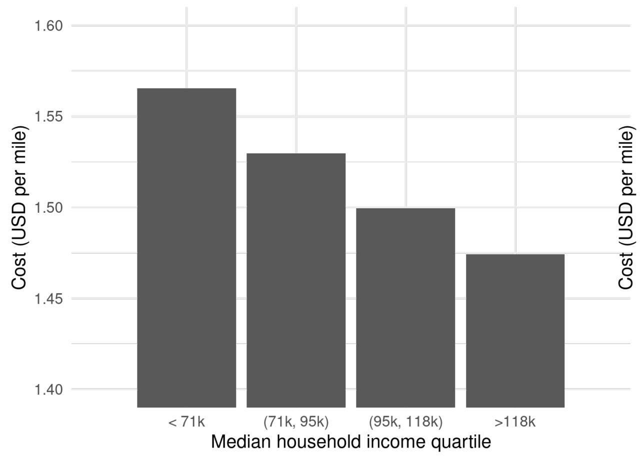

(a) New York-Jersey City Metro Area

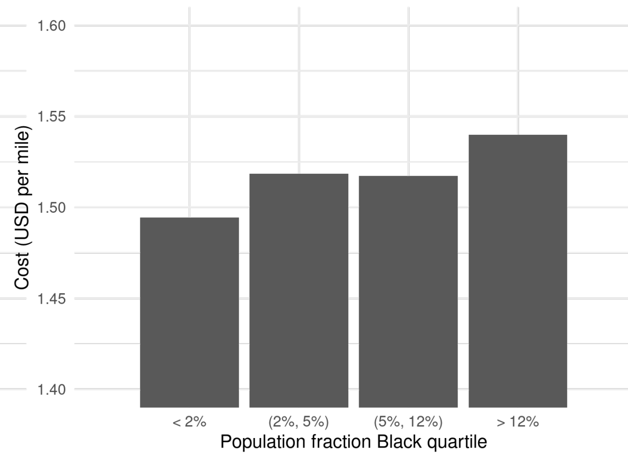

(b) New York-Jersey City Metro Area

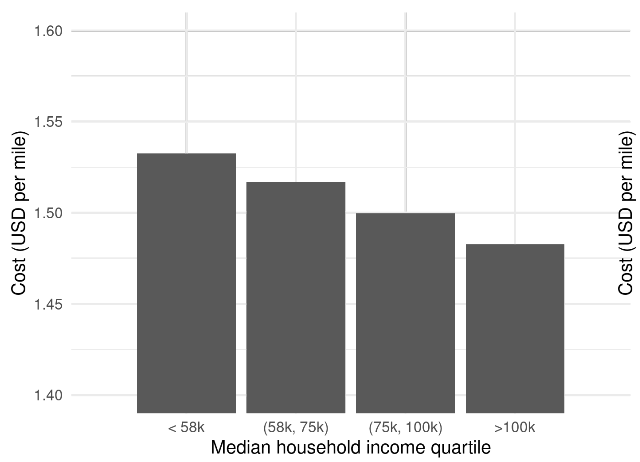

(c) Entire USA

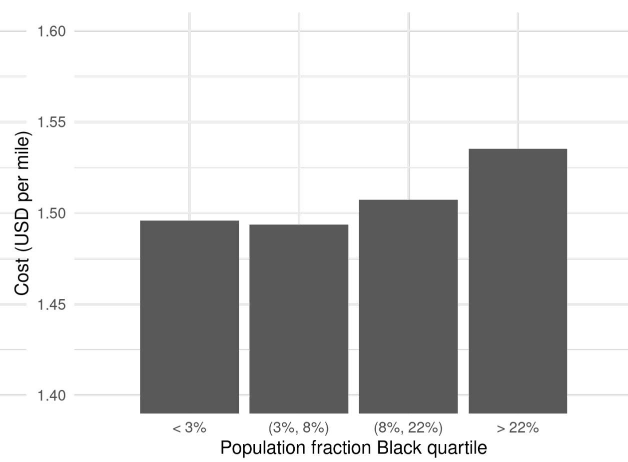

(d) Entire USA

Note: This figure shows the average per mile cost of local roads within quantile bins of income and the percentage of the population that is Black in panels (a)/(c) and (b)/(d), respectively. Panels (a) and (b) give results for the 221 towns our data covers in the New York-Jersey City MSA. Panels (c) and (d) give results for the entire USA after first residualizing by MSA fixed effects, then adding back the mean..

Table 9 examines the town-level correlates of roughness, including reported government spending on roads taken from the 2017 Annual Survey of State and Local Government Finances (ASLGF). The ASLGF is a voluntary survey that collects data on local government revenues, debt, and expenditures on services, including road related expenditures. The full survey is conducted in 5-year increments; in 2017, roughly 90% of all local governments responded.[31]

[31] We match 96% of towns to the ASLGF. For the remaining towns, we fill in the missing value with the average

Local roads are primarily maintained by local governments – state governments own only 4% of urban local roads. When counties and townships report road spending in the ASLGF, we distribute county-level spending to towns based on overlapping area. We use only towns with a population of at least 5,000 and for which we have Uber data on at least 10% of local road segments. Appendix Table A.12 shows that results are robust to different cut-offs.

We regress road roughness on town median income and racial and ethnic makeup, adding in town expenditure on roads as well as other town level demographic variables. Table 9 shows these results. Regression (1) shows a coefficient of $-.067$ on the log of income. Regression (2) includes the town's expenditures, which has a *positive* coefficient, presumably reflecting the fact that places with worse roads spend more. The median per capita expenditure on roads is $116 USD. In Appendix Table A.13, we regress spending on town characteristics to see which types of towns spend more. Spending per capita is increasing in median income, but there is no relationship between *total spending* and income, as wealthier towns are less populated.

Regression (3) includes a larger set of town-level covariates, including distance to the MSA's city center, population and land area. More distant towns have better roads, as do towns with lower population levels. The coefficient on log income drops to $-.043$ when we control for these variables, suggesting that a significant fraction of the infrastructure inequality is associated with richer people living further away from the city center in lower density locales.

Regression (4) shows a coefficient of .109 on the share of the population this is Black. This estimated coefficient changes little when we control for town-level expenditure in regression (5), and only falls to .084 controlling for our wider set of covariates in regression (6). The town-level correlation between race and road roughness is not explained by spending levels or density levels or proximity to the city center. Regression (7) does show that the coefficient on race drops to .073 when we control for income. Appendix Table A.13 shows that road expenditure per capita is weakly lower in towns with a higher share of Black households.

These results suggest that some fraction of the connection between race and roughness is associated with income and some fraction of the connection between income and roughness is associated with race. But as we have no clear source of exogenous variation for either variable, we cannot divide the impact of the two variables in any definitive manner. Richer people live in lower density areas and towns, which more distance to the city center, and that appears to explain some portion of the link between income and roughness. While Black households are also less suburban, controlling for density and proximity to the city center does little to reduce the coefficient on fraction Black.

spending. Appendix Figure A.16 shows a histogram of 2017 road related expenditure by local governments in our final sample. While the median expenditure is 3.2 million, the tails are wide, with both a significant mass at zero and New York City reporting 2.6 billion.

Table 9: Town Roughness and Characteristics

	Dependent variable:						
	Cost (USD per mile)						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Ln median income	-0.067^{***}	-0.071^{***}	-0.043^{***}				-0.040^{***}
	(0.009)	(0.009)	(0.008)				(0.011)
Fraction Black				0.109^{***}	0.113^{***}	0.084^{***}	0.073^{***}
				(0.016)	(0.016)	(0.015)	(0.019)
Fraction Hispanic				0.071^{***}	0.079^{***}	0.009	0.030
				(0.017)	(0.017)	(0.018)	(0.020)
Fraction Asian				−0.018	−0.008	−0.047	−0.023
				(0.033)	(0.033)	(0.030)	(0.033)
Ln expenditure per capita		0.006^{***}	0.007^{***}		0.005^{**}	0.005^{**}	
		(0.002)	(0.002)		(0.002)	(0.002)	
Miles to CBD			-0.035^{***}			-0.033^{***}	
			(0.004)			(0.005)	
Ln population			0.030^{***}			0.037^{***}	
			(0.006)			(0.006)	
Ln employment			−0.001			0.005	
			(0.004)			(0.004)	
Ln area (miles^2)			-0.029^{***}			-0.041^{***}	
			(0.006)			(0.006)	
Fraction drive to work			-0.209^{***}			-0.176^{***}	
			(0.032)			(0.033)	
Ln avg Uber usage			−0.010			−0.009	
			(0.010)			(0.010)	
MSA Fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1,235	1,235	1,235	1,235	1,235	1,235	1,235
Adjusted R^2	0.206	0.210	0.369	0.209	0.213	0.378	0.217

Note: This table shows coefficients from regressing average local roughness on town on log median characteristics of the town. We limit the data to towns with a population > 5,000, and where our coverage of local roads is greater than 10%. The outcome is cost per mile. The CBD location is defined as city hall location for the largest city in the MSA. Missing or zero town expenditure on roads is replaced with the mean expenditure, and an indicator variable is included for towns with imputed data. Robust standard errors in parentheses, *p<0.1; **p<0.05; ***p<0.01.

6 Cities' Road Repair Decisions

6.1 Local Road Resurfacing Targeting in Eleven Large Cities

How do cities decide which roads to repair? To find out, we collected road-segment level data on actual repair decisions for eleven large cities. We also surveyed over 100 towns and cities on their road data collection and repair strategies.

We collected data on road resurfacing that happened between August, 2021 (the period covered by our road roughness data) and January 1, 2023. We searched the top 25 cities by population for public data for our period. We were able to collect data from the following cities: Charlotte, NC, Chicago, IL, Columbus, OH, Dallas, TX, Denver, CO, Los Angeles, CA, New York City, NY, Phaldelphia, PA, Portland, OR, San Diego, CA, San Jose, CA. Data details and city-specific idiosyncrasies are described in Appendix section C.12.

For each city, we remove preventative work, such as slurry seals, as well as repairs on arterial roads,[32] and aggregate all data at the level of small grid cells.[33] With each grid cell we compute (1) the total length of local roads, (2) the total length of local roads repaved, (3) Uber road roughness cost of the segments in the grid, and (4) interpolated demographic variables from the American Community Survey (population, income, share of Black residents). We do not have Uber data everywhere. For example, 92% of grids within New York City have some Uber coverage, but only 37% do in Dallas. Figure A.18 shows a map of coverage of Dallas, where covered areas tend to be in the urban core. Since we cannot include grids where we have no Uber data, the analysis therefore asks whether, within the set of roads in the area of the city covered by Ubers, worse roads are resurfaced first.

Figure 12 shows repairs by roughness. For each city, we group grids into deciles by average road roughness, computed using our Uber data. The bar height indicates the share of road length within the decile bin that gets repaired. First, the plots show substantial variation in the *amount* of repairs undertaken by different cities. For example, New York City reported 731 miles of total repairs (including arterial roads), while Portland reported only 36 miles.[34]

The plots for New York City, Denver, and Philadelphia, show a clear correlation between resurfacing work and road roughness. However, the overall effect is small. For example, in New York City, the resurfacing share around 6-7% for the roads in grids with the smoothest roads at baseline, and around 14% for roads in grids with the roughest roads. These results suggest that resurfacing work is partly targeted to rough roads, yet this only explains a small share of which road segments are selected for repairs. We fail to see systematic patterns for the other cities.

[32] We classify repair segments as arterial roads if the line-segment provided by the city is contained in a 50 meter buffer of the network of OSM arterial roads.

[33] Grid cells are coordinate degrees rounded to .005. In New York City, this is approximately equal to .4km.

[34] We contacted the public work departments of Charlotte, Columbus, Portland, and San Diego, the four cities with the fewest reported repairs. We received no response from Columbus and San Diego, but Portland and Charlotte confirmed their data quality, and Portland noted that the jurisdiction was facing significant funding issues.

Figure 12: Local Road Roughness and Resurfacing Decisions in Eleven Cities

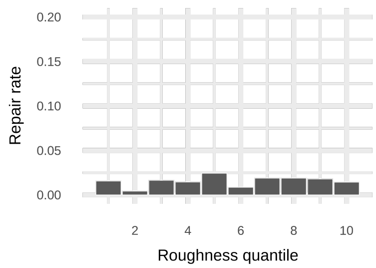

(a) Charlotte, NC

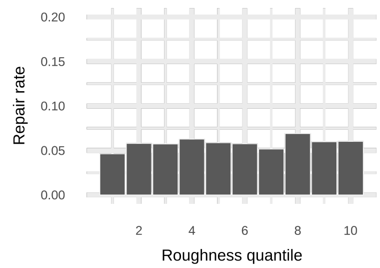

(b) Chicago, IL

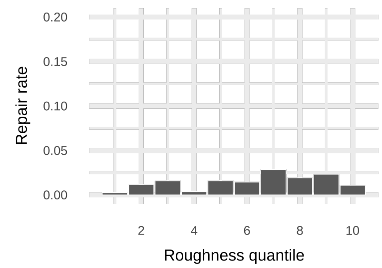

(c) Columbus, OH

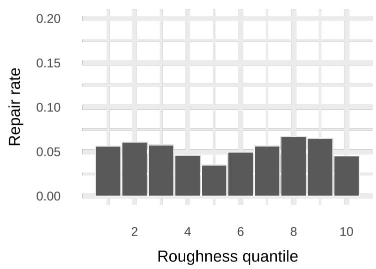

(d) Dallas, TX

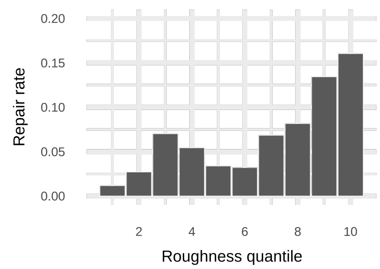

(e) Denver, CO

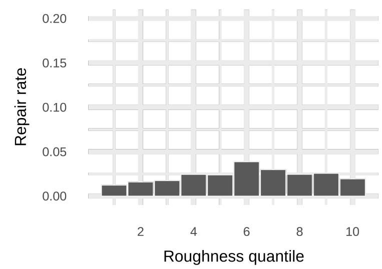

(f) Los Angeles, CA

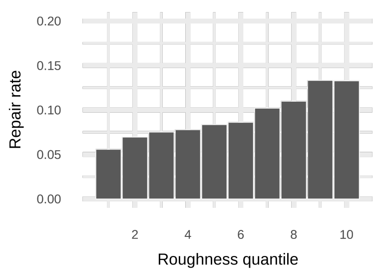

(g) New York City, NY

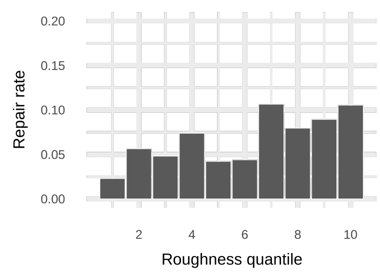

(h) Philadelphia, PA

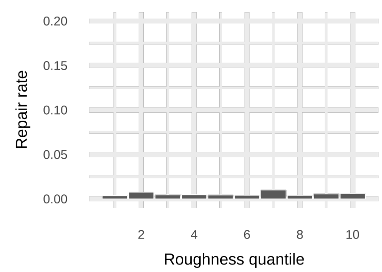

(i) Portland, OR

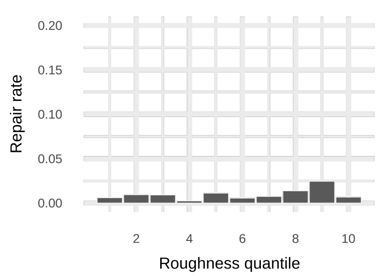

(j) San Diego, CA

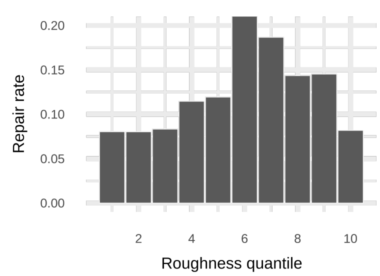

(k) San Jose, CA

Note: This figure shows the share of local road length repaired within each decile of grid cell roughness cost. Arterial and highways are excluded. Roughness cost is at the grid cell level, and is the average over the segments with Uber data within the cell.

Three theories could explain the weak correlation between roughness and repaving. First, there may be other unmeasured benefits of repaving particular roads first, such as high levels of usage. Second, repaving decisions may reflect political favoritism, perhaps towards the wealthy. Third, the local public capacity to target resources may be limited. We explore these hypotheses in this section and with the policymaker survey that follows.

What other factors explain road resurfacing decisions? Cities may prioritize high-traffic areas, as well as certain types of neighborhoods, speaking to the first two hypotheses above. Table A.14 reports linear regressions where the outcome is the repair rate in the grid. Uber road roughness predicts the road repair rate in New York City and Denver, but in none of the other cities. Higher population areas generally receive more repairs (excluding Denver), and higher income areas are also more likely to have a positive relationship with repair targeting (excluding San Diego). There is substantial variation across city for the other covariates. For example, repairs in Philadelphia are positively correlated with distance from downtown and negatively correlated with the fraction of households that are Black, while repairs in San Jose are positively correlated with the fraction of households that are Black.

Our measure of roughness might fail to predict repairs if there is selection bias due to incomplete coverage. In Table A.15, we present analogous tables where we remove grids with less than 20% and 30% coverage by Uber. In these results, New York and Denver continue to have a detectable relationship, and depending on the cutoff, San Diego and Philadelphia, also have statistically significant positive slopes, although the slope is very small for San Diego.

We test for the possibility that measurement error is driving the null results for cities other than New York, where we have much less coverage, by creating a sub-sample for New York City where we throw away observations until its distribution of coverage looks similar that of Dallas. We then run the same regressions as in Table A.14 on this sub-sample. The coefficient on roughness falls by a third but remains highly statistically significant. Details and results are reported in Appendix Table A.16.

6.1.1 A model of prioritizing road resurfacing

In this section, we estimate what share of a city's road resurfacing decisions can be explained by simple strategies that minimize the costs of roughness to drivers. We consider two benchmark strategies: prioritizing the road segments in worse condition, and prioritizing based on the product of segment traffic and cost reduction due to the repair, which may capture total cost savings from repairs. To implement the second strategy, we compute for each road segment r

$$ValueRepair_r = \left(c_r^0 - c_r^1\right) \times Traffic_r$$

We use the transformed measure of roughness (in dollars) defined in Section 4, for current cost c_r^0, and we set the cost after repair c_r^1 to be the 10% percentile of the cost distribution.

The measure $ValueRepair_r$ can be interpreted approximately as the marginal benefit to drivers from improving segment r. In general, the value of repaving a segment also depends on the availability of close substitute routes that drivers can take to avoid that segment. When drivers optimize their routes, the direct effect on current users – ignoring switchers – captures the first-order value of repaving.[35] Under this interpretation, $ValueRepair_r$ is also approximately the value of repairing segment r in a competitive equilibrium where road roughness determines road-based travel and trade costs, drivers choose optimal routes, and economic activity depends on travel costs (Hulten, 1978). This measure does not include external costs such as higher crash rates due to smoother, faster roads, as discussed in Section 4.5.2. Our measure of road segment traffic $Traffic_r$ is the number of Uber trips over that road segment, which may not represent overall road traffic patterns.

For both targeting strategies, we rank all segments in the city in decreasing order of value from repaving, and we assume that the planner repairs the same amount of total length of road as in reality, working down this list. To measure the weight that city repair decisions put on our counterfactual policies, we estimate the following linear regression

$$r_i = \alpha + \lambda r_i^* + \epsilon_i,$$

where r_i is the share of road length in grid i that is repaired in our data, and r_i^* is the share of road length repaired under either of our counterfactual policies. The coefficient λ measures the degree to which the counterfactual policy on average predicts the government's behavior.

