NBER WORKING PAPER SERIES

INFRASTRUCTURE INEQUALITY:
WHO PAYS THE COST OF ROAD ROUGHNESS?

Lindsey Currier
Edward L. Glaeser
Gabriel E. Kreindler

Working Paper 31981
http://www.nber.org/papers/w31981

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 2023

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, Alex Min, Taryn O'Connor, Michael Pak, Sam Patterson, Leo Saenger, 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. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

At least one co-author has disclosed additional relationships of potential relevance for this research. Further information is available online at http://www.nber.org/papers/w31981

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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2023 by Lindsey Currier, Edward L. Glaeser, and Gabriel E. Kreindler. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Infrastructure Inequality: Who Pays the Cost of Road Roughness?
Lindsey Currier, Edward L. Glaeser, and Gabriel E. Kreindler
NBER Working Paper No. 31981
December 2023
JEL No. L92,O18,R41,R42

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, and we find 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 the roughness of the median local road in the US generates welfare losses to drivers of at least 31 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 $318 per year more in driving pain if they live in a predominantly Black neighborhood than in a predominantly White neighborhood. Road roughness modestly predicts subsequent road resurfacing in New York City, but not in three other cities, which suggests that repaving is only weakly targeted towards damaged roads. Surveys from 120 towns and cities across the US suggest many reasons why resurfacing seems to be weakly targeted.

Lindsey Currier
Department of Economics
Littauer Center
Harvard University
Cambridge, MA 02138
United States
lcurrier@g.harvard.edu

Edward L. Glaeser
Department of Economics
315A Littauer Center
Harvard University
Cambridge, MA 02138
and NBER
eglaeser@harvard.edu

Gabriel E. Kreindler
Department of Economics
Harvard University
1805 Cambridge Street
Cambridge, MA 02138
and NBER
gkreindler@fas.harvard.edu

A data appendix is available at http://www.nber.org/data-appendix/w31981

1 Introduction

How bad are America's roads? 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) and estimates that America must spent over one trillion dollars over the next decade to rescue our "failing transportation infrastructure" (ASCE, 2020). Yet 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. That data does not cover the local roads that both make up a huge part of the driving experience and seem likely to contain far more of America's potholes.[1] Existing empirical work on US roads, which builds on previous theoretical analyses of optimal road quality for highways (Small and Winston, 1988), also largely focuses on highways because of data availability.

In this paper, we measure the roughness of America's local roads and highways using vertical accelerometer smartphone data provided by Uber.[2] This data is similar to the vertical acceleration measures of roughness used by the Department of Transportation, but the D.O.T. has the luxury of sending its drivers out on off-peak hours and requiring them to maintain a constant speed. A 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 observe significant variation in speed over the same road segment. We use this variation to estimate and recover the entire roughness-speed relationship at the road segment level. Equipped with this estimated relationship, we can compute predicted roughness at any given speed, which forms our primary measure of road quality. We estimate road roughness for over five million road segments in America that cover half a million kilometers and over one thousand towns in urban areas.

Section 3 presents our attempts to validate this data. The correlation between our measure and the D.O.T. measures of road roughness is 0.5-0.7 for highways and 0.7 for 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. We find a much lower (.24) correlation between our measure and the Pavement Condition Index (PCI) used in New York City, but our measure is far more strongly correlated than the PC is with traffic speeds. We also show that 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 vertical motion. When roughness is anticipated, drivers will slow down to reduce their vertical motion. If

[1] 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.

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

we assume a value of time, this behavioral response shows the willingness to pay to avoid roughness. Hence, our key empirical goal is to estimate this impact of salient, and hopefully anticipated, road roughness on driving speed.

We use median road segment speed from the Uber data as our measure of driving speed, and we focus on local roads, where roughness is more extreme and there are natural salient jumps in road roughness. This work builds on the analysis by Bock et al. (2021), who use an instrumental variable based on altitude and proximity to aquifers to show that roughness leads to lower driver speeds on highways in California. We first focus on discontinuities at 1,449 town borders, where we have sufficient Uber roughness data coverage on local roads near the border. We estimate an elasticity of speed with respect to roughness of around -0.3. We estimate the non-linear specification implied by our model and use our estimates to calculate the cost of road roughness. We also look at speed changes after repaving events in Chicago, as an alternate source of variation in road roughness. Comparing roughness and speed several months before and after these events, on treated road segments and control segments, we estimate a lower elasticity of speed with respect to roughness of -0.14, which may indicate lower awareness of variation in road roughness due to resurfacing. The repaving estimates actually imply a larger cost of road roughness, but we use the more conservative (lower) cost estimates from the border discontinuities in the rest of the paper.

The estimated model implies that the total costs of driving on rough roads are approximately 43% higher than the costs of travel time alone. If travel time's value is 15 USD per hour, then a road segment with median roughness costs 0.31 USD per mile more than a road with zero roughness. One standard deviation increase in roughness leads to 0.23 USD/mile higher costs. We use this figure to scale our roughness measure for local roads in the rest of the paper.

This cost depends directly on the assumed value of travel time (15 USD/hour), the estimated ratio of total costs to time costs (1.43), and the estimated relationship between chosen speed and road roughness (an elasticity of 0.3).[3] An advantage of our framework is that it allows readers who prefer other assumptions to adjust the cost of roughness numbers in the rest of our paper.

We supplement this exercise with more direct cost measures. We correlate our town-level roughness measures for Massachusetts with vehicle inspections results over the course of a year. Towns with rougher roads have a higher index of inspection failures, but this could just reflect car quality and maintenance levels in those towns, since we also find higher failure rates for damage to windshield wipers or mirrors, which are unrelated to road quality. We find that towns with rougher roads have fewer fatalities, even after controlling for population, income, and measures of employment and commute trips to work, which suggests that road roughness may generate benefits external to the driver who is choosing the speed. 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

[3]In the model, the speed-roughness relationship is non-linear, so the elasticity we cite here is an approximation.

our other controls, then combining the mortality-roughness and the speed-roughness relationships yields an implied elasticity of fatal crashes with respect to speed of about 1, which is lower than in much of the speed-mortality literature.[4]

Section 5 of the paper then uses the cost of roughness converted to dollars per mile and provides facts about the basic patterns of road roughness across towns and neighborhoods within the 100 largest metropolitan areas. We estimate jurisdiction level regressions, controlling for metropolitan area fixed effects. Cities and towns with higher 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 town, the average cost of roughness on local roads increases by 10.6 cents per road mile, or 318 dollars per year for a household that drivers 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. The coefficients on race at the neighborhood level within towns are considerably weaker than the coefficients on race within entire MSAs, but the relationship persists. As the share of Black households in a census tract increases from zero to 100%, the roughness-related costs of roads to drivers goes up by 1.6 cents per mile, or 48 dollars per year for a household that drives 3,000 miles on local roads. Income is also strongly correlated with road roughness at the neighborhood level. 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 4 large cities (New York City, Dallas, Portland, Columbus), where we have detailed data on road repaving events that occur after our road roughness measure. Our baseline measure of local road roughness only weakly predicts which road segments get resurfaced in New York City, and 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 of traffic levels (as measured by Uber use) and road quality. The observed repaving patterns in New York City are compatible with 10-12% 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.

To better understand that data, we ran a survey of 19 towns in the greater Boston area and 101 towns nationwide to learn how they asses road quality and how they target road repairs. Typically, towns have a measure of visual road roughness, and they take several factors into account when deciding which roads to resurface, including road conditions, but also traffic intensity, upcoming

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

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 are possibly 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. The value of better targeting may be quite important 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. Both GPS receivers and accelerometer sensors are ubiquitous in modern smartphones. Smartphone accelerometer readings are extremely precise compared to state-of-the-art acceleration sensors (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."[5] Smartphones record acceleration along three axes aligned with the physical device.[6] In order to identify sharp horizontal acceleration events such as harsh breaking, Uber created a proprietary algorithm to re-align these axes to the vehicle. This algorithm is aided significantly by the fact that drivers 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

[5] 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.

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

vertical acceleration) as our primitive measure to construct road roughness. Currently, we ignore the potential bias arising from sloped roads.

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.

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

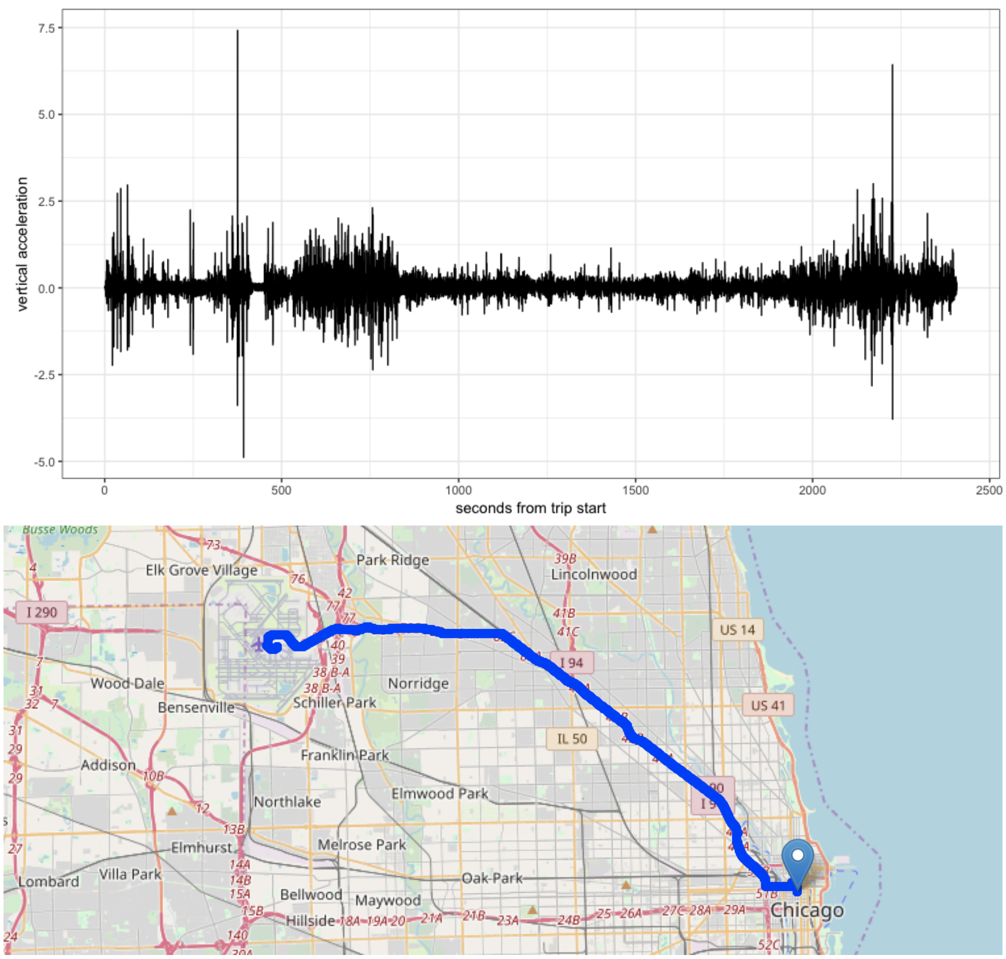

Note: These graphs show raw measures of vertical acceleration during a single Uber trip. The Z-acceleration measure is less variable during the middle portion of the trip, along highways, consistent with smoother pavement.

We begin with anoynmous 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, as well as April 2018, the latter of which we use only for testing and validation.

We first match raw GPS coordinates to a basemap of roads. 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, we aggregate to the level of an Uber trip by road segment by taking the standard deviation of acceleration. This leaves us with roughly 2 TB of data after aggregation.

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. We use the terms "vertical acceleration" 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. As drivers will drive at lower speed on segments with lower road quality, a failure to correct for driver speed will artificially compress the heterogeneity in road roughness.

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 to support the hypothesis that non-linearities in the relationship between speed and acceleration are impactful. For example, Appendix Figure A.1 shows that, for a single highly traveled segment in Chicago, measured vertical acceleration is close to linear in speed. To test for driver-specific road segment slopes, we use instances of the same driver going over the same segment repeatedly, and estimate the link between speed and vertical acceleration for the segment using within-driver variation in speed. We can then compare the driver-specific segment slopes with the overall segment slope, excluding that driver. Appendix Figure A.2 shows the results for the same highly travelled segment as Appendix Figure A.1. This graph again supports the view the driver slopes are typically similar to the overall segment slope.

Estimation. We estimate equation (2.1) separately for each Uber city.[7] The main challenge for estimation is the size of our data. For computational tractability we 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. 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).[8] The fact that we include driver fixed effects in equation (2.1) means that we are essentially comparing road segments holding the quality of the vehicle fixed. Combined with our approach to estimate the entire speed-roughness relationship, this method brings us close to a standardized measurement of road roughness.

