Urban form and driving: Evidence from US cities[§]

Gilles Duranton[*]
University of Pennsylvania
Matthew A. Turner[‡]
Brown University

11 January 2016

Abstract: We estimate the effect of urban form on driving. We match the best available travel survey for the US to spatially disaggregated national maps that describe population density and demographics, sectoral employment and land cover, among other things. We develop a novel approach to the sorting problem that follows from an intuitive definition of sorting and an assumption of imperfect mobility. We address the endogeneity problem by relying on measures of subterranean geology as sources of quasi-random variation in urban form. The data suggest that increases in density cause small decreases in individual driving. However, because densification policies must generally decrease population in source regions, this means that such policies can be expected to cause only tiny decreases in aggregate driving. This suggests that urban planning is unlikely to be a cost effective policy response to traffic congestion, automobile related carbon emissions, or other automobile related pollution.

Key words: urban form, vehicle-kilometers traveled, congestion.

JEL classification: L91, R41

[§]Financial support from the Canadian Social Science and Humanities Research Council and from the Sustainable Prosperity Network is gratefully acknowledged by both authors. Turner gratefully acknowledges the financial support and hospitality of the Property and Environment Research Center. We are also grateful to seminar audiences at the University of Arizona, Brown, the University of Houston, Kentucky, PERC and Yale for helpful comments and to Andy Foster for a particularly helpful suggestion about our model. This research could not have been conducted without the help of our research assistants, Tanner Regan, Nicholas Gendron-Carrier, Prottoy Akbar and Rebbeca Lindstrom.

[*]Wharton School, University of Pennsylvania, 3620 Locust Walk, Philadelphia, PA 19104, USA (e-mail: duranton@wharton.upenn.edu; website: https://real-estate.wharton.upenn.edu/profile/21470/).

[‡]Department of Economics, Box B, Brown University, Providence, RI 02912, USA (e-mail: matthew_turner@brown.edu; website: http://www.econ.brown.edu/fac/Matthew_Turner/), also affiliated with PSTC and S4 at Brown University and with the University of Toronto Department of Economics.

1. Introduction

We estimate the effect of urban form on driving. To conduct our analysis we exploit the best available travel survey data for the US and match them to spatially disaggregated national maps that describe, among other things, population density and demographics, sectoral employment, and land cover. Causal identification of the relationship between urban form and driving faces two primary obstacles, household sorting and the endogeneity of density. People with particular preferences for driving may sort into areas of particular density and unobserved factors correlated with urban form may also affect driving behavior. We develop a novel approach to the sorting problem that follows from an intuitive definition of sorting and an assumption of imperfect residential mobility. To address the omitted variables problem we rely on a variety of measures of subterranean geology. These measures predict the characteristics of surface development and are plausible as sources of quasi-random variation in urban form.

We find that urban density has a small causal effect on individual driving. In most of our estimations 'urban density' is the density of residents and jobs within a 10 kilometers radius of where a driver lives. We find that the elasticity of vehicle kilometers traveled (VKT) with respect to this measure of density is between -7% and -10%. This result is not sensitive to the particular measure of density, but is sensitive to the scale at which we measure density. Residents and employment more than 10 kilometers from a driver's residence do not have a measurable effect on driving behavior, nor do other measures of urban form. Point estimates suggest that households that prefer not to drive sort into denser neighborhoods, and that denser locations often have unobserved characteristics that increase the value of travel. However, both effects are economically and statistically small.

These estimates allow us to estimate the effects of hypothetical densification policies on aggregate driving. The decile of US population living at the lowest density occupies about 83% of the area of the continental US, while the population of the highest density decile occupies about 0.2%. If we consolidate the entire decile of population living at the lowest density into an area the same size as that occupied by the densest decile of population, an 867 fold increase in density for this part of the population, we achieve about a 5% decrease in aggregate driving. Draconian as it is, this example probably overstates the effect of densification on driving. More plausible densification policies create areas where density increases by decreasing density in other areas.

Provided they remain inhabited, the latter areas experience an *increase* in per capita driving while the denser areas experience a decrease. Our estimates suggest that such policies also cause only modest decreases in aggregate driving. A comparison of the effects of densification policies with what is known about the effects of gasoline taxes and congestion prices suggests that densification policies are unlikely to be a cost effective way to reduce aggregate driving or traffic congestion.

Our econometric framework derives from a simple theoretical description of the utility derived from driving in different landscapes and our econometric estimates allow us to recover the structural parameters of this model. This leads to two further conclusions. First, to the extent that we are able to check, our results are consistent with related results in the literature. Second, that the changes in the utility derived from driving appear to be small, and in particular seem consistent with our finding that sorting does not primarily explain the relationship between urban form and driving.

Our results are of interest for a number of reasons. First, traffic congestion is an important economic problem and land use change is a widely proposed policy response to this problem. For example: in a State of Arizona Department of Transportation professional paper, Kuzmyak (2012) concludes that "greater adherence to smart growth principles of compact, mixed-land use,..., may result in important reductions in average trip lengths and VMT [vehicle-miles traveled] demand on local and regional roads"; the US Department of Transportation states that "[t]ransportation demand is reduced when residential and commercial uses are planned to be within close proximity to each other...";[1] while the Brookings Institute asserts that to relieve traffic congestion "[w]e need to make places more efficient by joining up transportation with the housing, real estate, commercial, and industrial decisions it drives".[2] Our results provide a basis for evaluating such claims.

More generally, the hypothesis that changes in urban form have economically important effects on driving behavior is the subject of a large literature in urban planning, as we discuss in more detail below. This literature is based almost entirely on cross-sectional associations, typically calculated from small samples describing small areas. That is, the current empirical literature does not provide a foundation that allows policy makers to use urban planing to affect driving with any confidence of achieving their desired outcome. We improve on the existing literature by providing

[1] `http://www.fhwa.dot.gov/planning/processes/land_use/land_use_tools/page02.cfm#toc380582783`, September 17, 2015.

[2] `http://www.brookings.edu/blogs/the-avenue/posts/2015/08/27-urban-traffic-congestion-puentes?rssid=The+Avenue`, September 17, 2015.

plausibly causal estimates. We also exploit better data than has previously been available. This permits us to test different measures of urban form against each other and to investigate the scale at which urban form affects driving. Hence, we can provide some insight into which of the many correlated characteristics of urban form influence driving and which do not.

Urban planning also plays a prominent role in policy discussions of carbon abatement. The Fourth Assessment Report of the IPCC discusses land use as a potential policy to reduce the demand for automobile travel (e.g., section 5.5.1.1 of Intergovernmental Panel on Climate Change, 2007), the more recent Fifth Assessment suggests that "[u]rban Densification in the USA over about 50 years could reduce fuel use by 9-16%" (table 8.3, Intergovernmental Panel on Climate Change, 2014), and California's Senate Bill 375 (September 7,2006) asserts that "it will be necessary to achieve significant additional greenhouse gas reductions from changed land use patterns and improved transportation". Because driving accounts for a large share of carbon emissions, our analysis helps to inform these sorts of policies by providing causal estimates of how particular changes to urban form affect driving behavior.

Finally, the value of real estate holdings by US households nears 25 trillion dollars, while annual household expenditures on driving are above a trillion dollars.[3] Given the magnitude of the allocations involved, understanding how the spatial configuration of structures affects driving behavior is of intrinsic interest. More simply, much of the world's population lives in cities and the construction and operation of these cities is costly. Understanding how to better organize cities is clearly a policy question of the first order.

2. Literature

Three strands of literature are relevant to our inquiry. The first is the large literature on the relationship between urban form and driving. The second investigates the relationship between the characteristics of a place and behavior. The third examines the extent to which unobserved attributes of places affect the way that cities develop in these places.

[3]Board of Governors of the Federal Reserve System (2015) and U.S. Department of Labor, Bureau of Labor Statistics (2014). The first number arguably understates the aggregate value of the US property stocks as it concerns only real estate holdings by households. The number for transportation only concerns nominal expenditure on transportation, at least 95% of which is road transportation, and ignores the value of time spent in vehicle.

Urban form and driving

The relationship between urban form and driving is 'one of the most heavily researched subjects in urban planning' (Ewing and Cervero, 2010) and searching for the phrase 'urban form and driving' yields about 950,000 links on Google Scholar. The problem has also received attention from economists.

This literature is the subject of several surveys, including Ewing and Cervero (2010), Handy (2005), Cao, Mokhtarian, and Handy (2009), Ewing and Cervero (2001), and Boarnet (2011). The literature is overwhelmingly based on cross-sectional regressions. Most of the papers surveyed in Ewing and Cervero (2010) rely on samples of fewer than 1,000 people or households, though Boarnet (2011) describes a few studies based on samples approaching 10,000. Typically, these samples cover small geographic areas. Bento, Cropper, Mobarak, and Vinha (2005) is an exception in both regards. It is based on the 2002 wave of the NHTS and therefore exploits a national sample of about 22,000 households, although they restrict attention to urban form variables measured at the level of the MSA.

Research typically revolves around estimating the effect on driving behavior of the 'three D's' proposed in Cervero and Kockelman (1997); 'Density', 'Diversity' and 'Design'. That is: the density of residents or employment; the diversity of activity, in particular the extent to which residential and other uses are mixed; and, usually, characteristics of various transportation networks. Our data will allow us to investigate two of these three, the density and diversity of neighborhood population and economic activity, and to touch on the third, the characteristics of the neighborhood street network.

The possibility that an individual or household's location choice may depend on their predisposition to travel is widely recognized and Cao *et al.* (2009) survey the econometric techniques that have been applied to the problem. In short, the literature has yet to identify a good source of random or quasi-random variation in neighborhood choice. To the extent that the literature implements instrumental variables estimations to deal with sorting, it relies on variables such as race or housing stock age that seem unlikely to satisfy the relevant exogeneity condition and are subject to the conceptual problem we describe in section 4. Panel data sets are almost unknown and those that are available describe small areas and samples. In this light, the approach to the problem of sorting that we develop below is an advance.

The possibility that the neighborhood characteristics of interest may be correlated with unob-

served characteristics that affect driving, our 'endogenous density' problem, has been addressed only by Blaudin de Thé and Lafourcade (2015).

Although the relationship between urban form and total travel distance (e.g., Bento *et al.*, 2005, Brownstone and Golob, 2009) and the journey to work (e.g., Gordon, Kumar, and Richardson, 1989, Giuliano and Small, 1993, Glaeser and Kahn, 2004) have been a primary focus of this literature, the literature has also investigated the relationship between urban form and other travel outcomes, including pedestrian trips and energy consumption (e.g., Brownstone and Golob, 2009, Glaeser and Kahn, 2008, Blaudin de Thé and Lafourcade, 2015).

Places and behavior

Differences in individual outcomes across locations are widely observed, and determining whether these differences reflect causal effects of the location or the sorting of different types of people is a pervasive problem in economics. The justly famous 'Moving To Opportunity' experiment induced a random assignment of poor households to move to nicer neighborhoods than they would otherwise have chosen. The effects of this move for teenage children on educational attainment and economic outcomes are small while the effects for children under the age of 13 appear to be large (Kling, Ludwig, and Katz, 2005, Chetty, Hendren, and Katz, 2015). These effects are quite different from what we expect from a cross-sectional analysis where outcomes for individuals are usually strongly positively correlated across space (see Ioannides and Topa, 2010, for a survey).

The urban economics literature has devoted considerable effort to investigating the relationship between city size and wages. Combes, Duranton, and Gobillon (2008) find that about half of this effect is accounted for by basic demographic controls and unobserved individual traits and half is causal. Eid, Overman, Puga, and Turner (2008) find that all of the cross-sectional relationship between obesity and neighborhood characteristics can be accounted for by individual fixed effects. Dahl (2002) finds that cross-sectional estimates of the returns to education are biased up by the tendency of educated individuals to migrate to states where the returns to education are high. Currie and Walker (2011) find that automobile pollution has a causal effect on the health of residents in neighborhoods exposed to pollution and do not find evidence that this reflects the sorting of unhealthy residents into polluted neighborhoods.

Summing up, cross-sectional differences in outcomes across locations are sometimes due to sorting of people on the basis of observable or unobservable characteristics, but are also sometimes

due to causal effects of locations on people. Thus, concern about the role of sorting in determining the cross-sectional relationship between urban form and driving is well founded, but we have little basis for predicting the importance of sorting for our particular problem.

In addition, we note that extant approaches dealing with sorting all rely on strong identification assumptions. Some of the work cited above (e.g. Combes *et al.*, 2008, Eid *et al.*, 2008) relies on panel data and assumes that mobility is exogenous. Alternatively, the literature often assumes that sorting takes place for some choices (or particular spatial scales) but not others. For instance, Evans, Oates, and Schwab (1992) assume no sorting across cities but sorting within cities whereas Bayer, Ross, and Topa (2008) assume sorting across neighbourhoods but not across blocks within neighbourhoods. The approach we develop below relies instead on imperfect mobility.

Endogeneity of infrastructure

There is an active literature investigating the role of transportation infrastructure on the way that cities develop. Baum-Snow (2007) is a pioneering contribution to this literature and investigates the role that the interstate highway system played in the decentralization of US cities between 1950 and 1990. To address the possibility that highways were assigned to cities that would otherwise have decentralized, Baum-Snow (2007) relies on an early plan of the interstate system as a source of quasi-random variation. Duranton and Turner (2012) use a similar methodology to investigate the extent to which interstate highways caused population and employment growth in US cities. This literature is now large and is surveyed in Redding and Turner (2015). Often, but not always, this literature finds evidence that the assignment of infrastructure to cities is not random. For example, the results in both Baum-Snow (2007) and Duranton and Turner (2012) suggest that interstate highways are disproportionately assigned to US cities that grow less slowly than would be predicted from observable characteristics.

The literature on the effects of infrastructure is concerned with city level outcomes such as population growth or decentralization. The present inquiry is, for the most part, concerned with a smaller spatial scale. It is also the first to consider the possibility that neighborhood characteristics related to density, diversity and design may be correlated with unobserved characteristics that affect driving. Our solution to this problem involves instrumenting for urban form with underground geology, the pervasiveness of aquifers, and earthquake and landslide risk. Although our use of these instruments in this context is novel, the idea derives from Rosenthal and Strange

(2008). They use bedrock characteristics as an external predictor of population density because deep bedrock usually makes construction more expensive and limits the intensity of development.

3. A simple model of urban form and driving

To illuminate the inference problems that our empirical investigation must overcome, we first present a simple model of equilibrium driving behavior. Consistent with the regressions below, we focus on total travel distance by households. This is (arguably) the measure of travel that has received the most attention from both the academic literature and policy makers because it maps fairly directly into congestion, local pollution, and carbon emissions.

Consider a location with unit area and population density X. A resident with income W derives utility from the consumption of differentiated varieties and the numéraire good C,

$$U = C + \theta\,\delta\,\left(\int_{i=1}^{N} Q(i)\mathrm{d}i\right)^{\rho}, \tag{1}$$

where θ is a resident-specific term, δ is a location-specific term, and $0 < \rho < 1$. To consume a differentiated variety, the resident must make a dedicated trip. The cost of a unit of variety i is $\tau D(i)$ where $D(i)$ is the travel distance to variety i and τ parameterizes the cost of travel.[4] We imagine that restaurants and movie theaters as well as local recreational amenities such as parks or museums would each constitute a 'variety' in this context.