Table 10 reports estimates of λ.[36] In Panel A r_i^* is the worst-first counterfactual, and Panel B r_i^* is the use-weighted welfare optimizing counterfactual. New York City's repair decisions are consistent with a weight of 9-16% on model-implied targeting, and the rest due to orthogonal factors. We also find positive correlations in at least one panel for Columbus, Denver, Philadelphia, Charlotte, and Chicago, but no statistically significant relationships for the other five cities. Figure A.19 visualizes these results by plotting each city's actual targeting (from Figure 12) next to the implied repair rates from these counterfactual policies. Both policies target significantly more roads in the top decile of roughness than real cities do, and significantly fewer roads in the bottom four deciles of roughness. However, both policies have also broad coverage across grid roughness deciles, highlighting that there is substantial heterogeneity in road segment condition within grids.

Our measure of road traffic uses Uber trips may not be representative, especially if Uber riders are unusually wealthy. In Figure A.20, we re-estimate the counterfactual policy for New York City after first residualizing Uber trips by the median household income per grid. Controlling for income flattens the use-weighted cost minimizing counterfactual slightly, but it still suggests a policy much more tightly correlated with roughness than we see with actual repaving.

[35] While repairing a road segment is a large change locally, it is small for a typical trip. In our data, the median road segment length is less than 50 meters, much shorter than typical trips.

[36] Table A.17 reports results varying the coverage thresholds. The effects are similar.

Table 10: Repair Targeting versus Counterfactual Repair Targeting

Panel A: worst first

	Dependent variable:										
	Repair rate (percent road-miles repaved)										
	NYC	Dallas	Columbus	Portland	LA	San Jose	Denver	Philadelphia	Charlotte	San Diego	Chicago
Worst first	0.094***	0.030	0.00002	−0.016	−0.026	−0.003	0.164**	0.099	0.009	0.041	0.004
	(0.021)	(0.053)	(0.079)	(0.039)	(0.025)	(0.098)	(0.077)	(0.062)	(0.028)	(0.025)	(0.037)
Ln road miles	0.012***	0.038***	−0.018***	0.002	−0.005**	0.034*	−0.004	−0.005	0.006**	0.005**	−0.033***
	(0.003)	(0.007)	(0.004)	(0.001)	(0.002)	(0.017)	(0.011)	(0.004)	(0.003)	(0.002)	(0.005)
Constant	0.075***	0.119***	0.021***	0.003**	0.027***	0.100***	0.058***	0.071***	0.016***	0.010***	0.085***
	(0.004)	(0.008)	(0.005)	(0.001)	(0.002)	(0.021)	(0.012)	(0.007)	(0.002)	(0.002)	(0.004)
Observations	2,786	568	618	233	2,644	202	433	1,113	753	848	1,876
Adjusted R^2	0.011	0.047	0.023	−0.001	0.002	0.009	0.008	0.001	0.004	0.007	0.028

Panel B: use weighted

	Dependent variable:										
	Repair rate (percent road-miles repaved)										
	NYC	Dallas	Columbus	Portland	LA	San Jose	Denver	Philadelphia	Charlotte	San Diego	Chicago
Use-weighted	0.163***	−0.101	0.194*	−0.033	−0.047	−0.071	−0.068	0.161***	0.140***	−0.025	0.092***
	(0.023)	(0.104)	(0.099)	(0.075)	(0.032)	(0.084)	(0.083)	(0.058)	(0.053)	(0.047)	(0.032)
Ln road miles	0.011***	0.038***	−0.019***	0.002	−0.004**	0.031*	−0.011	−0.006	0.005**	0.005**	−0.033***
	(0.003)	(0.007)	(0.004)	(0.001)	(0.002)	(0.017)	(0.011)	(0.004)	(0.003)	(0.002)	(0.004)
Constant	0.071***	0.123***	0.019***	0.003**	0.027***	0.107***	0.073***	0.068***	0.014***	0.011***	0.081***
	(0.004)	(0.008)	(0.005)	(0.001)	(0.002)	(0.021)	(0.011)	(0.007)	(0.002)	(0.002)	(0.004)
Observations	2,786	568	618	233	2,644	202	433	1,113	753	848	1,876
Adjusted R^2	0.021	0.048	0.029	−0.001	0.002	0.012	−0.001	0.006	0.013	0.005	0.033

Note: This table shows coefficients from a linear regression of the share of roads miles in a grid that are repaved on the share repaved in one of two counterfactual policies: worst first and a use-weighted cost minimizing policy. Because the share will be mechanically extreme for grids with little Uber coverage in the counterfactual models, we drop grids with coverage under 20%.

6.2 An Administrative Survey of Road Maintenance Strategies Across the US

To better understand road maintenance strategies, we surveyed two samples of towns. First, we randomly sampled towns with Uber roughness data across the US. Second, we concentrated on towns in Massachusetts that also had Uber coverage. For each town, we contacted the Department of Public Works director or official email, or a city engineer by email. We sent follow-up emails and calls to increase the response rate, with the Massachusetts sample receiving additional follow-ups. The response rate for the national sample and Massachusetts towns is 16% and 73%, respectively. More details on data collection are included in Appendix C.13, which also examines the differences between towns that responded quickly and towns that responded only after repeated calls in the Massachusetts sample. These differences appear mild, which makes us hopeful that our US sample is also reasonably representative despite the low response rate.

Table 11 presents results on the resurfacing that actually takes place. Table 12 presents results on how towns decide which roads they should resurface. Our goal is to understand why road repaving may not be targeting the roughest roads.

The first panel of Table 11 shows respondents' assessments of what share of roads that need repaving actually get repaving. Only 20% of the Massachusetts sample and 32% of the national sample think that more than one-half of the roads that needed resurfacing get resurfacing. These numbers are low, especially given that we would expect respondents to highlight their own efficiency. We also find that 32% of the Massachusetts sample and 62% of the national sample report that less than 30% of roads that need repaving are repaved. Even if municipalities follow some objective for determining which roads "need" repaving, it seems quite plausible public works departments then select roads according to other criteria unrelated to road roughness when they among those roads. Yet, the numbers in Table 11 suggest spending scarce funds optimally is especially important given towns' inability to maintain their full road network.

The second panel in Table 11 shows that departments are allocating most of their resources to resurfacing, rather than preventative work. Only 20% of the Massachusetts sample and 31% of the national sample are spending more than 30% of their budgets on preventative work. Consequently, it seems reasonable to think that theses departments are focusing overwhelmingly on the task of rebuilding already highly deteriorated roads.

Table 11: City Road Maintenance and Repair Strategies

	Massachusetts	Rest of US
Panel A: Percent of Roads that Receive the Resurfacing they Require		
> 90%	0.08	0.06
70 - 90%	0.04	0.12
50 - 70%	0.08	0.14
30 - 50%	0.40	0.04
< 30%	0.32	0.62
Unsure	0.08	0.01
Panel B: Spending on Resurfacing vs Preventative Maintenance		
> 90% on resurfacing	0.12	0.21
70 - 90% on resurfacing	0.68	0.47
50 - 70% on resurfacing	0.16	0.14
50% on resurfacing	0.00	0.02
30 - 50% on resurfacing	0.00	0.04
< 30% on resurfacing	0.00	0.10
Unsure	0.04	0.01
Observations	25	96

Note: This table reports summary statistics for bureaucrat responses to our surveys.

Table 12 reports how towns gather information about road roughness. Panel A shows that, in both the US and Massachusetts samples, the modal source of information about road quality is a survey of roads performed by a private company; 80% and 59% of the Massachusetts and national sample, respectively. A large share also uses an in-house survey (48% and 55%, respectively), and "engineer's discretion", which indicates reliance on professional knowledge.

We also asked about two non-professional sources of information about road quality: elected officials and local citizens. While both sources of information could accurately reflect road quality, relying on these sources of information also leaves open the possibility that repaving is targeted towards more politically important or influential neighborhoods.[37] In Massachusetts, 4% of towns acknowledged a role for elected officials in targeting repaving. In the national sample, the figure was 16% of the national sample. Slightly more than one-fourth of the Massachusetts and 40% of the national sample base their assessments on input from ordinary residents. Appendix Figure A.21 shows a bar chart of use of these four information sources by town population quartiles. Use of the two non-professional sources show U-shaped relationships in population size. In contrast, use of a private firm is strictly increasing in town population.

The second panel in Table 12 shows that both samples are sharply bifurcated in the rate in which they survey roads. In both samples, almost one-half of towns survey more than 90% of their roads annually and yet about 40% of towns survey less than 50% of their roads annually. More than one-fifth of towns in both samples survey less than 30% of roads in any given year.

Panel C focuses on determining the need for repaving. Panel D repeats the exercise for prioritization, because if towns are unable to repave all of the roads that they think need repaving, then prioritization rules will determine repaving in practice. Formulas appear to be regularly used; 82% (88%) of the US (Massachusetts) sample report using a formula to determine roads that need resurfacing. 50% (32%) of US (Massachusetts) towns report their formula is more holistic than just road quality.[38] The percent of towns using a formula drops to 66% (68%) when we consider prioritization.

Upcoming utility work plays a particularly important role in determining repaving schedules, largely because utility work typically involves tearing up roads and so previous repaving effort is wasted. 60% (76%) of national (Massachusetts) respondents said that utility work helped determine need and 46% (44%) said that utility work determined prioritization. Only about one-half of both samples mentioned traffic utilization as a determinant of repaving need and prioritization. Transportation expert feedback, citizen complaint and accessibility play smaller roles in shaping assessment of repaving need and prioritization. About one-third of the national sample notes that elected official input shapes need assessment and prioritization. The role of elected officials appears to be smaller in Massachusetts.

These results help make sense of America's rough roads and of the poor targeting of roughness that we see in a few cities. Towns only resurface a small portion of the roads that they think need to be resurfaced. While formulas play a large role in determining need, especially in Massachusetts,

[37] Rizzo et al. (2021) show that local road investment in Italian cities follows electoral cycles.

[38] There are towns which said yes to using only road quality based formulas and more holistic formulas, which suggested that they either have two formulas or misunderstood the question.

they play a smaller role in determining prioritization. Other forces, especially the influence of elected officials in the national sample, help to determine which roads actually get resurfaced among the set of roads that need resurfacing. These results are compatible with our previous findings that suggested that very smooth roads are less likely to be resurfaced, but that there was little targeting among rougher roads.

Table 12: City Road Maintenance and Repair Strategies

	Massachusetts	Rest of US
Panel A: Method(s) to determine the state of roads		
Road survey (in-house)	0.48	0.55
Road survey (contracted to private firm)	0.80	0.59
Elected officials	0.04	0.16
Engineering discretion	0.44	0.51
Reporting by residents	0.28	0.40
Panel B: Percent of Roads Surveyed each Year		
> 90%	0.44	0.46
70 - 90%	0.00	0.04
50 - 70%	0.08	0.07
30 - 50%	0.16	0.14
< 30%	0.24	0.21
Unsure	0.00	0.03
Panel C: Criteria to Determine which Roads need Resurfacing		
A formula only considering road conditions	0.72	0.51
A formula considering other selected factors	0.32	0.50
Upcoming utility work	0.76	0.60
Traffic intensity	0.48	0.53
Citizen complaints	0.28	0.51
Transportation expert feedback	0.16	0.16
Accessibility	0.16	0.12
Elected official input	0.12	0.26
Panel D: Criteria to Determine which Roads are Prioritized		
A formula only considering road conditions	0.48	0.38
A formula considering other selected factors	0.28	0.42
Upcoming utility work	0.44	0.46
Traffic intensity	0.56	0.43
Citizen complaints	0.20	0.35
Transportation expert feedback	0.12	0.12
Accessibility	0.12	0.07
Elected official input	0.20	0.29
Observations	25	96

Note: This table reports summary statistics for bureaucrat responses to our survey. Questions in panels A, C and D accept multiple answers.

7 Conclusion

In this paper, we used vertical accelerometer data from Uber to measure road roughness across America's local roads. This data appears to correlate well with Department of Transportation IRI data, when such data is available, and it shows the same regional patterns as IRI data. The coasts are much rougher than the interior of the country. Big cities and colder winter places also have rougher local roads, just as they have rougher highways.

The central empirical exercise of the paper is to measure the willingness to pay for smoother roads, by looking at how speeds adjust to salient road roughness. The crucial assumption is that drivers care about lost time and experienced bumpiness, which is a function of both speed and road roughness. By focusing on breaks in road roughness at town borders, we estimate that a one standard deviation increase in roughness is associated with a 33 cents per mile increase in the cost of driving. This cost scales with the opportunity cost of time, which we fix at 15 dollars per mile.

We then use this cost estimate to calculate the costs that poor roads impose on minority and low-income households, whom our data suggests systemically live with worse infrastructure. A typical household living in a neighborhood that is 100 percent Black and that drives 3,000 miles per year on local roads will experience 450 dollars in harm per year because of rough local roads relative to a typical household that lives in a 100 percent White neighborhood. A doubling of town income, such as going from the 20th to the 80th percentile of towns in our sample, is associated with a 140-dollar-per-year decrease in pain due to road roughness. In both cases, the bulk of the effect occurs across rather than within towns. Within towns, the neighborhood-level link between race or income and roughness is stronger in rich towns and weaker in large towns.

Finally, we look at the correlation between roughness and road repaving in eleven American cities. We find the road roughness and usage seem to play some role in the repaving practices of New York, Denver, Philadelphia, and Charlotte, but not in the other cities. We also surveyed departments of public works and found that many of these departments do not have the resources that they need to appropriately repave. Our results suggest that within a class of roads that need some work, departments may be doing a weak job targeting repaving.

Measuring the costs of road roughness can help policymakers decide on the optimal amount of overall road repaving. These measures can also help to better target road repaving. The cost of road roughness should also inform decisions about the quality of roads when they are built. We hope that future work will improve our measures and will do more to integrate measures of road roughness into the broader problem of optimal repaving and infrastructure provision more broadly. We particularly hope that future work will do more to estimate the external effects of rough roads that are not included in our measure of willingness to pay with time for a smoother ride.

References

Akbar, Prottoy A, Victor Couture, Gilles Duranton, and Adam Storeygard, "The fast, the slow, and the congested: Urban transportation in rich and poor countries," Technical Report, National Bureau of Economic Research 2023.

Aleadelat, Waleed, Khaled Ksaibati, Cameron HG Wright, and Promothes Saha, "Evaluation of pavement roughness using an android-based smartphone," *Journal of Transportation Engineering, Part B: Pavements*, 2018, *144* (3), 04018033.

Apte, Joshua S, Kyle P Messier, Shahzad Gani, Michael Brauer, Thomas W Kirchstetter, Melissa M Lunden, Julian D Marshall, Christopher J Portier, Roel CH Vermeulen, and Steven P Hamburg, "High-resolution air pollution mapping with Google street view cars: exploiting big data," *Environmental science & technology*, 2017, *51* (12), 6999–7008.

ASCE, "Infrastructure Report Card: Report Card History," Technical Report 2017. https://www.infrastructurereportcard.org/making-the-grade/report-card-history/.

_ , "Failure to Act: Current Investment Trends in our Surface Transportation Infrastructure," Technical Report 2020. https://www.asce.org/uploadedFiles/Issues_and_Advocacy/Infrastructure/Content_Pieces/Failure-to-Act-Surface-Transpo-Preliminary-Final.pdf.

Ashenfelter, Orley and Michael Greenstone, "Using mandated speed limits to measure the value of a statistical life," *Journal of political Economy*, 2004, *112* (S1), S226–S267.

Benthem, Arthur Van, "What is the optimal speed limit on freeways?," *Journal of Public Economics*, 2015, *124*, 44–62.

Bock, Margaret, Alexander Cardazzi, and Brad R Humphreys, "Where the rubber meets the road: Pavement damage reduces traffic safety and speed," Technical Report, National Bureau of Economic Research 2021.

Brooks, Leah and Zachary Liscow, "Infrastructure costs," *American Economic Journal: Applied Economics*, 2023, *15* (2), 1–30.

Brueckner, Jan K, Jacques-François Thisse, and Yves Zenou, "Why is central Paris rich and downtown Detroit poor?: An amenity-based theory," *European economic review*, 1999, *43* (1), 91–107.

Correia, Sergio, "REGHDFE: Stata module to perform linear or instrumental-variable regression absorbing any number of high-dimensional fixed effects," Statistical Software Components, Boston College Department of Economics July 2014.

Cutler, David M, Edward L Glaeser, and Jacob L Vigdor, "The rise and decline of the American ghetto," *Journal of political economy*, 1999, *107* (3), 455–506.

DeAngelo, Gregory and Benjamin Hansen, "Life and death in the fast lane: Police enforcement and traffic fatalities," *American Economic Journal: Economic Policy*, 2014, *6* (2), 231–257.

Duranton, Gilles, Geetika Nagpal, and Matthew Turner, "Transportation Infrastructure in the US," *working paper*, 2020.

Elvik, Rune, "Speed and road safety: synthesis of evidence from evaluation studies," *Transportation Research Record*, 2005, *1908* (1), 59–69.

_ , Anna Vadeby, Tove Hels, and Ingrid Van Schagen, "Updated estimates of the relationship between speed and road safety at the aggregate and individual levels," *Accident Analysis & Prevention*, 2019, *123*, 114–122.

Fu, Ellen, Lyndsey Rolheiser, Christopher Severen et al., "The problem has existed over endless years: Racialized difference in commuting, 1980–2019," *Journal of Urban Economics*, 2023, p. 103542.

Glaeser, Edward L, Scott Duke Kominers, Michael Luca, and Nikhil Naik, "Big data and big cities: The promises and limitations of improved measures of urban life," *Economic Inquiry*, 2018, *56* (1), 114–137.

Gobillon, Laurent, Harris Selod, and Yves Zenou, "The mechanisms of spatial mismatch," *Urban studies*, 2007, *44* (12), 2401–2427.

Government Accountability Office, "Analysis of Available Data Could Better Ensure Equitable Pavement Condition," Report to Congressional Committees GAO-22-104578, Government Accountability Office 2022.

Grouios, George, Efthymios Ziagkas, Andreas Loukovitis, Konstantinos Chatzinikolaou, and Eirini Koidou, "Accelerometers in Our Pocket: Does Smartphone Accelerometer Technology Provide Accurate Data?," *Sensors*, 2022, *23* (1), 192.

Hershfield, David M, "Freeze-thaw cycles, potholes, and the winter of 1977–78," *Journal of Applied Meteorology (1962-1982)*, 1979, pp. 1003–1007.

Hulten, Charles R, "Growth accounting with intermediate inputs," *The Review of Economic Studies*, 1978, *45* (3), 511–518.

Rizzo, Leonzio, Massimiliano Ferraresi, and Riccardo Secomandi, "Electoral incentives, investment in roads, and safety on local roads," Working Papers 20210710, University of Ferrara, Department of Economics December 2021.

Rubin, Donald B, "The bayesian bootstrap," *The annals of statistics*, 1981, pp. 130–134.

Sayers, M. W., T. D. Gillespie, C. A. V. Queiroz, and University of Michigan, "International road roughness experiment: Establishing methods for correlation and a calibration standard for measurements," World Bank 1986.

Small, Kenneth A and Clifford Winston, "Optimal highway durability," *The American economic review*, 1988, *78* (3), 560–569.

Winston, Clifford and Quentin Karpilow, *Autonomous Vehicles: The Road to Economic Growth?*, Brookings Institution Press, 2020.

Appendix - For Online Publication

Infrastructure Inequality: Who Pays the Cost of Road Roughness?