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.[9] 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.

[7]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.

[8]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.

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

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. The road roughness measure also exhibits significant heterogeneity within Chicago. 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.[10] While our primary interest in cracks, potholes, and other road pavement irregularities that arise from deterioration, deliberately chosen aspects of bumpiness will also be captured by any measure that captures vertical acceleration. It is appropriate to include that chosen bumpiness when we estimate the link between roughness and travel speeds. However, when we measure the relationship between bumpiness and either income or race, we anticipate that chosen bumpiness will bias relationships downward since we anticipate that chosen roughness is 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 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.

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

Figure 2: Predicted Roughness in Cook County, IL

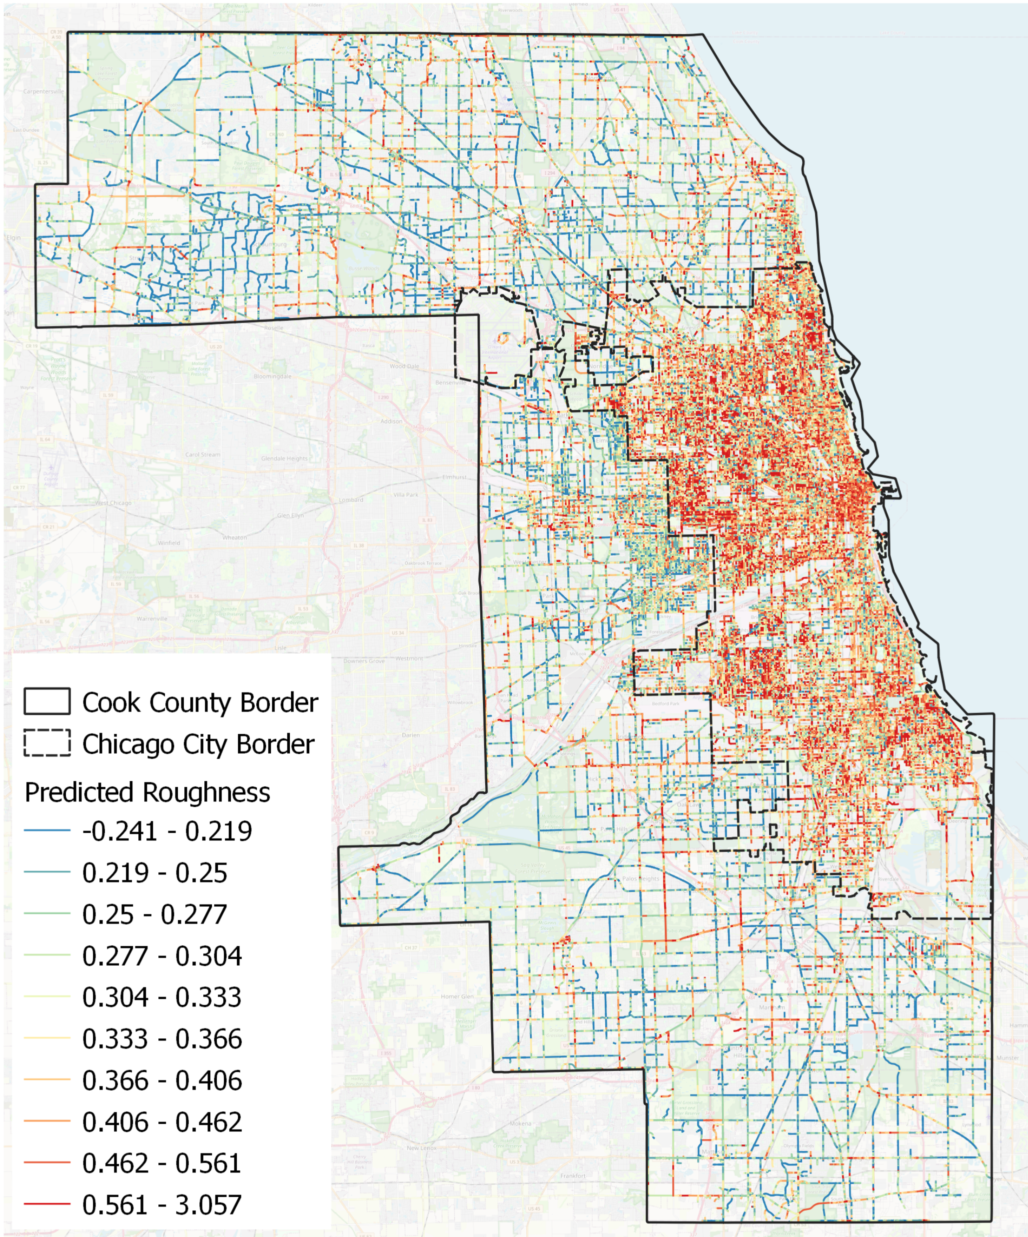

Note: This map plots predicted road roughness for all local and arterial road segments in Cook County. Colors correspond to deciles of the predicted roughness distribution. The city boundary of Chicago is shown in the dashed black line. Figure A.6 shows an analogous map for highway roughness.

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

One concern is that our road segment estimates may be very noisy. Despite having a large amount of data, we are estimating over 16.4 million parameters: an intercept and a slope for 8.2 million directed road segments across the US. To understand whether this is an issue in practice, we perform the following cross-validation exercise in Chicago, where we have multiple months of data. We estimate predicted roughness using data from March, 2021 and then repeat the exercise and estimate predicted roughness using a completely separate data set from April, 2021. Note that during this period road roughness may change in reality due to deterioration or road work, making the comparison a lower bound on the reliability of our measure. Table A.3 reports regression results of predicted roughness from April on the same measure from March. This yields 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. These results establish that our measure of road roughness captures a persistent characteristic of the road segment rather than estimation noise.

The key empirical exercise needed to estimate the model of driver costs that we will set up is to measure the impact of road roughness on driver speed. This raises two potential concerns. First, given that we use speeds when we estimate equation (2.1), if our linear specification is misspecified and the coefficient estimates contain some bias from speed, putting speed again on the left hand side will also lead to biased estimates on the relationship between roughness and speed. Second, our estimates of the impact of roughness on speed may be biased in the presence of road traffic congestion, which limits drivers' ability to choose their speed. We perform the following split-sample exercise to check these issues. Focusing again on 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, and we do this exercise 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 statistically significantly different at conventional levels, 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.

Our Uber data only covers a subset of all road segments in the US, and that may affect exercises we do later in this paper, where we use our segment-level estimates to study differences between places in the US. Using the Uber data, we obtain fixed effect estimates for 5.6 million undirected road segments, which represents 14.8% of all the 38.1 million undirected road segments in our entire base map. The road segments with fixed effects cover a total of 516,138 kilometers of road, or 6.7% of the entire road network. We have estimates for 8,211 towns or designated places in the US, out of 29,526 in total. Table A.1 reports coverage broken down by road type. It also shows that when we weight by town population, segment coverage is 12-14% for local roads, and 63-75% for arterials and highways, depending on whether we use the number of segments or total segment length. Figure A.4 shows that our measure's coverage increases rapidly in the population level.

The key concern with partial coverage is that our sample of segments may not be representative in terms of road quality of all road segments in a given area. This would happen if the road segments that have large Uber coverage are systematically more (or less) rough compared to all road segments in the area. While we cannot test this hypothesis for road segments without Uber data, we can 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 is encouraging. This suggests that road segment selection into the Uber sample is not biasing our results.

2.3 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 five 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. Second, we compare the geographic patterns for the Uber measure and for IRI data. Third, 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. Fourth, in the city of Chicago we test whether our measure of road roughness is significantly higher on road segments that intersect railroad tracks. Fifth, 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.

3.1 Correlation with Highway IRI in Cook County and across the U.S.

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.[11] This data is generally produced by state employees who drive vans with laser sensors and onboard computers at fixed speeds. 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.

A further difficulty with matching our data with the governmental IRI measures is that, 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. We report results from two exercises. First, nationwide, we divide the US into square grid cells with lengths of 1, 10 and 100 kilometers. Second, in Cook county, Illinois, we perform a "fuzzy match" between our OSM road segments and the IRI basemap road segments. Appendices C.4 and C.5 have more details on data processing.

The coarsest grid cells are the easiest to visualize: Figure 3 shows a scatter plot of IRI roughness over Uber Roughness for national 10,000 km^2 cells. Figure 4 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.

[11] 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.

Figure 3: Uber Roughness and DOT's International Roughness Index (IRI)

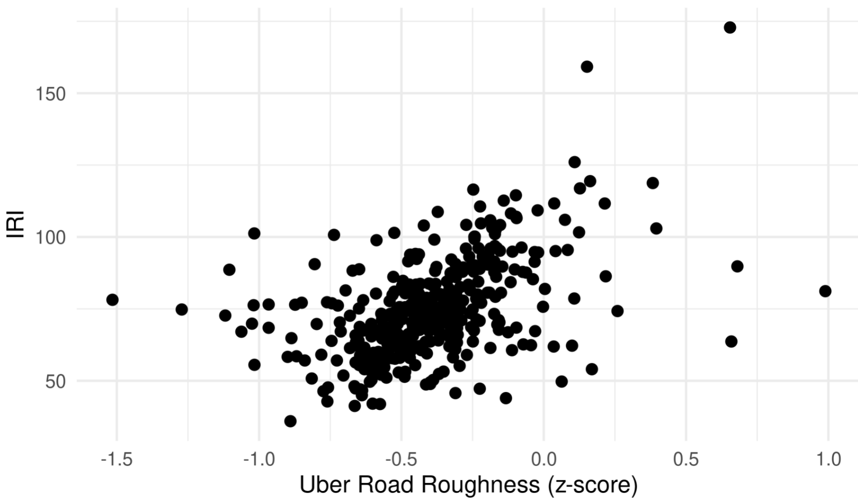

Note: This is a scatter plot at the level of 100 km × 100 km grid cells of highway roughness measured in two ways. On the X axis, we use the Uber data and predicted roughness at 48mph, from August 2021. We then convert roughness to units of standard deviations. On the Y axis we use the DOT IRI data for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

Figure 4: Uber Roughness and DOT IRI in Massachusetts

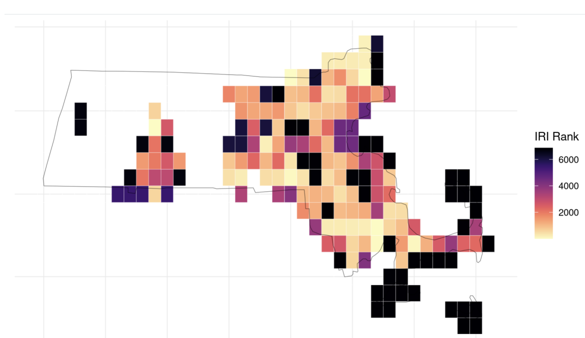

(a) Ranked road quality by IRI

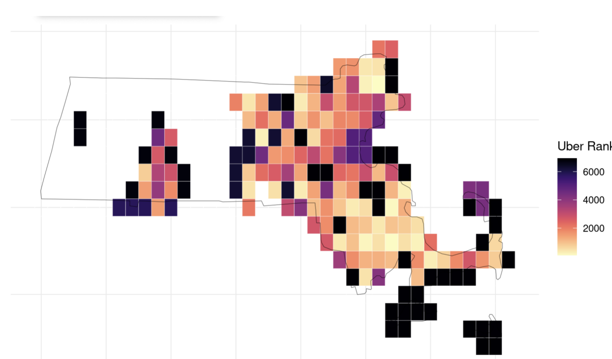

(b) Ranked road quality by Uber

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 1, 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 1: 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.

We now turn to our segment-level analysis for Cook County. We use 2018 IRI data directly published by the Illinois Department of Transportation, which is more extensive and 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 different sample of Uber data in Cook County from March, 2018.

Table 1, Panel (b) show the correlations between our road roughness measures and the IRI measure for highways and arterial roads in Chicago. The correlation coefficients between IRI roughness and our measures are .74 and .69 for highways and arterial roads respectively. Figure 5 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 temporal mismatch between the two data sources.