Residence in a location requires the consumption of a unit of housing at price P_h. The budget constraint of a resident is thus $C + P_h + \int_{i=1}^{N} \tau D(i) Q(i)\mathrm{d}i = W$. To keep the problem tractable we assume that: (i) there are 'enough' varieties so that residents never consume the full set of available varieties, (ii) varieties can only be consumed in unit quantity $Q(i) = 1$, and (iii) varieties are symmetrically located around the resident so that $D(i) = D$ for all varieties i.[5] The budget constraint simplifies to $C + P_h + N\tau D = W$. Next, we can substitute this budget constraint into the utility function and simplify to obtain

$$U(N) = W - P_h + \theta\delta\, N^{\rho} - N\tau D\,. \tag{2}$$

[4]We impose an 'iceberg' (multiplicative) specification for travel costs to keep the consumer program tractable. This type of specification is extremely standard to model trade in goods (Head and Mayer, 2014). Its gravity implications also appear to describe commuting patterns extremely well (Ahlfeldt, Redding, Sturm, and Wolf, 2016).

[5]Besides imposing convenient functional forms, our simple model also ignores many common features of travel such as the possibility of chaining trips. In addition, we do not explicitly deal with commutes and other work-related trips. Some of these complications are addressed in our regressions below. Our priority is to develop a tractable framework to underpin our regressions and to highlight the key econometric challenges that we face.

Assuming income is high enough, the maximization of utility with respect to the number of varieties implies that

$$N = \left(\frac{\rho\theta\delta}{\tau D}\right)^{\frac{1}{1-\rho}} . \tag{3}$$

This expression indicates that residents take more trips if they have a greater taste for differentiated varieties, θ. For instance, some residents may enjoy dining out more than others. More generally, θ captures an individual resident's propensity to travel. The number of trips also increases with δ. For instance, a neighborhood near a nice beach may generate more trips than a neighborhood near a dirty beach. Our model can capture this by assigning one location a higher value of δ. Residents also make more trips when they are cheaper. This can occur because the cost of travel, τ, is lower or because trip distance, D, is shorter. In turn, differences in τ and D across locations arise as locations differ in how congested they are and in how compact they are. Finally, the number of trips increases with ρ, which measures the (opposite of the) concavity of the utility function with respect to differentiated varieties.

We are ultimately interested in how travel distance relates to density around a resident. Total travel distance is given by,

$$Y \equiv N\,D = \left(\frac{\rho\theta\delta}{\tau}\right)^{\frac{1}{1-\rho}} \left(\frac{1}{D}\right)^{\frac{\rho}{1-\rho}} , \tag{4}$$

where the last equality results from the use of equation (3). Like the number of trips, travel distance also increases with θ and δ and decreases with the unit cost of travel τ and trip distance D. The latter effect arises because the demand for trips is elastic with respect to trip distance.

The effect of density at a location on the demand for travel is ambiguous. A higher density reduces trip distance through greater accessibility. In turn, this reduces travel distance for a given number of trips but it also makes trips cheaper and thus elicits more trips. In addition, a higher density increases the unit cost of travel through more congestion. The net effect of improved accessibility and increased congestion on travel distance is ambiguous.

More specifically, to model the reduction in travel distance per trip that comes with greater population density, we assume

$$D = X^{-\zeta} , \tag{5}$$

where we refer to ζ as the accessibility elasticity.[6]

[6] As accessibility improves residents face both more and closer options. Our formulation reflects this tradeoff, albeit in a simple, reduced-form manner. See Couture (2014) for micro-foundations.

Turning to congestion, we expect it to depend on aggregate travel in the location. To capture this stylized fact in our model, suppose that travel costs are

$$\tau = \left(X\overline{Y}\right)^{\phi}, \tag{6}$$

where $\overline{Y}$ is mean travel distance and ϕ measures the elasticity of travel cost per unit with respect to aggregate travel. We refer to ϕ as the congestion elasticity. Note that, by construction, individuals do not account for their impact on τ. Thus, the equilibrium we describe below will not be optimal. There will be too much driving. It follows that, even if changes in urban form reduce congestion and increase utility, they do not remove the need for congestion pricing. We return to this point below.

After defining $\overline{\theta} = \left[\frac{1}{X}\int_{j\in X}\theta(j)^{1/(1-\rho)}\mathrm{d}j\right]^{1-\rho}$, an index of the preferences of residents in a location, inserting equations (5) and (6) into (4) implies

$$Y = \theta^{\frac{1}{1-\rho}}\left(\frac{\rho\delta}{\overline{\theta}^{\frac{1}{1-\rho}}}\right)^{\frac{1}{1-\rho+\phi}} X^{-\frac{\phi-\zeta\rho}{1-\rho+\phi}}, \tag{7}$$

after simplifications.

Substituting equations (4)-(7) into the utility function (2) leads to:

$$\begin{aligned} U &= W - P_h + (1-\rho)\rho^{\frac{1}{1-\rho}}\left(\frac{\theta\delta}{\tau D}\right)^{\frac{1}{1-\rho}} \\ &= W - P_h + (1-\rho)\rho^{\frac{\rho}{1-\rho}-\phi}\delta^{\frac{1}{1-\rho}-\phi}\theta^{\frac{1}{1-\rho}}\overline{\theta}^{\frac{\phi}{1-\rho}\frac{\rho-\phi}{1-\rho+\phi}}X^{\zeta-\phi+\phi\frac{\phi-\zeta\rho}{1-\rho+\phi}}. \end{aligned} \tag{8}$$

We draw four conclusions from equations (7) and (8). First, if the congestion elasticity, ϕ, is larger than the product of the accessibility elasticity, ζ, and the utility term, ρ, then travel distance decreases with population density. Two forces are at play. Travel distance increases with population density because of improved accessibility. This increase in travel distance also depends on how much the consumption of differentiated goods that require travel is valued in utility terms. At the same time, the cost of travelling also increases with density because of rising congestion. It is only when $\phi > \zeta\rho$ that travel distance declines with population density.

Second, even if we assume no other cost or benefit from density, the effect of density on equilibrium utility is ambiguous. The coefficient associated with density in equation (8), $\zeta - \phi + \phi\frac{\phi-\zeta\rho}{1-\rho+\phi}$ is involved because of the feedback of aggregate travel distance into individual driving that occurs through congestion in equation (6). However, the term in τD in the first line of equation (8) makes it clear that utility increases with density when it reduces trip distance D more than it increases

unit travel cost τ. For $\rho < 1$, utility increases with density when the accessibility externality is large enough, $\zeta > \frac{\phi}{1+\phi}$.

We note that it is only when the exponent on X is positive for utility in equation (8) and the exponent on X for travel distance is negative in equation (7) that travel declines while utility increases with density. These two conditions require $\frac{\phi}{\rho} > \zeta > \frac{\phi}{1+\phi}$. That is, the accessibility externality, ζ, must be large enough for utility to increase with density but not too large to avoid an increase in travel with density. Hence, it is only in a particular parameter region that the model both predicts a widely conjectured empirical relationship and satisfies a necessary condition to rationalize policies to increase population density.

Third, it is easy to see from equation (8) that $\frac{\partial^2 U}{\partial \theta \partial X} \geq 0$ when $\frac{\partial U}{\partial X} \geq 0$. In words, there is a positive complementarity between the propensity to take trips and population density when utility increases with population density. In this case, residents with a greater propensity to make trips benefit more from a higher population density than residents with a smaller θ. In section 7, we extend our model to solve for the location choices of residents. In this extension, we show that the single-crossing condition implied by this complementarity between the propensity to take trips and density leads to the perfect sorting of residents across locations of different density. More specifically, residents with a greater propensity to make trips choose to locate in denser locations. The opposite form of sorting occurs when utility decreases with density. Hence, in general, we expect a non-zero correlation between the propensity to make trips, θ, and population density, X, to be a feature of our data.

In addition, it is also easy to see that in general, $\frac{\partial^2 U}{\partial \delta \partial X} \neq 0$. Hence, we should also expect a non-zero correlation between how beneficial trips are in a location, δ, and population density, X.

If residential sorting is perfect in equilibrium, then must have $\overline{\theta} = \theta$. In fact, we expect sorting to be less precise than this, and our econometric model relies on the fact that residential mobility is imperfect. To describe such a process parsimoniously, we instead suppose that $\theta = \overline{\theta}\nu$, where ν is an error term.[7] Using this relationship in equation (7) and taking logs then gives,

$$y = \frac{\log \rho}{1 - \rho + \phi} - \frac{\phi - \zeta\rho}{1 - \rho + \phi} x + \epsilon \,, \tag{9}$$

[7]In regressions reported in table 2, we will see that our data do not allow us to separately identify the effects of individual and neighborhood average demographic characteristics on household driving. This suggests that ν is small relative to θ.

where

$$\epsilon = \frac{\log \delta}{1-\rho+\phi} + \frac{\rho-\phi}{(1-\rho+\phi)(1-\rho)} \log \theta + \frac{1}{(1-\rho+\phi)(1-\rho)} \log \nu \,. \tag{10}$$

and $y \equiv \log Y$ and $x \equiv \log X$.

Equation (9) describes a regression of driving on urban form. This regression, typically conducted with cross-sectional survey data, forms the basis of the large literature described in section 2. Because local gains from trips, δ, and the propensity to make trips, θ, are not observed, they enter the error term. Given their expected correlation with population density, the estimated coefficient of x is potentially biased. The sorting of travellers and the endogeneity of density are the two main identification challenges we face in our empirical work below.[8]

4. Econometric model

We would like to estimate the relationship between urban form and driving behavior. We begin by considering the problem of sorting and then turn to the problem of endogenous urban form.

Each person (household) is assigned to a geographic unit. As we discuss below, these will usually be regular grid cells approximately 1 kilometer square. For each such unit we construct measures of urban form, usually a measure of density, which we also discuss below. Let i index individuals and j index residential locations. We are interested in explaining how driving behavior y_{ij} varies with urban form. More specifically, we are interested in knowing how the driving behavior of a randomly selected person or household changes when we change urban form in or around their residential location.

Let x_j^0 denote the urban form variable of interest for geographic unit j at an initial period (density in the model above), usually around 1990 and let x_j^1 denote the urban form variable of interest usually around 2010, contemporaneous to y. Define $\Delta x_j = x_j^1 - x_j^0$. We observe both contemporaneous and historical descriptions of urban form at each location, but we observe each driver only once.

Suppose that driving for each person is described by the following equation,

$$y_{ij} = \theta_{ij} + \beta x_j + \delta_j, \tag{11}$$

[8]The direction of the bias is ambiguous because, as pointed out above, the correlation between density and the taste parameter θ can be positive or negative, depending on parameter values.

so that observed driving for each person is determined by an individual specific intercept, θ_{ij}, a location specific intercept, δ_j, and the urban form in person i's location j, x_j. The parameter of interest, β, measures the effect of local urban form on distance travelled.

We note that this is equivalent to the equilibrium driving equation (9) derived above, where, in a slight abuse of notation, we renormalize θ_{ij} and δ_j to improve legibility. Importantly, in both equations (9) and (11) individual taste parameters and location specific effects enter only through the intercept. They do not lead to individual or neighborhood level differences in β, the rate at which individuals change their behavior in response to density. This simplifies our econometric task considerably and we appeal to theoretical analysis above to justify this restriction. This assumption also finds some empirical support in our results: we perform our main regression on many different subsamples and do not find any measurable differences in β across samples.

Given equation (11), our two main inference problems are that people do not choose their locations at random and that observed and unobserved attributes of urban form are correlated and may affect driving. We address each problem in turn.

To begin, suppose that individual specific intercepts are not observed, but are drawn from the real interval Θ, let w denote observable individual characteristics related to location choice and let the distribution of individual types at each location j be determined by

$$\theta_{ij} = \alpha_0 + \alpha_1 x_j + \alpha_2 w_{ij} + \mu_{ij} \tag{12}$$

and $E(x_j \mu_{ij}) = 0$. That is, the assignment of types to location j depends on urban form, on observable individual characteristics, and on unobserved individual characteristics. If $\beta > 0$, then we expect drivers with a larger θ to sort into neighborhoods with a larger x and conversely. As μ increases, residents derive more utility from trips for reasons unrelated to x.

Using both equation (12) and (11), we have that

$$\begin{aligned} y_{ij} &= (\alpha_0 + \alpha_1 x_j + \alpha_2 w_{ij} + \mu_{ij}) + \beta x_j + \delta_j \\ &= \alpha_0 + (\alpha_1 + \beta) x_j + \alpha_2 w_{ij} + \epsilon_{ij}, \end{aligned} \tag{13}$$

where $\epsilon = \mu + \delta$. Thus, if $\alpha_1 \neq 0$ or $E(\epsilon_j x_j) \neq 0$, OLS estimates of β will be biased.

Our approach to this sorting problem relies on an assumption of imperfect mobility. We now consider two time periods $t = 0$ and $t = 1$ and suppose that at $t = 0$ all agents match to locations as described above. At $t = 1$ a randomly selected fraction, s_j, of these residents relocates and is

replaced by agents who sort on the basis of current conditions. With these assumptions in place, for a location where $x_j^1 = x_j^0 + \Delta x_j$, expected driving at $t = 1$ is

$$\begin{aligned} y_{ij}^1 &= (1 - s_j)\left[(\alpha_0 + \alpha_1 x_j^0 + \alpha_2 w_{ij} + \mu_{ij}) + \beta x_j^1 + \delta_j\right] \\ &\quad + s_j\left[(\alpha_0 + \alpha_1 x_j^1 + \alpha_2 w_{ij} + \mu_{ij}) + \beta x_j^1 + \delta_j\right] \\ &= \alpha_0 + (\alpha_1 + \beta) x_j^0 + \alpha_1 s_j \Delta x_j + \beta \Delta x_j + \alpha_2 w_{ij} + \epsilon_{ij} \\ &= A_0 + A_1 x_j^0 + A_2 s_j \Delta x_j + A_3 \Delta x_j + \alpha_2 w_{ij} + \epsilon_{ij}. \end{aligned} \tag{14}$$

In fact, we will not observe s_j directly. Instead, we observe characteristics that vary systematically with the mobility rate, e.g., driver age or mean housing tenure in the driver's home cell. To understand how this allows similar tests, denote our mobility proxy by $\tilde{s}$ and suppose that mobility varies with $\tilde{s}$ according to $s = g(\tilde{s})$. Taking a linear approximation, we have $s = \gamma_1 \tilde{s}$, where $\gamma_1 \neq 0$ is assumed. Substituting this expression for s into (14) we see that the coefficient on $\tilde{s}\Delta x$ is $\alpha\gamma_1$. Substituting into (14) gives

$$y_{ij}^1 = A_0 + A_1 x_j^0 + A_2 \gamma_1 \tilde{s}_j \Delta x_j + A_3 \Delta x_j + A_4 w_{ij} + \epsilon_j. \tag{15}$$

Equation (14) suggests two parametric tests of the importance of sorting. First, the difference between the coefficients of x^0 and Δx is α_1. This is the parameter that describes how the unobserved individual propensity to drive varies with urban form in equation (12). Since $\alpha_1 = A_1 - A_3$, we can reject the hypothesis that $\alpha_1 = 0$ by rejecting the hypothesis that $A_1 = A_3$. Second, we can reject the hypothesis that $\alpha_1 = 0$ by rejecting the hypothesis that $A_2\gamma_1 = 0$. In fact, our estimates will generally indicate the $A_2\gamma_1$ is tiny and not significantly different from zero. However, because this test compounds two structural coefficients, we regard it as less informative than tests based on the difference $A_1 - A_3$. Although they are imprecise, point estimates in our preferred specification suggest that $\alpha_1 < 0$ and is about one sixth the magnitude of β. That is, individuals with smaller propensity to drive move to dense places, but this sorting is most likely makes only a modest contribution to the observed relationship between urban form and driving.