Lindsey Currier, Edward L. Glaeser and Gabriel E. Kreindler

December 15, 2023

Contents

A Appendix Figures

Figure A.1: Linear and Local Linear Speed Slopes

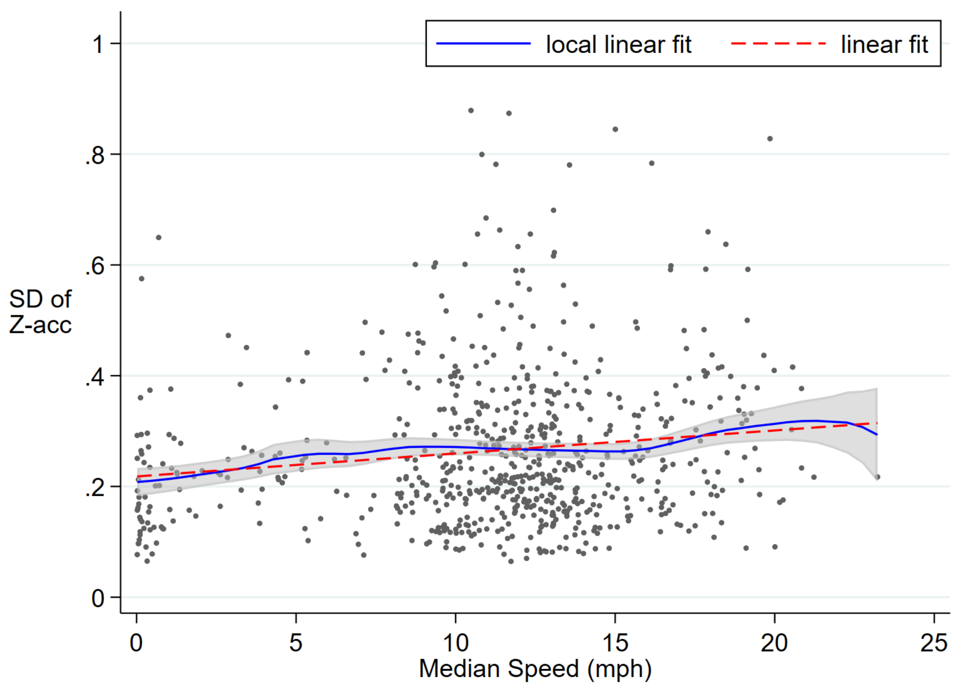

Note: This graph shows the relationship between the standard deviation of Z-acceleration and median speed. Each observation is at the driver-segment level. The sample is the road segment with the largest number of observations in April 2018 (after we keep no more than 33 observations per road segment per day). The blue, solid line is a local linear fit with 95% confidence intervals in gray. The red, dashed line is the linear fit. This graph suggests that the relationship between speed and the standard deviation of vertical acceleration is approximately linear.

Figure A.2: Segment- and Driver-specific Speed Slopes

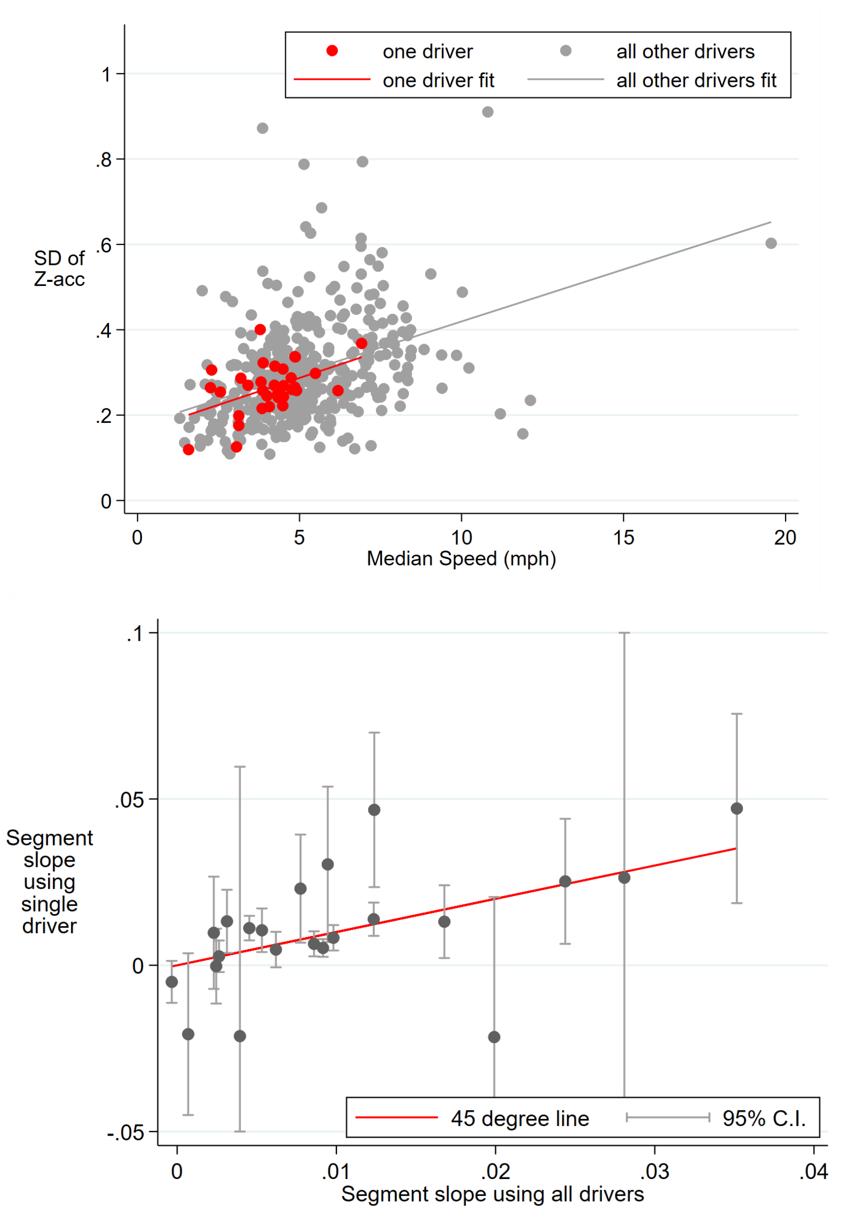

Note: These graphs show that the relationship between median speed and standard deviation of Z-acceleration is for the most part not driver-specific within segment. In the top graph, we select the driver with the most observations over any given segment. The red points represent the 34 observations for the selected driver and segment (linear fit in red) and the gray points are all the other observations for that segment (linear fit in gray). (Results are virtually identical when dropping the outlier.) In the bottom graph, we select all segments with at least 20 observations from the same driver. For each such segment, we compute the segment-specific slope excluding that driver (X axis), and the segment-driver-specific slope using only the observations for that driver (Y axis). The gray lines represent 95% confidence intervals of the segment-driver-specific slope (censored at -0.05 and $+0.1$). The 45 degree line in red falls within nearly all the confidence intervals of the segment-driver-specific slopes.

Figure A.3: Cross-Validation of Segment-Level Uber Road Roughness (Chicago)

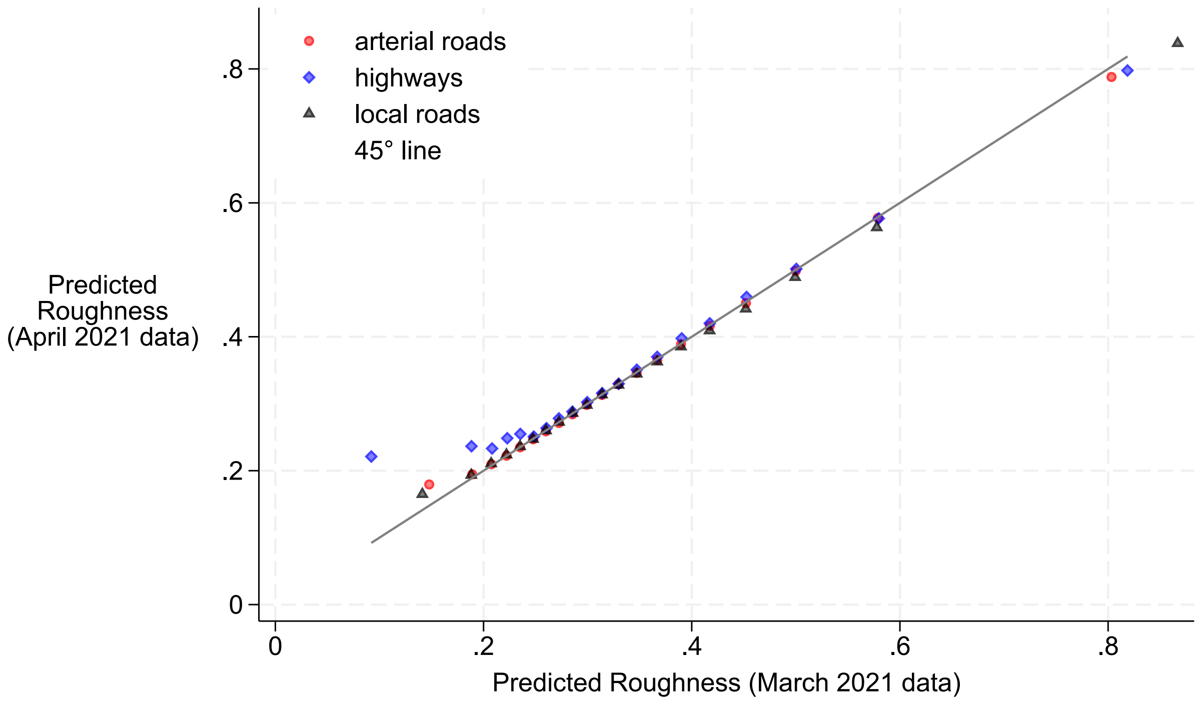

Note: This graph shows the binned scatter plot of predicted roughness at the road segment level in Chicago, using only March 2021 data (X axis), and using only April 2021 data (Y axis). Predicted roughness is computed using 20mph for local roads, 32mph for arterial roads, and 48mph for highways. See Table A.3 for additional results..

Figure A.4: Uber Measure Coverage and Population (Town level)

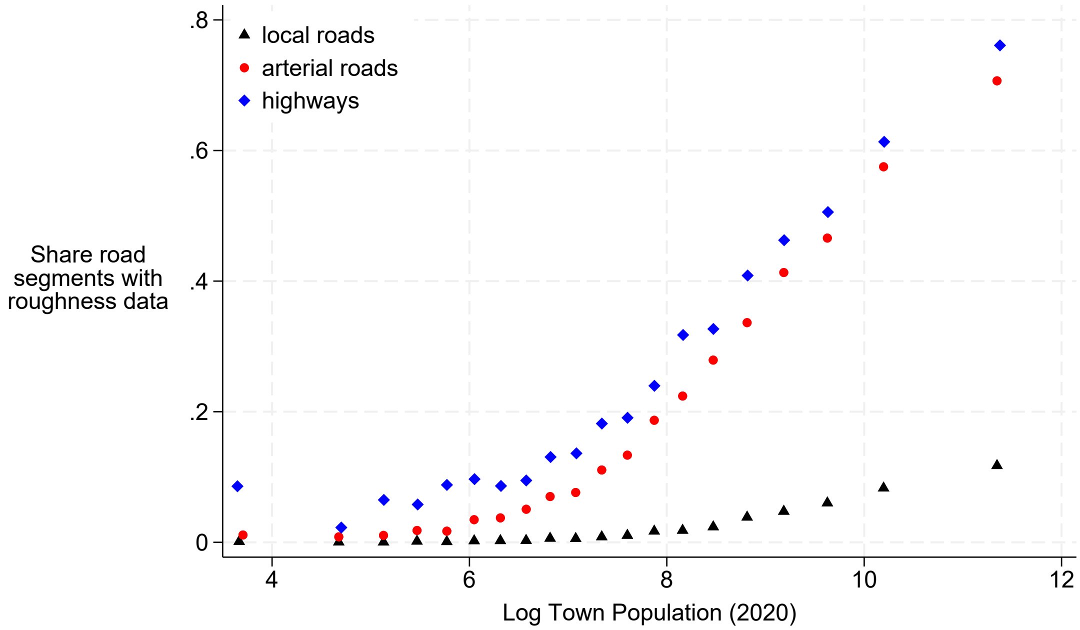

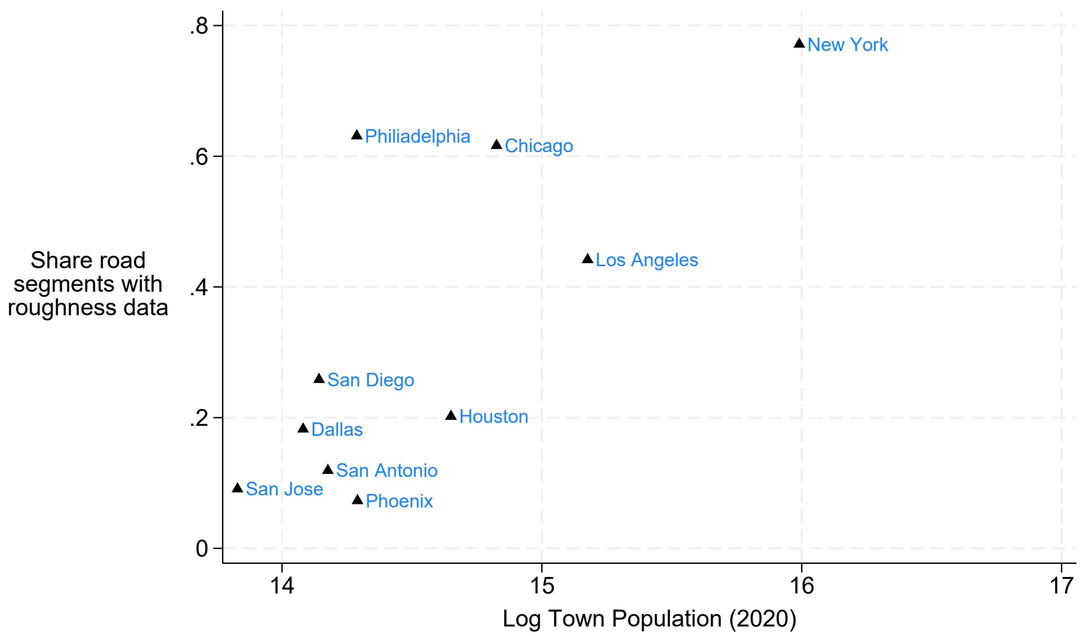

Note: The top graph shows a binned scatter plot of the share of road segments with Uber road roughness data at the town level, as a function of log town population, by road type. The bottom graph displays coverage for local roads for the census places with census population above 1 million. See Table A.1 for additional results.

Figure A.5: Number of Uber Trips per Segment and Road Roughness (Local Roads)

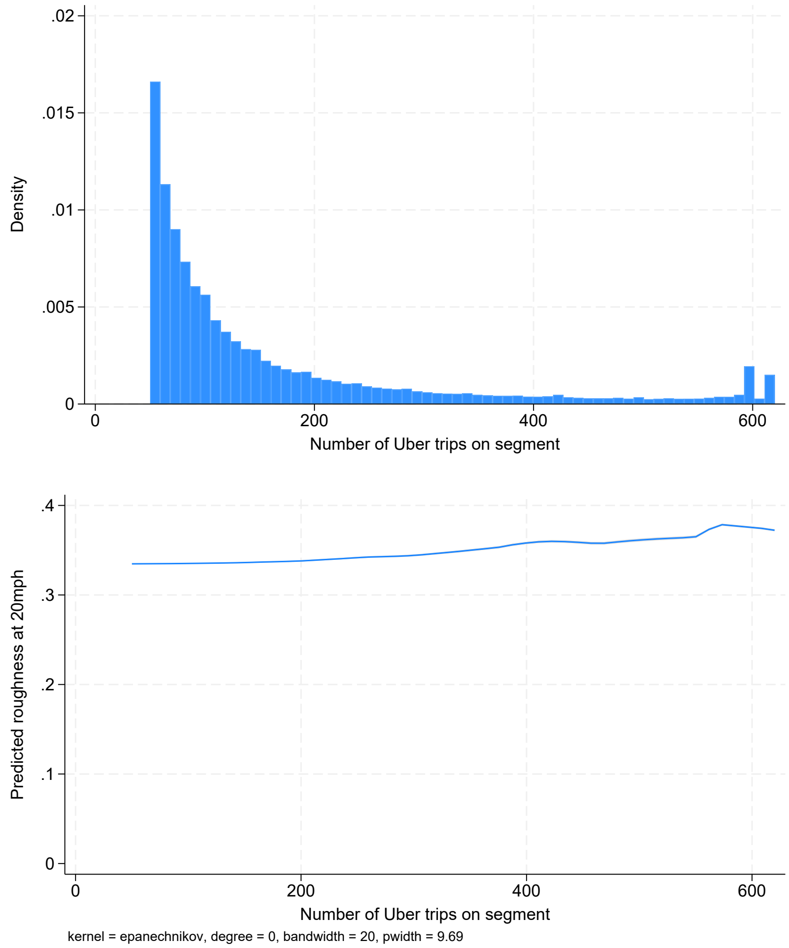

Note: The top graph shows the histogram of number of Uber trips per segment (censored at 620, which corresponds to 20 trip per day in our one month of data) in our sample of 1,829,526 local road segments. The bottom graph plots a locally linear regression of predicted road roughness at 20 mph versus trips per segment..