Figure 5: Uber Roughness and IRI Segments

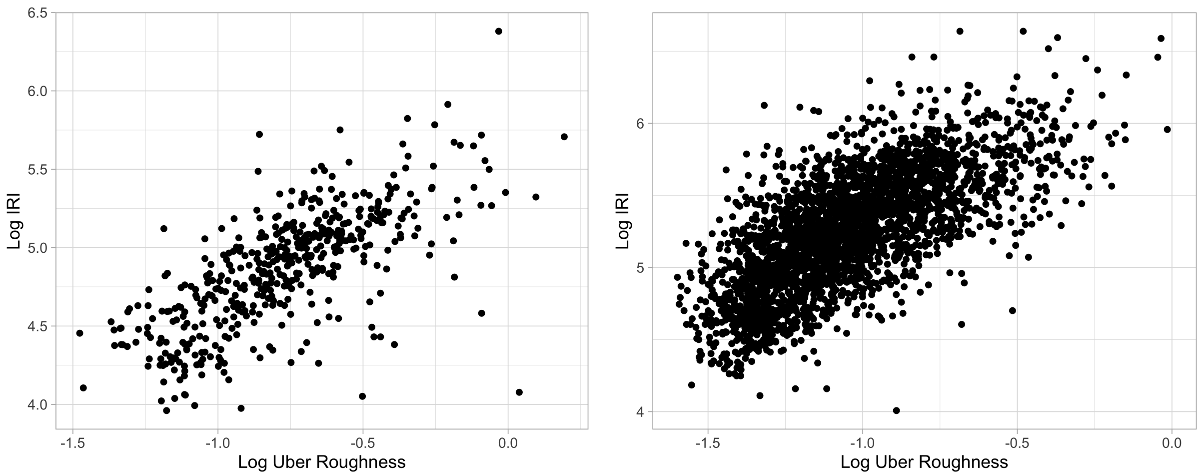

(a) Highway Roads (b) Arterial Roads

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.

3.2 Geographic Patterns

Our next validation exercise is to look at whether our data shows the same regional patterns as the D.O.T. roughness data for highways and arterial roads. Table 3 shows the importance of regional factors, which may be related to weather and soil quality, in the Uber data. The first regression shows that 18 percent of variation in highway roughness across towns can be explained by metropolitan area (MSA) fixed effects. The second regression finds that a larger amount, 31 percent, of arterial roughness is explained by MSA fixed effects. The third regression shows that only 9 percent of the variation in the roughness of local roads can be explained by MSA fixed effects. There is far more variation in road quality within metropolitan areas than across metropolitan areas, and this is especially true for local roads which are typically maintained by local governments.

Table 3: 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.

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 4. 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 4 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. All 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). Cold winters mean that groundwater freezes, and as it freezes the ground expands and 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, but the coefficient is not statistically distinct from zero in four of the five regressions. 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), with the former statistically significant at the 10% level.

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 coefficient on highways and arterials are $-.25$ 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. These estimates seem reasonably close to us. 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.

The final row looks at the coefficient on the logarithm of metropolitan area population. In all four specifications, a 1 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 coefficents are .36 and .38. For arterials, the two coefficients are .48 and .40. The coefficient in the local roads regression is substantially higher.

Table 4 both highlights broad regional patterns, and shows similarities and differences between the two data sets. Every coefficient had the same sign with the two data sets, and in many cases, the magnitudes were also similar. The largest differences occurred with arterial roads and the weather variables.

Table 4: Road Roughness and Climate

	Road Roughness (z-score)				
	highway	arterial	local road	highway	arterial
		Uber		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 largests 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 a dummy 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.

3.3 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,[12] 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.11 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.[13] 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.

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

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

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

within a cell. Regressions of speed on the two measures of roughness are reported in Table 5.

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 5: 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.

3.4 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.

We begin by analyzing a single road segment in Chicago and using the full richness of our data to map experienced roughness at precise locations within the road segment. Figure A.12 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 displays strong systematic patterns on different sections of the road segment, including around the railway crossing.

We next assemble a data set of all the local road segments with railway crossings in Cook County, and nearby and comparable control segments. We begin from the Illinois Department of Transportation data set of the universe of grade (street-level) railway crossings. We match each crossing to the nearest road segment in our data, cross-checking that the segment intersects an active railway based on Open Street Maps. For the control group, we next 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.

Road segments with railways are significantly rougher compared to nearby control segments (Figure 6). Predicted road roughness at 20 mph is 63% higher on these road segments, with a p-value < 0.001. This is equivalent to 1.27 standard deviations of the outcome the control distribution. 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 .97 mph lower (p-value 0.06) or 0.21 standard deviations lower, from a basis of 15.4 mph on the nearby control segments. These results are based on averaging over the entire road segment, so the differences narrowly around the railway crossing are likely to be even starker. Overall, these results show that the Uber roughness data picks up real differences in road quality.

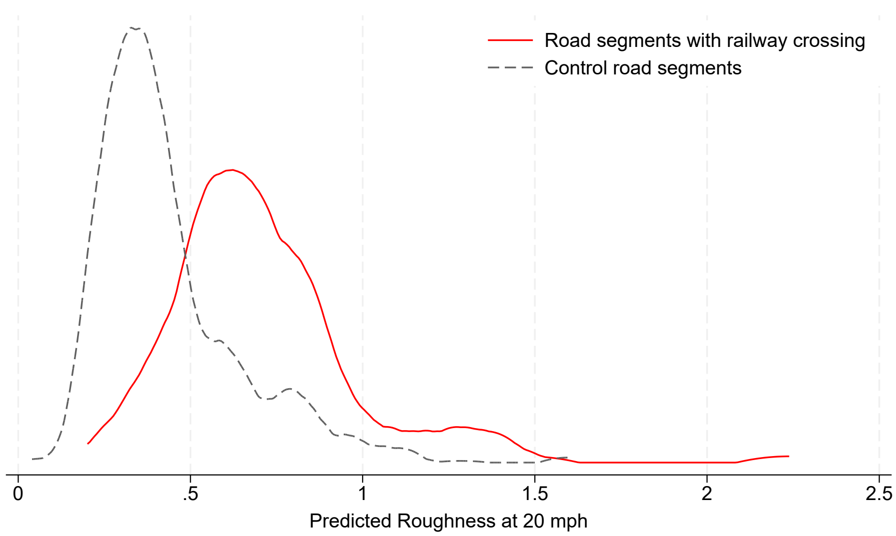

Figure 6: Roughness on Roads with Railway Crossings

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).

3.5 Impact of Resurfacing in Chicago

Our final validation exercise looks at the impact of road resurfacing on road roughness and driving speeds measured with the Uber data. 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 on the data and how we process it.

We analyze all 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. We use this donut strategy to avoid falsely tagging repaved or partially repaved road segments as control, while keeping control segments as similar as possible along unobservable dimensions. We restrict the sample to local roads. 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 7 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. Our outcome is the raw standard deviation of the vertical acceleration. Recall that this measure combines the road's intrinsic roughness with drivers' endogenous choice of speed. Using this measure as outcome will tend to attenuate differences in roughness because drivers slow down to compensate for high roughness, and vice-versa. Nevertheless, using this outcome allows us to have high temporal resolution.

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)). These results give a first indication that road resurfacing affects both roughness and driver speeds. 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.

[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.

Figure 7: Experienced Roughness and Speed on Repaved Road Segments

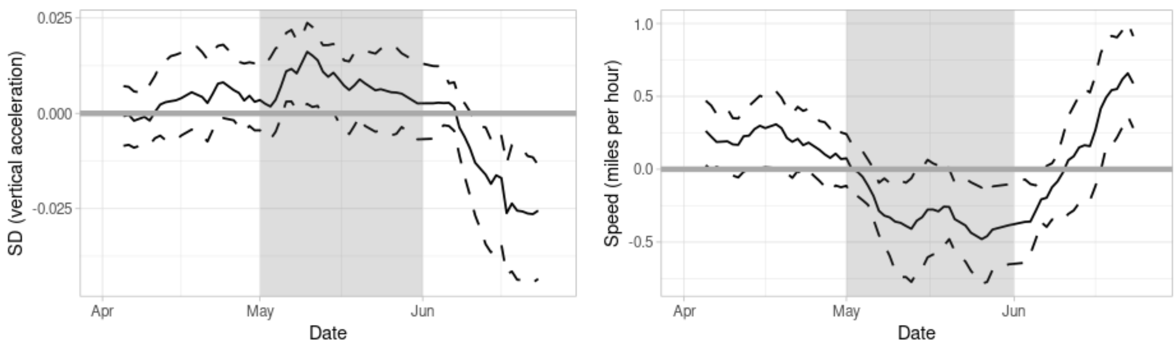

(a) Experienced Roughness

(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 treated segments. We plot point-wise 95% confidence bands based on bootstrapping the entire procedure at the level of clusters of road repair events.

4 The Speed Costs of Rough Roads

4.1 A Model of Driver Optimal Road Segment Speed

In this section, we present a simple 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. The structure of the model, in turn, allows us to estimate the direct benefit by looking at the shape of the relationship between speed and roughness. The total cost of rough roads from our model simplifies to a multiple of the time costs. Readers who are skeptical about the assumptions needed to estimate the model's parameters and who have their own estimate of the relationship between total costs and time costs can simply multiply our cost estimates accordingly.

Our model of driving costs assumes that total costs are an additive function of time and speed-related costs. The latter include both roughness-related damage or discomfort and other costs of driving fast, which we proxy by driving speed relative to the speed limit. The model predicts that at the optimal speed, time costs are a fixed multiple of speed-related costs and that multiple depends on the elasticity of optimal speed to speed costs. We estimate this elasticity parameter using two salient sources of variation in road roughness: town borders and repaving events. 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.

4.1.1 Model Setup and Optimal Speed

We consider the optimization decision of a driver on a 1 mile-long road segment r. The utility cost of driving on segment r during trip i is the sum of the cost of driving time and speed costs:

$$c_{ri}(s) = \underbrace{\frac{v_i}{s}}_{\text{time cost}} + \underbrace{\frac{\beta_{ri}}{\sigma}\left(\gamma_r s + \theta \frac{s}{s_r^{\text{LIM}}}\right)^{\sigma}}_{\text{speed cost}}. \tag{4.1}$$

where v_i is the value of (driving) time in dollars per hour, and s is the chosen speed in miles per hour. Speed costs are a constant-elasticity transformation of the sum of an experienced roughness term, $\gamma_r s$, and other costs of speed (e.g. the increased probability of an accident). The term β_{ri} captures segment, driver, and trip-level variation in the cost of driving, and acts as a speed cost shock, and we assume it is given by

$$\log(\beta_{ri}) = -(1+\sigma)(B + B_r + \xi_{ri}), \tag{4.2}$$

where B is a constant, B_r is a mean-zero road segment-specific speed cost term, and ξ_{ri} is a mean-zero random shock at the level of trip i and road segment r. The term B_r captures systematic factors that impact the cost of speed on a road segment. We return to this term when discussing endogeneity issues when estimating the impact of roughness on speed.

The term $\gamma_r s$ captures experienced roughness, which is linear in speed s with a slope γ_r.[15] We assume that other speed costs, $\theta \frac{s}{s_r^{\text{LIM}}}$, scale linearly with speed and are inversely proportional to the speed limit. This captures the facts that driving over the speed limit increases the probability of receiving a fine and that speed limits are a proxy for other costs of driving fast, such as the risk of accidents. The parameter θ determines the relative importance of other speed costs.

The driver observes the road roughness slope parameter γ_r, the speed limit s_r^{LIM}, and the cost parameter β_{ri}, and chooses the speed that minimizes cost. The first-order and negative second-order conditions imply that optimal speed is positive, interior, and given by

$$(s_{ri}^*)^{-(1+\sigma)} = \frac{\beta_{ri}}{v_i}\Big(\gamma_r + \frac{\theta}{s_r^{\text{LIM}}}\Big)^{\sigma}. \tag{4.3}$$

Drivers choose to go slower on roads where experienced roughness increases faster with speed (as captured by γ_r), where the speed limit is lower, and where the speed cost term β_{ri} is higher. The elasticity σ controls the shape of the relationship between road segment characteristics and optimal speed.[16]

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

[16] In the limit as σ tends to zero, optimal speed does not depend on roughness, while if σ is negative, optimal speed

Estimating the roughness-speed relationship. Substituting β_{ri} from (4.2) into (4.3) and taking logs yields:

$$\ln(s_{ri}^*) = \frac{1}{\sigma+1}\log(v_i) + B - \frac{\sigma}{\sigma+1}\log\Big(\gamma_r + \frac{\theta}{s_r^{\text{LIM}}}\Big) + (B_r + \xi_{ri}).$$

This equation captures how optimal speed depends on the roughness slope γ_r, speed limit s_r^{LIM}, and the speed cost term B_r. We estimate this equation, given data on speeds s_{ri} and an estimate γ_r of the roughness slope, with

$$\ln(s_{ri}) = \alpha - \frac{\sigma}{\sigma+1}\log\Big(\gamma_r + \frac{\theta}{s_r^{\text{LIM}}}\Big) + \epsilon_{ri}, \tag{4.4}$$

where α is our estimate of $\frac{1}{\sigma+1}\log(v_i) + B$ and the residual ϵ_{ri} captures the sum $B_r + \xi_{ri}$.