This methodology requires two comments. Identification rests on the assumption that as urban form changes, so do the characteristics of the marginal resident. Not only does this seem like a reasonable hypothesis, it also a common sense definition of 'sorting'. While we express the intuition precisely and in particular functional forms, the underlying intuition seems unrestrictive. Second, as we have described it, sorting affects only residents moving to a location, not those

moving away from it. More realistically, we might expect a non-random sample of people to move from a location, and in the case of an increase in density, they should value density less highly than the average current resident, who in turn should value density less highly the average arrival. Generalizing our framework to describe this intuition precisely is straightforward and leads to an indistinguishable empirical strategy.

While the estimation described in equation (15) addresses the problem of sorting by unobserved individual characteristics, it does not address the possibility of omitted location variables correlated with urban form and driving. For example, municipal snow removal may be systematically worse in dense areas and affect driving. To address this problem, we consider the system of equations,

$$y_{ij} = \theta_i + \beta x_j + \delta_j \,, \tag{16}$$

$$x_j = \gamma_0 + \gamma_1 z_j + \eta_j \,. \tag{17}$$

In the context of this system, our omitted variables problem may be stated as $E(x_j \delta_j) \neq 0$. We resolve this problem by relying on instrumental variables estimation. As the system above suggests, this requires an instrument that predicts urban form but that does not otherwise affect driving, or more formally, that $\gamma_1 \neq 0$ and $E(z_j \delta_j) = 0$. In our empirical work, we rely on various measures of subterranean geology as instrumental variables. As we will see, these measures are important determinants of urban form and it is difficult to imagine other channels through which they could affect driving behavior than by affecting the urban form.

Although this is a standard instrumental variables estimation, in our context, it requires two comments. First, we should not expect our instrumental variables estimation to resolve the problem of sorting. To see this, let $\widehat{x}_j = \gamma_0 + \gamma_1 z_j$ and rewrite equation (13) using (17) as,

$$\begin{aligned} y_{ij} &= \alpha_0 + (\alpha_1 + \beta)(\widehat{x}_j + \eta_j) + \epsilon_j \,, \\ &= \alpha_0 + (\alpha_1 + \beta)\widehat{x}_j + ((\alpha_1 + \beta)\eta_j) + \epsilon_j) \,. \end{aligned}$$

That is, as long as residents sort on the component of the urban form predicted by underground geology in the same way as they sort on the residual component, the instrumental variables regression does not lead to unbiased estimates of β. Thus, instrumental variables estimation can solve the problem of unobserved local characteristics, but it cannot solve the problem of unobserved individual characteristics.

Table 1: Descriptive statistics for NHTS households, MSA sample

Variable	Mean	Std. Dev.	5th percentile	95th percentile	Observations
Vehicles km travelled (VKT)	37,022	29,826	4,459	87,906	99,875
log VKT	10.17	1.01	8.40	11.38	99,875
Annual VKT	33,014	29,766	3,645	82,620	93,602
Odometer VKT	33,123	24,647	6,388	74,483	71,742
Household daily VKT	73.2	66.8	6.5	208.1	83,313
Household daily travel minutes	98.7	70.0	17	234	83,313
Household daily speed	42.6	38.8	13.9	75.6	83,313
Distance to work	22.5	34.4	1.6	61.6	95,532
10-km density	1,072	1,559	44.9	3,222	99,875
log 10-km density	6.30	1.31	3.81	8.08	99,875
10-km population density	755	1,027	34.7	2,211	99,875
10-km share developed (%)	4.40	5.61	0.07	15.5	99,875

Notes: Authors calculations for 2006-2011. Distances are measured in kilometers and monetary values in current American dollars. Household age is mean age for the adult members of the household. Household daily VKT, travel time, and speed are computed for all households with positive travel by summing all trips across the surveyed members of the household. Household speed is computed by dividing VKT by travel time for each household and averaging across households. 'Density' refers to the sum of jobs and residents unless it is qualified by employment or residential population. All densities are reported per square kilometer.

In light of the intuition above, we would ideally implement our instrumental variables strategy in the context of equation (15) which explicitly accounts for sorting. In practice, our instruments are not able to predict changes in urban form, only levels. Thus, in spite of its theoretical appeal, this strategy is beyond the reach of our data. With this said, the data strongly suggest that neither sorting nor omitted variables cause economically important biases in our estimates, so we can reasonably conjecture that allowing these two biases to interact would also be unimportant.

5. Data

Our analysis requires three main types of data; household and individual level travel behavior, a description of urban form for each household, and finally, a description of subterranean geology. To implement our response to the sorting problem, we require panel data describing urban form, but only cross-sectional travel data.

We also require a way of matching survey respondents to landscapes. To accomplish this, we construct a regular grid of 990 meter cells covering the entire continental US. Each household is matched to the cell with its centroid closest to the centroid of the household's census block group. We will refer to this cell as an individual or household's 'home cell', and in a slight abuse of

language, describe cells as having an area of one square kilometer. We convert all data describing urban form to this resolution as described below. With this data structure in place, we can construct urban form measures for each household on the basis of arbitrary geographies by averaging over the relevant sets of grid cells. In particular, we can examine the square kilometer surrounding each household by reporting the characteristics of its home cell, we can average over all cells within 10 kilometers of the home cell or over all cells lying in the same MSA.

Data on individual travel behavior come from the 2008-2009 National Household Transportation Surveys (NHTS).[9] The NHTS survey reports several measures of total annual driving for each household or individual in a nationally representative sample of households. Our main dependent variable is household annual vehicle kilometers travelled (VKT) and is reported in the first row of table 1.[10] This measure of household annual mileage is computed by the survey administrators, 'bestmiles', and is their preferred measure. In robustness checks, we consider four other measures of individual and household driving distance, stated annual vehicle kilometers traveled, a reported odometer measure of kilometers traveled, individual daily kilometers traveled on the survey day, and distance to work.

Table 1 reports descriptive statistics for several measures of driving from the NHTS. The three measures of total household driving have sample means of 37,022, 33,014 and 33,123 kilometers over slightly different samples of households. Except where noted otherwise, we restrict attention to households and individuals who live in MSAs.[11] Aggregating individual VKT and travel time at the household level implies that households travel 73.2 kilometers in 98.7 minutes at an average speed of 42.6 kilometers per hour.[12] Individual distance to work is 22.5 kilometers. These values reflect the sample of household members who filled out a travel diary reporting positive travel and those who reported driving to work. Because households often consist of more than one member, each NHTS survey describes more individuals than households.

[9] U.S. Department of Transportation, Federal Highway Administration (2009).

[10] Our initial NHTS sample contains 150,147 households of whom we can locate 149,638 on our grid. We have a positive measure of vehicle kilometers traveled for 136,530 households. After restricting our sample to those observations for which we have a full set of household and individual characteristics, we are left with 126,203 households, 99,875 of whom live in an MSA as defined by their 1999 definitions.

[11] This is purely for expositional convenience. It allows us to include MSA indicator variables in our regressions without changing our sample.

[12] This is an average across households. Dividing aggregate VKT by aggregate travel time implies a speed of 44.5 kilometers per hour. Couture, Duranton, and Turner (2014) report a mean speed per trip of 38.5 kilometers per hour. The differences between those numbers are due to the fact that shorter trips are slower. Averaging across trip gives them a greater weight than averaging total travel across households. In turn, a household average will also weight shorter trips more does the ratio of aggregate distance to aggregate travel time.

The NHTS survey reports household and individual demographics. These demographic variables provide a description of household race, size, income, educational attainment, and homeownership status. Mean household income is $71,257 and the average over households of the average age of household adults is 53.5 years. We also note that nearly 90% of households in our sample are homeowners.

Urban form data are more complicated. To measure the share of developed land cover, we rely on the 1992, 2002 and 2006 National Land Cover Data (NLCD).[13] While the NLCD reports many land cover classifications, we sum the urban classes in each year to measure the share of urban cover in each grid cell. Table 1 reports descriptive statistics for our sample. For an average survey respondent, 4.40% of the land area within 10 kilometers of their home cell is in urban cover in 2006.[14]

To assign 2000 census data to our grid cells, we distribute block group data to our grid cells using an area weighting based on a geocoded map of 2000 census block groups. We perform a similar exercise for 1990 and 2010.[15] With this correspondence between block groups and grid cells in place, we are able to assign any block group variable reported in the 1990, 2000 or 2010 census and in the American Community Survey (ACS) to our grid.[16] All urban form variables involving demographic characteristics are computed on this basis. Table 1 reports that for an average survey respondent, the average residential density within a radial distance of 10 kilometers of their home cell is 755 per square kilometer.

Using ACS and census tabulations, we also measure a number of other local characteristics such as an average length of tenure of 10.3 years and a renter share of 26.0%. We use these variables in estimations below and note that there is some variation across households in the mobility and tenure rates of their neighborhoods.

Employment data are based on zipcode business patterns. These data report both aggregate and sectoral employment by zipcode. We assign these data to our grid on the basis of zipcode

[13]United States Geological Survey (2000), United States Geological Survey (2011*a*) and United States Geological Survey (2011*b*).

[14]Note that all densities for rings around a survey respondent's home are normalized by the number of grid cells for which we have population and employment information. This prevents us from underestimating density for households who live by the sea, a lake, or uninhabitable terrain.

[15]The particular census maps we use are: Environmental Systems Research Institute (1998*a*), Environmental Systems Research Institute (2004), U.S. Department of Commerce, U.S. Census Bureau, Geography Division (2010).

[16]Sources for these data are: Missouri Census Data Center (1990), Missouri Census Data Center (2000), Missouri Census Data Center (2010) and National Historical Geographic Information System (2010).

maps using the same procedure that we use for census data.[17] We use zipcode business patterns for the years closest to the NHTS survey years, and to reduce measurement errors, average over the nominal year of the survey and the preceding year.

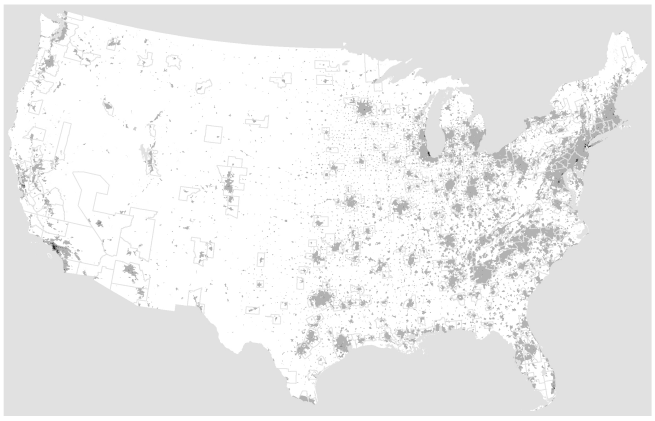
For much of our analysis, we use the total number of people living or working within 10 kilometers of each survey respondent to measure urban form and call this measure '10-kilometer density'. We sometimes also work with the corresponding measure based only on the household's home cell and call this measure '1-kilometer density'. When the scale of analysis is clear, we sometimes refer to these quantities as 'density'. Table 1 reports that for an average household survey respondent, the 10-kilometer density is 1,072 per square kilometer. People and jobs tend to be denser nearer survey respondents' homes, 1-kilometer density is 1,513.

Panel (a) of figure 1 illustrates the way that people in the US are exposed to our measure of 1-kilometer density. In this map, the white area contains the 10% of the US population living at the lowest density. This region is about 5.8 million square kilometers and 83% of the land area of the continental US. On average, the about 30 million people living in this region have 6.25 people or jobs in their home cell. The barely visible black areas in this map contain the 10% of the US population living at the highest densities. This area is less than 1,5000 square kilometers and about 0.2% of the land area of the continental US. On average, residents of these areas share their home cells with about 5,421 other people and workers. That is, the decile of US population living at the highest densities lives at densities about 870 times higher than the lowest density decile. The medium gray area in this figure houses the residual 80% of the population.

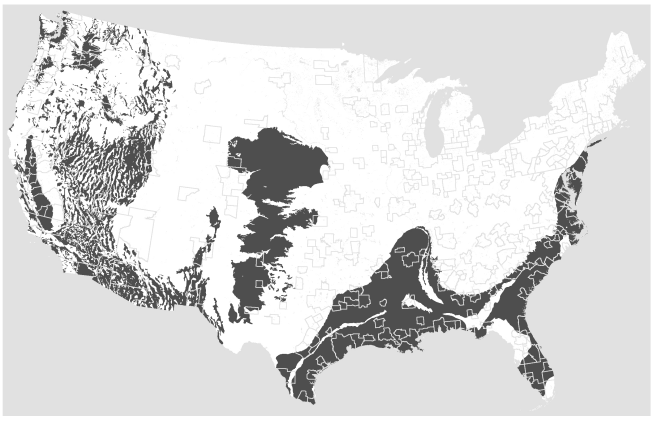
In our instrumental variables estimation, we rely on variables constructed from United States Geological Survey (2001, 2003, 2005). United States Geological Survey (2003) describes the incidence of aquifers in the continental US. Using this map, we determine which grid cells overlay consolidated or semi-consolidated aquifers. Panel (b) of figure 1 illustrates these pixels. Burchfield, Overman, Puga, and Turner (2006) find that an MSA level index of aquifer prevalence is a good predictor of an aggregate measure of urban form. We will also find aquifers are good predictors of local density. Usefully, the map indicates that aquifers are broadly distributed across the landscape so that instrumental-variables estimates will not be driven by variation within particular small regions. United States Geological Survey (2005) describes a measure of earthquake intensity that

[17]Sources for our zipcode maps are: Environmental Systems Research Institute (1998*b*), U.S. Department of Commerce, U.S. Census Bureau, Geography Division (2010).

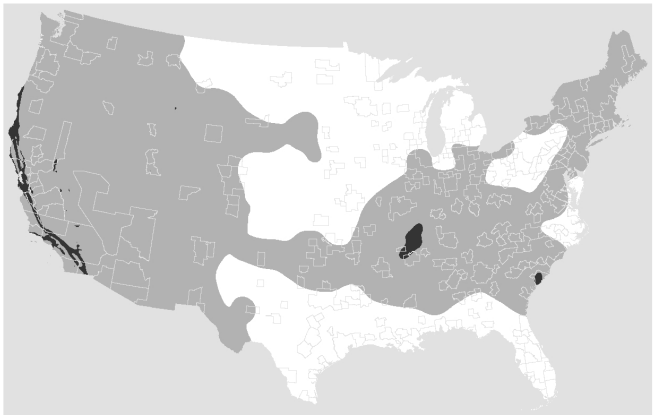
Figure 1: Maps.