Figure A.6: Predicted Roughness at 48 mph in Cook County, IL

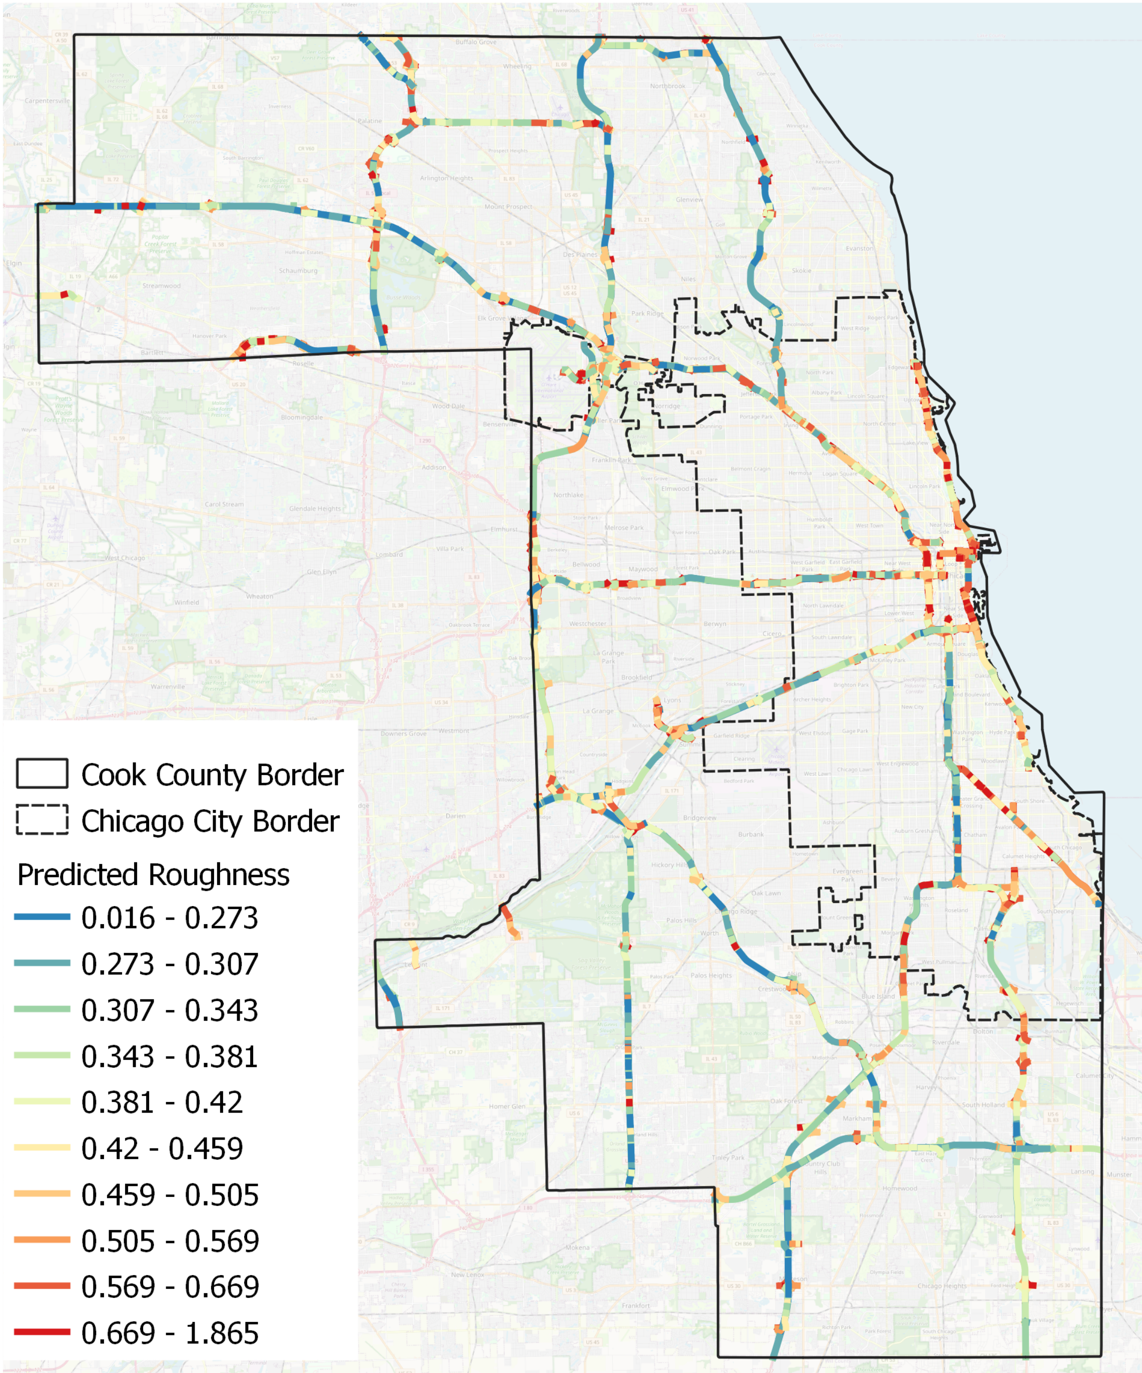

Note: This map plots predicted road roughness for all highway road segments in Cook County. Colors correspond to deciles of the roughness distribution at 48mph.

Figure A.7: Brick Roadways in the Village of Wilmette, IL

Note: The top panel highlights the high predicted roughness in the village of Wilmette in New Trier Township in Cook County, IL. (Colors correspond to deciles of the roughness distribution.) The bottom panel shows brick roadway in a Google Street View image (©Google).

Figure A.8: Local Road Roughness at MSA level (top 100 MSAs)

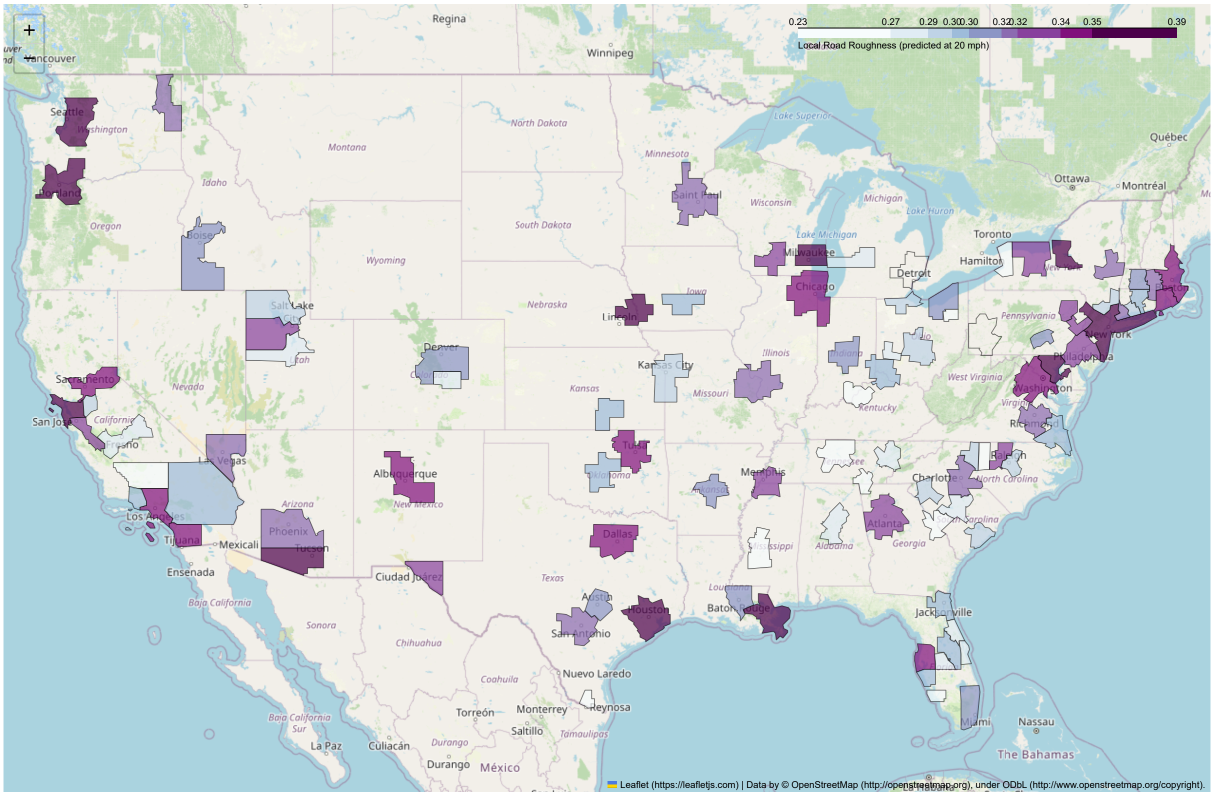

Note: This map plots the average road roughness on local roads for the top 100 MSA by population in our data. MSA-level average road roughness is winsorized at 2.5%..

Figure A.9: Arterial Road Roughness at MSA level (top 100 MSAs)

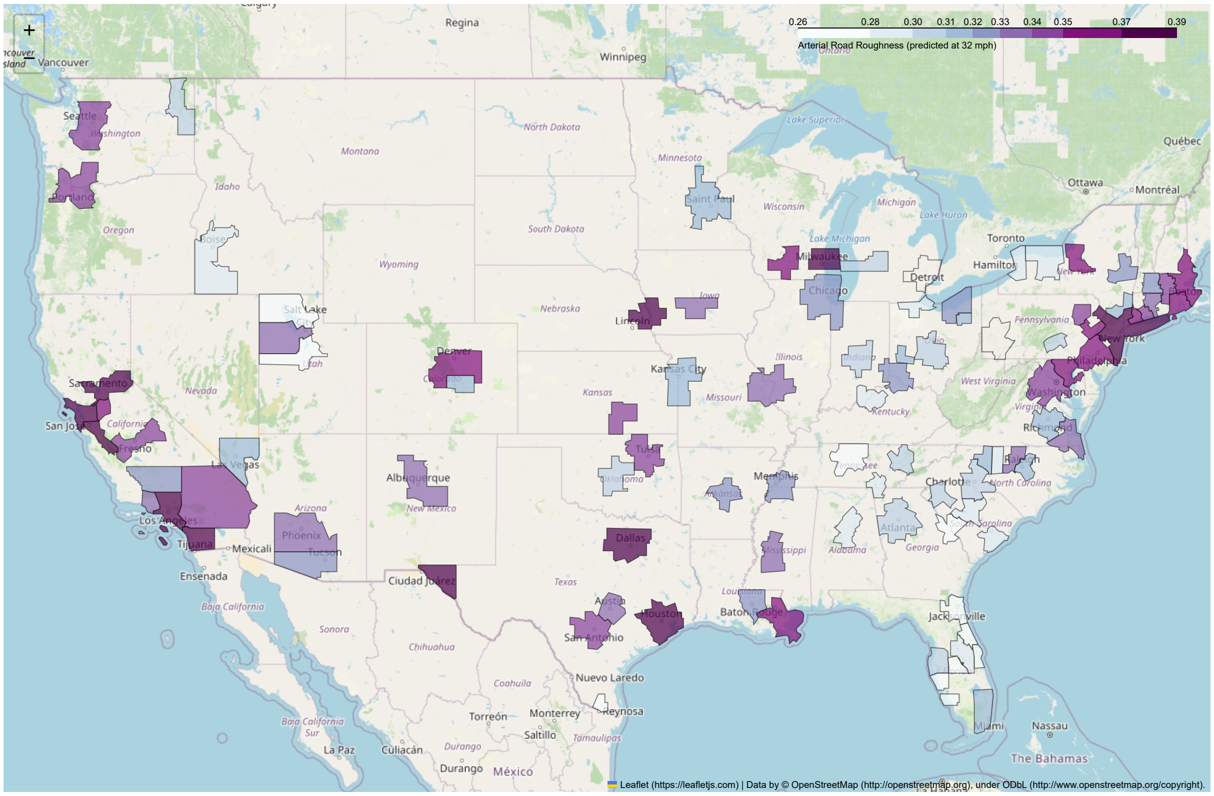

Note: This map plots the average road roughness on arterial roads for the top 100 MSA by population in our data. MSA-level average road roughness is winsorized at 2.5%..

Figure A.10: Highway Roughness at MSA level (top 100 MSAs)

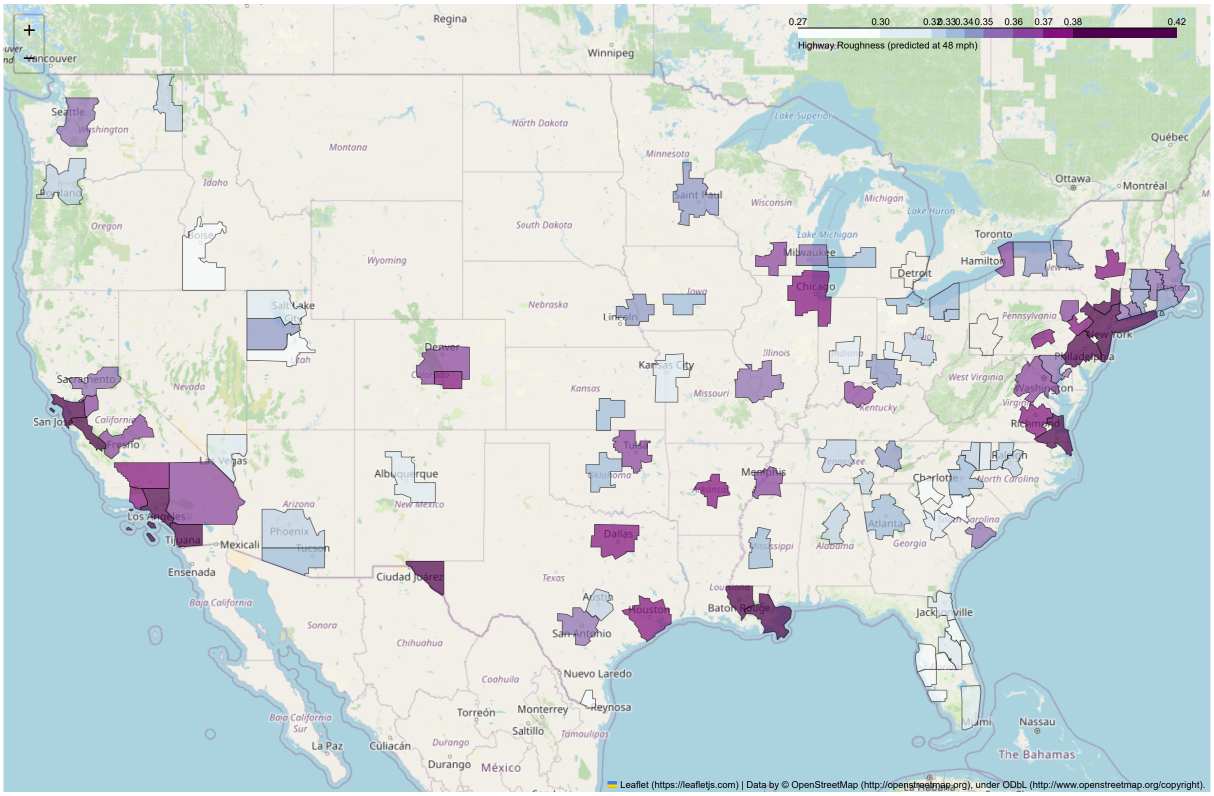

Note: This map plots the average road roughness on highways for the top 100 MSA by population in our data. MSA-level average road roughness is winsorized at 2.5%..

Figure A.11: Uber Data Has Signal: Railroad Crossing

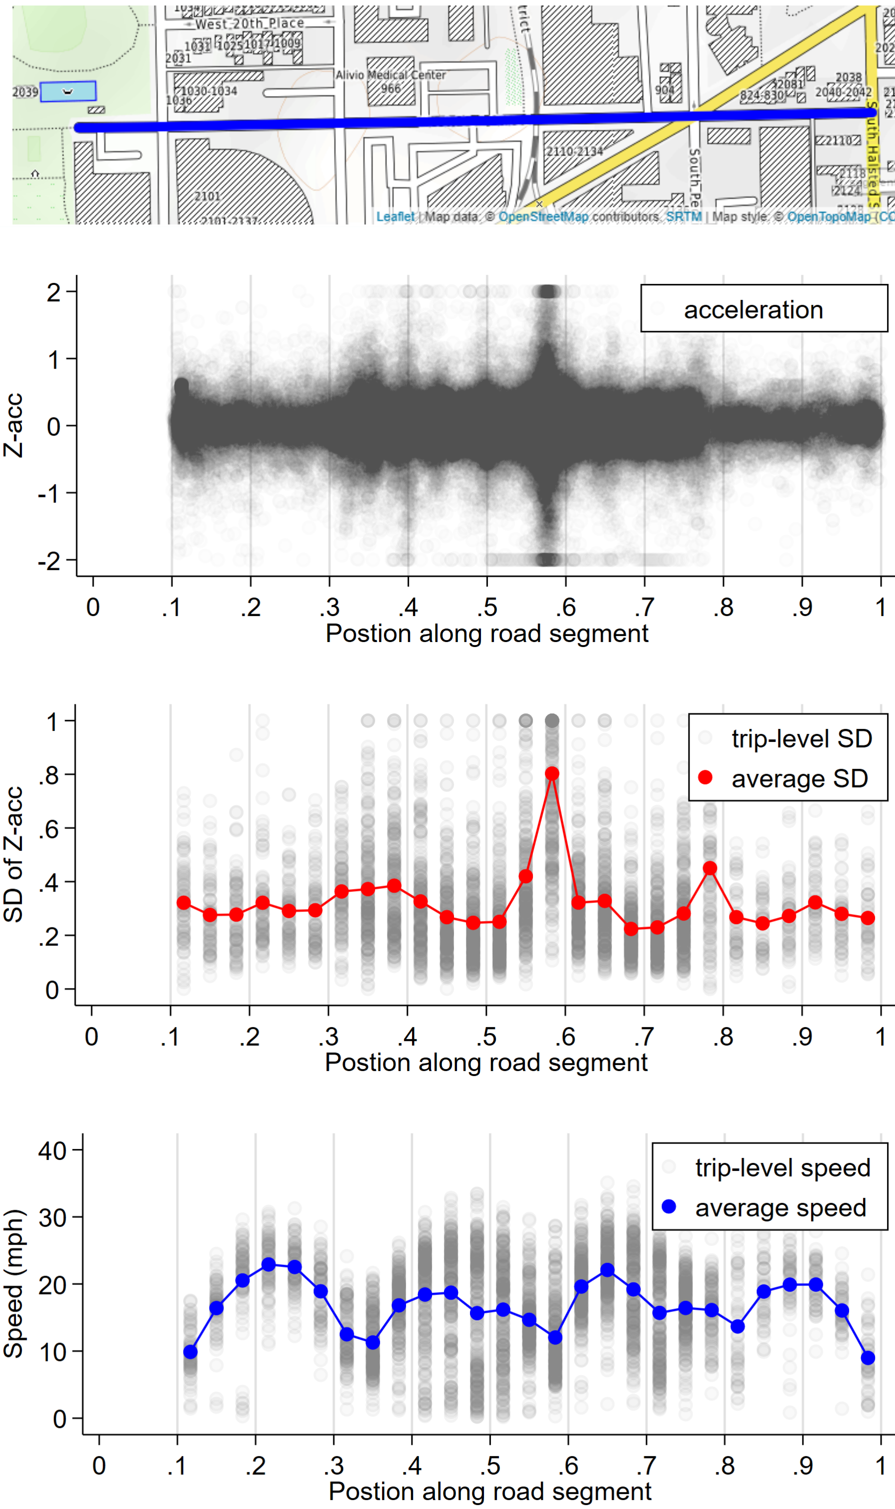

Note: These graphs show vertical acceleration, its standard deviation, and speed, for all trips covering a given road segment with a railroad crossing in Chicago (shown in blue in the map in the top panel)..

Figure A.12: Predicted Roughness and Speed around the Chicago border

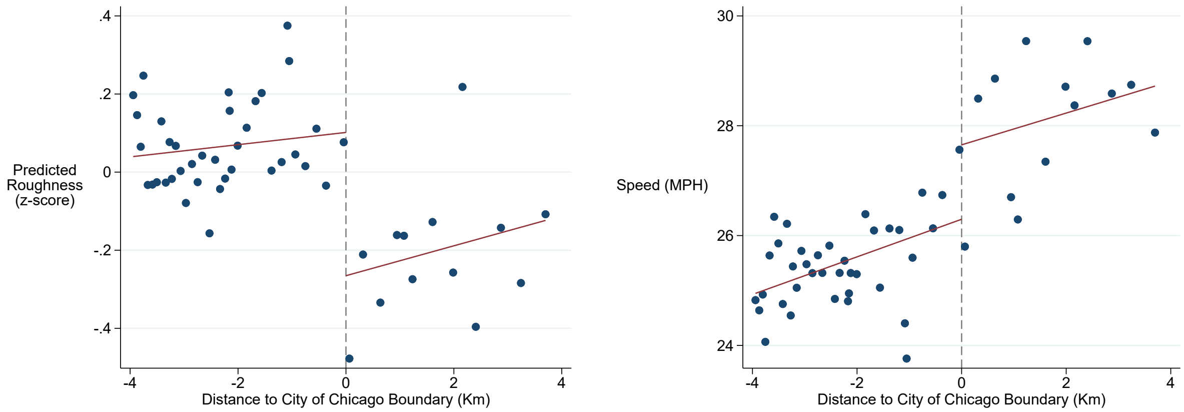

Panel (A) Predicted Roughness

Panel (B) Speed

Note: Version of Figure 7 for artery road segments..