The main challenge to estimation is that the error term ϵ_{ri} includes the road segment costs B_r that may be correlated with roughness γ_r. To address this, we will use quasi-random variation in the roughness slope γ_r. Section 4.2 describes our two separate approaches for identification, using road resurfacing events, and town border comparisons.

Total costs of road roughness. The total user cost of driving at the optimal speed is

$$c_{ri}^* = \frac{v_i}{s_{ri}^*}\Big(1 + \frac{1}{\sigma}\Big). \tag{4.5}$$

The effect of road roughness γ_r on total user costs has two key components. First, roughness affects optimal speed s_{ri}^* and hence the value to time spent driving v_i/s_{ri}^*. This is the relationship that we will estimate in the data, using quasi-random and anticipated variation in road roughness. This effect is stronger for higher σ.

Second, total user costs are a multiple of the value of time spent driving. In our model, this multiplicative factor is fixed and only depends on σ, which is estimated when we link speed with road roughness. The factor is decreasing in σ, which means that when we observe drivers decrease their speed more in response to rougher roads, the time costs of roughness account for a higher share of total costs.[17]

We define the total cost of roughness as the total cost of driving over a road segment of a certain roughness minus the cost of driving over the same road segment where $\gamma_r = 0$, or

$$C_{ri} = \Big(\frac{v_i}{s_{ri}^*} - \frac{v_i}{s_{ri}^0}\Big)\Big(1 + \frac{1}{\sigma}\Big),$$

is increasing in roughness.

[17]The overall effect of roughness on total costs has an ambiguous relationship with σ when we hold fixed the optimal speed at the median value of road roughness, as motivated by our two empirical settings. See Appendix D.1 for details.

where s_{ri}^0 is the optimal speed implied by (4.3) with $\gamma_r = 0$.

To standardize analysis, we will evaluate this 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.[18] We will evaluate C at the median cost shock $\overline{\beta}_{ri}$ and median speed limit ($\overline{s}^{\text{LIM}} = 25\text{mph}$).

4.2 Empirical Results: US-wide Town Border Discontinuities

To calculate the costs of roughness, 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.[19]

To illustrate our approach and provide a first round of evidence of this channel, we first show results from the border of Chicago. Figure 8 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. These differences are large, consistent with substantially worse roads in Chicago relative to the neighboring wealthier towns. The changes are sharp around the boundary and remain relatively stable away from the boundary.

[18] 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.

[19] 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.

Figure 8: Predicted Roughness and Speed around the Chicago border

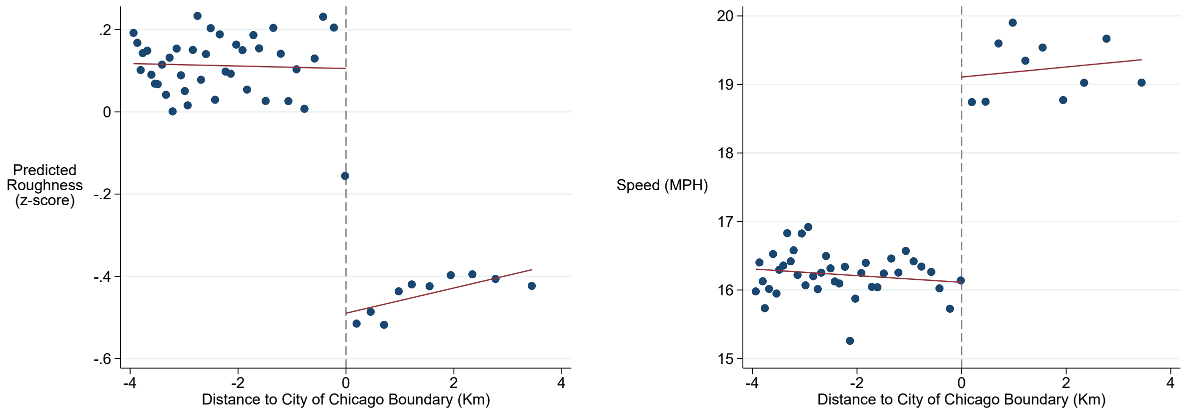

Panel (A) Predicted Roughness

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.13 repeats the exercise for arterial roads..

4.2.1 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, controlling for speed limits. 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 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. Our sample includes all the local road segments within 1000 meters to the border.[20] This leaves us with a sample of 1,257 towns covering 1,449 border pairs and 76 Uber cities. We follow equation (4.4), and use data at the level of each town border pair b, side of the border $j \in \{0, 1\}$, and road segment r, to estimate

$$\log(s_{bjr}) = \psi_b - \frac{\sigma}{\sigma+1} \log\left(\gamma_r + \frac{\theta}{s_{bj}^{\text{LIM}}}\right) + \epsilon_{bjr}, \tag{4.6}$$

where s_{bjr} is the median speed on road segment r measured in the Uber data, and ψ_b are border pair fixed effects, which means we only use variation within town border areas. We measure s_{bj}^{LIM} using the median speed limit in the town on side j of border pair b, based on the speed limit data

[20]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).

from HERE.com, a mapping data platform.[21]

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.[22] 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 equation (4.6) using IV-GMM and instruments for each border pair interacted with one of the sides $j = 1$. (The sides are defined arbitrarily for each border.) Our moments are $\mathbb{E}_r \epsilon_{bjr} 1(b = b', j = 1) = 0$ for all border pairs b'.[23]

The exclusion restrictions we impose require that for each border pair, the speed cost terms $B_r = B_{bjr}$ from equation (4.2) are in expectation equal on the two sides of the border. This amounts to assuming that in the narrow 1km band around each border, the only systematic variation in log speed on the $j = 1$ side versus the $j = 0$ side of the border is due to the term in (4.6) that depends on road roughness and speed limits.

We also estimate a specification with flexible controls for distance from a road segment to the border and with separate coefficients for each side of each border pair:

$$\log(s_{bjr}) = \mu_b - \frac{\sigma}{\sigma + 1} \log\left(\gamma_r + \frac{\theta}{s_{bj}^{\text{LIM}}}\right) + \xi_b Dist_{br} + \zeta_b Dist_{br} \cdot 1(j = 1) + \epsilon_{bjr}. \tag{4.7}$$

In this case, we also allow the speed cost term B_{bjr} to vary with the distance from the town border:

$$B_{bjr} = \xi_b^B Dist_{br} + \zeta_b^B Dist_{br} \cdot 1(j = 1) + \epsilon_{bjr}^B.$$

The exclusion restrictions in this specification require that for each border pair, the cost shock residuals ϵ_{bjr}^B are in expectation equal on the two sides of the border.

Table 6 reports our IV-GMM estimates and confidence intervals constructed using a Bayesian bootstrap procedure at the level of border pairs (Rubin, 1981).[24] In addition to the coefficient $-\frac{\sigma}{\sigma+1}$ from equation (4.6), we also report the term $1 + \frac{1}{\sigma}$, which is the multiplicative factor on pure

[21] For data details, see Appendix C.

[22] 20mph is approximately the median speed for local roads in the sample.

[23] Technically, we also need moments $\mathbb{E}_r \epsilon_{bjr} 1(b = b') = 0$ for every border pair b', in order to pin down the border fixed effects $\psi_{b'}$. In practice, for every assumed σ and θ, we compute residuals $\epsilon_{bjr}^1 = \log(s_{bjr}) + \frac{\sigma}{\sigma+1} \log(\gamma_r + \frac{\theta}{s_{bj}^{\text{LIM}}})$ and regress ϵ_{bjr}^1 on border dummies to obtain residuals ϵ_{bjr}^2, using which we compute the moment conditions $\mathbb{E}_r \epsilon_{bjr}^2 1(b = b', j = 1) = 0$ for each b'. We use the identity weight matrix to obtain the GMM objective function.

[24] We also estimate log-linear approximations in Appendix Table A.5. These results also show transparently that the log speed limit is a predictor for driving speed, but its effect is orthogonal to the speed impact due to town border variation in roughness. The headline elasticity of speed with respect to road roughness is -0.3 in columns (7-8). This is the number we cite in the introduction.

time costs in the expression for total user costs from (4.5).

In column 1 we estimate the model using cross-town variation in roughness, using town dummies as instruments.[25] The speed cost term — itself a composite of road roughness and inverse speed limits – has a negative, statistically significant effect on the speed chosen by Uber drivers. The magnitude implies an elasticity σ of 4.4, which implies that total user costs are 1.23 times larger than the value of time.

We next estimate equation (4.6) in column 2. These estimates use variation in speeds and road segment speed slopes up to 1km on each side of the town border, and control for the effect of town-level median speed limit on each side of the border. The effect of the speed cost term remains negative but falls in magnitude, implying a smaller elasticity and larger total cost factor of 2.42.

Column 3 and 4 focus on variation closer to the border. In column 3, we restrict to observations within 500 meters of the border. The coefficient on log speed costs is -0.63 and remains statistically distinct from zero.

The last column reports the estimates from the regression discontinuity specification (4.7), our preferred specification, where we again include all road segments up to 1000 meters from the border, controlling flexibly for distance to the border and estimating separate coefficients on each side of each border pair. The log speed cost term coefficient is -0.7 and the coefficient on the speed limit cost is 0.33. These results imply a total cost multiplier of 1.43. In other words, the total costs of roughness are 43% higher than the value of travel time costs implied by the change in speed at different roughness levels.

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.

[25] The sample includes local road segments with above median number of observations, restricted to towns that enter the town border sample, where we have speed limit data. Using data at the level of road segment r in town k in Uber city $c(k)$, we estimate the following equation using IV-GMM using town-level dummies as instruments:

$$\log(s_{kr}) = \mu_{c(k)} - \frac{\sigma}{\sigma+1}\log\left(\gamma_r + \frac{\theta}{s_k^{\mathrm{LIM}}}\right) + \epsilon_{kr}. \tag{4.8}$$

Table 6: The Impact of Roughness on Driver Speed on Local Roads at Town Borders (IV-GMM)

	Dependent variable:			
	Log speed (mph)			
	(1)	(2)	(3)	(4)
Log speed costs $\frac{\sigma}{\sigma+1}$	-0.82	-0.41	-0.63	-0.70
	[-1.21 , -0.61]	[-0.64 , -0.26]	[-0.88 , -0.41]	[-0.91 , -0.45]
Inverse speed limit θ	0.05	0.16	0.35	0.33
	[0.00 , 0.18]	[0.04 , 0.34]	[0.16 , 0.58]	[0.19 , 0.40]
Total user cost factor $1+\frac{1}{\sigma}$	1.23	2.42	1.59	1.43
	[0.83 , 1.65]	[1.57 , 3.85]	[1.14 , 2.45]	[1.10 , 2.24]
Sample:	Town	Borders	Borders	Borders
Sample restriction:		$<$ 1km	$<$ 500m	$<$ 1km
Instruments (dummies):	Town	Border side	Border side	Border side
Uber city FE	Yes			
Border pair FE		Yes	Yes	Yes
Distance to border controls				Yes
Uber cities	90	76	72	76
Towns	1,516	1,257	1,209	1,257
Border pairs		1,449	1,285	1,449
Observations	566,046	219,225	109,097	219,225

Note: This table reports IV-GMM estimates of the model-implied non-linear regression equations (4.8) (column 1), (4.6) (columns 2 and 3), and (4.7) (column 4). In column 1, we use town dummies as instruments, while in columns 2-4 we use town border pair side dummies as instruments. Column 4 includes controls for the distance to the border, with a separate coefficient for each border and each side. In parentheses, 95% confidence intervals based on bootstrap at the town level (column 1) or border pair level (columns 2-4).

4.2.2 Total user costs of roughness.

We now use the estimated model and equation (4.5) 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 construct total costs, we use the estimate of σ from column 4 in Table 6, which is our preferred empirical specification. 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.[26]

Figure 9 shows the implied relationship between the roughness slope γ_r and cost. We include the histogram of γ_r across road segments in the border sample for reference. The solid blue line plots total user costs, which include the cost of time $\frac{v_i}{s^*_{ri}}$ as well as the speed costs due to roughness and due to speed limits. The dashed black line plots the hypothetical cost of driving 1 mile at zero roughness when $\gamma_r = 0$.