(a) Density deciles of population



(b) Aquifers



(c) Earthquake intensity

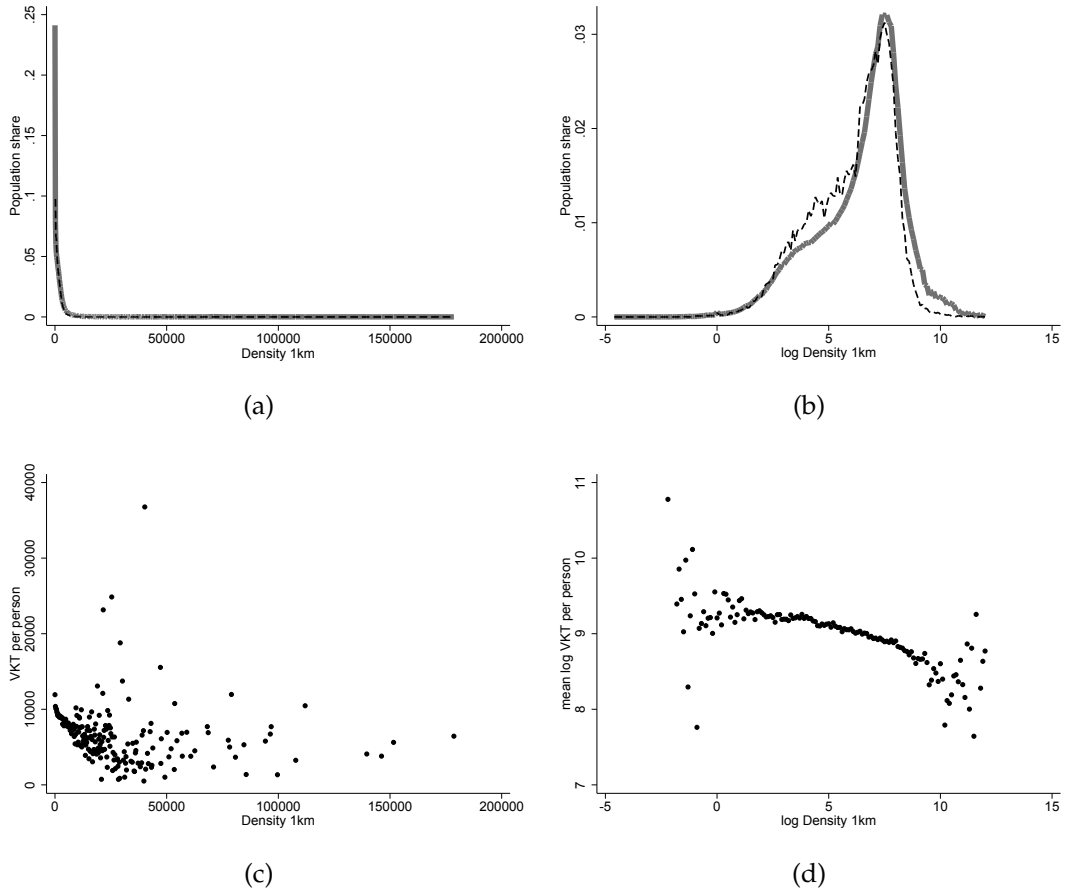


(d) Landslide risk

Notes: Panel(a): White indicates the area inhabited by people living in the bottom decile of density. Black indicates the area inhabited by people living in the top decile of density. Gray indicates the area inhabited by the 80% of the people living at intermediate densities. Panel (b): Gray indicates areas overlying unconsolidated or semi-consolidated aquifers and white indicates the absence of such aquifers. Panel (c): Darker gray indicates areas subject to larger earthquakes. Panel (d): Darker gray indicates areas subject to higher landslide risk. 2000 MSA boundaries shown in light gray in all four maps.

ranges from 0 to 18. We consolidate to three categories; low, medium or high earthquake exposure. Panel (c) of figure 1 illustrates these regions. Areas of high earthquake intensity are dark. United States Geological Survey (2001) describes landslide susceptibility. The source data contains six categories, which we consolidate to low, medium and high risk. Panel (d) of figure 1 illustrates high risk areas in dark gray, medium risk areas in light gray and low risk areas in white. Like the aquifers map, neither landslide nor earthquake risk are concentrated in small geographic areas so that instrumental variables estimates based on these variables are not driven by small regions of the country.

Figure 2: Density and VKT plots



Notes: All figures are histograms. Panel (a) describes the distribution of population conditional on density. The dashed black line describes this distribution for people surveyed by the NHTS. Here, we calculate number of NHTS people in each cell and then take the share of the total NHTS population living in cells of given densities. The heavy gray line provides the corresponding information for census population, i.e., for the whole contiguous continental US. Panel (b) is like panel (a), but density is in log scale. In panel (c) we first calculate mean per person VKT for each home cell. We then calculate mean VKT conditional on density as density varies. Panel (d) is like panel (c) but plots mean log per person VKT for each cell, conditional on log density.

Figure 2 presents four histograms that illustrate the variation in our data that drives our main estimation results. The left panels are identical to the right, except that axes on the right are logarithmic and on the left they are levels. The two top panels each present two probability distribution functions, the fine dashed black line for NHTS sample population and the heavy gray line for census population. In the left panel, the two distributions are indistinguishable, and as we might expect from panel (a) of figure 1, heavily skewed to the left. In the logarithmic version of panel (b) we see a distribution with a mode around 8, which converting from logs to levels, is density of about 3,000. The distribution of census and NHTS people diverges slightly at high densities. This may either

confirm that the NHTS tends to sample nearby households in less dense locations (U.S. Department of Transportation, Federal Highway Administration, 2009), or may reflect the fact that we exclude from our sample households for whom no household driving information is available and who live at much higher densities. We return to this issue below.

The bottom panels of figure 2 illustrate mean per person driving as a function of density.[18] The left panel shows a clear downward trend in driving at low levels of driving, but becomes disorganized at higher levels. The logarithmic scale right panel shows a clear log linear trend at intermediate levels of density. On the basis of the two bottom panels of figure 2, we rely on multiplicative rather than additive regression equations.

A few comments about these plots are required. First, the graphs in figure 2 are based on the whole sample of the NHTS for which we record household VKT, not the MSA only sample on which we base most of our regressions and table 1. Dropping the non-msa observations to be consistent with our regressions does not qualitatively change the figures. Second, while the far right part of panel (d) is noisy, it suggests that the effect of density on driving may increase at very high levels of density. This region of the graph describes a very small area of the country. Only about 260 cells have log(density)$>$ 10, and of these, fully two thirds are in the New York MSA. While we will look for this sort of non-linearity in our regressions, such results should be regarded with caution. These tiny areas of the country differ from the rest for many reasons besides their high density.

Finally, in some of our results we investigate the relationship between road density and driving. To describe the US road network, we rely on the 2007 National Highway Planning Network map (Federal Highway Administration, 2005). This map is part of the federal government's efforts to track roads that it helps to maintain or build. It describes all interstate highways and most state highways and arterial roads in urbanized areas. To construct measures of road density, for each grid cell containing a survey respondent, we construct disks of radius 5, 10 and 25 kilometers centered on this cell. For each such disk, we then calculate kilometers of each type of road network in that disk. In addition to these data we also use the PRISM gridded climate data (PRISM Climate Group at Oregon State University, 2012*a*,*b*) to measure temperature and precipitation in each grid cell.

[18] We calculate per person driving from the count of household members and household measures of driving.

6. Results

We proceed in steps. First, we present OLS results showing the relationship between our preferred measures of driving and urban form, household VKT and the density of residents and jobs within 10 kilometers. Second, we verify that these relationships are robust to different measures of driving and to the scale at which we calculate the urban form variable. Third, we consider the problems of sorting and endogeneity. Finally, we investigate other measures of urban form.

A OLS *estimations*

Table 2 reports the results of OLS regressions of driving on urban form in US MSAs. Our unit of observation is a household described by the 2008 NHTS. In every column, our dependent variable is the log of household VKT, reported in the second row of table 1. In all specifications, our measure of urban form is the log of 10-kilometer density, also as described by table 1.

In column 1, we regress log annual household VKT on the log of density to find an elasticity of -8.7%. Households in locations with a 10% higher density drive 0.87% less and a one-standard deviation increase in density within 10 kilometers is associated with a 0.11 standard deviation decrease in VKT. At the sample mean, this represents about 3,300 kilometers annually. Because the estimated coefficient of density is stable across specifications, these magnitudes are relevant to most of the tables presented below.

In column 2, we add household characteristics to our specification and estimate a slightly larger effect of density on household VKT with an elasticity of -9.8%. White and asian households drive about 2% more. Female households drive less. The coefficient of -0.26 implies that a single female is predicted to drive 23% less on average than a single male. Large households also drive more, but not proportionately so. The coefficients on log household size and the indicator for one-person households show that two-person households will drive about 30% more than one-person households. We also observe that VKT is concave in age. At age 20, an extra year of age is associated with 2% more driving. Then, VKT peaks around the age of 45 before declining. The elasticity of VKT with respect to income is large at around 26%. VKT increases with education (which is coded 1 to 5) for low levels of educational achievement and then decreases for the most educated households. Because the coefficients on households' characteristics are stable across specifications, we do not report or discuss them for subsequent tables.

Table 2: Driving and density, baseline OLS estimations

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
log 10-km density	-0.087[a]	-0.098[a]	-0.089[a]	-0.093[a]	-0.091[a]	-0.12[a]	-0.082[a]	-0.075[a]
	(0.0024)	(0.0020)	(0.0083)	(0.0089)	(0.013)	(0.0075)	(0.0051)	(0.0050)
White/Asian		0.019[b]			0.020[c]		0.024[b]	0.020[b]
		(0.0088)			(0.0100)		(0.0094)	(0.0090)
Female		-0.26[a]			-0.26[a]		-0.26[a]	-0.26[a]
		(0.010)			(0.012)		(0.011)	(0.010)
log household size		0.49[a]			0.49[a]		0.49[a]	0.49[a]
		(0.011)			(0.011)		(0.011)	(0.012)
Single		-0.24[a]			-0.24[a]		-0.24[a]	-0.24[a]
		(0.012)			(0.013)		(0.013)	(0.013)
Age		0.045[a]			0.044[a]		0.044[a]	0.044[a]
		(0.0010)			(0.00099)		(0.00098)	(0.0011)
Age^2 (/1000))		-0.51[a]			-0.50[a]		-0.50[a]	-0.50[a]
		(0.0098)			(0.0095)		(0.0097)	(0.010)
log income		0.26[a]			0.25[a]		0.25[a]	0.25[a]
		(0.0043)			(0.0054)		(0.0053)	(0.0050)
Education		0.10[a]			0.095[a]		0.092[a]	0.091[a]
		(0.016)			(0.014)		(0.014)	(0.016)
Education^2		-0.014[a]			-0.012[a]		-0.012[a]	-0.011[a]
		(0.0024)			(0.0020)		(0.0020)	(0.0024)
log precipitation			-0.015		0.051		-0.13[b]	-0.060
			(0.053)		(0.042)		(0.057)	(0.079)
log precipitation sd			0.025		-0.036		0.11[c]	0.090
			(0.069)		(0.049)		(0.063)	(0.076)
Temperature			0.062[b]		0.049[b]		-0.080	-0.021
			(0.028)		(0.024)		(0.060)	(0.039)
Temperature sd			-0.043[b]		-0.033[c]		0.052	0.014
			(0.019)		(0.017)		(0.040)	(0.027)
Share higher educ.				-0.50[a]	-0.18		-0.25[a]	-0.30[a]
				(0.072)	(0.11)		(0.078)	(0.071)
Share higher educ^2				-0.020	-0.20[b]		-0.21[b]	-0.12
				(0.080)	(0.092)		(0.091)	(0.073)
log local income				0.68[a]	0.18[a]		0.23[a]	0.22[a]
				(0.014)	(0.033)		(0.014)	(0.015)
R^2	0.01	0.36	0.01	0.05	0.37	0.02	0.37	0.36
Observations	99,875	99,875	99,875	99,875	99,875	99,875	99,875	99,874

Notes: The dependent variables is log household VKT in all columns. All regressions include a constant including 275 MSA fixed effects in columns 6 and 7 and 837 county fixed effects in column 8. Robust standard errors in parentheses, clustered by MSA in columns 3-7, and by county in column 8. [a], [b], [c]: significant at 1%, 5%, 10%.

In column 3, we consider geographic characteristics. Relative to column 1, the coefficient on density changes little. The results of this column indicate that VKT is higher where temperature is on average higher and varies less over the year. We find no significant effect of precipitation or its variation over the year. In other specifications like in column 7, we sometimes find that VKT is higher in places with less precipitation and more variation over the year.

In column 4, we consider neighborhood socio-economic characteristics. We find that driving declines with the share of university educated workers and increases with average local income. Because richer neighbourhoods are also on average denser, the coefficient on density also increases marginally in magnitude relative to the one estimated in column 1. In column 5, we consider all the controls together and estimate an elasticity of VKT with respect to density of -9.1%. Relative to column 4, we note that the magnitudes of neighbourhood characteristics drop sharply and lose significance. This is unsurprising. Richer and more educated households tend to live in richer and more educated neighbourhoods and the resulting co-linearity makes it difficult to separately estimate the effects of household and neighborhood income.

In column 6, we return to the specification of column 1 but also include a fixed effect for each Metropolitan Statistical Area (MSA). Estimating the elasticity of VKT with respect to density within MSAs yields a coefficient larger in magnitude relative to column 1. This is because richer and more educated households, both drive more and tend to locate in denser MSAs. Consistent with this, including all the household, geographic, and neighbourhood characteristics in column 7 gives a coefficient on density close to that of column 5 and, in spite of the extra fixed-effects, does not improve the fit of the regression. This specification is our benchmark OLS specification.[19] Finally, column 8 introduces a fixed effect for each of the 837 counties where metropolitan households are located. At -0.075, the coefficient on local density is marginally lower but statistically indistinguishable from our preferred coefficient in column 7 or from the coefficient obtained in column 1, the simplest estimation.

Our choice of explanatory variables in table 2 controls for obvious determinants of household travel, like household demographics or the geography of where they live. We also include controls for the neighborhood socio-economic characteristics, in spite of the fact they are may be correlated with urban form and capture some of its effect. Given our concern about the sorting of household

[19]In alternative specifications we also used the distance to the CBD as explanatory variable. Adding it in log to the specification of column 7 makes the coefficient on density marginally smaller in absolute value at -0.074. The elasticity of VKT with respect to distance to the CBD is small at 0.015.