Figure A.13: Total User Cost of Road Roughness: Simple and Full Model

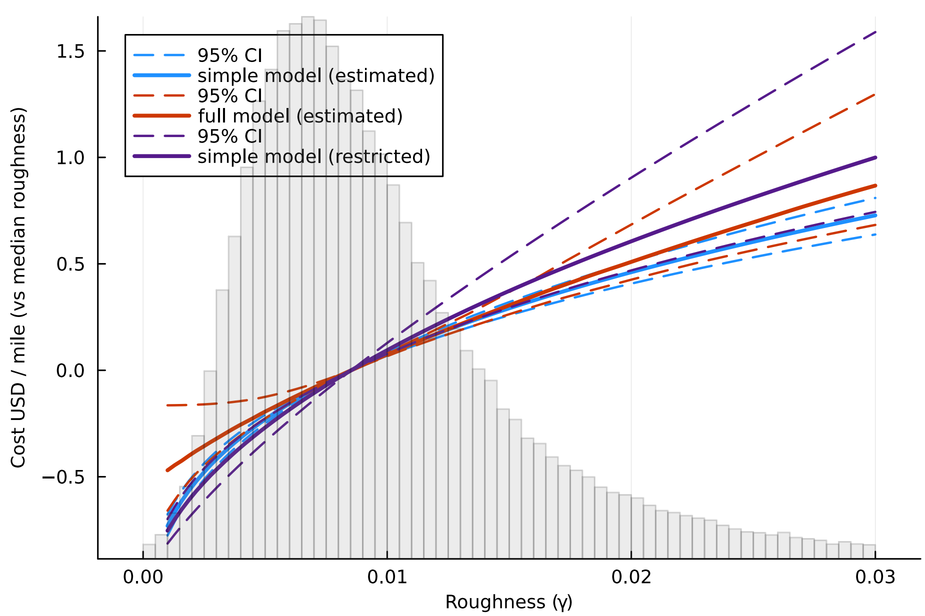

Note: This graph replicates Figure 8 for the full and simple models. The estimated simple model (Table A.8, column 2) in blue, the estimated full model (Table A.8, column 1) in red. In purple, we use the full model estimates but set $K = 0$. The costs from this model are steeper than the full model, as predicted by equation (4.3).

Figure A.14: Total User Cost of Road Roughness: Sensitivity to Roughness Elasticity

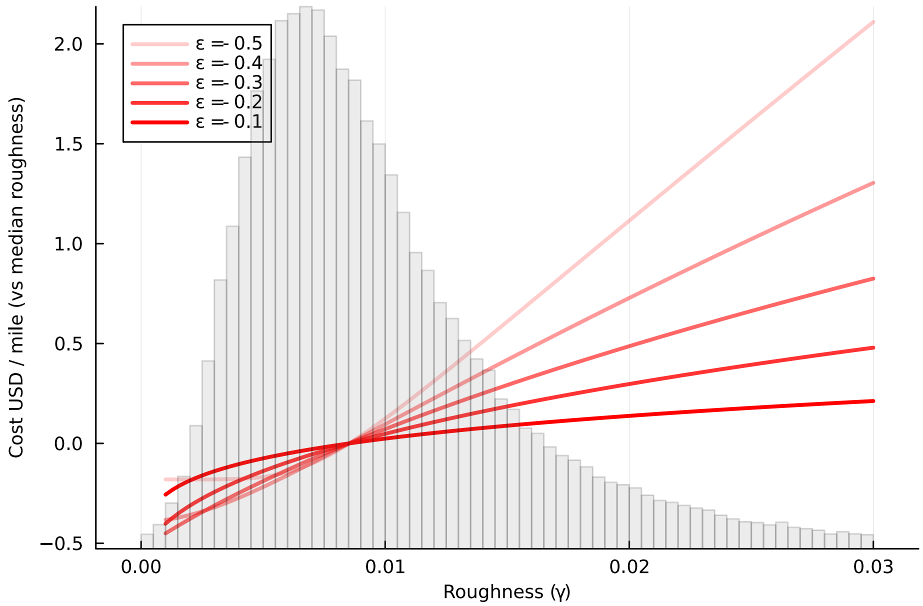

Note: This graph replicates Figure 8 where we re-estimate the model varying the elasticity of speed with respect roughness (equal to −0.31 in the main estimation).

Figure A.15: Road Quality in the New York-Jersey City Metro Area

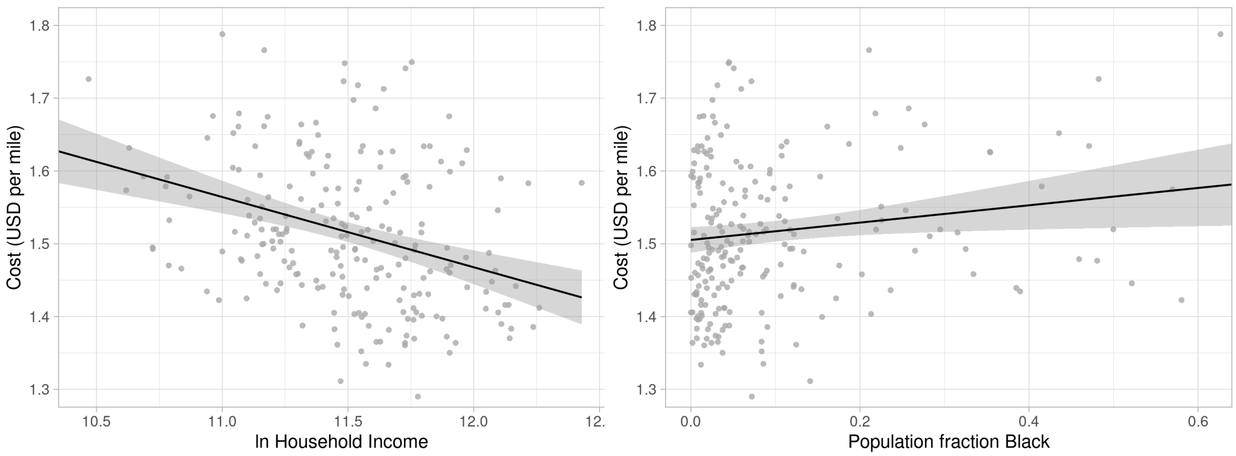

(a) (b)

Note: This figure shows the average per mile cost of local roads by log household income and percentage of the population that is Black in panels (a) and (b), respectively, for towns in the New York-Jersey City MSA..

Figure A.16: Road Expenditure by Local Goverments

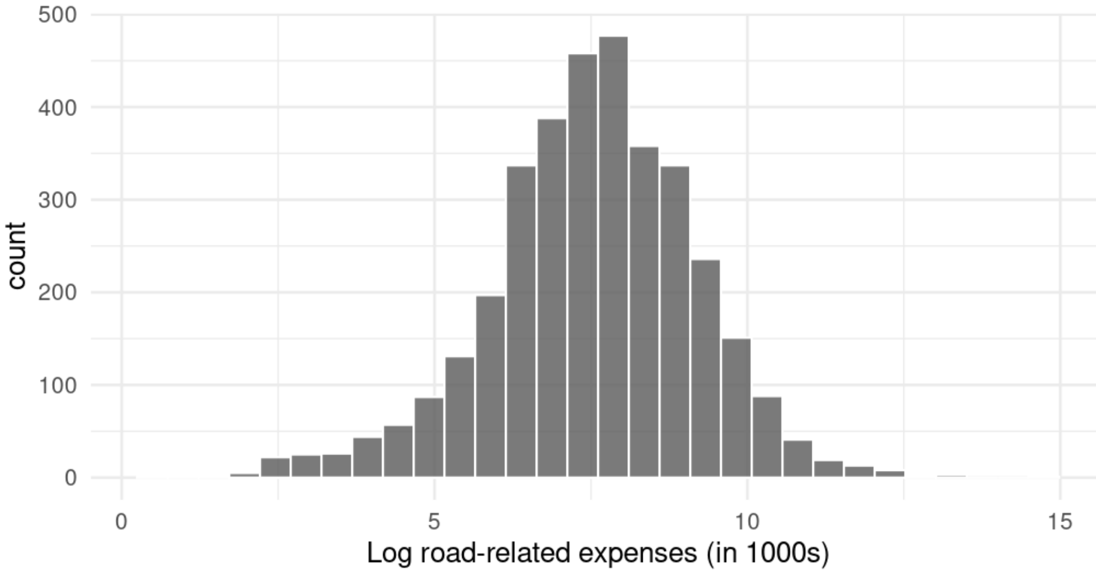

(a)

Note: This figure shows a histogram of 2017 spending on local roads for the towns/cities/CDPs matched to the ASLGF with positive spending. Spending at the country level is assigned to places based on the share of overlapping land area. We show results under a log(x) transformation because our data contains a long right tail. For example, NYC reported 2.5 billion, Chicago reported 670 million, Los Angeles 550 million, and Seattle 430 million. The median town in the analysis reported spending 2 million.

Figure A.17: Dallas grids with greater than zero Uber coverage on local roads

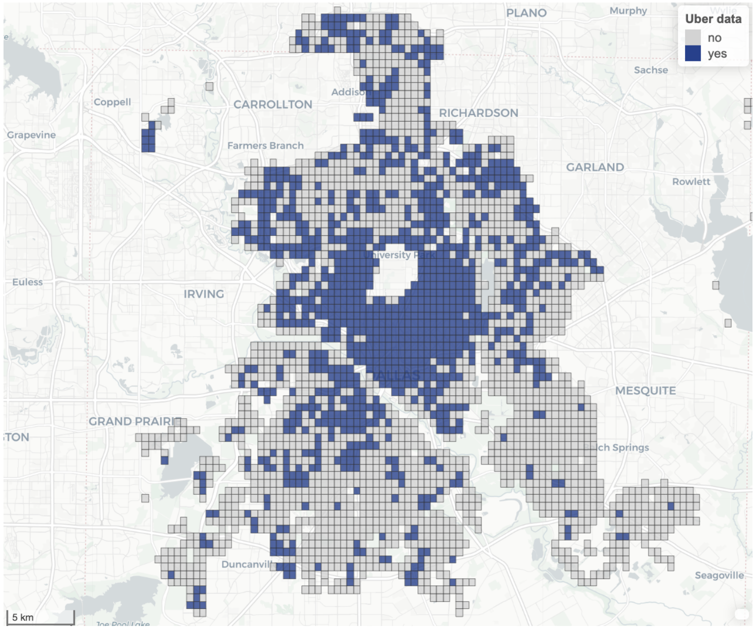

Figure A.18

Note: This figure shows the percent of girds in Dallas that have some Uber coverage, i.e. for which we have a measure of local road roughness. 37% of grids within the city boundary have some Uber coverage. Coverage varies by city; for example New York City has 93% coverage, Columbus has 51%, and Portland has 43%.

Figure A.19: Actual and Model-Implied Targeting in Eleven Cities

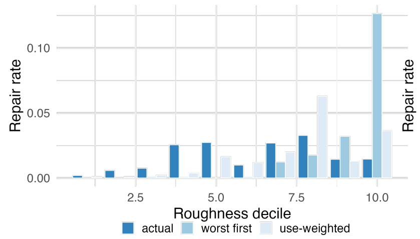

(a) Charlotte, NC

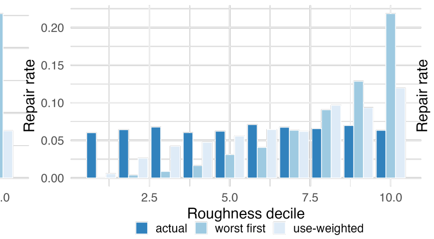

(b) Chicago, IL

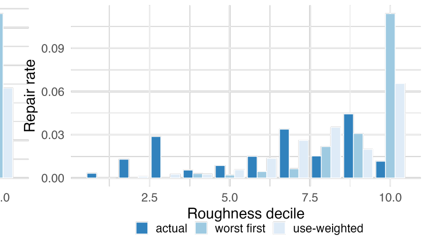

(c) Columbus, OH

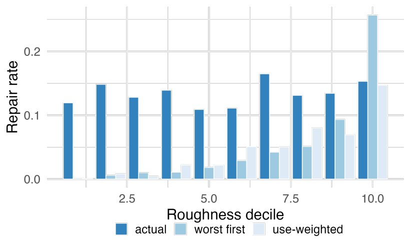

(d) Dallas, TX

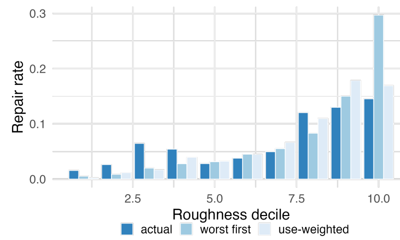

(e) Denver, CO

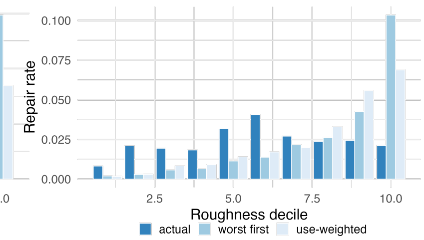

(f) Los Angeles, CA

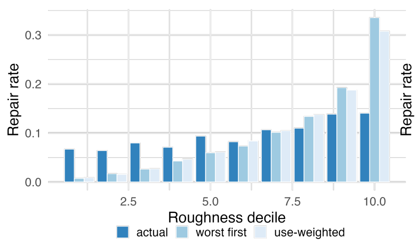

(g) New York City, NY

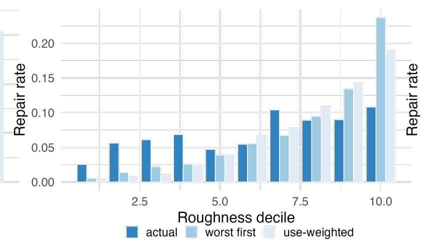

(h) Philadelphia, PA

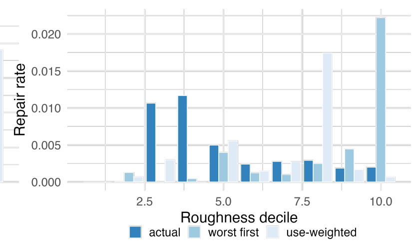

(i) Portland, OR

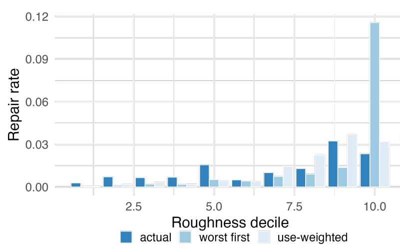

(j) San Diego, California

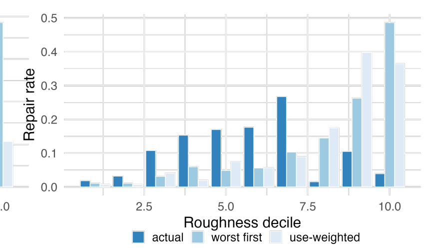

(k) San Jose, CA

Note: This figure shows the percent of roads that are repaired within each decile of roughness cost in the data (dark blue), and for counterfactual policies that prioritize roads based on baseline roughness (medium blue), and based on roughness with road use weights (light blue). Grid cells with Uber data on less than 20% of local roads are dropped.

Figure A.20: NYC Roughness Targeting Controlling for Income

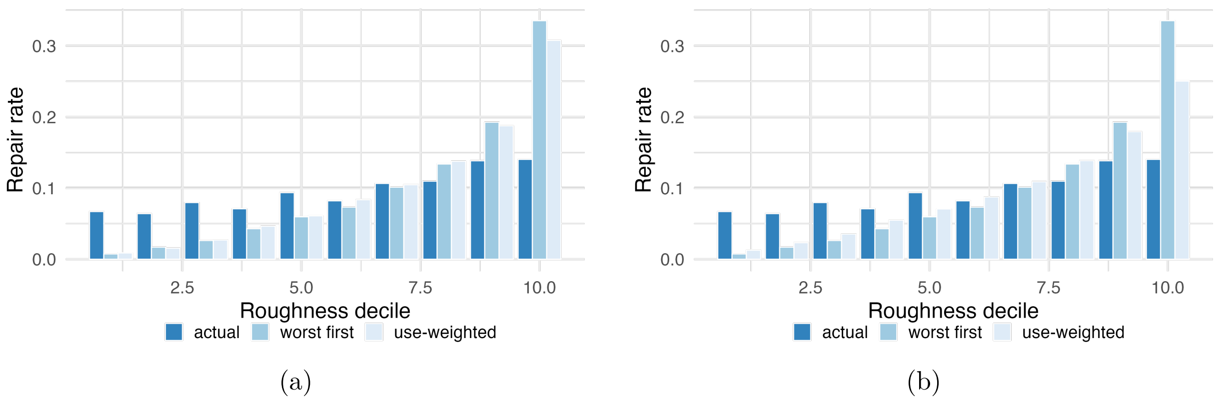

Note: This figure shows the percent of roads that are repaired within each decile of roughness cost. Panel (a) repeats the graph for New York City from Figure A.19, while Panel (b) shows the same results where the road usage measure (Uber traffic) has been residualized by income in the use-weighted cost minimizing counterfactual. For both graphs, arterial and highways are excluded. Income is at the grid level and imputed from ACS Census tract household medians. Roughness cost is at the grid cell level, and is the average over the segments with Uber data within the cell. Grid cells with Uber data on less than 20% of local roads are dropped.

Figure A.21: Road quality information sources and town population

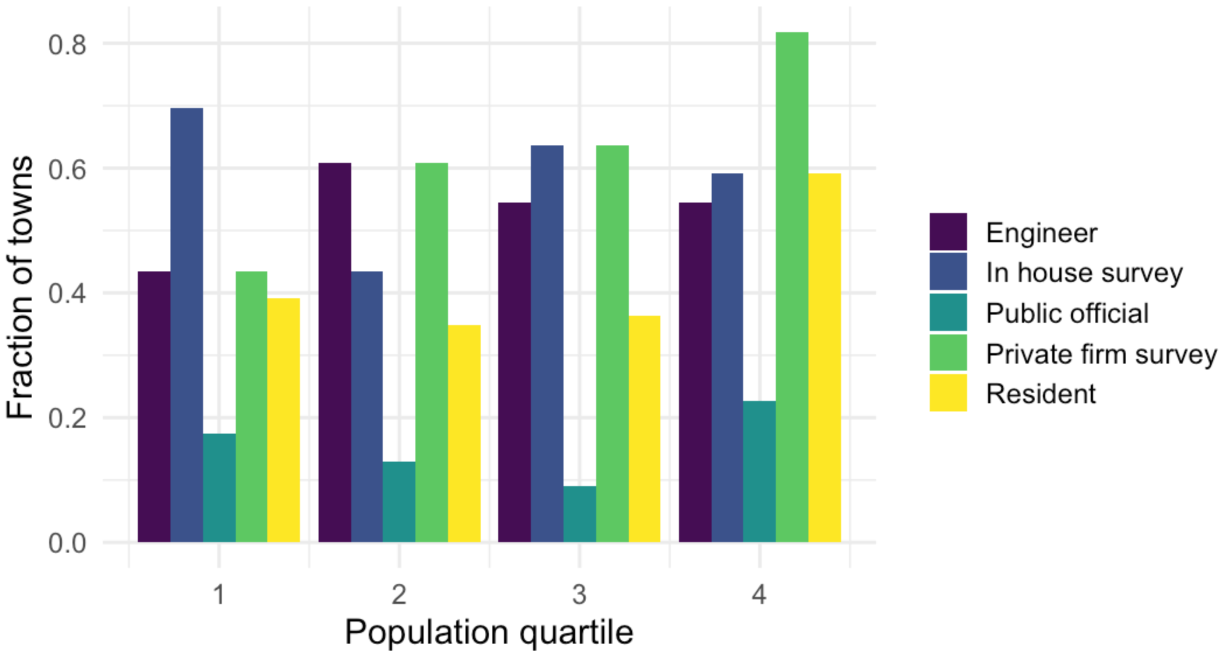

Note: This figure shows the percent of towns within each population quartile across towns in the sample that use various information sources to learn about quality of their roads. There are 90 towns in the sample. The population ranges in each quartile are (777,22972), (22972,44960), (44960,129880), and (129880,1508083).