How costly is road roughness? A road segment with the median roughness slope has a total user cost of 1.05 USD per mile, while a road segment with one standard deviation larger roughness slope has a total user costs of 1.27 USD per mile.[27] Hence, a one standard deviation of road roughness increases user costs by 0.23 USD/mile, with bootstrapped 95% confidence interval of [0.20, 0.30] USD/mile. The counterfactual with no roughness is 0.74 USD per mile, which means that the for the median level of road roughness in our sample, the cost of roughness is 0.31 USD per mile.

Figure 9: Total User Cost of Roughness vs No Roughness (Local Roads)

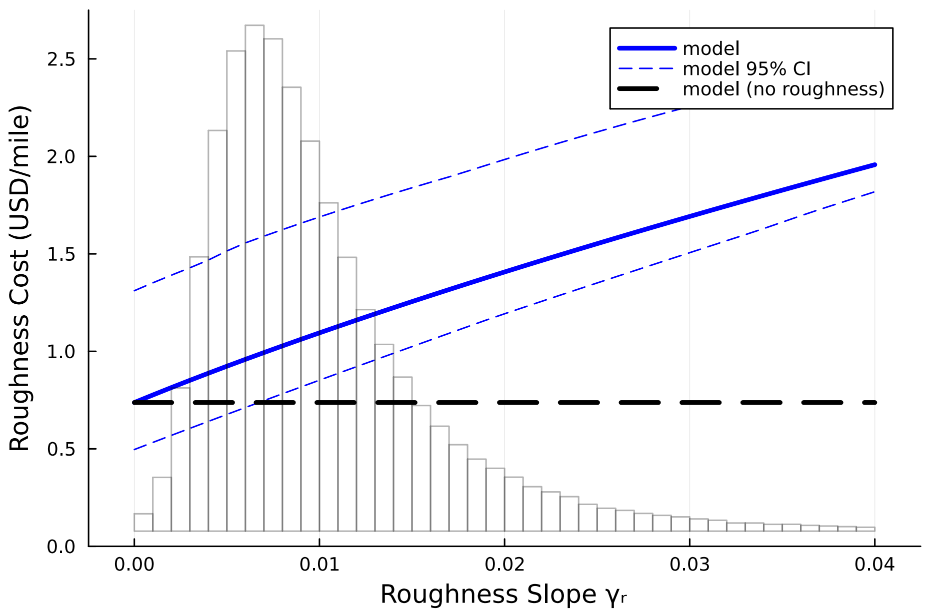

Note: This graph plots the total user costs of roughness per mile of road, with and without roughness, from equation (4.5). We compute optimal speed at different road roughness slopes γ_r in the model using the estimates from column (4) in Table 6, at the median speed limit (25 mph), and at the median value of the speed shock β_{ri}. The histogram of the road segment slope γ_r in the border pair sample is displayed in the background. The dashed blue lines indicate point-wise 95% confidence intervals from bootstrapping the entire estimation procedure at the level of town border pairs.

[26] In Appendix Figure A.14 we show that results are similar using actual speeds at different levels of road roughness.

[27] These numbers are directly proportional to the value of time spent driving and the factor $1 + \frac{1}{\sigma}$. In our preferred specification this is equal to 1.43.

4.3 Empirical Results: The Impact of Road Resurfacing

The second source of variation in road quality that we use comes from road resurfacing in the city of Chicago. Here, we exploit over-time variation within road segment that comes from resurfacing work. Our data allows us to net out time trends in a differences-in-difference framework using other nearby local road segments that were not resurfaced as controls. This approach allows us to control for unobserved, time-invariant factors at the road segment level that may be correlated with roughness. Relative to changes in roughness around town borders, this variation in roughness is highly localized, with resurfacing events covering an area between less than a city block up to several blocks. Drivers may change their speed less in response to variation in roughness at this level relative to town borders.

We introduced the resurfacing events in Section 3.5. 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:

$$\log(s_{rpt}) = \psi Post_t + \phi_p + \delta Repaved_r - \frac{\sigma}{\sigma+1} \log\left(\gamma_{rpt} + \frac{\theta}{s_r^{\text{LIM}}}\right) + \epsilon_{rpt}, \tag{4.9}$$

where $Post_t$ is a dummy for $t = 1$, ϕ_p are fixed effects for all segments that we associate to a given repaving event p, both the repaved segment and the control segments, $Repaved_r$ is a dummy for whether the segment r was repaved, and s_r^{LIM} is the speed limit at the road segment level. We estimate this equation using the interaction variable $Repaved_r \times Post_t$ as instrument. The exclusion restriction requires that speed cost factors do not change differentially over time for repaved relative to the control road segments.

While we have data on the speed limit for each road segment in our sample, speed limits do not predict actual speeds in a linear regression.[28] A possible explanation is that we are making comparisons within the city of Chicago and only among local (residential) roads, where there is less variation in speed limits compared to our town border analysis. The speed limit at the road segment level is likely to be measured with noise, which may attenuate this coefficient. (In the town border analysis, we used the median speed limit over an entire town.) For our benchmark specification, we use the value that we estimated in our town border analysis, $\theta = 0.33$. As a robustness check, we also report results using a much smaller $\theta = 0.01$ (Appendix Table A.8).

We report the estimation results in Table 7. Rougher roads lead to slower driving speeds. The coefficient of -0.28 on the log speed cost term is statistically significantly different from zero. The

[28] See Appendix Table A.7. The linear specification also produces the elasticity of speed with respect to roughness of -0.14 that we cite in the introduction.

variation comes only from road repaving events. This effect is smaller than the one what we found in Table 6 when we used variation in road roughness at town borders. As road resurfacing covers short stretches of road, drivers may be less likely to adjust speed at this scale compared to when they enter a different town. It may also take a longer than a few months for drivers to learn about these changes.

The costs of road roughness computed using these estimates are roughly twice as large as those we estimated based on the town border analysis. Our estimates imply that the median cost of roughness is 0.62 USD, and the cost of 1 SD of road roughness is 0.44 USD. Appendix Figure A.15 replicates Figure 9 with the data and estimates from Chicago local roads.[29]

In the rest of this paper, we will use the more conservative (lower) costs of road roughness for local roads based on the town border analysis in the previous section. Readers with a different view of the impact of road roughness on optimal driver speed, or of the relationship between time costs and total costs (including road roughness), may easily adjust our estimates accordingly.

[29]It may appear counter-intuitive that a *smaller* responsiveness of speed to roughness leads to a *larger* cost of higher road roughness. The cost of higher road roughness combines two effects. First, roughness affects optimal driver speed, and this effect is lower for Chicago resurfacing events compared to the town border analysis. Second, our model implies that total costs are a fixed multiplicative factor of time costs, and this factor is decreasing in σ. The overall effect of increasing roughness, holding speed at median roughness constant, is in general ambiguous (see Appendix D.1).

Table 7: The Impact of Roughness on Driver Speed on Local Roads (Chicago Repaving Events)

	Dependent variable:
	Log speed (mph)
	(1)
Log speed costs $\frac{\sigma}{\sigma+1}$	-0.28
	[-0.36, -0.21]
Inverse speed limit θ	0.33
	[0.33, 0.33]
Repaved road	-0.01
	[-0.02, 0.01]
Post repaving	-0.01
	[-0.02, -0.01]
Welfare factor $1+\frac{1}{\sigma}$	3.61
	[2.76, 4.76]
Repair events	611
Observations	19,878

Note: This table reports estimates of the model-implied non-linear regression equation (4.9). The analysis covers 611 road repaving events between May and mid-July 2021 in Chicago. The sample also includes control local road segments between 100 and 200 meters away from each repaved segment, and two time periods, April and August 2021. We fix the coefficient on inverse speed limit to $\theta = 0.33$, based on Table 6, column 4. Appendix Table A.8 reports results using $\theta = 0.01$. We use an interaction of post repaving (August) and treated road segment as instrument. In parentheses, 95% confidence intervals based on bootstrap at the repaving event level .

4.4 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. Our model, like all such models, imperfectly reflects reality. To obtain a more direct measure of costs, we look at damage found in Massachusetts' vehicle inspections. We test whether there is more evidence of damage to road-related car parts, such as tires and suspensions, in towns with rougher roads.

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 cannot test for all of the possible externalities from road roughness, but we will look at the link between road roughness and fatal vehicle crashes.

4.4.1 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 A.9. 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 10 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).

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):

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

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.[30]

[30] Appendix Table A.9 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.

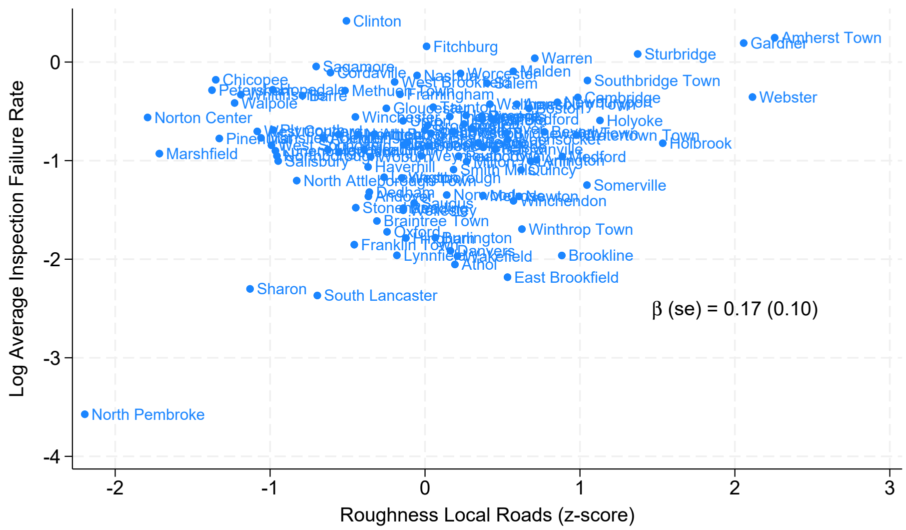

Figure 10: Inspection Failure Rate vs Local Road Roughness (Main Indicators)

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.

4.4.2 Other Costs: Vehicle Crashes

The biggest possible upside to rougher roads is that slower speeds might reduce the number of collisions, or make collisions less deadly. We now turn to the relationship between road roughness and traffic fatalities. Our results are only suggestive, as there may be omitted variables correlated with both roughness and fatalities. For example, we do not have good estimates of vehicle miles traveled.

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:[31]

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

where $Fatalities_i^r$ is the number of fatalities in 2021 in town i on road type r, Z_i^r denotes the

[31] 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.

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.

Table 8: 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.

The first six regressions of Table 8 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 ($\log Pop_i$) and median household income ($\log Y_i$), 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 ($\log Pop_i^{\text{drive}}$), and total local employment ($\log Empl_i$) 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 -.6 for local, arterial and highway roads respectively. The fact that controlling for these imperfect measures of 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.

Another way of gauging the plausibility of these results is to combine our estimates of the links between roughness and speed 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 estimates coefficients change little with or without our two road traffic controls, which may be explained by the fact that most of the road segments in our data are relatively uncongested.[32]

The coefficients of roughness on speed would suggest our coefficients of fatalities on roughness if the elasticity of fatilities with respect to speed was approximately one. In fact, most external estimates estimate a much larger imact 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.[33]

The impact we find of road roughness on traffic fatalities is not unreasonably large given the impact of road roughness on speed that we have documented. In fact, the 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. If this were the case, then smoother roads enable faster speeds, which inevitably means more accidents, but also reduces accidents based on road roughness.

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

[32]The semi-elasticity with respect to roughness on local roads is in a similar range as the estimates that we get using border discontinuities in Section 4.2 (see Appendix Table A.6).

[33]Ashenfelter and Greenstone (2004) shows that a speed limit increase led to a 4% increase in speeds and 35% increase in fatalities, which means an elasticity of around 9. See also DeAngelo and Hansen (2014) and Van Benthem (2015).

on the costs imposed by road roughness on minorities and low-income Americans. We first focus on the tract level variation, and ask whether road quality is worse in poorer neighborhoods, and in neighborhoods 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 for 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. A congressional report prepared by the GAO in 2022 used highway road pavement data to show 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). Our data allows us to analyze these patterns on all types of roads, including local roads, which carry an important share of traffic. 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 (Brueckner et al., 1999).

5.1 Rough Roads, Income and Race

The model and estimation laid out 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 to drivers from driving slower, 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 4.2.2.

Our main analysis is at the Census tract level,[34] but we also report results for the localities that actually administer local roads.[35] 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.

[34] Census tracts are a subdivision of counties and contain roughly 4,000 people per tract.

[35] 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.

The median tract in the sample has roughness data for 13% of its local road segments.