Table 3: Robustness of baseline OLS estimations to measures of travel

Dependent: variable:	(1) stated km	(2) odometer km	(3) ind. day km	(4) dist. to work	(5) ind. day minutes	(6) speed	(7) number of trips	(8) mean trip distance
log 10-km density	-0.11[a]	-0.095[a]	-0.13[a]	-0.18[a]	-0.026[a]	-0.11[a]	0.014[a]	-0.15[a]
	(0.0054)	(0.0055)	(0.0066)	(0.0097)	(0.0036)	(0.0040)	(0.0020)	(0.0061)
R^2	0.42	0.43	0.18	0.11	0.12	0.14	0.33	0.10
Observations	93,602	71,742	83,313	86,387	85,996	82,849	83,313	83,313

Notes: All regressions include controls for household demographics, geography, local socio-economic characteristics, and 275 MSA fixed effects. Robust standard errors clustered by MSAs in parentheses. [a], [b], [c]: significant at 1%, 5%, 10%. The dependent variables and explanatory variables of interest are in log in all columns except for the number of trips in column 7. Demographic controls include a white/asian indicator, log income, log household size, a single indicator, age, age squared, gender, education, and education squared. Geographic controls include average precipitation and its standard deviation, and average temperature and its standard deviation. Local socio-economic controls include the share of residents with higher education and its square and log local income.

on the basis of unobserved tastes for driving, we prefer the larger set of control variables.[20] As it turns out, once we control for basic household demographics, including further controls does not measurably affect the coefficient of urban form.

B *Robustness to measure of driving and urban form*

In table 3 we assess the stability of the results of table 2 as we vary our dependent variable. In each column of this table we estimate a specification similar to that of column 7 of table 2, with controls for households demographics, neighbourhood socio-economic characteristics, and geography as well as a full set of MSA fixed effects. In column 1, we replace our preferred measure of VKT with a stated measure of VKT. We find a density elasticity of -11% instead of -8.2%. Measuring VKT through odometer readings by households in column 2, we estimate a density elasticity of -9.5%. Using a measure of daily VKT for individual drivers aggregated at the household level in column 3, the elasticity is again slightly larger at -13%. Using instead, distance to work in column 4 yields an even larger elasticity of -18%.

These elasticities for alternative measures of kilometers traveled are estimated on slightly different samples of households. In supplemental results we restrict attention to the about 37,000

[20] We experimented with many characteristics and included all those that are 'often' significant in the preliminary regressions we estimated. For instance, we include an indicator variables for households that are white or asian. As can be seen in table 2 below, this variable is often significant but the magnitude of its effects is small. We grouped white and asian households because differences between them were minimal. Similarly we grouped all other minorities together because the differences between them were also minimal.

households for whom we observe our preferred measure of travel and the four alternatives from columns 1-4 of table 3, we find the following elasticities: -9.2% for our preferred measures of travel, -12% for stated miles, -11% for odometer miles, -14% for daily travel, and -18% for distance to work. The differences from the corresponding elasticities reported in tables 2 and 3 are small.

We find some differences across different measures of travel, but note that these differences are small and that these measures are conceptually distinct. For instance, daily VKT is measured at the individual level whereas odometer VKT is measured for vehicles regardless of the number of household members who travel. Distance to work is more sensitive to local density. This is not surprising because commutes often take place when congestion is at its worst. Importantly, commutes represent 27% of household VKT and the density elasticity is -18% for commute distance. Hence, commutes account for $(0.27 \times 0.18)/0.092 \approx 53\%$ of the density elasticity of -9.2% that we estimate for all travel.

In column 5, our dependent variable is a measure of travel time that corresponds to kilometers traveled in column 3. For this measure of travel time, we estimate an elasticity of -2.6%, much lower than for travel distance. In column 6, we use travel speed as the dependent variable and estimate an elasticity of -11%.[21] Although residents in denser locations travel fewer kilometers, their travel time is only marginally lower because travel is slower. In column 7, we use the number of trips as the dependent variable and estimate a small positive density elasticity of 1.4%. Finally, in column 8, we estimate an elasticity of mean trip distance to 10-kilometer density of -15%. This shows that the lower VKT of residents in denser locations is exclusively explained by shorter trips not by fewer trips. If anything, residents of denser locations tend to travel more often.

In table 4, we assess the stability of the results of table 2 as we vary our explanatory variable of interest. In column 1, we use 1-kilometer density to measure urban form instead of 10-kilometer density. Relative to the -8.2% elasticity we estimate with 10-kilometer density, the estimate here is modestly lower at -6.7%. Columns 2 and 3 use the number of residents and the number of jobs within a 10 kilometer radius to estimate comparable elasticities. We estimate a smaller elasticity in column 4 when using the share of developed land within a 10 kilometer radius as measure of urban form. We return to these measures below when we consider several measures of urban form in the same regression.

[21] Although our approach is very different from that developed in Couture *et al.* (2014), they estimate a comparable elasticity of travel speed with respect to population of -13% across the largest 100 US MSAs.

Table 4: Robustness of baseline OLS estimations to measure of density

Sample restriction	(1) None	(2) None	(3) None	(4) None	(5) No NY	(6) No high density	(7) Non-MSA HH	(8) No high-VKT HH
Urban form:	1-km density	10-km pop. den.	10-km emp. den.	10-km land cover	10-km density	10-km density	10-km density	10-km density
	-0.067[a]	-0.083[a]	-0.065[a]	-0.055[a]	-0.080[a]	-0.078[a]	-0.081[a]	-0.067[a]
	(0.0036)	(0.0052)	(0.0046)	(0.0033)	(0.0045)	(0.0035)	(0.0046)	(0.0050)
R^2	0.37	0.37	0.37	0.36	0.37	0.37	0.38	0.34
Observations	99,875	99,875	99,870	99,423	94,970	74,864	26,328	90,662

Notes: All regressions include controls for household demographics, geography, local socio-economic characteristics, and 275 MSA fixed effects. Robust standard errors clustered by MSAs in parentheses. [a], [b], [c]: significant at 1%, 5%, 10%. The dependent variables and explanatory variables of interest are in log in all columns. Demographic controls include a white/asian indicator, log income, log household size, a single indicator, age, age squared, gender, education, and education squared. Geographic controls include average precipitation and its standard deviation, and average temperature and its standard deviation. Local socio-economic controls include the share of residents with higher education and its square and log local income.

In columns 5 to 8, we consider various sample restrictions to confirm that our results are not driven by a small subgroup of locations or drivers. In column 5, we estimate our preferred OLS estimation of column 7 of table 2 without the New York MSA. Although travel behavior in New York is dramatically different from the rest of the country in many ways, surprisingly, the elasticity of VKT with respect to density is unchanged when we exclude it. In results not reported here, we estimate the same specification for only the New York MSA and obtain an elasticity of -14%. In column 6, we eliminate all observations in the top density quartile and still estimate an elasticity of -7.8%. In column 7, we consider only the non-MSA residents who are excluded from most of our specifications and estimate an elasticity of -8.1%. Finally, in column 8, we eliminate the 10% of households with the highest VKT. Collectively, these households are responsible for more than 20% of aggregate VKT. As these high VKT households are disproportionately located in low density areas, we are bound to estimate a lower elasticity of VKT with respect to density. We do but, interestingly, the change is modest. We still estimate an elasticity of -6.7%.

Overall, these results suggest that our findings are broadly consistent across a variety of measures of driving and landscape, but that particular measures of driving may be more or less sensitive to urban form.

C *Sorting*

We now turn attention to the possibility that households and individuals in dense areas are different from those in less dense areas, and that it is the difference in people rather than the difference in urban form that causes the difference in driving behavior across neighborhoods.

Prior investigations of urban form and driving often consider the possibility of sorting on unobservables in the context of (Heckman) selection models estimated without an exclusion restriction. That is, they resolve the problem of sorting on unobservables entirely by functional form restrictions. Appendix table 11 reports the results of a series of such regressions. In the first five columns of this table, we consider the selection into driving of NHTS households. These estimates are very close to those in our preferred specifications. When we include demographic, neighborhood, and geographic controls, the resulting estimate of the elasticity of VKT with respect to the density of jobs and residents within 10 kilometers is -8.3%. Comparing this to our benchmark estimate in column 7 of table 2 shows that this change of functional form and accounting for non-drivers is not important.

We also report some results where we allow households to select into density. We do this in the interests of completeness since, in absence of an exclusion restriction, it is unclear how these estimations should be interpreted. We nonetheless note that unless we focus again on the top density decile, the elasticities we estimate for VKT are not dramatically different from our OLS estimates.

As in table 2, the existing literature also controls for observable individual characteristics. Intuitively, if such controls change the estimate of the coefficient of interest, then we worry that other unobserved variables might also be important. We see in table 2 that this does not occur. Oster (2015) refines this intuition and points out that observed control variables do not generally inform us about the importance of unobserved controls unless the observed controls improve the R^2 of the regression. In addition, Oster (2015) provides a parametric test for bias caused by sorting on unobservables, conditional on an assumption about the extent to which unobserved controls are 'like' observed controls. Performing this test on the regressions of columns 2 and 5 suggests that unobserved controls must satisfy an implausible condition in order to bias our estimates while columns 3 and 4 are uninformative about this issue.

Our primary strategy for addressing the possibility of sorting, however, revolves around variants of equation (15). Consistent with the discussion of section 3, we proxy for the mobility rate in

Table 5: Selection and mobility using information about local mobility measured through the tenure length of local residents

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Period	90 to 10	90 to 10	90 to 10	90 to 10	90 to 10	90 to 10	00 to 10	00 to 10
Household sample	All	All	All	Big Δ	Small Δ	<50	All	All
Initial log 10-km density	-0.080[a]	-0.075[a]	-0.046[a]	-0.084[a]	-0.069[a]	-0.080[a]	-0.076[a]	-0.076[a]
	(0.0052)	(0.0055)	(0.014)	(0.0073)	(0.0060)	(0.0072)	(0.0055)	(0.0054)
Δ log 10-km density	-0.12[a]	-0.063[b]	-0.030	-0.061	-0.056	-0.033	-0.014	-0.014
	(0.025)	(0.029)	(0.033)	(0.054)	(0.060)	(0.035)	(0.043)	(0.044)
Mobility $\times\Delta$ log density	0.0077[b]	-0.0033	-0.00059	0.0014	-0.0048	0.0023	-0.00014	-0.00014
	(0.0034)	(0.0036)	(0.0038)	(0.0067)	(0.0072)	(0.0044)	(0.0057)	(0.0057)
Mobility rate		-0.0099[a]	-0.040[a]	-0.0056	-0.016[a]	-0.010[a]	-0.010[a]	-0.010[a]
		(0.0029)	(0.013)	(0.0043)	(0.0044)	(0.0033)	(0.0027)	(0.0027)
Mobility $\times$ log density			0.0026[b]					
			(0.0011)					
Past Δ log 10-km density								-0.0017
								(0.022)
F-test 1 p-value	0.061	0.0020	0.24	0.82	0.44	0.13	0.0011	0.0055
F-test 2 p-value	0.073	0.65	0.54	0.71	0.88	0.16	0.14	0.15
R^2	0.37	0.37	0.37	0.36	0.37	0.26	0.37	0.37
Observations	99,875	99,875	99,875	46,942	48,939	39,253	99,875	99,875
Number of MSA	275	275	275	263	272	275	275	275

Notes: The dependent variable is log household VKT in all columns. Mobility is measured as - average tenure length in of residents of the same home cell. All regressions are estimated with OLS and include 275 MSA fixed effects with demographic controls (a white/asian indicator, log income, log household size, a single indicator, age, age squared, gender, education, and education squared), geographic controls (average precipitation and its standard deviation, and average temperature and its standard deviation) and local socio-economic controls (the share of residents with higher education and its square and log local income). Robust standard errors clustered by MSA in parentheses. [a], [b], [c]: significant at 1%, 5%, 10%. F-test 1 is a joint test of the equality of the coefficients on initial log 10-km density and Δ log 10-km density and of the coefficient on mobility $\times\Delta$ log density being zero. F-test 2 is a test of the equality of the coefficients on Initial log 10-km density and Δ log 10-km density.

a given neighborhood with the mean tenure of a resident in the survey respondent's home cell.[22] We multiply by minus one so that increases in our proxy correspond to increases in mobility.

Equation (15) offers two parametric tests of sorting. One of these tests involves the coefficient of the interaction of a mobility proxy with the change in urban form, and the second involves the difference between the coefficients of the level and of the change in the measure of urban form.

All of the specifications in table 5 contain these three terms. In addition to the controls from our preferred specification in column 7 of table 2 (household, neighbourhood and geographic

[22]Our information on residential tenure comes from the ACS block group data (National Historical Geographic Information System, 2010). We impute this variable to grid cells as described in section 5.

characteristics, and MSA fixed effects), column 1 also controls for the 1990 level of the density within 10 kilometers, the change in this measure between 1990 and 2010, and the interaction of the change in density with minus one times mean tenure. In order to address the possibility that driving behavior varies with tenure, column 2 also controls for the level of the mobility proxy. This specification closely approximates equation (15) and is our preferred specification. Column 3 also controls for the mobility rate interacted with the initial level of density. Column 4 restricts attention to bottom and top quartile of density growth in a 10-kilometer radius (excluding the top and bottom percentiles). Column 5 considers the complementary sample of households located in locations at the second and third quartile of density change. Column 6 restricts attention to survey respondents with household age below 50. Columns 7 and 8 consider the ten year periods from 1990 to 2000 and from 2000 to 2010.

In every case, we find that the coefficient on initial density and that on the change in density are statistically close. Except for columns 1 and 8, the coefficient on the change in density is less than a standard deviation apart from the coefficient on density. With this said, point estimates are different and, from equation (15), the difference between these two coefficients is α_1, our measure of selection. Hence, while we cannot, in a majority of cases, reject the hypothesis that $\alpha_1 = 0$, point estimates suggest it is negative. In our preferred specification in column 2, we have $\alpha_1 = -0.075 - (-0.063) = -0.012$ and $\beta = -0.063$, so that sorting accounts for about one sixth of the effect of density on driving. In equation (15), we can also implement the second test for $\alpha_1 = 0$ by rejecting the hypothesis that the coefficient of Mobility $\times$ log(10-kilometer density)$= 0$. In every specification, we see that this coefficient is small, precisely estimated and usually indistinguishable from zero. This also suggests that α_1 is small.

We note that, in the same spirit as equation (15), we can also compare the coefficient on initial density in high-mobility locations (column 4) and low-mobility locations (column 5). In addition, we can even compare the coefficient obtained when estimating our preferred specification on a sample of more mobile residents (those below 50 as in column 6) to the overall sample in column 2. In both cases, the differences are close to zero and the coefficients are precisely estimated.

More generally, and in light of the hypothesis tests developed in section 3, table 5 suggests that as density changes the driving behavior of people who leave is not statistically distinguishable from that of the people who arrive. With this said, point estimates indicate a modest amount of sorting.

In the remainder of this section, we report a number of robustness tests for this result. First, appendix table 12 replicates our preferred estimation from column 2 of table 5 under various sample restrictions, with a purely residential population based measure of density, and using alternative dependent variables. These results are consistent with our findings so far. Excluding high-density locations or high-VKT households makes no difference. Focusing more narrowly on more mobile households in locations facing greater changes in population or on less mobile households in more stable locations yields elasticities of VKT with respect to density that are of the same magnitude. Using only population instead of population and employment to measure density makes no difference. We also confirm the results of table 3. That is, the elasticity of daily VKT is slightly larger than the annual measure, the elasticity of travel time is close to zero, and this difference is still explained by the difference in travel speed.