B Appendix Tables

Table A.1: Uber Roughness Data Coverage of Road Network

	(1) With FE	(2) Total	(3) Coverage Share	(4) Coverage Share (Pop. Weighted)
Panel A. Number of Segments				
Local roads	1,268,601	27,112,992	0.047	0.14
Arterial roads	3,850,221	9,905,938	0.39	0.63
Highways	523,538	1,147,462	0.46	0.73
Panel B. Segment Length (kilometers)				
Local roads	89,274	5,607,603	0.016	0.12
Arterial roads	296,290	1,701,138	0.17	0.63
Highways	130,575	399,187	0.33	0.75

Note: This table reports raw and population-weighted coverage of the Uber road roughness data. To construct it, we start with the universe of road segments. The first three columns report the number of segments and total length with Uber road roughness fixed effects (column 1), in total (column 2), and their ratio (column 3). Column (4) restricts to observations within towns (census places) and reports the share with fixed effects using town population weights.

Table A.2: Number of Uber Trips per Segment and Road Roughness

	Predicted roughness at 20mph			
	(1)	(2)	(3)	(4)
Log trips per segment	0.0116***	0.0025***	0.0013	-0.0008
	(0.0002)	(0.0002)	(0.0008)	(0.0008)
Constant	0.2838***	0.3279***	0.3292***	0.3384***
	(0.0008)	(0.0009)	(0.0035)	(0.0035)
Uber city FE		Yes		Yes
Sample: below median trips			Yes	Yes
SD of the outcome	0.16	0.16	0.17	0.17
Observations	1,829,526	1,829,524	914,586	914,584
R2	0.00	0.04	0.00	0.04

Note: This table reports the correlation between log number of Uber trips and predicted road roughness at 20mph for all local road segments in our data. The data covers 240 Uber cities. Only road segments with at least 50 trips per segment are included. Column 3 and 4 restrict to the sample of road segments with below-median number of trips per segment (103 trips). Robust standard errors in parentheses, *p<0.1; **p<0.05; ***p<0.01.

Table A.3: Cross-Validation of Segment-Level Uber Road Roughness and Speed (Chicago)

	Predicted Roughness			Log Speed		
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Outcome Using April Data						
Predicted Roughness (March data)	0.91***	0.97***	0.93***			
	(0.01)	(0.00)	(0.00)			
Log Speed (March data)				1.00***	0.99***	0.99***
				(0.00)	(0.00)	(0.00)
Road type	highway	arterial	local	highway	arterial	local
Observations	8,823	225,035	156,144	8,823	225,035	156,144
Adj R2	0.82	0.90	0.84	0.99	0.99	0.98
Panel B: Outcome Using August Data						
Predicted Roughness (March data)	0.90***	0.94***	0.90***			
	(0.01)	(0.00)	(0.00)			
Log Speed (March data)				0.99***	0.99***	0.98***
				(0.00)	(0.00)	(0.00)
Road type	highway	arterial	local	highway	arterial	local
Observations	8,816	225,742	159,632	8,816	225,742	159,632
Adj R2	0.80	0.86	0.80	0.99	0.98	0.97

Note: This table reports cross-validation results in Chicago for predicted roughness and log median speed using March and April data (Panel A) and August data (Panel B). Each column reports results from a regression at the road segment level where the outcome is estimated only using the later data (April or August 2021), and the explanatory variable is estimated only using March 2021 data. Predicted roughness is computed using 20mph for local roads, 32mph for arterial roads, and 48mph for highways. Robust standard errors in parentheses, *p<0.1; **p<0.05; ***p<0.01 .

Table A.4: Split Sample Correlation Between Predicted Roughness and Speed (Chicago)

	Median Speed								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Train	Test	Off-peak	Train	Test	Off-peak	Train	Test	Off-peak
Predicted Roughness (Training data)	-29.0***	-29.4***	-29.3***	-27.6***	-30.5***	-30.4***	-20.5***	-21.2***	-21.1***
	(1.0)	(1.1)	(1.1)	(0.2)	(0.2)	(0.2)	(0.1)	(0.2)	(0.2)
Constant	59.9***	59.3***	59.3***	40.6***	40.8***	40.7***	26.9***	26.4***	26.3***
	(0.5)	(0.5)	(0.5)	(0.1)	(0.1)	(0.1)	(0.1)	(0.1)	(0.1)
Road type	highway	highway	highway	arterial	arterial	arterial	local	local	local
Observations	9,021	9,021	9,021	219,661	219,661	219,661	94,177	94,177	94,177
Adj R2	0.11	0.10	0.10	0.21	0.22	0.22	0.22	0.21	0.21

Note: This table reports same- and split-sample regressions of road segment median speed on predicted roughness, including using only off-peak speed measurement. This table uses data from August 2021. We first divide the data into a 75% sample used for estimation (training), and a 25% hold-out sample used for testing. We compute median road segment speed in both samples. We further compute median speed only during off-peak hours () in the testing sample. In all the regressions, we regress median speed on predicted roughness. In odd columns, we compute median speed in the same data that we use for estimating roughness. In odd columns, we use the hold-out sample to compute speed. Predicted roughness is computed using 20mph for local roads, 32mph for arterial roads, and 48mph for highways. Robust standard errors in parentheses, *p<0.1; **p<0.05; ***p<0.01 .

Table A.5: Town Roughness and MSA Fixed Effects

	Road Roughness (z-score)		
	(1)	(2)	(3)
	Highway	Arterial	Local
MSA FE	Yes	Yes	Yes
Observations	4,352	7,087	4,844
Adj. R2	0.18	0.31	0.09

Note: This table shows coefficients from regressing average road roughness in a town on MSA fixed effects.

Table A.6: The Impact of Roughness on Driver Speed on Local Roads at Town Borders

	Log speed (mph)							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Log roughness slope	-0.455***	-0.406***	-0.327***	-0.327***	-0.314***	-0.314***	-0.298***	-0.298***
	(0.020)	(0.013)	(0.015)	(0.015)	(0.017)	(0.017)	(0.020)	(0.020)
Log speed limit	0.237***		0.188***		0.192**		0.044	
	(0.037)		(0.053)		(0.063)		(0.022)	
Sample:	Town	Town	Borders	Borders	Borders	Borders	Borders	Borders
Sample restriction:			$<$ 1km	$<$ 1km	$<$ 500m	$<$ 500m	$<$ 1km	$<$ 1km
Uber City FE	Yes	Yes						
Border pair FE			Yes	Yes	Yes	Yes	Yes	Yes
Distance to border controls							Yes	Yes
Uber cities	83	170	73	73	72	72	73	73
Towns	1,509	3,752	1,245	1,245	1,209	1,209	1,245	1,245
Border pairs	0	0	1,366	1,366	1,285	1,285	1,366	1,366
N	1,509	3,752	2,732	2,732	2,570	2,570	2,732	2,732

Note: Additional results related to columns 1 and 2 in Table 2. In columns 7 and 8 we control for distance to the town border, linearly with a different slope for each border pair and each side of the border. Standard errors, adjusted for clustering at the Uber region level (columns 1-2) and border pair level (columns 3-8) in parentheses. *p$<$0.1; **p$<$0.05; ***p$<$0.01.

Table A.7: The Impact of Roughness on Driver Speed on Local Roads in Chicago using Repaving Events (linear IV)

	Log roughness slope		Log speed (mph)		
	(1)	(2)	(3)	(4)	(5)
Log roughness slope			-0.145***	-0.145***	-0.141***
			(0.026)	(0.026)	(0.030)
Repaved	0.206***	0.206***	-0.009	-0.009	
	(0.036)	(0.036)	(0.012)	(0.012)	
Post repaving	-0.025***	-0.025***	-0.014***	-0.014***	0.000
	(0.007)	(0.007)	(0.002)	(0.002)	(NaN)
Post repaving × Repaved	-0.191***	-0.191***			
	(0.020)	(0.020)			
Log speed limit	0.156*		0.010		
	(0.077)		(0.022)		
Repaving Event FE	Yes	Yes	Yes	Yes	Yes
Road Segment FE					Yes
Repaving Event × Post FE					Yes
Estimator	OLS	OLS	IV	IV	IV
Repaving Events	611	611	611	611	611
N	19,878	19,878	19,878	19,878	19,878

Note: The estimating equation is (4.7). Column 1 and 2 report the first stage, with and without controlling for the segment speed limit. In columns 3-5, log roughness slope is instrumented using the Post repaving × Repaved interaction. Standard errors, adjusted for clustering at the road repaving event level in parentheses. *p<0.1; **p<0.05; ***p<0.01.

Table A.8: Roughness Cost Models Parameter Estimates

	(1)	(2)
Bumpiness cost level β_1	2.75	1.69
	(40.9)	(0.11)
Bumpiness cost exponent β_2	1.06	0.47
	(0.69)	(0.04)
Other speed costs κ	0.39	
	(0.16)	
Estimator	CMD	CMD
N	2,330	2,330

Note: This table reports classic minimum distance estimation results from the full model based on (4.1) in column 1 and the simple model with $K = 0$ in column 2. We use the town border sample and specification from column 2 in Table 2. For inference, we propagate the uncertainty from the Bayesian bootstrap at the border pair level through the CMD estimation. Standard errors in parentheses.

Table A.9: Roughness Cost Sensitivity

Elasticity of speed wrt roughness	*Cost of* $+1SD$
-0.5	0.67
-0.4	0.45
-0.3	0.31
-0.2	0.20
-0.1	0.09

Note: This table reports the cost of increasing road roughness γ by one SD, starting from its median value. We re-estimate the full model five times varying the value of the elasticity of speed with respect to roughness moment (which is -0.31 in the main estimation).

Table A.10: Inspection Failure Rate and Road Roughness By Indicator

	(1) Brakes	(2) Front	(3) Suspensions	(4) Frame	(5) Muffler	(6) Bumper	(7) Tires	(8) Wipers	(9) Windshield	(10) Mirror
Local Roughness (z-score)	0.022	0.084	0.072	0.009	0.076	0.028	0.084	0.052	0.018	0.022
	(0.055)	(0.088)	(0.077)	(0.009)	(0.043)	(0.036)	(0.084)	(0.047)	(0.011)	(0.015)
Mean Outcome	0.18	0.97	0.70	0.03	0.45	0.21	1.10	0.47	0.11	0.06
Observations	926	926	926	926	926	926	926	926	926	926

Note: This table reports the correlation between inspection failure rates and local road roughness in the same town. The data is at the inspection station level. The outcome is the level of failures per inspection, for the respective category, namely brakes, front end, steering and suspension frame, muffler and exhaust system, bumpers/fenders/exterior sheet metal, tires, windshield wipers and cleaner, windshield, and rear view mirror. The last three indicators are the placebo indicators. Robust standard errors in parentheses, *p<0.1; **p<0.05; ***p<0.01.

Table A.11: Tract Coverage Robustness Check

	Dependent variable:					
	Cost (USD per mile)					
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: > 5% coverage						
Ln median income	-0.053^{***}	-0.069^{***}	-0.019^{***}			
	(0.002)	(0.002)	(0.002)			
Fraction Black				0.109^{***}	0.116^{***}	0.011^{**}
				(0.004)	(0.004)	(0.005)
Climate controls	Yes			Yes		
MSA Fixed effects		Yes			Yes	
Town Fixed effects			Yes			Yes
Observations	23,809	23,948	23,948	23,809	23,948	23,948
Panel B: > 10% coverage						
Ln median income	-0.042^{***}	-0.058^{***}	-0.013^{***}			
	(0.002)	(0.002)	(0.002)			
Fraction Black				0.086^{***}	0.094^{***}	0.005
				(0.005)	(0.005)	(0.005)
Climate controls	Yes			Yes		
MSA Fixed effects		Yes			Yes	
Town Fixed effects			Yes			Yes
Observations	18,520	18,634	18,634	18,520	18,634	18,634
Panel C: > 15% coverage						
Ln median income	-0.031^{***}	-0.046^{***}	-0.008^{***}			
	(0.003)	(0.003)	(0.002)			
Fraction Black				0.068^{***}	0.076^{***}	0.003
				(0.005)	(0.005)	(0.005)
Climate controls	Yes			Yes		
MSA Fixed effects		Yes			Yes	
Town Fixed effects			Yes			Yes
Observations	15,046	15,142	15,142	15,046	15,142	15,142
Climate controls	Yes			Yes		
MSA Fixed effects		Yes			Yes	
Town Fixed effects			Yes			Yes

Note: This table reports results from regressions following the specifications in Table 5 for income and the faction of households that are Black. Here, we adjust the data sample by removing any tract for which the percent of road segments we have roughness for is below a threshold. In the first panel, we remove tracts with less than 5% coverage of segment roughness; in the second, less than 10; in the final panel, less than 15%. The coefficients persist as we restrict the sample, but the magnitude shrinks noticeably. Given the evidence in Figure A.5 that there is no relationship between roughness and number of Uber observations, we interpret these results as suggestive that the relationship between roughness and income/race is strongest in tracts with the worst Uber coverage. Robust standard errors in parentheses, *p<0.1; **p<0.05; ***p<0.01.

Table A.12: Town Coverage Robustness Check

	Dependent variable:						
	Cost (USD per mile)						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: no threshold							
Ln median income	-0.064^{***}	-0.064^{***}	-0.040^{***}				-0.047^{***}
	(0.006)	(0.006)	(0.007)				(0.007)
Fraction Black				0.123^{***}	0.123^{***}	0.080^{***}	0.078^{***}
				(0.014)	(0.014)	(0.013)	(0.015)
Ln expenditure per capita		-0.0001	0.002		-0.002	0.001	
		(0.002)	(0.002)		(0.002)	(0.002)	
MSA Fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Additional controls			Yes			Yes	
Observations	3,142	3,142	3,142	3,142	3,142	3,142	3,142
Panel B: 5% coverage							
Ln median income	-0.078^{***}	-0.080^{***}	-0.053^{***}				-0.062^{***}
	(0.007)	(0.007)	(0.007)				(0.008)
Fraction Black				0.118^{***}	0.119^{***}	0.090^{***}	0.066^{***}
				(0.013)	(0.013)	(0.012)	(0.014)
Ln expenditure per capita		0.006^{***}	0.007^{***}		0.002	0.006^{***}	
		(0.002)	(0.002)		(0.002)	(0.002)	
MSA Fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Additional controls			Yes			Yes	
Observations	1,902	1,902	1,902	1,902	1,902	1,902	1,902
Panel C: 15% coverage							
Ln median income	-0.064^{***}	-0.069^{***}	-0.037^{***}				-0.047^{***}
	(0.010)	(0.010)	(0.010)				(0.011)
Fraction Black				0.097^{***}	0.098^{***}	0.091^{***}	0.064^{***}
				(0.017)	(0.017)	(0.015)	(0.018)
Ln expenditure per capita		0.008^{***}	0.008^{***}		0.004	0.007^{***}	
		(0.003)	(0.003)		(0.003)	(0.003)	
MSA Fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Additional controls			Yes			Yes	
Observations	866	866	866	866	866	866	866

Note: This table reports results from regressions following the specifications in Table 9, for income and the faction of households that are Black. Here, we adjust the data sample by removing any town for which the percent of road segments we have roughness for is below a threshold. In the first panel, we remove no towns; in the second, less than 5; in the final panel, less than 15% (in Table 9, we remove towns with less than 10% coverage). The coefficients persist as we restrict the sample, but the magnitude shrinks noticeably. Given the evidence in Figure A.5 that there is no relationship between roughness and number of Uber observations, we interpret these results as suggestive that the relationship between roughness and income/race is strongest in tracts with the worst Uber coverage. Robust standard errors in parentheses, *p<0.1; **p<0.05; ***p<0.01.

Table A.13: Road Expenditure Correlates

	Dependent variable:					
	Ln expenditure per capita		Ln expenditure			
	(1)	(2)	(3)	(4)	(5)	(6)
Ln median income	0.528***		0.045		0.068	
	(0.105)		(0.148)		(0.119)	
Fraction Black		−0.341*		0.184		−0.098
		(0.191)		(0.274)		(0.189)
Fraction Hispanic		−1.590***		−0.701**		−0.891***
		(0.205)		(0.294)		(0.239)
Fraction Asian		−1.783***		−0.615		−2.074***
		(0.403)		(0.577)		(0.384)
Ln miles of local road					0.205	0.013
					(0.171)	(0.166)
Ln population					0.362***	0.508***
					(0.116)	(0.111)
Miles to CBD					0.025	−0.009
					(0.056)	(0.057)
Ln employment					0.273***	0.310***
					(0.047)	(0.048)
Ln area (miles^2)					0.404***	0.414***
					(0.111)	(0.110)
Fraction drive to work					−2.187***	−1.550***
					(0.417)	(0.434)
MSA fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1,223	1,223	1,223	1,223	1,223	1,223
Adjusted R^2	0.215	0.239	0.157	0.161	0.638	0.648

Note: This table reports results from regressions of road expenditure on town level covariates. The first two columns have log expenditure per capita as the outcome. The last four columns have log expenditure as the outcome. Note that since we include log population as a covariate in the last two columns, using the per capita measure as the dependent variable would produce the same coefficients on all the covariates except log population. Expenditure per capita is each local government's reported direct expenditures on "construction and maintenance of roads, sidewalks and bridges; street lighting; snow removal; highway engineering, control, and safety." County and township spending are distributed to towns based on shared land area. Robust standard errors in parentheses, *p<0.1; **p<0.05; ***p<0.01.

Table A.14: Predictors of Road Repair Targeting

	Dependent variable:										
	Repair rate (percent road-miles repaved)										
	NYC	Dallas	Columbus	Portland	LA	San Jose	Denver	Philadelphia	Charlotte	San Diego	Chicago
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
Local roughness (z score)	0.023***	−0.004	0.001	0.001	0.004	−0.003	0.077***	0.021	−0.0001	0.003	0.002
	(0.006)	(0.007)	(0.003)	(0.002)	(0.002)	(0.017)	(0.016)	(0.014)	(0.002)	(0.002)	(0.004)
Ln road miles	0.002	0.034***	−0.016***	0.001	−0.004**	0.011	0.004	−0.014***	−0.002	0.003	−0.036***
	(0.004)	(0.006)	(0.003)	(0.002)	(0.002)	(0.015)	(0.010)	(0.005)	(0.003)	(0.002)	(0.004)
Ln population	0.009**	−0.003	0.009**	0.0002	0.002	0.039**	−0.017**	0.017**	0.006	0.003*	0.025***
	(0.004)	(0.006)	(0.004)	(0.002)	(0.001)	(0.019)	(0.009)	(0.007)	(0.004)	(0.001)	(0.005)
Ln miles to CBD	−0.042***	−1.286	−0.042	2.007*	1.095*	3.045	8.103*	0.545***	−0.186	−2.704**	−0.446
	(0.008)	(1.739)	(0.309)	(1.029)	(0.629)	(11.660)	(4.618)	(0.125)	(0.208)	(1.083)	(0.650)
Fraction Black	−0.031***	−0.029	−0.002	0.015	0.026**	2.051***	0.133	−0.074***	0.004	−0.045*	0.003
	(0.011)	(0.026)	(0.017)	(0.024)	(0.011)	(0.564)	(0.120)	(0.021)	(0.012)	(0.024)	(0.010)
Ln median income	0.002	0.012**	0.032***	−0.003	0.00004	−0.042	−0.015	0.002	0.009	−0.006***	0.024***
	(0.003)	(0.006)	(0.009)	(0.004)	(0.002)	(0.038)	(0.022)	(0.009)	(0.005)	(0.002)	(0.006)
Ln Uber volume	−0.004	−0.025**	−0.0002	−0.001	−0.0001	−0.109***	0.007	0.027**	0.013***	0.003	−0.005
	(0.006)	(0.010)	(0.005)	(0.003)	(0.003)	(0.034)	(0.016)	(0.011)	(0.004)	(0.003)	(0.007)
Fraction drive to work	0.011	0.191**	−0.182***	0.012	0.014	0.194	−0.017	0.079*	−0.045*	−0.008	−0.020
	(0.021)	(0.084)	(0.033)	(0.010)	(0.013)	(0.226)	(0.068)	(0.044)	(0.026)	(0.019)	(0.023)
Constant	0.120**	9.543	0.062	−16.289*	−8.817*	−24.112	−61.828*	−2.732***	1.071	21.755**	2.738
	(0.052)	(12.962)	(2.017)	(8.345)	(5.058)	(94.526)	(35.506)	(0.563)	(1.366)	(8.692)	(4.455)
Observations	2,918	1,086	1,249	619	3,531	530	766	1,209	1,486	1,456	2,128
Adjusted R^2	0.040	0.052	0.050	−0.003	0.003	0.048	0.029	0.042	0.008	0.013	0.057

Note: *p<0.1; **p<0.05; ***p<0.01

Note: This table shows coefficients from a linear regression of the fraction of a grid that gets repaved on covariates that may impact repair targeting. We then include two proxies for traffic volume, Uber traffic and the fraction of workers that drive to work.