There is certainly mismeasurement of any aggregate geography in our data, but the available evidence seems to suggest that our small sample are not biased, which they could be if Uber drivers typically selected smoother roads or if additional Uber use made roads rougher. 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 we cannot rule out the possibility that roads never used by Uber drivers are either much rougher or much smoother. Appendix Tables A.10 and A.12 show results where we restrict the sample by requiring tracts or towns to meet several coverage thresholds.

Figure 11 visualizes our resulting tract-level data for Chicago and New York City. The two cities show quite different patterns. 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.

Figure 11: Road Roughness by Census Tract

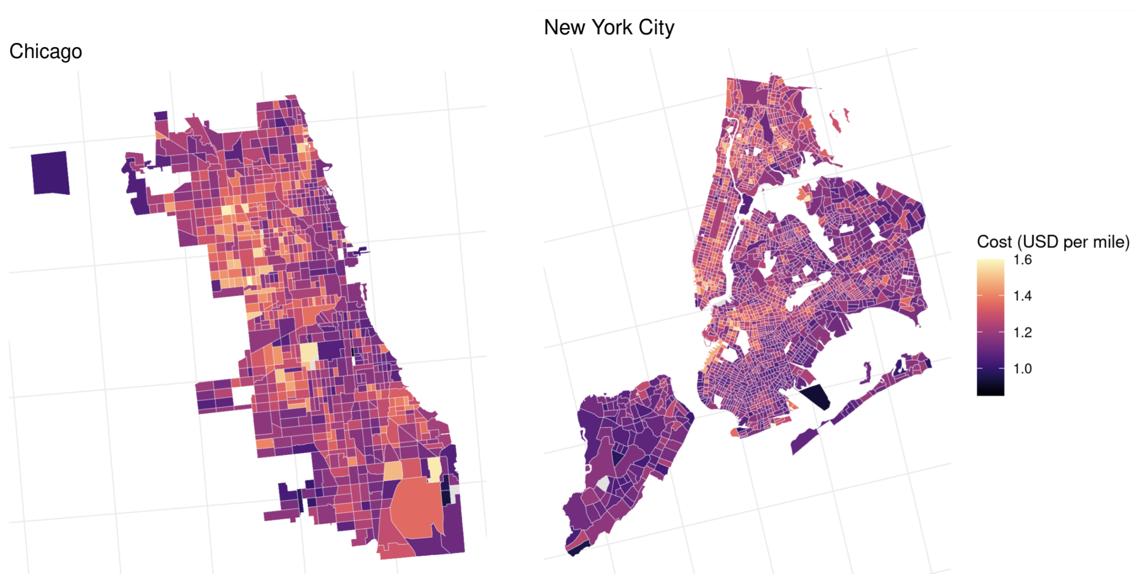

Note: These maps show costs of local road roughness at the Census tract level for Chicago and New York City. To construct them, for each local road segment we compute the cost per mile using the model from Section 4, and take the average at the tract level.

Table 9 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 $-.043$ in regression (1) means that as income increases by 1 log point, the cost per mile of local roads drops by 4.3 cents. If a household drivers 3,300 miles per year on local roads, then a one log point increase in income is associated with 142 dollars less harm from bad roads per year.[36] In the second regression, which controls for metropolitan area fixed effects, the coefficient increases in magnitude to $-.052$. 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 172 dollars less harm as its income increases by one log point.

In the third regression, we control for town fixed effects. In this case, the coefficient declines in magnitude to $-.017$. 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 shouldn't be surprising. We expect richer areas to spend more on road repaving and perhaps also have access to better expertise in road maintenance or higher quality initial construction. As towns own the majority of local roads, the tax base should impact upkeep. 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. The coefficient of .106 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 10.6 cents per mile. A household that drives 3,000 miles annually on local roads would pay an extra roughness cost of 318 dollars per year.[37] The second regression shows that controlling for metropolitan area fixed effects does not change the coefficient at all. 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.

Regression (6) includes town fixed effects and the coefficient on fraction Black drops to .016. 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

[36] 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.

[37] Black households drive an average of 19,386 miles per year, based on 2017 NHTS, which implies that Black households on average drive 2,966 miles on local roads and 8,607 on local and arterial roads. See footnote 36 for details.

that the large gaps in road quality across the US are better understood by sorting into different towns, then 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 9: Tract Roughness, Income, and Race

	Dependent variable:					
	Cost (USD per mile)					
	(1)	(2)	(3)	(4)	(5)	(6)
ln median income	-0.043^{***}	-0.052^{***}	-0.017^{***}			
	(0.001)	(0.001)	(0.002)			
fraction Black				0.106^{***}	0.104^{***}	0.016^{***}
				(0.003)	(0.003)	(0.004)
fraction Hispanic				0.090^{***}	0.077^{***}	0.031^{***}
				(0.003)	(0.004)	(0.004)
fraction Asian				0.085^{***}	0.038^{***}	-0.034^{***}
				(0.007)	(0.007)	(0.008)
Climate controls	Yes			Yes		
MSA Fixed effects		Yes			Yes	
Town Fixed effects			Yes			Yes
Observations	32,967	33,134	33,134	32,967	33,134	33,134
Adjusted R^2	0.053	0.144	0.346	0.073	0.144	0.346

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). 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.

Table 10 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) in Table 10 reproduce columns (3) and (6) in Table 9 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 (2) and (5) 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 10 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. However, when we control for income and race jointly in regression (7), it is the income coefficient that remains significant.

In Table 11, we look at heterogeneity in the within town relationship between roughness and income/race. For race, we focus on the fraction of the population that is Black. 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 fraction Black covariates. The results suggest the within town relationship between income and better roads is stronger in less populated towns and in higher income towns. The within town relationship between tract share Black and worse roads is stronger in smaller towns and in higher income towns.

Table 10: Tract Roughness and Tract Characteristics

	Dependent variable:						
	Cost (USD per mile)						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
ln median income	−0.017***	−0.007***	−0.005***				−0.015***
	(0.002)	(0.002)	(0.002)				(0.002)
fraction Black				0.016***	0.016***	0.013***	−0.003
				(0.004)	(0.004)	(0.004)	(0.004)
fraction Hispanic				0.031***	0.015***	0.015***	0.011**
				(0.004)	(0.004)	(0.004)	(0.005)
fraction Asian				−0.034***	−0.042***	−0.040***	−0.044***
				(0.008)	(0.008)	(0.008)	(0.008)
ln miles to CBD		−0.009***	−0.008***		−0.008***	−0.006***	
		(0.002)	(0.002)		(0.002)	(0.002)	
ln population		0.008***	0.009***		0.008***	0.009***	
		(0.002)	(0.002)		(0.002)	(0.002)	
ln employment		0.005***	0.006***		0.006***	0.007***	
		(0.001)	(0.001)		(0.001)	(0.001)	
ln area (miles^2)		−0.019***	−0.020***		−0.020***	−0.020***	
		(0.001)	(0.001)		(0.001)	(0.001)	
fraction drive to work			−0.056***			−0.059***	
			(0.007)			(0.007)	
ln avg Uber usage			−0.021***			−0.021***	
			(0.002)			(0.002)	
Town fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	33,134	33,134	33,134	33,134	33,134	33,134	33,134
Adjusted R^2	0.346	0.357	0.360	0.346	0.358	0.362	0.347

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.

Table 11: Tract Roughness and Tract Characteristics, Town Heterogeneity

	Dependent variable:							
	Cost (USD per mile)							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
ln tract income	-0.056^{***}	-0.018^{***}	-0.015^{***}	-0.045^{***}				
	(0.010)	(0.004)	(0.002)	(0.011)				
ln tract income × high town population	0.040^{***}			0.037^{***}				
	(0.010)			(0.010)				
ln tract income × high town fraction Black		0.001		-0.008				
		(0.005)		(0.005)				
ln tract income × high town income			-0.015^{***}	-0.015^{***}				
			(0.004)	(0.004)				
tract fraction Black					0.136^{***}	0.053	0.010^{***}	0.081
					(0.030)	(0.044)	(0.004)	(0.052)
tract fraction Black × high town population					-0.122^{***}			-0.101^{***}
					(0.030)			(0.031)
tract fraction Black × high town fraction Black						-0.037		0.029
						(0.045)		(0.045)
tract fraction Black × high town income							0.095^{***}	0.090^{***}
							(0.015)	(0.015)
Town fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	31,419	31,419	31,419	31,419	31,543	31,543	31,543	31,543
Adjusted R^2	0.337	0.337	0.337	0.338	0.334	0.334	0.335	0.335

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 dummy 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. All regressions include town fixed effects. 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 12 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 7 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 12: Road Quality by Town Income and Percent Black

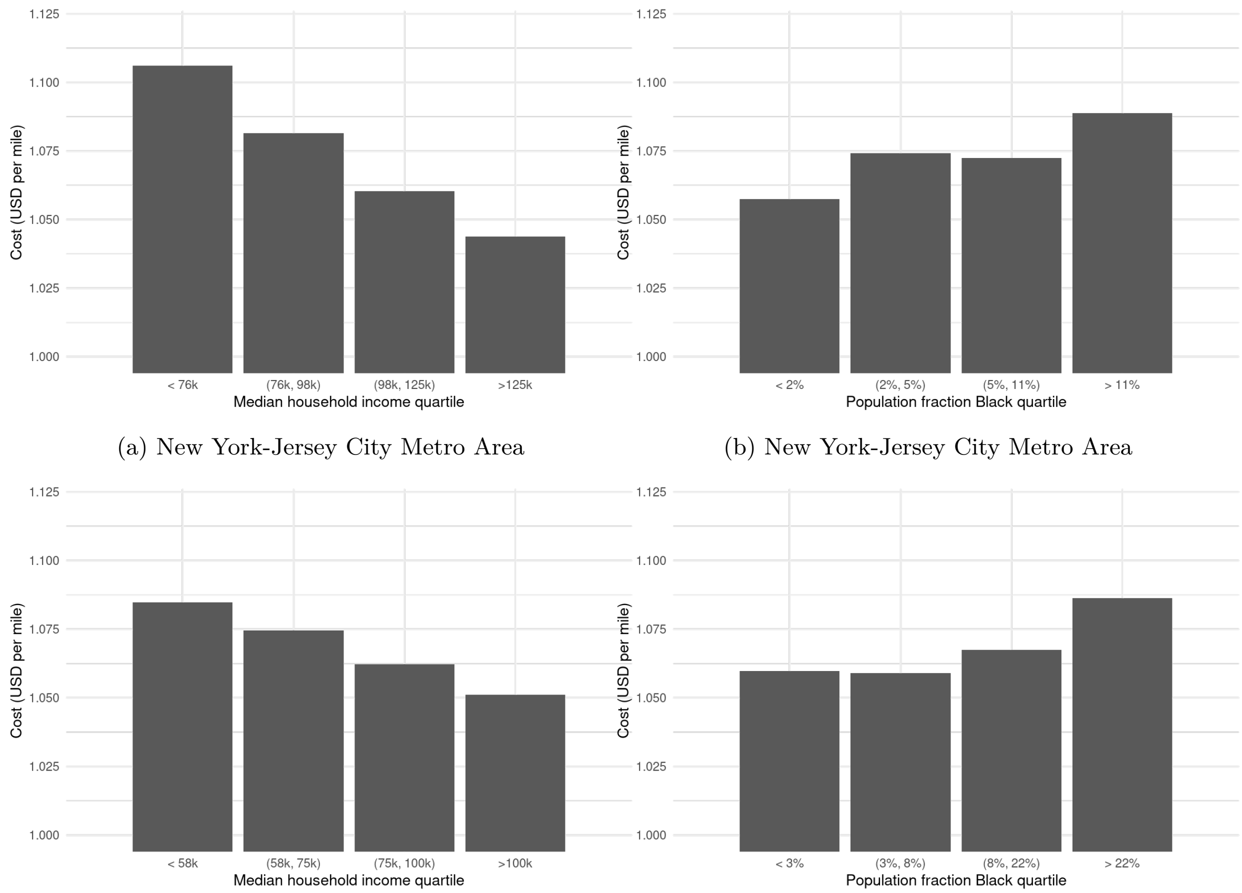

(a) New York-Jersey City Metro Area

(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..

We refine this analysis in Table 12. In addition to the ACS and Zip Code Business Patterns data, we use reported government spending on roads 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.[38]

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

One obvious reason higher income towns have better roads may be that they spend more. Local roads are primarily maintained by local governments – the federal government owns only .6% of urban roads and state governments own 14% of urban roads, but only 4% of urban local roads. When counties and townships report road spending in the ASLGF, we distribute this to towns based on overlapping area. Consequently, our expenditure variable should capture nearly all spending on local roads for each town. The median per capita expenditure on roads is $116 USD. 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 12 shows these results. Regression (1) shows a coefficient of $-.046$ on the log of income. Regression (2) includes the town's expenditures, which comes in with a positive coefficient. Here we expect to capture two effects. First, a negative causal effect of expenditure on road roughness. Second, towns with rougher roads spend more. Overall, the small positive correlation we observe is consistent with the second effect dominating. Towns with rougher roads spend more, but this does not fully offset the roughness differences. In Appendix Table A.11, 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. Further, wealthier towns have more roads per person, so they do not have greater spending per mile. Spending alone (without accounting for degradation rates or initial quality) does not explain why higher income towns have better roads.