Second, appendix table 13 presents a series of regressions that are identical to those of table 5, except that we proxy for the mobility rate with the share of renters in the cell of the survey residents. These results are qualitatively similar to those of table 5 except that the interaction terms are marginally larger and are estimated somewhat less precisely. In spite of this, these results suggest the same conclusion as does table 5. That is, as the landscape changes, the driving behavior of arrivals is like that of those who leave.

We next use age as a proxy for mobility. However, given that the relationship between age and residential mobility is unlikely to be linear, we use a vector of decadal age dummies to describe the age of drivers. Then, consistent with the intuition developed in equation (15), we interact these indicators with changes in landscape. We include these interactions in regressions that also contain the 1990 level of urban form and changes in urban form. Table 6 reports these results. Column 1 includes only the log level and change of density within 10 kilometers of a survey respondent's home cell, along with an extensive set of control variables. Column 2 includes the interaction terms. Columns 3-8 repeat column 1 on a variety of subsamples. The results of this table are striking. In every specification the coefficient of the level and change in urban form are statistically indistinguishable and coefficients do not vary across specifications. This does not allow us to reject $\alpha_1 = 0$ in equation (15), and as above, in most specifications point estimates suggest that α_1 is a small negative number.

Table 7 reports the interaction terms for column 2 of table 6. On the basis of equation (15) the coefficients of the last set of interaction terms, the interaction of decade of life with change in

Table 6: Sorting on age OLS estimations

Household sample	(1) All	(2) All	(3) <50	(4) >60	(5) Big Δ	(6) Small Δ	(7) Big Δ <50	(8) Small Δ >60
log 10-km density 1990	-0.082[a]	-0.085[a]	-0.086[a]	-0.074[a]	-0.087[a]	-0.080[a]	-0.087[a]	-0.073[a]
	(0.0053)	(0.0063)	(0.0073)	(0.0053)	(0.0068)	(0.0068)	(0.0086)	(0.0076)
Δ_{90-10} log 10-km density	-0.071[a]	-0.068[a]	-0.080[a]	-0.058[b]	-0.091[a]	-0.093	-0.092[a]	-0.13
	(0.013)	(0.019)	(0.022)	(0.025)	(0.019)	(0.057)	(0.027)	(0.096)
Controls:								
Demographics	Y	Y	Y	Y	Y	Y	Y	Y
Geography	Y	Y	Y	Y	Y	Y	Y	Y
Local socio-econ.	Y	Y	Y	Y	Y	Y	Y	Y
Decade indicators	N	Y	N	N	N	N	N	N
Decade $\times$ log density	N	Y	N	N	N	N	N	N
Decade $\times\Delta$ log density	N	Y	N	N	N	N	N	N
F-test 1 p-value	.	0.0028	.	.	.	.	.	.
F-test 2 p-value	0.31	0.32	0.71	0.51	0.80	0.81	0.81	0.52
R^2	0.37	0.37	0.26	0.26	0.36	0.37	0.25	0.27
Observations	99,875	99,875	39,253	40,421	46,942	48,939	18,710	19,980
Number of MSA	275	275	274	274	263	272	247	257

Notes: All regressions include MSA fixed effects. Robust standard errors clustered by MSA in parentheses. [a], [b], [c]: significant at 1%, 5%, 10%. The dependent variables and explanatory variables of interest are in log in all columns. Demographic controls include a white/asian indicator, log income, log household size, a single indicator, age, age squared, gender, education, and education squared. Geographic controls include average precipitation and its standard deviation, and average temperature and its standard deviation. Local socio-economic controls include the share of residents with higher education and its square and log local income. When decade effects are introduced, households in their 40s are used as reference. See table 7 for the detailed results of column 2. F-test 1 is a joint test of the equality of the coefficients on log 10-km density in 1990 and Δ log 10-km density and of the coefficients on decade indicators interacted with Δ log density all being zero. F-test 2 is a test of the equality of the coefficients on Initial log 10-km density and Δ log 10-km density.

Table 7: Detailed results for column 2 of table 6

Age	20-29	30-39	40-49 (ref.)	50-59	60-69	>70
Decade indicators	-0.057	0.087	0	0.034	-0.22[a]	-0.13
	(0.098)	(0.10)		(0.075)	(0.082)	(0.087)
Decade $\times$ log 10-km density 1990	0.0033	-0.0073	0	-0.0039	0.014[b]	0.0035
	(0.0079)	(0.0085)		(0.0060)	(0.0068)	(0.0067)
Decade $\times$ Δ_{90-10} log 10-km density	-0.010	-0.0097	0	0.027	0.022	-0.053
	(0.022)	(0.020)		(0.020)	(0.028)	(0.036)

Notes: This table reports the coefficients on decades of age, interactions between decades of age and log 10-km density in 1990, and interactions between decades of age and log density changes between 1990 and 2010.

urban form, should inform us about α_1 for that subgroup. We see that these coefficients are all small relative to the effect of density on driving and are indistinguishable from zero. We note that

the table includes a complete set of interaction as controls. We are concerned that driving behavior may vary by age or that relationship between driving and density was different in places with different initial demographics.

We have now completed five distinct tests of the role of sorting. In our OLS results, we control for observable characteristics. We find these controls have only a tiny effect on our estimates of the effect of density on driving and the more formal test of Oster (2015) indicates that unobservables are unlikely to bias our estimates. We also develop a parametric test for the role of sorting and implement it using three different proxies for the mobility rate of residents. In each case, we find little support for the idea that sorting is an important determinant of the relationship between density and driving. Finally, we report the result of Heckman type selection corrections. These estimates also suggest that the relationship between urban form and driving does not reflect sorting of individuals across places on the basis of their propensity to drive.

D *Endogeneity*

Table 8 reports the results of a series of instrumental variables estimations. These regressions are all variants of equation (16) in which we rely on permutations of three instruments. These instruments measure the share of the 10-kilometer disk surrounding a respondent's home cell that overlays an aquifer that can provide residential water. This variable is well known to predict urban form (Burchfield *et al.*, 2006). In addition, we construct variables measuring a respondent's exposure to earthquakes and landslides. These variables have a remarkably strong ability to predict surface employment and residential density, and it is not easy to see how they might influence driving through any other channel.

In column 1 we present an instrumental variables regression using our aquifers instrument but do not include other controls. In the second column, we add MSA indicators and the same long list of controls that we use in column 7 of table 2. In the subsequent columns we experiment with the different instruments and with permutations of these instruments. The coefficient of density is stable across specifications. In every case our instruments are not weak according to conventional tests, and in regressions including more than one instrument, we comfortably pass over-id tests.

Most importantly, coefficient estimates are statistically indistinguishable from those in our table of OLS estimations. This suggests that omitted variables correlated with driving and urban form

Table 8: IV regressions

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
log 10-km density	-0.13[b]	-0.100[c]	-0.12[a]	-0.076[a]	-0.080[a]	-0.075[c]	-0.069[a]	-0.075[a]
	(0.060)	(0.054)	(0.032)	(0.026)	(0.025)	(0.041)	(0.025)	(0.024)
Controls	N	Y	Y	Y	Y	Y	Y	Y
MSA effects	N	Y	Y	Y	Y	Y	Y	Y
Instruments:								
Aquifers	Y	Y	N	N	Y	Y	N	Y
Earthquakes	N	N	18	N	N	3	3	3
Landslides	N	N	N	Y	Y	N	Y	Y
Overidentification p-value	.	.	0.14	0.27	0.39	0.60	0.38	0.45
First-stage statistic	202	32.6	24.2	81.3	82.4	29.5	87.6	83.8
Observations	99,874	99,874	99,874	99,874	99,874	99,874	99,874	99,874
Number of MSA	275	275	275	275	275	275	275	275

Notes: All regressions TSLS regressions with a constant. Controls are demographic controls (a white/asian indicator, log income, log household size, a single indicator, age, age squared, gender, education, and education squared), geographic controls (average precipitation and its standard deviation, and average temperature and its standard deviation) and local socio-economic controls (the share of residents with higher education and its square and log local income). In column 3, we use all 18 values of earthquake intensity as dummy variables. In columns 6 to 8, we group them into three groups (intensity below 2, between 3 and 14, and above 15.) [a], [b], [c]: significant at 1%, 5%, 10%. Robust standard errors clustered by county in parentheses. Clustering is by county to have a sufficient number of clusters to compute robust covariance matrices more reliably than when clustering by MSA. The dependent variables and explanatory variables of interest are in log in all columns.

are not causing economically important bias in our estimates of the relationship between urban form and driving.[23]

E *Other landscape variables*

On the basis of our work so far, it appears that neither sorting nor omitted variables cause bias in OLS estimates. Given this, we now turn to an investigation of the effects of different measures and spatial scales of urban form on driving using OLS regressions.

Tables 9 reports results for a series of 'horse races' between measures of urban form. Three main conclusions emerge from this table. The first is that, although population and employment appear to play a role in explaining VKT, once we use our preferred measure of density of both jobs and residents within 10 kilometers, other measures such as the ratio of jobs to residents have no

[23]Including measures of topography in our results does not change the coefficients on landscape variables in either OLS or IV results. However, it does change our first stage. In particular, our underground geology variables do not generally pass weak instrument tests if we include topographical variables as controls.

Table 9: Driving and urban form, extended OLS estimations

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
log 10-km density	-0.082[a]	0.075		-0.082[a]		-0.085[a]	-0.090[a]	-0.088[a]
	(0.0051)	(0.063)		(0.0055)		(0.0065)	(0.0060)	(0.0061)
sq. log 10-km density		-0.0070[b]						
		(0.0029)						
log 10-km land cover			-0.0051					
			(0.0064)					
log 10-km population			-0.050[a]					
			(0.013)					
log 10-km employment			-0.024[a]					
			(0.0068)					
log 10-km job ratio				-0.0020				
				(0.0066)				
log 1-km density					-0.026[a]			
					(0.0066)			
log 1-to-5 km density					-0.044[a]			
					(0.0099)			
log 5-to-10 km density					-0.0042			
					(0.0078)			
log 10-to-25 km density					-0.0091			
					(0.0080)			
log 1-km roads						-0.00086[c]		
						(0.00048)		
log 5-km roads						-0.00068		
						(0.0011)		
log 10-km roads						-0.0019		
						(0.0025)		
log 25-km roads						0.021[b]		0.018[b]
						(0.0090)		(0.0082)
log 25-km arterials							0.021[a]	
							(0.0057)	
log 25-km highways							0.0012	
							(0.0011)	
R^2	0.37	0.37	0.37	0.37	0.37	0.37	0.37	0.37
Observations	99,875	99,875	99,423	99,870	99,861	99,875	99,875	99,875
Number of MSA	275	275	275	275	275	275	275	275

Notes: The dependent variable is log household VKT in all columns. All regressions are OLS regressions with MSA fixed effects. Controls are demographic controls (a white/asian indicator, log income, log household size, a single indicator, age, age squared, gender, education, and education squared), geographic controls (average precipitation and its standard deviation, and average temperature and its standard deviation) and local socio-economic controls (the share of residents with higher education and its square and log local income). Robust standard errors clustered by MSA in parentheses. [a], [b], [c]: significant at 1%, 5%, 10%.

measurable effect on VKT despite small standard errors. This conclusion holds more broadly than for the specifications we reported here. We experimented extensively with measures of job vs. residential locations. The effect we estimate for our preferred measure of 10-kilometer density is robust to the inclusion of many alternative measures of urban form and none of these alternative measures of urban form appears to systematically affect VKT.

Our second conclusion is that most of the effect of density takes place within 10 kilometers of a household's place of residence. The results of column 5 are even suggestive that it is density within 5 kilometers that is most important. In spite of this, when used 'alone' 10-kilometer density is often more precisely estimated than 5-kilometer density, and so we rely more heavily on 10-kilometer density in reported results. In regressions not reported here, we have also used the MSA fixed effects estimated in our preferred OLS regression from column 7 of table 2 and regressed them on variables that describe MSAs. We found no effect of MSA population, area, education, income, or geography. We also found no effect of measures of MSA employment concentration, residential concentration, and mismatch between jobs and residents. We found weak effects for some measures of segregation and the share of manufacturing employment. As we experimented with a large number of MSA characteristics, we expect the coefficients of a small proportion of them to be significant. We interpret this large majority of insignificant coefficients as an absence of MSA effects after controlling for the characteristics of households and their immediate landscape. This absence of metropolitan effect is consistent with the fact that the insertion of MSA fixed effects in table 2 does not improve the R^2 of the regressions. That most of the effect of the landscape on VKT should take place within a reasonably short range may not be surprising given that mean trip distance is slightly less than 13 kilometers in our data.

Our third conclusion concerns roads. We estimate a small positive association between roads within a 25-kilometer radius of a household's place of residence and household VKT. We acknowledge that roads may be simultaneously determined with VKT. This said, we note that the small effects of roads that we estimate are conditional on density and many other variables. Such small effects are not inconsistent with new major arterials and highways eliciting a lot of traffic as households may choose to locate closer to roads (Baum-Snow, 2007).

7. Discussion

A *Using the model: driving and welfare*

Consider first a situation with no sorting, that is $\alpha_1 = 0$. We also find no evidence of an important role played by unobserved local characteristics. That is δ is uncorrelated with X after controlling for local geography and socio-economic characteristics. These two features imply that the coefficient β estimated from a regression of log VKT on log density identifies $-\frac{\phi-\zeta\rho}{1-\rho+\phi}$ as per equation (7) of our model. The OLS estimate of β in column 7 of table 2 is -0.082. Taking alternative measures of travel distance, the first 3 columns of table 3 estimate slightly larger magnitudes for β between -0.095 and -0.13. To keep matters simple, we retain -0.1. Hence

$$\beta = -\frac{\phi - \zeta\rho}{1-\rho+\phi} = -0.1\,. \tag{18}$$

By dividing travel distance in equation (7) by mean trip distance in equation (5), we obtain the number of trips. Hence, regressing the log number of trips on log density provides an estimate of $\zeta + \beta$. Column 7 of table 3 provides such an estimate. It suggests that $\zeta = -\beta + 0.014 = 0.114$.[24]

From equations (6) and (7), regressing speed – an inverse measure of travel cost – on density, provides an estimate of $(1+\beta)\phi$. Hence, $\beta = -0.1$ and the estimated density elasticity of speed of -0.107 in column 6 of table 3 imply $\phi = 0.119$. Knowing β, ζ, and ϕ, it is now easy from equation (18) to provide a value for ρ, the concavity of utility: $\rho = 0.5$.

We note that the implied values of ζ, and ϕ are not sensitive to the exact choice of β. By contrast, the implied value of ρ is sensitive to β. A value of -0.09 for β implies $\rho > 1$ whereas a value of 0.11 implies $\rho < 0$. This is because the value of ρ in equation (18) results from dividing a small numerator by a small denominator.