.

Table A.15: Repair Targeting with Varying Coverage Thresholds

Panel A: > 20% coverage

	Dependent variable: Repair rate (percent road-miles repaved)										
	NYC	Dallas	Columbus	Portland	LA	San Jose	Denver	Philadelphia	Charlotte	San Diego	Chicago
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
Local roughness (z score)	0.026***	−0.011	0.004	−0.003	0.005	−0.002	0.082***	0.025	0.002	0.009***	0.003
	(0.006)	(0.011)	(0.008)	(0.002)	(0.004)	(0.030)	(0.025)	(0.016)	(0.004)	(0.003)	(0.005)
Ln road miles	0.004	0.036***	−0.023***	0.002	−0.005***	0.019	0.003	−0.013**	0.002	0.005*	−0.036***
	(0.004)	(0.007)	(0.005)	(0.002)	(0.002)	(0.019)	(0.012)	(0.005)	(0.003)	(0.002)	(0.004)
Ln population	0.008**	0.005	0.013**	−0.00005	0.002	0.030	−0.003	0.013*	0.009*	0.001	0.025***
	(0.004)	(0.008)	(0.007)	(0.001)	(0.002)	(0.026)	(0.011)	(0.008)	(0.005)	(0.002)	(0.005)
Ln miles to CBD	−0.040***	4.745*	−0.104	0.781	0.456	−2.815	5.911	0.529***	−0.159	−5.910***	−1.118
	(0.008)	(2.715)	(0.602)	(1.361)	(0.823)	(19.393)	(7.009)	(0.141)	(0.258)	(1.738)	(0.744)
Fraction Black	−0.032***	−0.045	0.022	−0.006	0.024*	2.712***	0.113	−0.072***	−0.015	−0.079*	0.0001
	(0.011)	(0.041)	(0.029)	(0.035)	(0.012)	(0.728)	(0.179)	(0.022)	(0.015)	(0.042)	(0.010)
Ln median income	0.002	0.007	0.051***	−0.001	0.001	−0.037	−0.055	0.004	0.004	−0.005	0.020***
	(0.003)	(0.007)	(0.015)	(0.005)	(0.002)	(0.057)	(0.035)	(0.010)	(0.006)	(0.003)	(0.006)
Ln Uber volume	−0.009	−0.042***	0.011	−0.003	−0.002	−0.073	−0.004	0.027**	0.016***	0.001	−0.003
	(0.007)	(0.014)	(0.009)	(0.004)	(0.003)	(0.045)	(0.021)	(0.012)	(0.005)	(0.004)	(0.007)
Fraction drive to work	0.009	0.170	−0.234***	0.00001	0.014	0.514*	0.061	0.092*	−0.052*	−0.038	−0.013
	(0.023)	(0.108)	(0.052)	(0.012)	(0.015)	(0.274)	(0.090)	(0.049)	(0.029)	(0.027)	(0.024)
Constant	0.143**	−35.268*	0.201	−6.316	−3.680	22.949	−44.654	−2.668***	0.943	47.530***	7.374
	(0.056)	(20.241)	(3.921)	(11.041)	(6.621)	(157.143)	(53.902)	(0.637)	(1.691)	(13.951)	(5.108)
Observations	2,732	568	618	233	2,635	202	427	1,089	737	848	1,870
Adjusted R^2	0.039	0.070	0.072	−0.019	0.004	0.080	0.026	0.034	0.028	0.023	0.057

Panel B: > 30% coverage

	NYC	Dallas	Columbus	Portland	LA	San Jose	Denver	Philadelphia	Charlotte	San Diego	Chicago
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
Local roughness (z score)	0.029***	−0.019	0.011	−0.003	0.001	−0.020	0.067**	0.029*	0.006	0.008*	0.004
	(0.007)	(0.014)	(0.012)	(0.002)	(0.004)	(0.030)	(0.028)	(0.017)	(0.006)	(0.004)	(0.005)
Ln road miles	0.002	0.042***	−0.032***	0.002*	−0.004*	0.017	0.003	−0.013**	0.003	0.006**	−0.036***
	(0.004)	(0.008)	(0.006)	(0.001)	(0.002)	(0.020)	(0.013)	(0.005)	(0.004)	(0.003)	(0.005)
ln population	0.009**	0.005	0.011	0.001	0.002	0.040	0.001	0.013	0.011	0.001	0.026***
	(0.004)	(0.009)	(0.008)	(0.001)	(0.002)	(0.029)	(0.013)	(0.008)	(0.007)	(0.003)	(0.006)
Ln miles to CBD	−0.042***	6.230*	0.076	2.244*	0.777	10.113	6.987	0.524***	−0.320	−6.870***	−1.487*
	(0.008)	(3.399)	(0.872)	(1.248)	(0.949)	(21.757)	(7.986)	(0.153)	(0.374)	(2.143)	(0.799)
Fraction Black	−0.029***	−0.084*	0.017	−0.021	0.020	1.609*	0.084	−0.071***	−0.035	−0.050	−0.006
	(0.011)	(0.050)	(0.040)	(0.032)	(0.014)	(0.872)	(0.195)	(0.023)	(0.021)	(0.061)	(0.011)
Ln median income	0.002	0.008	0.054***	−0.001	0.001	−0.049	−0.053	0.005	0.0002	−0.003	0.018***
	(0.003)	(0.008)	(0.019)	(0.004)	(0.002)	(0.063)	(0.040)	(0.010)	(0.009)	(0.004)	(0.006)
Ln Uber volume	−0.011	−0.057***	0.022*	−0.001	0.0003	−0.086*	−0.005	0.023*	0.014*	−0.001	−0.005
	(0.008)	(0.017)	(0.012)	(0.003)	(0.004)	(0.046)	(0.023)	(0.014)	(0.008)	(0.005)	(0.008)
Fraction drive to work	0.012	0.181	−0.273***	0.008	0.022	0.220	0.159	0.100*	−0.040	−0.046	−0.016
	(0.023)	(0.124)	(0.067)	(0.011)	(0.016)	(0.302)	(0.107)	(0.051)	(0.040)	(0.033)	(0.026)
Constant	0.152**	−46.253*	−0.997	−18.233*	−6.267	−81.467	−53.017	−2.629***	2.026	55.236***	9.933*
	(0.060)	(25.346)	(5.662)	(10.123)	(7.637)	(176.254)	(61.379)	(0.691)	(2.448)	(17.199)	(5.485)
Observations	2,596	427	396	160	2,140	136	350	993	453	618	1,693
Adjusted R^2	0.042	0.100	0.103	0.014	0.001	0.054	0.013	0.030	0.032	0.022	0.059

Note: *p<0.1; **p<0.05; ***p<0.01

Note: This table reports results from regressions following the specifications in Table A.14. Here, we vary the Uber coverage threshold under which we remove any grid from the sample. In the first panel, we remove grids with less than 20% coverage of segment roughness; in the second, we remove grids with less than 30% coverage of segment roughness.

Table A.16: NYC Repair Targeting with Sub-Sample

	Dependent variable:	
	Repair rate (percent road-miles repaved) NYC	
	(1)	(2)
Local roughness (z score)	0.015^{***}	0.015^{***}
	(0.005)	(0.005)
Ln road miles	0.004	0.002
	(0.004)	(0.005)
Ln population	0.007^{**}	0.011^{***}
	(0.003)	(0.004)
Ln miles to CBD	-0.040^{***}	-0.048^{***}
	(0.006)	(0.008)
Fraction Black	-0.031^{***}	-0.028^{***}
	(0.011)	(0.011)
Ln median income	0.005^{*}	0.002
	(0.003)	(0.003)
Ln Uber volume		-0.005
		(0.006)
Fraction drive to work		0.025
		(0.022)
Constant	0.080^{**}	0.121^{**}
	(0.038)	(0.052)
Observations	2,708	2,708
Adjusted R^2	0.043	0.043

Note: This table reports results from regressions following the specifications in Table A.14 for New York City, but where the data has been subsampled to match the distribution of Uber coverage across grids in Dallas. Specifically, for Dallas and New York City, we group grids into 5 coverage quantiles (conditional on >0 Uber data) and get the average coverage within each quantile. For each NYC grid g we get the hypothetical length of road it would need to have Uber coverage on to match the mean Dallas coverage in the same quantile. We define $p_g = \min(\text{the hypothetical length/the true Uber segment length}, 1)$ Then we keep or discard each Uber segment in g with probability p_g.

Table A.17: Repair Targeting vs. Counterfactual Targeting with Varying Coverage Thresholds

Panel A: > 0 coverage

	Dependent variable: Repair rate (percent road-miles repaved)										
	NYC	Dallas	Columbus	Portland	LA	San Jose	Denver	Philadelphia	Charlotte	San Diego	Chicago
Use-weighted	0.164***	−0.043	0.077	−0.008	−0.045	−0.101*	0.0003	0.183***	0.097**	−0.006	0.107***
	(0.023)	(0.033)	(0.054)	(0.025)	(0.028)	(0.055)	(0.068)	(0.055)	(0.041)	(0.043)	(0.031)
Ln road miles	0.010***	0.037***	−0.015***	0.001	−0.003*	0.020	−0.007	−0.006	−0.0004	0.004**	−0.034***
	(0.003)	(0.006)	(0.003)	(0.001)	(0.002)	(0.014)	(0.009)	(0.004)	(0.002)	(0.002)	(0.004)
Constant	0.070***	0.129***	0.020***	0.005***	0.025***	0.122***	0.070***	0.063***	0.014***	0.008***	0.080***
	(0.004)	(0.005)	(0.003)	(0.001)	(0.001)	(0.013)	(0.008)	(0.006)	(0.002)	(0.001)	(0.004)
Observations	2,974	1,086	1,249	619	3,541	530	774	1,234	1,504	1,456	2,134
Adjusted R^2	0.021	0.040	0.020	−0.002	0.001	0.009	−0.002	0.008	0.002	0.002	0.034

Panel B: > 30% coverage

	Dependent variable: Repair rate (percent road-miles repaved)										
	NYC	Dallas	Columbus	Portland	LA	San Jose	Denver	Philadelphia	Charlotte	San Diego	Chicago
Use-weighted	0.169***	−0.164	0.288**	−0.032	−0.055	−0.065	−0.086	0.151**	0.105	−0.025	0.088***
	(0.023)	(0.118)	(0.125)	(0.059)	(0.035)	(0.082)	(0.086)	(0.060)	(0.064)	(0.055)	(0.034)
Ln road miles	0.010***	0.042***	−0.024***	0.002*	−0.003	0.026	−0.009	−0.006	0.008**	0.006**	−0.034***
	(0.003)	(0.008)	(0.006)	(0.001)	(0.002)	(0.017)	(0.012)	(0.005)	(0.004)	(0.002)	(0.005)
Constant	0.070***	0.131***	0.013*	0.003**	0.028***	0.092***	0.078***	0.072***	0.021***	0.012***	0.083***
	(0.004)	(0.010)	(0.007)	(0.001)	(0.002)	(0.024)	(0.013)	(0.007)	(0.004)	(0.003)	(0.004)
Observations	2,650	427	396	160	2,148	136	354	1,013	468	618	1,699
Adjusted R^2	0.023	0.060	0.046	0.012	0.002	0.010	−0.001	0.005	0.014	0.006	0.034

Note: *p<0.1; **p<0.05; ***p<0.01

Note: This table reports results from regressions following the specifications in Table 10, but where we adjust the data sample by removing any grid for which the percent of road segments we have roughness for is below a different threshold. In Table 10, the threshold is 20%. In the first panel, we remove no grids; in the second, we remove grids with less than 30% coverage of segment roughness. There is a trade-off to restricting the sample. Requiring higher coverage reduces measurement error, but also changes the sample if there is selection in where Uber riders travel.

C Data Appendix

C.1 Geographic data.

We obtain boundaries for census tracts, counties, towns (Census places), CBSAs, and states, as well as geographic crosswalks between them, from the US Census' TIGER/Line Shapefiles.

C.2 Census data, 2015-2019.

We use data on population, income, race, and ethnicity from the 2019 American Community Survey 5-Year estimates (2015-2019) at the Census tract and Census place level. We assign tracts to towns (Census Designated Places) using crosswalks from the Census. We join this with data from the 2018 Zip Codes Business Patterns. Zip Code Tabulation Areas are matched to towns by share of overlapping area.

C.3 Climate data for the entire United States, 2016-2022.

We use county-level monthly maximum temperatures, minimum temperatures, and precipitation rates from the National Oceanic and Atmospheric Administration (NOAA). For each county, we take a five-year average of each outcome.

C.4 IRI for the entire United States, 2017-2018 (highways).

We use International Roughness Index (IRI) data provided by the federal Department of Transport's Highway Monitoring Performance System (HPMS). The latest available data as of 2023 is from 2017-2018. We download the data from https://www.fhwa.dot.gov/policyinformation/hpms/shapefiles.cfm. The data includes a shapefile of road segments and IRI measurement. We include all roads in the National Highway System for IRI data, while for our Uber data, we only those roads that our data classifies as "highways".

C.5 International Roughness Index (IRI) in Cook county, 2018 (arterial roads, highways).

We use data on segment-level road roughness measured as the International Roughness Index (IRI) for Cook County provided by the Illinois Department of Transportation (IDOT) and collected as part of the Transportation Asset Management Plan (TAMP). Road quality information is collected 1) to prioritize highway rehabilitation needs, and 2) for incorporation into the Highway Performance Monitoring System (HPMS), the Federal Highway Administration's (FHWA) database used to report on the national road network and apportion federal funding to states for transportation needs. The IRI, along with other road measures, is collected by vans with laser sensors and on-board computers. This data is primarily available for highways and arterial roads, with very sparse coverage of residential roads.

Data on the road network is published with respect to a basemap, which provides the coordinates for each road segment on a model of the Earth. Different basemaps differ in the precise location of roads. Therefore, comparing across datas ets requires matching segments. We match OSM basemap road segments in Chicago to the TAMP basemap road segments based on the distance separating them, a fuzzy name match of the road names, and the angle generated by the intersection of tangent lines at the nearest point between the segments.

For each OSM segment, we assign it an IRI roughness level equal to the weighted average of overlapping TAMP segments, where the weights are the length of the overlap between segments divided by the length of the OSM segment. We also calculate a match quality measure that depends on how much TAMP segments extend away from the OSM segment, as in these cases the IRI measure for a given OSM segment may be particularly noisy.

C.6 Pavement Condition Index (PCI) in New York City, 2021.

We use data on the PCI on New York City streets provided by NYC OpenData. We limit our analysis to PCI estimates collected between June and October 2021.

C.7 Railway crossings at grade in Chicago.

We obtained the universe of at-grade railway crossings in Illinois from 2017 from the Illinois Department of Transportation, from https://gis-idot.opendata.arcgis.com/ (last accessed June 26, 2023).

C.8 Road resurfacing in Chicago, 2021.

We collected data on road resurfacing in Chicago. This data is based on moratoriums on street work. The City of Chicago imposes an increased permit fee on anyone who wishes to excavate a given street within five years of its last resurfacing. The city consequently provides the start and end address of the moratorium section of road and the moratorium start and end date. We infer the approximate resurfacing completion date by subtracting five years from the expiration date.[39] As we discuss when we return to resurfacing, these inferences are likely to contain some spatial and temporal noise.

[39] There are some concerns with the validity of the provided dates. Although the public website says the moratorium length is 7 years for resurfacing, discussions with the Chicago Department of Transportation confirmed that the correct length is 5 years. This is, indeed, the modal length between start and end dates in the data. However, there are also a number of odd lengths, possibly due to the start date corresponding to the data entry date rather than the construction date. For this reason, we follow the advice of the Chicago DOT and infer completion date as 5 years prior to the moratorium expiration date.

For this analysis, we compile the set of road resurfacing moratoriums completed in Chicago from May through mid-July, 2021. We match the repaired sections of road to our data by using the Google Geocoding API to turn addresses to coordinates, and an algorithm based on distance and a fuzzy street name match. These road segments are part of the treated or repaved group.

C.9 Speed limit data, 2023.

We collect road-segment level data on official posted speed limits for a random sample of road segments used in our town border analysis and for all road segments in our Chicago road resurfacing analysis. We use the HERE.com API. (HERE.com is a mapping data platform.) For the town border analysis, we compute the median speed limit at the town level.

C.10 Vehicle inspection failure rates data for Massachusetts, 2021-2022.

We obtained data on non-commercial vehicle inspection failure rates in Massachusetts from the Massachusetts Department of Transport (MassDOT/Registry of Motor Vehicle). We obtain data at the level of each inspection location (station) with the total number of inspections, average vehicle age, and total number of failures by category for 10 categories, for the period May 2021 to April 2022. We pre-selected seven "main" forms of inspection failure that we hypothesized could be linked to vehicle damage due to rough roads: brakes, front end, steering and suspension frame, muffler and exhaust system, bumpers/fenders/exterior sheet metal, and tires. We also pre-selected three "placebo" types of failure that we thought unlikely to be related to rough roads: windshield wipers and cleaner, windshield, and rearview mirror.

C.11 Vehicle Crash Fatalities for the entire United States, 2021.

We use data on the universe of vehicle crash fatalities from the Fatality Analysis Reporting System (FARS) published by the National Highway Traffic Safety Administration (NHTSA). We accessed data for 2021 from https://www.nhtsa.gov/file-downloads?p=nhtsa/downloads/FARS/.

C.12 Road resurfacing in a sample of cities, Sep 2021 - Dec 2022.

We collect data on resurfacing in Charlotte, NC, Chicago, IL, Columbus, OH, Dallas, TX, Denver, CO, Los Angeles, CA, New York City, NY, Philadelphia, PA, Portland, OR, San Diego, CA, San Jose, CA. For each city, except Chicago, we download the location of repairs from an arcGIS map published on each respective city's public website. We use data from September 2021 through December 2022 for New York, Dallas, and Portland, and data from only 2022 for Columbus and San Jose, which do not publish more precise data on times. Some cities publish all historical roadwork, some only resurfacing, and some publish moratorium roads. Whether or not the data comes from repairs or moratoriums is presented for each city below. To remove highway and arterial repairs, we drop any repaired segments that are contained by a 25-meter buffer of OSM basemap arterial and highway roads. The third annd fourth columns of the table present the original repair miles and the miles after restricting to local roads. Road segments are then assigned to the .4km x .4km grid cell they are in. The fifth columns gives the number of local repair miles in grids that we also have Uber data for. The final column gives the share of all grids in the city with Uber data.