Regression (3) includes a larger set of town-level covariates, including distance to the city center, population and land area. More distant towns have better roads, as do towns with lower population levels. The coefficient on the logarithm of town income drops to $-.029$ 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 .073 on the share of the population this is Black. This variable changes little when we control for town-level expenditure in regression (5). Even controlling for our wider set of covariates in regression (6) only leads the coefficient to fall by 20 percent to .056. The tendency of towns with a higher share of the population that is Black 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 .048 when we control for income. Appendix Table A.11 shows that road expenditure per capita is weakly lower in towns with a higher share of Black households. Overall spending levels are not significantly different in such towns, although they are significantly lower in towns with large Hispanic or Asian populations.

---

spending. Appendix Figure A.17 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.

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. The correlation that relates to income is somewhat more explicable than the correlation that relates to race. 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.

Table 12: Town Roughness and Characteristics

	Dependent variable:						
	Cost (USD per mile)						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
ln median income	-0.046^{***}	-0.049^{***}	-0.029^{***}				-0.029^{***}
	(0.006)	(0.006)	(0.006)				(0.008)
fraction Black				0.073^{***}	0.076^{***}	0.056^{***}	0.048^{***}
				(0.011)	(0.011)	(0.010)	(0.013)
fraction Hispanic				0.050^{***}	0.055^{***}	0.006	0.021
				(0.012)	(0.012)	(0.012)	(0.014)
fraction Asian				-0.010	-0.004	-0.031	-0.013
				(0.023)	(0.023)	(0.021)	(0.022)
ln expenditure per capita		0.004^{**}	0.004^{***}		0.003^{**}	0.003^{**}	
		(0.002)	(0.002)		(0.002)	(0.002)	
miles to CBD			-0.024^{***}			-0.022^{***}	
			(0.003)			(0.003)	
ln population			0.021^{***}			0.026^{***}	
			(0.004)			(0.004)	
ln employment			-0.0001			0.004	
			(0.003)			(0.003)	
ln area (miles^2)			-0.021^{***}			-0.029^{***}	
			(0.004)			(0.004)	
fraction drive to work			-0.140^{***}			-0.118^{***}	
			(0.022)			(0.023)	
ln avg Uber usage			-0.006			-0.006	
			(0.007)			(0.007)	
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	0.209	0.368	0.207	0.210	0.376	0.215

Note: This table shows coefficients from regressing average local roughness on town on log median charactistics 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. Missing or zero town expenditure on roads is replaced with the mean expenditure, and a dummy 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 Four Large Cities

How do cities decide which roads to repair? To find out, we collected road-segment level data on actual repair decisions for four large cities, and also fielded a survey of over 100 towns and cities on their road data collection, targeting and repair strategies. We discuss the survey results in the next section and focus on the first exercise here.

We collected data on road resurfacing that happened between August, 2021 (the period covered by our road roughness data) and January 1, 2023. We can thus use our measure of road roughness as a baseline, and study how later resurfacing decisions depend on roughness, together with other factors.

We focus on the cities of New York, Dallas, Columbus, and Portland. These cities have publicly available data on road resurfacing after August 2021, including a GIS map of the location of the repaving. There is large variation in the number of repairs by city; excluding preventative work like slurry seals, New York City reported 553 miles of repairs over the period and Dallas reported 368 miles, while Columbus reported 64 miles of repairs and Portland reported only 25 miles.[39] We also remove repairs on arterial roads and on highways, which reduces the number of repairs on average by 27%.

We focus on local road segments and aggregate all data at the level of small grid cells.[40] For each grid cell we compute (1) the total length of local roads within, (2) the total length of local roads repaved, (3) various statistics of Uber road roughness cost of the segments in the grid, and (4) interpolated demographic variables from the ACS (population, income, share of Black residents). Further details are in Appendix C.9.

Figure 13 shows the results for repairs by roughness for the four cities. For each city, we place grids into deciles of the average grid level road roughness, computed using our Uber data. The bar height indicates the share of road length within the decile bin that gets repaired. 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. While 93% of grids within the boundary of New York City have some Uber coverage, Dallas, Columbus, and Portland have 37%, 51%, and 43% respectively. The areas covered tend to be the urban core of the city; for example, a map of coverage of Dallas is given in Appendix Figure A.19.

In New York City, around 10% of the entire local road network was resurfaced during our study period. This share is lower, 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

[39]We contacted the public works departments of Portland and Columbus to confirm their online maps are not missing data. Portland confirmed that the number of repairs is likely accurate, as the jurisdiction is facing significant funding issues.

[40]The grid cells are coordinate degrees rounded to .005. In New York City, this is approximately equal to .4km by .4km squares.

resurfacing work is partly targeted to rough roads, yet this only explains a small share of which road segments are selected for repairs.

No such pattern holds for Dallas. The repair rate is lower overall, around 5%, and it does not increase systematically with our measure of road roughness. Columbus and Portland made remarkably few repairs over this period, and we also fail to see systematic patterns.

Figure 13: Local Road Roughness and Resurfacing Decisions in Four Cities

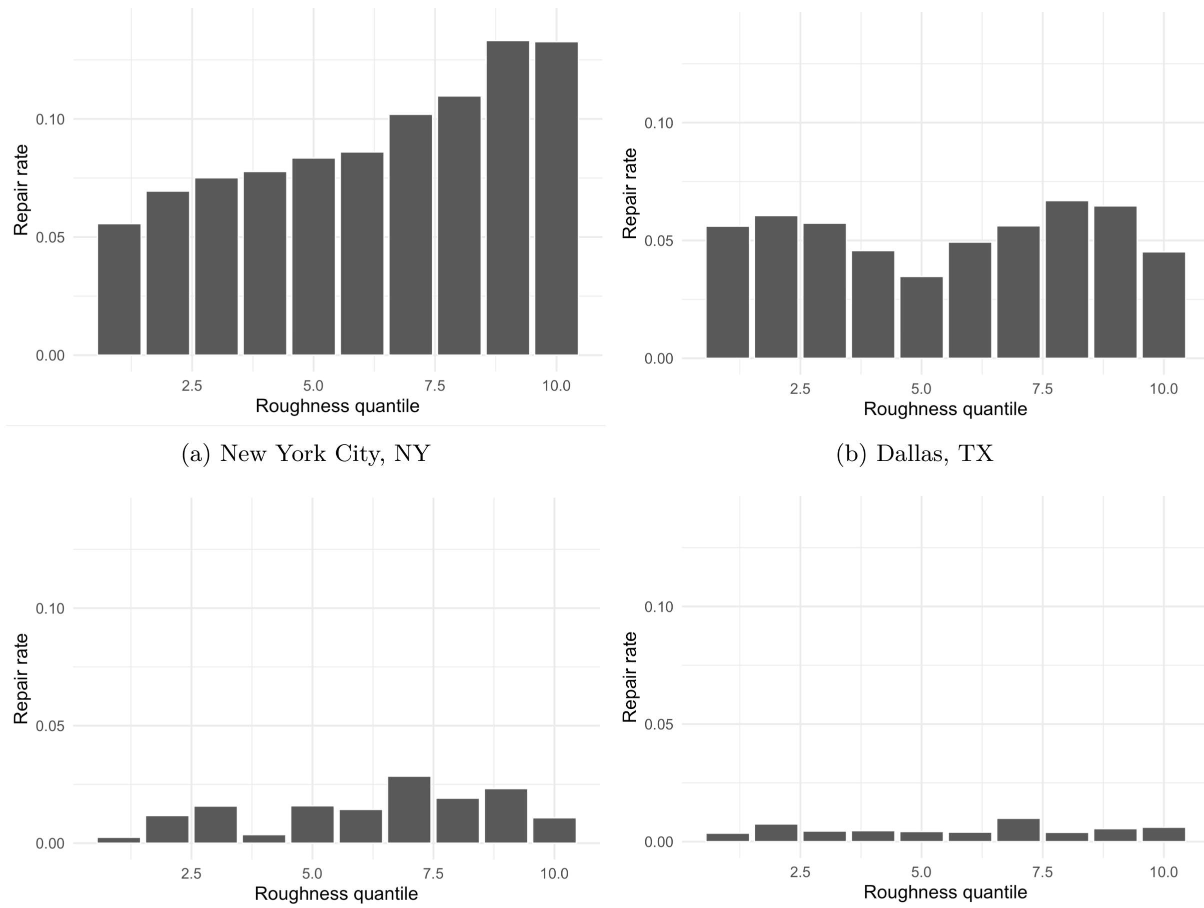

(a) New York City, NY

(b) Dallas, TX

(c) Columbus, OH

(d) Portland, OR

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.

6.1.1 Predictors of road resurfacing

What other factors explain road resurfacing decisions? Cities may prioritize high-traffic areas, as well as direct repairs in certain types of neighborhoods. We first analyze empirically the correlates of road repairs in the four cities, and then turn to comparing city decisions with those implied by our Uber roughness data and a simple model of cost-minimizing resurfacing decisions.

Table 13 reports linear regressions where the outcome is the repair rate in the grid. In odd columns, we include geographic and road network covariates at the grid level.[41]

Uber road roughness predicts the road repair rate in NYC but not in the other cities. Repairs target more populous areas in NYC and Columbus, and more central areas in NYC and Dallas. Areas with high Black population get fewer repairs in NYC, while richer areas get slightly more repairs in Dallas and Columbus. We then include two proxies for traffic volume, Uber traffic and the fraction of workers that drive to work. The share of grid residents who drive to work has a strong correlation with repairs in Dallas but a negative correlation in Columbus.

Since we have significantly better coverage of New York City than the other cities, a key concern with the null results we identify for Dallas, Columbus and Portland is that measurement error is driving the results. We test for this possibility 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 13 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.13.

[41] In Appendix Table A.14, we present analogous tables where we remove grids with less than 20% and 30% coverage by Uber. There is a trade-off in removing grids, as low coverage grids will have more measurement error in average roughness and lead the counterfactual targeting strategies to over-target low coverage grids, but restricting to higher coverage grids may change the sample due to selection.

Table 13: Predictors of Road Repair Targeting

	Dependent variable:							
	Repair rate (percent road-miles repaved)							
	NYC		Dallas		Columbus		Portland	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
local roughness (z score)	0.022***	0.023***	−0.002	−0.005	0.003	0.001	0.001	0.001
	(0.005)	(0.006)	(0.007)	(0.007)	(0.003)	(0.003)	(0.002)	(0.002)
ln road miles	0.003	0.002	0.036***	0.031***	−0.017***	−0.016***	0.001	0.001
	(0.004)	(0.004)	(0.006)	(0.006)	(0.003)	(0.003)	(0.002)	(0.002)
ln population	0.007**	0.009**	−0.004	−0.002	0.010**	0.008**	−0.00001	0.00002
	(0.003)	(0.004)	(0.006)	(0.006)	(0.004)	(0.004)	(0.002)	(0.002)
ln miles to CBD	−0.037***	−0.042***	0.005	−0.018*	−0.016***	−0.003	0.002	0.005
	(0.006)	(0.008)	(0.009)	(0.010)	(0.005)	(0.006)	(0.002)	(0.004)
fraction Black	−0.033***	−0.031***	−0.030	−0.030	−0.009	−0.0002	0.017	0.016
	(0.011)	(0.011)	(0.025)	(0.026)	(0.013)	(0.013)	(0.024)	(0.024)
ln median income	0.003	0.002	0.010*	0.014**	0.018**	0.033***	0.002	0.003
	(0.003)	(0.003)	(0.006)	(0.006)	(0.008)	(0.008)	(0.004)	(0.004)
ln Uber volume		−0.004		−0.030***		−0.0003		−0.0001
		(0.006)		(0.011)		(0.005)		(0.003)
fraction drive to work		0.011		0.261***		−0.174***		−0.013
		(0.021)		(0.094)		(0.036)		(0.016)
Constant	0.092***	0.120**	0.036	−0.067	−0.201**	−0.223**	−0.022	−0.031
	(0.035)	(0.052)	(0.068)	(0.134)	(0.091)	(0.094)	(0.044)	(0.050)
Observations	2,918	2,918	1,086	1,086	1,249	1,249	619	619
Adjusted R^2	0.041	0.040	0.040	0.054	0.035	0.051	−0.005	−0.007

Note: This table shows coefficients from a linear regression of the fraction of a grid that gets repaved on covariates.
.