We now consider situations with sorting. We assume above that $\theta = \bar{\theta}\nu$. Using this in equation (7) and the parameterization of sorting in equation (12) implies that the density elasticity of VKT is now $\frac{\alpha_1(1-\rho)}{1-\rho+\phi} + \beta = -0.1$. We can obtain an estimate of the sorting term $\frac{\alpha_1(1-\rho)}{1-\rho+\phi}$ from the difference between the coefficient on density and that on the change in density in table 5. Our preferred estimate from column 2 of table 5 indicates $\frac{\alpha_1(1-\rho)}{1-\rho+\phi} = -0.012$. Hence we have $\beta = -0.112$ when we consider sorting instead of -0.1 when we do not. From equations (4) and (7), the density

[24] We can also obtain a value $\zeta + \beta$ by taking the difference between the density elasticity of mean trip distance in column 8 of table 3 and the density elasticity of daily VKT in column 3 of the same table. These two measures are directly comparable as they rely on the same measure of VKT. They imply that the sum $\zeta + \beta$ is very much the same: $0.151 - 0.134 = 0.017$ instead of 0.014 when this quantity is estimated directly in column 7 of table 3.

elasticity of the number of trips is now $\frac{\alpha_1(1-\rho)}{1-\rho+\phi} + \beta + \zeta$. This implies $\zeta = 0.126$. From equations (6) and (7), the density elasticity of speed is now $\left(1 + \beta + \frac{\alpha_1(1-\rho)}{1-\rho+\phi}\right)\phi$ which leaves the value of $\phi = 0.119$ unchanged as sorting affect travel distance and travel time in the same manner and thus disappears when estimating speed as function of density.[25]

While there is a large literature that estimates congestion effects through traffic flows and traffic speed (Small and Verhoef, 2007), most of it is concerned with estimating effect of the (endogenous) number of vehicles on traffic speed for a particular segment of roads or groups of road segments. Attempts to measure congestion for an area depending on its population are extremely rare. Couture *et al.* (2014) estimate the effect of MSA vehicle travel time on a measure of MSA speed and find an elasticity of -0.13 for the largest 100 US MSAs. With the caveat that MSA population and the density of residents and workers are different objects, we nonetheless note that this estimated value of ϕ of 0.13 in Couture *et al.* (2014) is very similar to our implied value of 0.119 despite a very different methodology.[26]

We know of no alternative estimate of the accessibility elasticity ζ in the literature that could be directly compared with ours. Couture (2014) estimates the (constant) elasticity of substitution between restaurants using a logit model of travel demand. His framework imposes a constant trip time, which is consistent with the extremely small elasticity we estimate in table 3. His estimates of the elasticity of substitution are about nine which are consistent with accessibility benefits associated with the number of restaurants of about (9/8-1=)0.125. Although this comparison is somewhat of a stretch, this value is remarkably close to our estimate of $\zeta = 0.126$ with sorting.

Although they do not explicitly model travel behaviour, Ahlfeldt *et al.* (2016) estimates the consumption benefits from greater population density with a structural model that they implement using detailed data for the city of Berlin. Their structural model estimates an elasticity of block-level amenities with respect to a discounted measure of nearby residential density of 14%. This measure of consumption spillover is probably best interpreted as a measure of the importance of accessibility to nearby amenities and goods. To repeat, this does not directly correspond to our measure of accessibility ζ but it is nonetheless suggestive of a similar magnitude.

[25] We note that ρ is no longer separately identified from α_1 in this context unless further assumptions are being made.

[26] Geroliminis and Daganzo (2008) attempt to measure speed-flow relationships for larger spatial units. They estimate elasticities of speed with respect to the number of vehicles that are much larger in magnitude, in the order of -0.5. A strong distinction needs to be made between the number of vehicles at a particular point in time and population density. Put differently, the estimates of Geroliminis and Daganzo (2008) are for 'peak-hour' congestion which represents only a small fraction of all kilometers driven.

As another check on the consistency of our results, we can return to equilibrium utility as given by equation (8). As made clear by this equation, the elasticity of utility with respect to density is $\zeta - \phi + \phi\frac{\phi-\zeta\rho}{1-\rho+\phi} \equiv b$. Using our implicit value of the congestion elasticity ϕ of 0.119 and our implicit value of the accessibility elasticity ζ of 0.126, the elasticity of utility with respect to density remains between 0 and 0.02 as the value of the concavity parameter ρ goes from 0 to 1. This implies that equilibrium utility is fairly insensitive to density. In turn, this is consistent with little sorting being detected in our data.

To further assess the quantitative effects of this potentially weak sorting on our estimations, we need to specify our model further. We assume that the price of housing P_h increases with density X: $P_h = X^\gamma$. We can then use the utility function (8) and solve for the optimal choice of density depending on the propensity to travel θ. In equilibrium, perfect sorting occurs and $\overline{\theta} = \theta$.[27] It is easy to show that in equilibrium the elasticity of θ with respect to density X is $(1-\rho)(\gamma - b)\frac{1-\rho+\phi}{1-\rho+\phi(1+\rho-\phi)}$. Then the bias $\zeta + \beta$ caused by sorting in our OLS estimation can be directly computed from equation (10). It is equal to $\alpha_1 = \frac{(\gamma-b)(\rho-\phi)}{1-\rho+\phi(1+\rho-\phi)}$. To quantify this bias we need an estimate of γ. Combes, Duranton, and Gobillon (2012) estimate an elasticity of housing costs with respect to urban density of about 4% in France. Their approach is replicated on US data by Albouy and Ehrlich (2013) who also estimate a similar value of about 4%. For $\gamma = 0.04$, $\zeta = 0.126$, $\phi = 0.119$, and $\rho = 0.5$, we can compute a value of 0.013 for α_1, which is extremely close to our preferred estimated value of $\alpha_1 = 0.014$ used above to start these computations.[28]

The simple model proposed in section 3 was first used to derive an empirical specification and discuss identification concerns. In this section, we connect our empirical results to our model. This leads to three conclusions. Our empirical results imply estimates of the accessibility and congestion elasticities which are consistent with previous literature. In line with our empirical results about the weakness of sorting, these structural parameters also imply weak effects of density on utility. In turn, a reasonable parametrisation of the housing costs associated with greater density combined with our structural estimates of the accessibility and congestion elasticities predict nearly exactly the amount of sorting we estimate.

[27] Adding idiosyncratic preferences for locations, it would be easy to generate the type of imperfect sorting assumed in our empirical work above.

[28] The elasticity of unit housing prices is arguably larger than 0.04 but we expect residents to substitute away from housing as it becomes more expensive. But even if we ignore that housing is only about 25% of household expenditure, the bias remains moderate for $\gamma = 0.08$ which would still suggest an elasticity of driving with respect to density smaller in magnitude than -0.20.

Table 10: **Driving and population by density decile**.

(1) Decile	(2) Area share (%)	(3) Density	(4) VKT pp	(5) Area VKT share (%)	(6) Area VKT (10^9 km)
Continental US:					
1	83.26%	6.25	19,329	12.2%	620.86
10	0.21%	5,421.04	12,497	7.9%	401.44
MSA only:					
5	1.94%	816.19	15,175	9.9%	391.14
9	0.82%	2,741.09	13,779	9.0%	355.14

Notes: Top panel describes first and tenth density deciles of US population. 2010 census and 2008 NHTS populations are 3.08 million and 321,000, the total area of the continental US is 7.03 million km^2, and total NHTS VKT is 5.08 trillion kilometers. The bottom panel describes the fifth and ninth density deciles of MSA population. Census and NHTS MSA populations are 2.47 million and 257,000, the total area of the continental US MSAs is 1.66m km^2, and total NHTS VKT is 3.94 trillion kilometers.

B *Some simple general equilibrium implications*

Table 10 describes the way that people, driving and density are distributed. The top panel describes the continental US and the bottom panel restricts attention to the approximately 20% of area and 80% of population within MSA boundaries, the sample on which our regressions are primarily based. The rows of each panel describe 'density deciles of population'. For example, the first row of the top panel describes the 10% of the NHTS population living in the least dense parts of the US while the second row describes the 10% of the NHTS population living in the densest parts of the country. These are the subsets of the population that live in the white and black regions of figure 1a.

Calculating density deciles of population requires calculating threshold values of density that divide the NHTS population into tenths. The second column of the table describes the share of land area occupied at the densities intermediate between these thresholds. For example, 83% of US land area is occupied at lower densities than the threshold density for the bottom density decile of the NHTS population. Moving across columns to the right; the average density of these pixels is 6.25 people or jobs per square kilometer, the average travel per person for NHTS people living in these pixels is 19,329 vehicle-kilometers, and the population of this decile accounts for 12% of all driving in the NHTS, as measured by household odometer readings. Finally, column 7 gives aggregate driving in each decile in millions of km per year.

Table 10 permits calculations to assess the impact of policies to change density on aggregate

driving. For example, consider a policy which relocates the bottom density decile of US population into an area whose density is equal to the average density occupied by the top density decile of population. To implement such a policy we would take population dispersed across 83% of the country's land area and settle them in about 0.2% of the country's land area, concentrating the population resident in the white area of the map in figure 1a into an area the size of the barely visible black area. From column 3 of table 10, this involves an 867 fold increase in density, from 6.25 to 5,421. From our estimate in column 7 of table 2, this results in about a $(1-(5421/6.25)^{0.082}=)43\%$ decrease in driving for this decile of the population. Since this decile of population accounts for about 12% of total driving, this gives about a 5.1% decrease in aggregate driving.

It is also useful to consider a more plausible densification policy. For example, a policy that moves 1% of the MSA population and employment from the area inhabited by the fifth population decile of density to the ninth. Although only 1% of the population moves, the entire 20% of the MSA population in the source and destination regions experiences a change in density. Thus, to calculate the aggregate change in VKT we must calculate the change in aggregate driving for three groups; the 9% of the population that stays in region five, the 10% of population initially in region nine and the 1% of the population that moves from region five to region nine. The group that stays in region five is initially responsible for 0.9×391.14 billion VKT. They experience a 10% reduction in density. Using our preferred density elasticity of -0.082, this causes driving to increase by a factor of 1.0087, an increase of 3.06 billion VKT. The population of region nine initially drives 355.14 billion VKT. They experience a 10% increase in density which causes their driving to decrease by a factor of 0.9922, for a total decrease of 2.77 billion VKT. Finally, the group of movers initially accounts for 0.1×391.14 billion VKT. Their residential density increases by a factor of 3.69 from the initial level in region five, 861.14, to the final level in region nine, $1.1 \times 2{,}741.09$. This causes their driving to decrease by a factor of 0.90, for a total decrease of 3.91 billion VKT. Summing, this relocation causes a change in aggregate driving of $3.06 - 2.77 - 3.91 = -3.62$ billion VKT. Since aggregate driving is about 5.08 trillion VKT, this is a decrease of 0.07%.

Note that the decrease in driving by the initial residents of region nine is slightly smaller than the increase in region five. This reflects the fact that the population of region five initially drove more than those in region nine. Thus, this densification policy results in a small net increase in driving by the stationary population. This partly offsets the decrease in driving by the people who move. This result appears to be general and suggests that the sort of extreme relocation policy

we considered first will have a slightly larger impact on driving than we might expect from more realistic densification policies.

This hypothetical reallocation involved 1% of the total MSA population, about 2.5m people. It is difficult to assess the costs of such a policy. However, at a minimum, it requires the abandonment of about 1% of the total MSA housing stock. With a 1 trillion dollars annual expenditure on road transportation, a 0.07% decline is a 700 million dollars annual savings. In addition, according to Parry, Walls, and Harrington (2007), the external cost associated with car driving is about seven cents per kilometer. The gain coming from less congestion, fewer accidents, and less pollution from a 0.07% reduction in driving is thus about 25 million dollars. On the other hand, relocating 2.5 million people implies foregoing about a million houses. At 200,000 dollars per unit, this is a 200 billion dollars stock of housing that is abandoned. Using a 5% interest rate, the annual cost is about 10 billion dollars, which is an order of magnitude larger than the transportation gain.

C *Densification vs. gas taxes and congestion pricing*

Assessing the wisdom of using urban planning to manage traffic requires that we evaluate the effects of urban form on driving, as we do above, and also that we compare urban planning to other policies that we might use to manage driving, gasoline taxes and congestion pricing in particular.

There is a large literature on the relationship between gasoline prices and consumption. Hughes, Knittel, and Sperling (2015) survey this literature, while Coglianese, Davis, Kilian, and Stock (2015) provide recent contributions to the literature. Broadly, the short run price elasticity of gasoline appears to be about 0.3, while the long run elasticity is estimated at about 0.8. Hughes *et al.* (2015) find evidence that this short run elasticity has declined, to around 0.1, over the last decade. In the short run, we expect that gasoline consumption and driving will move together closely. Over the long run it is less clear: in response to higher gasoline prices we expect consumers to use more fuel efficient cars and to organize their lives so that they involve less driving. Since these two adjustments operate in opposite directions, it is difficult to forecast how the long run response of driving to changes in gasoline prices will diverge from that of gasoline consumption.

With this said, on the basis of the available estimates of the gasoline price elasticity, it seems reasonable to imagine that the gasoline price of elasticity of driving is at least 0.1. In this case, a fifty percent increase in gasoline causes a 5% reduction in total driving. This is about the same decrease in aggregate driving as was accomplished by the extreme relocation policy described above, but

it is accomplished with price variation that is well within the range of prices experienced in the US over the past five years. If the objective of policy is to reduce aggregate driving, it is hard to imagine that gasoline taxes do not accomplish this objective at a lower cost than urban planning.

Congestion pricing schemes are also used as a tool to manage traffic in urban areas and involve time of day, area specific road tolls. These programs are relatively recent so our understanding of their effects is based on just a few cities and roads. The three best known examples of congestion pricing are London, Singapore and Stockholm, although Murray (2012) lists a number of others, most often highway segments in California and Minnesota.

The London congestion charge began in 2003 and required the payment of about 8 USD to enter central London, an area of about 22 square kilometers, during working hours. This policy led to a dramatic reduction in travel, about 34% for cars and 12% for all vehicle types, an increase in peak hour travel speeds from 14.3 to 16.7 kilometers per hour and a dramatic decrease in delay relative to free-flow travel speeds (Leape, 2006). The Singapore congestion charge began in 1975 and was about 2.5 USD per day. It converted from a paper-based to electronic enforcement system in 1999 with somewhat lower charges. At its beginning, this program was responsible for about a 45% reduction in peak area vehicle travel in the affected area and an increase in travel speeds from 19 to 36 kilometers per hour (Santos, 2005). Borjesson, Eliasson, Hugosson, and Brundell-Freij (2012) describe Stockholm's experience with congestion pricing. Begun in 2006, with a time of day charging that peaks at about 3 USD at peak hours and tapers to zero during off peak times, this program caused about 30% reduction in vehicles in the affected areas and a dramatic decrease in travel times.