We interpolate Census variables defined at the Census tract level to the grid cell level as follows. First, we assume that variables are uniformly distributed within each census tract, so we can use the area of the intersection between a tract and a grid cell. For "mean" variables, such as income, we take the weighted mean of the variable for the census tracts that intersect a given grid cell, with weights given by the area of the intersection. For aggregate variables, such as total population, we first apportion a Census tract's value to all of its intersections with grid cells, proportional to the areas of the intersections. Next, for each grid cell, we take the sum of the variable for the intersections.

Name	Data type	Total miles	Local miles	Miles in Uber grid	Fraction Uber grids
Charlotte	repairs	65.69	55.20	28.70	0.47
Chicago	moratoriums	283.11	247.84	236.57	0.88
Columbus	repairs	64.22	42.41	28.06	0.51
Dallas	repairs	368.46	180.96	75.22	0.37
Denver	moratoriums	121.82	108.72	75.84	0.54
LA	repairs	265.11	140.30	129.80	0.80
New York City	moratoriums	731.09	550.58	544.82	0.92
Philadelphia	moratoriums	249.47	173.60	172.49	0.86
Portland	moratoriums	36.10	18.52	7.02	0.43
San Diego	repairs	35.51	24.62	20.64	0.54
San Jose	repairs	208.22	195.81	92.50	0.35

C.13 Policymaker survey of road repair data and decisions.

The goal of this survey is to obtain data on road repaving strategies from municipalities across the US. The list of cities is derived from the 2010 Census and excludes Census Designated Places, as well as 15 towns for miscellaneous reasons (no contact information online, no department of public works, county responsible for road repaving, etc.). The sample includes a national sample, with a 16% response rate and a Massachusetts sample for which more follow-ups were conducted, achieving a 73% response rate. The contact protocol includes randomizing email send time, spacing 3 business days between emails and 4 days between emails and calls, and ensuring that every business day, 25 new cities and 25 previously contacted cities are emailed the survey.

We report here results corresponding to Tables 11 and 12 for cities that responded later to our survey in the Massachusetts sample (Table A.18).

Table A.18: Survey Responses by Late Response

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	In-house Survey	Private Survey	Officials	Engineering	Residents	Share Surveyed (%)		
Responds Late	-0.359*	-0.224	-0.083	-0.276	-0.263	-1.439		
	(0.195)	(0.160)	(0.080)	(0.199)	(0.179)	(15.012)		
Constant	0.667**	0.917**	0.083	0.583**	0.417**	62.273**		
	(0.140)	(0.116)	(0.058)	(0.144)	(0.129)	(10.843)		
Observations	25	25	25	25	25	23		
	Formula Road Conditions	Formula Other	Utility Work	Traffic	Citizen Complaints	Transportation Expert	Accessibility	Elected
Responds Late	-0.199	-0.263	-0.436*	-0.045	-0.096	-0.090	-0.090	0.064
	(0.204)	(0.179)	(0.186)	(0.207)	(0.166)	(0.134)	(0.134)	(0.166)
Constant	0.583**	0.417**	0.667**	0.583**	0.250*	0.167*	0.167*	0.167
	(0.147)	(0.129)	(0.134)	(0.149)	(0.120)	(0.097)	(0.097)	(0.120)
Observations	25	25	25	25	25	25	25	25
	Gets Required Resurfacing (%)	Resurfacing Spending Share (%)						
Responds Late	7.917*	2.652						
	(3.742)	(10.546)						
Constant	74.583**	38.182**						
	(2.646)	(7.618)						
Observations	24	23						

Note: These tables analyze heterogeneity of results from Tables 11 and 12. We sent reminder emails and calls to towns that did not fill out the survey. The variable "Responds Late" is an indicator for above median response time.

D Validation Appendix

D.1 Correlation with Highway IRI across the U.S.

In this section, we report validation results where we compare our road roughness measure with the IRI measure nationwide. The most recent FHWA data at the time of this study was collected between 2017 and 2018 while our data comes from 2021, and so we expect an imperfect correlation between the two measures.

Matching our data with the governmental IRI measure is difficult because unlike us, the FHWA does not use Open Street Map as a basemap. Consequently, the coordinates for the same roads differ. Also, the IRI is reported for road segments of different lengths than the OSM road segments. To match roads from the two data sources, we divide the US into square grid cells with lengths of 1, 10 and 100 kilometers. (Appendix C.4 has more details on data processing.)

The coarsest grid cells are the easiest to visualize: Figure A.22 shows a scatter plot of IRI roughness over Uber Roughness for national 10,000 km^2 cells. Figure A.23 plots the overall ranks for 10 km grids in Massachusetts. The basic pattern seems quite consistent. Roads in the core of the Boston metropolitan area are rougher than roads elsewhere in the state.

Figure A.22: Uber Roughness and DOT's International Roughness Index (IRI)

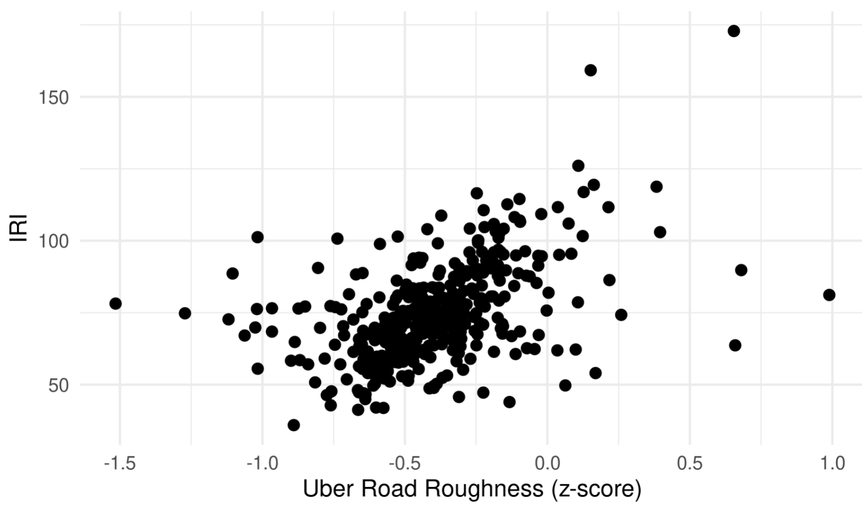

Note: This is a scatter plot at the level of 100 km × 100 km grid cells of highway roughness measured in two ways. On the X axis, we use the Uber data and predicted roughness at 48mph, from August 2021. We then convert roughness to units of standard deviations. On the Y axis we use the DOT IRI data for 2017-2018..

Figure A.23: Uber Roughness and DOT IRI in Massachusetts

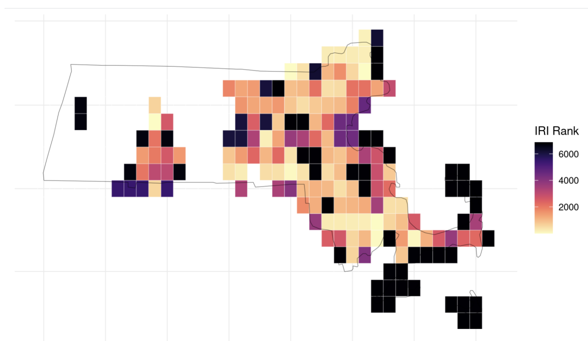

(a) Ranked road quality by IRI

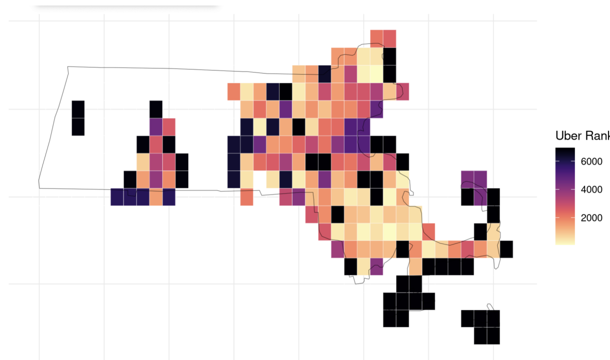

(b) Ranked road quality by Uber

Note: These maps show average highway roughness at the level of 10 km × 10 km grid cells in Massachusetts. Panel (a) uses IRI data from the DOT from 2017-2018, panel (b) uses Uber data. In both cases, the outcome is the rank of the average roughness in the grid in the nationwide distribution.

Table A.19, panel (a) shows Spearman's rank-order correlation for grid cells. The correlations range from .49 for the largest grid cells to .54 for the smallest. This panel also shows that the correlation between estimated highway lengths for the IRI roads and for our roads falls from .64 for the largest cells to .35 for the smallest cells. For the smaller grid cells, our roughness measure is more correlated with IRI than our highway lengths are correlated with the highway system highway lengths. This fact suggests that the imperfect correlation of the roughness measures is at least partially related to an imperfect match of the roads being measured.

Table A.19: Correlation between Uber Roughness and IRI

grid	roughness	length	sample size
10,000 km^2	0.493	0.635	421
100 km^2	0.537	0.530	7096
1 km^2	0.537	0.349	43746

(a) Grid cells, national

Class	Roughness	Sample size
Highway	0.735	654
Arterial	0.688	6919

(b) Segments, Cook County IL

Note: This table reports the Spearman rank correlation between IRI and Uber roughness. Panel (a) compares average for highways at the grid cell level, between the 2017-2018 DOT data and the 2021 Uber data, for three grid areas. The third column reports the correlation of total road length at the grid cell level. Panel (b) performs the same analysis at the OSM segment level, using 2018 IDOT data from Cook County, and 2018 Uber data. The segment matching procedure is described in Appendix C.5. We report the correlation of road roughness with IRI separately for the highways and arterial roads.

D.2 Correlation with New York City PCI

We next compare our measure to the NYC's Pavement Condition Index (PCI). The New York City Department of Transportation performs ongoing assessment of New York City local streets, rating the pavement quality on a scale from 1 to 10, with ratings 8 through 10 considered good. NYC measures a sample or roads every year and so contains some measurements from 2021, the primary year of our data collection.

There is a substantial literature linking PCI with IRI measures in different contexts. The difficulty with gleaning general facts about the correlations between these measures is that IRI is a far more standardized measure than PCI. ASTM International publishes a standard procedure,[40] but even that measure only "represents the collective judgement of pavement maintenance engineers." Moreover, the PCI measures used by individual jurisdictions can vary quite significantly from the ASTM standard.

The subjectivity and methodological variation with PCI measures helps explain why studies estimating the correlation between PCI and IRI vary widely. At one end of the estimates, Elhadidy et al., (2021) find that IRI could explain 99.2% of the out-of-sample variation in PCI. Suryoto et al. (2016) shows a correlation that is almost as high. At the other end of the estimates, Piryanosi and El-Diraby estimate an R-square of .31, although this increases to over .7 in some subsets of the data. Arhin et al. (2015) look at roads in the District of Columbia and estimate r-squared ranging from .53 for highways to .74 for local roads. Moreover, when PCI and IRI diverge, it is unclear which measure better captures driver discomfort and vehicle harm. However, in NYC, 91 percent of the PCI ratings for New York City local roads take on the integer values of 6, 7, 8 or 9, which suggests limited precision.

In New York City, inspectors perform a visual inspection of roads and rate them based on "overall condition, patching, and cracking." While more information on this procedure is not easily available, this strikes us as a fairly spartan set of variables.

We analyze PCI data collected between June and October 2021, within two months of the date of Uber data collection. Figure A.24 provides maps of road segment coverage for the PCI and Uber data in New York City. We aggregate the segment data at the level of 1 kilometer square grid cells, which leaves us with 617 1 km^2 grids with local road data.[41] The Spearman's rank-order correlation for the PCI and Uber roughness in NYC is .24. While this correlation is low, the difficulty of matching exact road segments makes comparison with previous work difficult. Nonetheless, it is positive and significant suggesting that both measures are getting at some common underlying attribute.

[40]See https://www.astm.org/d6433-20.html, last accessed August 2023. ASTM International was formerly "The American Society for Testing and Materials."

[41]We remove all grids where less than 50% of local roads have Uber observations.

Figure A.24: Uber and PCI data coverage on local roads in NYC

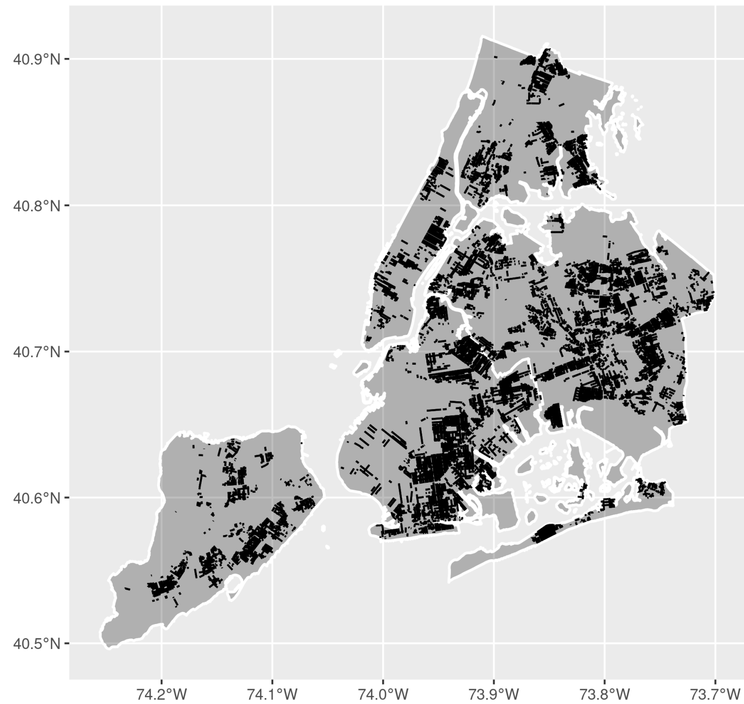

(a) PCI (June-October 2021)

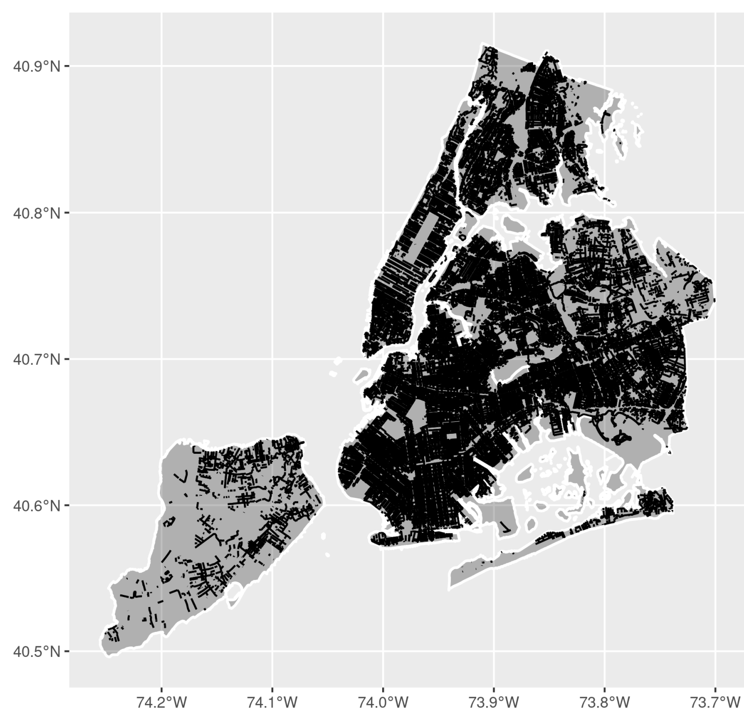

(b) Uber (August 2021)

Note: These figures show the local road segments with data for PCI and Uber roughness, respectively. PCI from New York City from June-October, 2021. Uber data is for August alone..

We test the relative predictive power of the city's PCI data with our Uber data, by using driving speed from the Uber data, averaged within grid cells. To avoid any potential mechanical correlation with Uber road roughness, we estimate speed and roughness from a split sample, where we use 75% of the sample to estimate road roughness profiles, and a separate 25% of the sample to measure speed. To avoid any possible impact of traffic congestion, our outcome measure is the 75th percentile of travel speeds along a given road segment. We then take the average of this measure within a cell. Regressions of speed on the two measures of roughness are reported in Table A.21.

Column (1) shows that Uber roughness strongly predicts slower drives and can explain 33% of the variation in speeds across cells. Column (2) show that PCI also predicts slower speeds, but with an R-squared below .01. Both measures are significant in column (3), but including Uber roughess adds .36 to the R-squared relative to the NYC PCI alone, while adding the NYC PCI adds .03 to the R-squared of the Uber roughness regression. The much greater power of the Uber data in predicting speeds is reassuring. However, the Uber roughness data has the advantage of corresponding to the exact same roads and month as the Uber speed data, while the PCI roughness data does not.

Table A.21: Driving Speed and Roughness Measures

	Dependent variable:		
	75th percentile of speed (mph)		
	(1)	(2)	(3)
Uber Rougness (z-score)	-5.152^{***}		-5.465^{***}
	(0.293)		(0.294)
NYC PCI		-0.145	-0.568^{***}
		(0.135)	(0.110)
Constant	19.832^{***}	19.265^{***}	24.146^{***}
	(0.141)	(1.007)	(0.847)
Observations	617	617	617
Adjusted R^2	0.333	0.0003	0.360

Note: This table reports the correlation between two measures of roughness and driving speed, at the level of $1km^2$ grid cells in New York City. The outcome is grid average of the 75th percentile of speed among a given road to avoid any bias from congestion. The mean "uncongested" speed is 18mph, while the mean speed is 15mph. The results are robust to replacing the 75th percentile with the mean speed outside the hours 5am to 9am and 4pm to 8pm; for example, the coefficients in the third column become -5.65, -0.67, and 21.52. Robust standard errors in parentheses, *p<0.1; **p<0.05; ***p<0.01.

E Model Appendix

E.1 Non-parametric Identification in Roughness Model

In this section we derive a non-parametric identification result for the marginal cost of roughness. We do this in the context of a slightly more general model, where the utility cost of driving over a segment is given by

$$cost_{ri}(s;\gamma_r) = \frac{v_i}{s} + L(s,\gamma_r), \tag{E.1}$$

where speed costs depend on the function L that is differentiable, increasing in both arguments and satisfies $L(0,\gamma) \equiv 0$.[42]

The first-order condition is

$$v_i(s_r^*)^{-2} = L_1(s_r^*, \gamma_r),$$

and the envelope theorem states that

$$\frac{dcost_{ri}^*}{d\gamma_r} = L_2(s_r^*, \gamma_r).$$

For identification we will assume that we know the following two comparative statics:

$$\frac{ds^*}{d\gamma} = \frac{-L_{12}}{L_{11} + 2v_i(s^*)^{-3}},$$
$$\frac{ds^*}{dv_i} = \frac{(s^*)^{-2}}{L_{11} + 2v_i(s^*)^{-3}}.$$

The first expression is the key empirical exercise in our paper – how optimal speed varies with road roughness. The second looks at how drivers change speed as the value of time increases. These two objects (evaluated at all values of (v_i, γ)) are sufficient for identifying the costs of roughness. To see this, note that from their ratio (and observed speed) we can estimate $L_{12}(s^*, \gamma)$. Note that by varying v we can obtain any optimal speed s^*, so we know L_{12} everywhere. We can then write

$$L_2(s,\gamma) = L_2(0,\gamma) + \int_0^{s^*} L_{12}(s,\gamma)ds$$

The first term is zero from our assumption that $L(0,\gamma) \equiv 0$. This concludes the proof that roughness costs are identified from knowledge of the two comparative statics with respect to road roughness and value of time. In the main paper, we do not have variation in the value of time to estimate the second comparative static. Instead, model identification follows from the parametric assumptions we make and from observing how driver speed depends on the speed limit.

[42]This is equivalent to writing the speed cost as $M(s\gamma, s)$ and assuming that M is continuous around $(0,0)$.