6.1.2 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 cost-minimization strategies. We consider two benchmark strategies for targeting road repairs. First, we consider a simple strategy of prioritizing the road segments in worse condition. Second, ordering by 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.[42] 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). However, this measure does not include external costs such as the possibility of higher crash rates due to smoother, faster roads, as we discussed in Section 4.4.2.

Our proxy for road segment traffic $Traffic_r$ is the number of Uber trips over that road segment. While this measure of traffic at the level of each local road segment is the best available to us in the four cities, it may not represent overall road traffic patterns. We return below to this issue of whether this traffic measure over-weights high-income neighborhoods.

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.

Figure 14 compares each city's actual targeting (in red) to the two hypothetical targeting mechanisms. In these counterfactual policies, repairs are more strongly increasing in average road roughness, compared to what we see in the data. They target significantly more roads in the top decile of roughness, and significantly fewer roads in the bottom four deciles of roughness. However, both policies have broad coverage across grid roughness deciles, highlighting that there is substantial heterogeneity in road segment condition within grids.

[42]While repairing a road segment is a significant event for that particular road segment, it is small for any driver if trips have many road segments. This is the case in our data, where the median road segment length is less than 50 meters, much shorter than typical trips.

Figure 14: Actual and Model-Implied Targeting in Four Cities

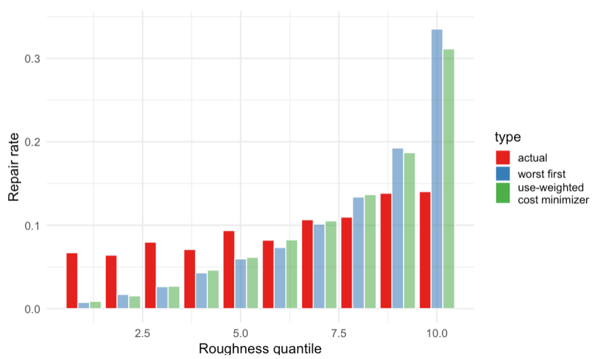

(a) New York City

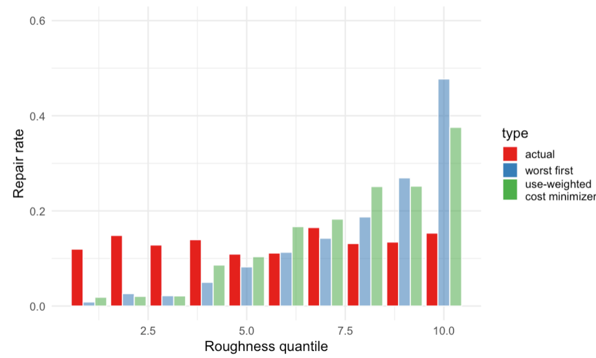

(b) Dallas

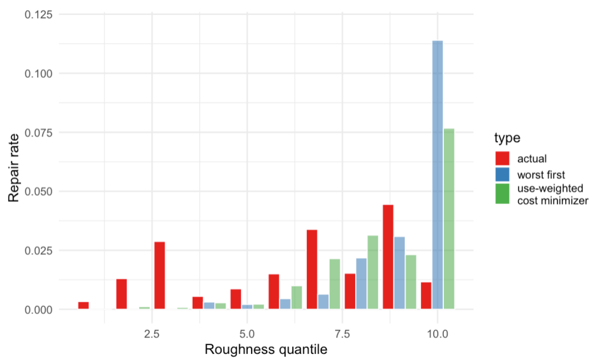

(c) Columbus

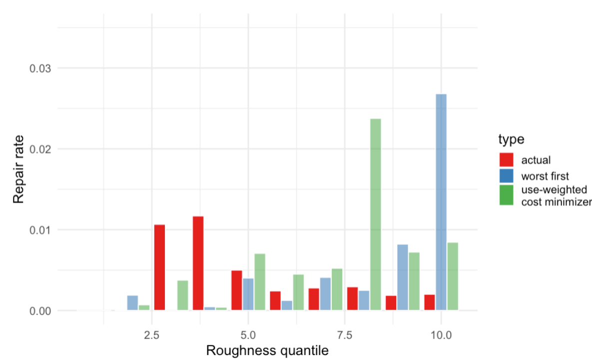

(d) Portland

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

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 14 reports estimates of λ.[43] New York City's repair decisions are consistent with a weight of 10-12% on model-implied targeting, and the rest due to orthogonal factors. For the other three cities, we find no evidence that repair decisions are statistically related to what our simple model

[43] Appendix Table A.15 reports results varying the coverage thresholds.

implies.

One concern with this analysis is that our measure of road traffic uses Uber trips, which may not be representative. For example, Uber ridership may be correlated with other variables, such as neighborhood income, which would bias the counterfactual policies towards targeting wealthy neighborhoods. In Appendix Figure A.20, we re-estimate the counterfactual policy after first residualizing Uber trips by the median household income per grid. Controlling for income flattens the use-weighted cost minimizing counterfactual somewhat, but it maintains a much steeper slope in roughness than actual repaving targeting shows.

Table 14: Repair Targeting versus Counterfactual Repair Targeting

	Dependent variable:							
	Repair rate (percent road-miles repaved)							
	NYC		Dallas		Columbus		Portland	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
worst first	0.094***		0.030		0.00002		−0.016	
	(0.021)		(0.053)		(0.079)		(0.039)	
use-weighted		0.159***		−0.128		0.069		−0.040
		(0.023)		(0.104)		(0.099)		(0.083)
ln road miles	0.012***	0.011***	0.038***	0.038***	−0.018***	−0.018***	0.002	0.002
	(0.003)	(0.003)	(0.007)	(0.007)	(0.004)	(0.004)	(0.001)	(0.001)
Constant	0.075***	0.071***	0.119***	0.124***	0.021***	0.020***	0.003**	0.003**
	(0.004)	(0.004)	(0.008)	(0.008)	(0.005)	(0.005)	(0.001)	(0.001)
Observations	2,786	2,786	568	568	618	618	233	233
Adjusted R^2	0.011	0.020	0.047	0.049	0.023	0.023	−0.001	−0.0005

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

Our surveys focused on two samples of towns with Uber coverage. First, we randomly sampled towns with Uber roughness data across the US. Second, we focused on 30 towns in Massachusetts. 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. The response rate for the national sample is 16.3%. The response rate for Massachusetts towns rises to 63.3%. More details on data collection are included in Appendix C.13, which also includes evidence on the differences between towns that responded quickly and towns that responded only after repeated calls. These differences seem mild, which makes us hopeful that our sample is reasonably representative of similar towns.

Table 15 presents results on the resurfacing that actually takes place. Table 16 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 15 shows respondents' assessments of what share of roads that need repaving actually get repaving. Only 21 percent of the Massachusetts sample and 29 percent of the national sample thought that more than one-half of the roads that needed resurfacing were getting resurfacing. These numbers seem low given that the respondents should have an interest to highlight their own efficiency. Further, 32 percent of the Massachusetts sample and 62 percent of the national sample report that less than 30 percent of roads that need repaving are repaved. These figures suggest that we may be seeing little targeting of resurfacing either because repaving plans doesn't target rough roads, or because plans require more repaving than the city can handle, and public works departments then select roads according to convenience or other criteria unrelated to road roughness.

The second panel in Table 15 shows that departments are allocating most of their resources to resurfacing. Only 21 percent of the Massachusetts sample and 33 percent of the national sample are spending more than 30 percent 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 15: City Road Maintenance and Repair Strategies

	Massachusetts	Rest of US
Panel A: Percent of Roads that Receive the Resurfacing they Require		
> 90%	0.11	0.05
70-90%	0.05	0.11
50-70%	0.05	0.13
30-50%	0.42	0.06
< 30%	0.32	0.62
Unsure	0.05	0.03
Panel B: Spending on Resurfacing vs Preventative Maintenance		
> 90% on resurfacing	0.11	0.20
70 - 90% on resurfacing	0.68	0.47
50 - 70% on resurfacing	0.16	0.18
50% on resurfacing	0.00	0.02
30 - 50% on resurfacing	0.00	0.04
< 30% on resurfacing	0.00	0.09
Unsure	0.05	0.01
Observations	19	101

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

The first panel of Table 16 shows how towns get information about road roughness. In both samples, the modal source of information about road quality was a survey of roads performed by a private company. Almost 90% of the towns we surveyed in Massachusetts and 62% of our national sample pay for such a survey. In the national sample, 59% of towns say that they use an in-house survey. The share of Massachusetts towns with an in-house survey is 42 percent, which is understandable given that almost all of these places are also using a professional surveying service. While it is less clear what towns meant when they claimed to use "engineer's discretion," the intent is clearly to rely on professional knowledge. More than one-half of the national sample and almost one-third of the Massachusetts sample claim to use such discretion.

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.[44] None of the Massachusetts towns admitted a role for elected officials, while 18% of the national sample said that politicians played some role in targeting repaving. Citizens' complaints were more important in both samples.

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

Slightly more than one-fourth of the Massachusetts and 41% of the national sample base their assessments on input from ordinary residents.

The second panel shows that both samples are sharply bifurcated in the rate in which they survey roads. In both samples, about one-half of towns survey more than 90% of their roads annually and about 40% of towns survey less than 50% of their roads annually. We suspect that some of these towns are surveying either one-half or one-third of their roads every year, although more than one-fifth of towns in both samples survey less than 30% of roads in any given year. As roads can deteriorate fairly quickly, this low cadence could mean that targeting is quite imperfect in those places.

The third are fourth panels ask about their practices in targeting repaving. Panel C focuses on assessing the need for repaving. Panel D asks about prioritizing repaving. We asked these questions separately, because if towns are unable to repave all of the roads that they think need repaving, then their prioritization rules will determine repaving in practice. Formulas appear to be regularly used in both determining need and prioritizing repaving. In the Massachusetts sample, 72% of towns used a formula based only on road quality in assessing need and 32% used a formula that also considers other factors. In the US sample, 55% use a formula based on road quality alone and 52% used a more holistic formula. There are towns which said yes to both types of formulas, which suggested that they either have two formulas or misunderstood the question. Formulas became somewhat less prevalent when determining prioritization of roads. In Massachusetts, 47% of towns used a formula based on road quality alone to prioritize repaving, and 26% of towns reported using a formula which included other factors. In the national sample, 41% of respondents reported using a road quality-based formula to prioritize and 44% suggested that they use a formula based on a broader set of criteria.

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. In Massachusetts, 84% of towns reported that upcoming utility worked determined whether there was a need for repaving, and 42% reported that utility work determined repaving prioritization. In the national sample, 63% of respondents said that utility work helped determine need and forty-eight% said that utility work determined prioritization. As the benefits of road smoothness depend largely on their use, traffic utilization would seem to be a crucial determinant of the value of repaving. Only about one-half of both samples mentioned traffic utilization as a determined 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 are only able to 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, they play a smaller role in determining prioritizion. 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 16: City Road Data Collection and Targeting

	Massachusetts	Rest of US
Panel A: Method(s) to determine the state of roads		
Road survey (in-house)	0.42	0.59
Road survey (contracted to private firm)	0.89	0.62
Elected officials	0.00	0.18
Engineering discretion	0.32	0.56
Reporting by residents	0.26	0.41
Panel B: Percent of Roads Surveyed each Year		
> 90%	0.58	0.46
70-90%	0.00	0.04
50-70%	0.00	0.08
30-50%	0.21	0.14
< 30%	0.21	0.24
Unsure	0.00	0.05
Panel C: Criteria to Determine which Roads need Resurfacing		
A formula only considering road conditions	0.74	0.55
A formula considering other selected factors	0.32	0.52
Upcoming utility work	0.84	0.63
Traffic intensity	0.42	0.57
Citizen complaints	0.00	0.01
Transportation expert feedback	0.11	0.19
Accessibility	0.11	0.16
Elected official input	0.05	0.30
Panel D: Criteria to Determine which Roads are Prioritized		
A formula only considering road conditions	0.47	0.41
A formula considering other selected factors	0.26	0.44
Upcoming utility work	0.42	0.48
Traffic intensity	0.47	0.48
Citizen complaints	0.00	0.00
Transportation expert feedback	0.05	0.15
Accessibility	0.05	0.10
Elected official input	0.16	0.33
Observations	19	101

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 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 23 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 318 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 142-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 New York, Portland, Dallas and Columbus. We find that rougher roads are more likely to be repaved in New York, 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.

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.