Relative to the marginal and uncertain reductions in driving that appear to result from densification policies, it is hard to imagine that congestion pricing is not a more cost effective way to reduce urban congestion than is urban planning.[29]

8. Conclusion

Urban density appears to have a small causal effect on driving. Our estimates of the density elasticity are generally between -7% and -10% and is about -8% in our preferred specification. The

[29] While congestion pricing appears to have dramatic effects on the volume and speed of travel, there is some debate over whether such programs are welfare improving. The central issue is that the demand for travel appears to be very elastic, so that deadweight loss from congestion is small, while the costs of implementing congestion pricing plans can be large. See Prudhomme and Bocarejo (2005) for a nice illustration of these issues, which are also discussed in Couture *et al.* (2014).

literature on this issue is large. Our estimates improve on those in the literature in four ways. First, we use better data. We are the first to use a data set as large as the NHTS to estimate the effect of urban form on driving using individual level landscape data. Second, we develop a parametric test for sorting. Although the literature has long been aware that cross-sectional differences in driving behavior across locations may reflect sorting, it has yet to develop a persuasive quasi-experimental design. Given this, our ability to test for sorting using cross-sectional travel survey data and panel landscape data is an advance. Third, we implement a quasi-experimental design for dealing with the possibility of endogenous determination of density. Specifically, we use subterranean geology to instrument for surface density. Fourth, our econometric model is motivated by a theoretical foundation. Ultimately, this means that we are able to recover the structural parameters governing the way that travel behavior responds to density. To the extent that we are able to check, these structural parameters appear to be consistent with related estimates in the literature. This structural model also highlights that, even if densification is welfare improving, it does not remove the need for congestion pricing. Whether neighborhoods are high density or low, without congestion pricing, drivers do not account for their contribution to congestion without an explicit pricing program.

Our estimates of the relationship of driving to urban form allow us to assess the cost effectiveness of densification as a policy response to excessive driving. These estimates suggest that urban form is not cost effective compared to explicit pricing programs. In particular, even concentrating the population residing in 83% of the area the continental US into an area of about 1500 square kilometers would result in only about a 5% decrease in aggregate driving, and this policy appears to describe the upper envelope of what densification policies can accomplish. On the other hand, existing estimates of the gasoline price elasticity of driving suggest that a similar decrease in driving would be accomplished with a gas tax that is no larger than gasoline price fluctuations observed over the past five to ten years. Congestion pricing programs appear to have even larger effects. In sum, while dense urban development may well be desirable because it provides a residential environment where people want to live and that make them work more productively (e.g., Rosenthal and Strange, 2008), it is probably more costly to manipulate driving behavior through densification policies than through congestion pricing or gasoline taxes.

References

Ahlfeldt, Gabriel M., Stephen J. Redding, Daniel M. Sturm, and Nikolaus Wolf. 2016. The economics of density: Evidence from the Berlin Wall. *Econometrica* 84(forthcoming).

Albouy, David and Gabriel Ehrlich. 2013. The distribution of urban land values: Evidence from market transactions. Processed, University of Illinois.

Baum-Snow, Nathaniel. 2007. Did highways cause suburbanization? *Quarterly Journal of Economics* 122(2): 775–805.

Bayer, Patrick, Stephen L. Ross, and Giorgio Topa. 2008. Place of work and place of residence: Informal hiring networks and labor market outcomes. *Journal of Political Economy* 116(6): 1150–1196.

Bento, Antonio, Maureen L. Cropper, Ahmed Mushfiq Mobarak, and Katja Vinha. 2005. The effects of urban spatial structure on travel demand in the United States. *Review of Economics and Statistics* 87(3): 466–478.

Blaudin de Thé, Camille and Miren Lafourcade. 2015. The carbon footprint of suburbanization: Evidence from French household data. Processed, Paris School of Economics.

Board of Governors of the Federal Reserve System. 2015. Z.1 financial accounts of the united states. flow of funds, balance sheets, and integrated macroeconomic accounts. second quarter 2015. www.federalreserve.gov/releases/z1/current/z1.pdf. Accessed: 2015-11-03.

Boarnet, Marlon G. 2011. A broader context for land use and travel behavior, and a research agenda. *Journal of the American Planning Association* 77(3): 197–213.

Borjesson, Maria, Jonas Eliasson, Muriel B. Hugosson, and Karin Brundell-Freij. 2012. The Stockholm congestion charges-5 years on. Effects, acceptability and lessons learnt. *Transport Policy* 20(1): 1–12.

Brownstone, David and Thomas F. Golob. 2009. The impact of residential density on vehicle usage and energy consumption. *Journal of Urban Economics* 65(1): 91–98.

Burchfield, Marcy, Henry G. Overman, Diego Puga, and Matthew A. Turner. 2006. Causes of sprawl: A portrait from space. *Quarterly Journal of Economics* 121(2): 587–634.

Cao, Xinyu, Patricia L. Mokhtarian, and Susan L. Handy. 2009. Examining the impacts of residential self-selection on travel behaviour: A focus on empirical findings. *Transport Reviews* 29(3): 359–395.

Cervero, Robert and Kara Kockelman. 1997. Travel demand and the 3Ds: Density, diversity, and design. *Transportation Research Part D: Transport and Environment* 2(3): 199–219.

Chetty, Raj, Nathaniel Hendren, and Larry Katz. 2015. The effects of exposure to better neighborhoods on children: New evidence from the moving to opportunity experiment. *American Economic Review* (forthcoming).

Coglianese, John, Lucas W. Davis, Lutz Kilian, and James H. Stock. 2015. Anticipation, tax avoidance, and the price elasticity of gasoline demand. NBER Working Paper No. 20980.

Combes, Pierre-Philippe, Gilles Duranton, and Laurent Gobillon. 2008. Spatial wage disparities: Sorting matters! *Journal of Urban Economics* 63(2): 723–742.

Combes, Pierre-Philippe, Gilles Duranton, and Laurent Gobillon. 2012. The costs of agglomeration: Land prices in French cities. Processed, Wharton School, University of Pennsylvania.

Couture, Victor. 2014. Valuing the consumption benefits of urban density. Processed, University of California Berkeley.

Couture, Victor, Gilles Duranton, and Matthew A. Turner. 2014. Speed. Processed, Wharton School, University of Pennsylvania.

Currie, Janet and Reed Walker. 2011. Traffic congestion and infant health: Evidence from E-Z Pass. *American Economic Journal: Applied Economics* 3(1): 65–90.

Dahl, Gordon B. 2002. Mobility and the return to education: Testing a Roy model with multiple markets. *Econometrica* 70(6): 2367–2420.

Duranton, Gilles and Matthew A. Turner. 2012. Urban growth and transportation. *Review of Economic Studies* 79(4): 1407–1440.

Eid, Jean, Henry G. Overman, Diego Puga, and Matthew A. Turner. 2008. Fat city: Questioning the relationship between urban sprawl and obesity. *Journal of Urban Economics* 63: 385–404.

Environmental Systems Research Institute. 1998*a*. ESRI Data and Maps 1998: Discs 3&4 - Block Group Boundaries. This map data was included with an ArcGIS 7 License.

Environmental Systems Research Institute. 1998*b*. ESRI Data and Maps 1998: Disc 2 - Zip Code Tabulation Areas. This map data was included with an ArcGIS 7 License.

Environmental Systems Research Institute. 2004. ESRI Data and Maps 2004: United States Disc - Block Group Boundaries. This map data was included with an ArcGIS 9 Licence.

Evans, William N., Wallace E. Oates, and Robert M. Schwab. 1992. Measuring peer group effects: A study of teenage behavior. *Journal of Political Economy* 100(5): 966–991.

Ewing, Reid and Robert Cervero. 2001. Travel and the built environment: A synthesis. *Transportation Research Record: Journal of the Transportation Research Board* 0(1780): 87–114.

Ewing, Reid and Robert Cervero. 2010. Travel and the built environment: A meta-analysis. *Journal of the American Planning Association* 76(3): 265–294.

Federal Highway Administration. 2005. National highway planning network. www.fhwa.dot.gov. Accessed: 2014-05-24.

Geroliminis, Nikolas and Carlos F. Daganzo. 2008. Existence of urban-scale macroscopic fundamental diagrams: Some experimental findings. *Transportation Research Part B: Methodological* 42(9): 759–770.

Giuliano, Genevieve and Kenneth A. Small. 1993. Is the journey to work explained by urban structure? *Urban studies* 30(9): 1485–1500.

Glaeser, Edward L. and Matthew E. Kahn. 2004. Sprawl and urban growth. In Vernon Henderson and Jacques-François Thisse (eds.) *Handbook of Regional and Urban Economics*, volume 4. Amsterdam: North-Holland, 2481–2527.

Glaeser, Edward L. and Matthew E. Kahn. 2008. The greenness of cities: Carbon dioxide emissions and urban development. *Journal of Urban Economics* 67(3): 404–418.

Gordon, Peter, Ajay Kumar, and Harry W. Richardson. 1989. The influence of metropolitan spatial structure on commuting time. *Journal of Urban Economics* 26(2): 138–151.

Handy, Susan. 2005. Smart growth and the transportation-land use connection: What does the research tell us? *International Regional Science Review* 28(2): 146–167.

Head, Keith and Thierry Mayer. 2014. Gravity equations: Workhorse, toolkit, and cookbook. In Gita Gopinath, Elhanan Helpman, and Kenneth Rogoff (eds.) *Handbook of International Economics*, volume 4. Amsterdam: North-Holland, 131–195.

Hughes, Jonathan E., Christopher R. Knittel, and Daniel Sperling. 2015. Evidence of a shift in the short-run price elasticity of gasoline demand. *The Energy Journal* 29(1): 113–134.

Intergovernmental Panel on Climate Change. 2007. *Working Group III contribution to fouth assessment report of the IPCC: Mitigation of Climate Change*. Cambridge, UK and New York, NY, USA: Cambridge University Press.

Intergovernmental Panel on Climate Change. 2014. *Working Group III contribution to fifth assessment report of the IPCC: Mitigation of Climate Change*. NY, New York, USA: Cambridge University Press.

Ioannides, Yannis M. and Giorgio Topa. 2010. Neighborhood effects: Accomplishments and looking beyond them. *Journal of Regional Science* 50(1): 343–362.

Kling, Jeffrey R., Jens Ludwig, and Lawrence F. Katz. 2005. Neighborhood effects on crime for female and male youth: Evidence from a randomized housing voucher experiment. *Quarterly Journal of Economics* 120(1): 87–130.

Kuzmyak, J. Richard. 2012. Land use and traffic congestion. Professional Paper 618, Arizona DOT and US DOT.

Leape, Jonathan. 2006. The London congestion charge. *Journal of Economic Perspectives* 20: 157–176.

Missouri Census Data Center. 1990. Block-group and county populations in 1990. mcdc.missouri.edu/websas/geocorr90.shtml. Accessed: 2014-01-14.

Missouri Census Data Center. 2000. Block-group and county populations in 2000. mcdc.missouri.edu/websas/geocorr2k.html. Accessed: 2014-01-14.

Missouri Census Data Center. 2010. Block-group and county populations in 2010. mcdc.missouri.edu/websas/geocorr12.shtml. Accessed: 2014-01-14.

Murray, Peter. 2012. Congestion pricing for roads: An overview of current best practice, and the economic and transport benefits for government. *Public Infrastrucutre Investment* 1(8): 34–42.

National Historical Geographic Information System. 2010. Block-group level demographic variables for 2006-2010 and 2007-2011. www.nhgis.org. Accessed: 2014-01-30.

Oster, Emily. 2015. Unobservable selection and coefficient stability. Processed, Brown University.

Parry, Ian W.H., Margaret Walls, and Winston Harrington. 2007. Automobile externalities and policies. *Journal of Economic Literature* 45(2): 1335–1353.

PRISM Climate Group at Oregon State University. 2012*a*. United States Average January Precipitation, 1981-2010. prism.oregonstate.edu. Accessed: 2014-04-23.

PRISM Climate Group at Oregon State University. 2012*b*. United States Average January Mean Temperature, 1981-2010. prism.oregonstate.edu. Accessed: 2014-04-23.

Prudhomme, Rémy and Juan Pablo Bocarejo. 2005. The London congestion charge: A tentative economic appraisal. *Transport Policy* 12(1): 279–287.

Redding, Stephen J. and Matthew A. Turner. 2015. Transportation costs and the spatial organization of economic activity. In Gilles Duranton, Vernon Henderson, and William C. Strange (eds.) *Handbook of Regional and Urban Economics*, volume 5B. Amsterdam: Elsevier, 1339–1398.

Rosenthal, Stuart S. and William C. Strange. 2008. The attenuation of human capital spillovers. *Journal of Urban Economics* 64(2): 373–389.

Santos, Georgina. 2005. Urban congestion charging: A comparison between London and Singapore. *Transport Reviews* 25(5): 511–534.

Small, Kenneth A. and Erik T. Verhoef. 2007. *The Economics of Urban Transportation*. New York (NY): Routledge.

United States Geological Survey. 2000. NLCD 1992 Land Cover. www.mrlc.gov. Accessed: 2013-12-19.

United States Geological Survey. 2001. Landslide Incidence and Susceptibility in the Conterminous United States. www.nationalatlas.gov/atlasftp.html. Accessed: 2014-05-20.

United States Geological Survey. 2003. Principal Aquifers of the 48 Conterminous United States. www.nationalatlas.gov/atlasftp.html. Accessed: 2014-05-20.

United States Geological Survey. 2005. Significant United States Earthquakes, 1568-2004. www.nationalatlas.gov/atlasftp.html. Accessed: 2014-05-20.

United States Geological Survey. 2011*a*. NLCD 2001 land cover version 2.0. www.mrlc.gov. Accessed: 2014-01-24.

United States Geological Survey. 2011*b*. NLCD 2006 Land Cover Version 2.0. www.mrlc.gov. Accessed: 2014-01-24.

U.S. Department of Commerce, U.S. Census Bureau, Geography Division. 2010. Block-group level US Census 2010 geography Shapefiles,.

U.S. Department of Commerce, U.S. Census Bureau, Geography Division. 2010. Tiger/line shapefile, 2010, 2010 nation, u.s., 2010 census 5-digit zip code tabulation area (zcta5) national. www.census.gov/geo/www/tiger. Accessed: 2014-09-06.

U.S. Department of Labor, Bureau of Labor Statistics. 2014. Consumer expendituresurvey. http://www.bls.gov/cex/csxcombined.htm. Accessed: 2015-11-03.

U.S. Department of Transportation, Federal Highway Administration. 2009. 2009 National Household Travel Survey. http://nhts.ornl.gov. Accessed: 2014-06-13.

Appendix A. Robustness checks

Table 11: Heckman selection models (one-step MLE)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Selection into:	Driving					High density		
Panel (a). Selection equation								
log 10-km density	-0.089[a]	-0.13[a]	-0.13[a]	-0.092[a]	-0.69[a]	-	-	-
	(0.0043)	(0.0058)	(0.0058)	(0.0064)	(0.033)	-	-	-
Top density decile				-0.29[a]				
				(0.020)				
Panel (b). VKT equation								
log 10-km density	-0.053[a]	-0.091[a]	-0.083[a]	-0.079[a]	-0.21[a]	-0.12[a]	-0.13[a]	-0.21[a]
	(0.0026)	(0.0023)	(0.0029)	(0.0029)	(0.035)	(0.0059)	(0.0064)	(0.022)
Top density decile				-0.053[a]				
				(0.012)				
Controls:								
Demographics	N	Y	Y	Y	Y	Y	Y	Y
Geography	N	Y	Y	Y	Y	N	Y	Y
Local socio-econ.	N	Y	Y	Y	Y	N	Y	Y
MSA fixed effects	N	N	Y	Y	Y	N	N	N
Observations	108,857	108,857	108,857	108,857	10,925	99,875	99,875	99,875

Notes: Results reported for the main regression using log household VKT as dependent variable. The selection equation regards selection into driving in columns 1-5, selection into above median density in columns 6-7, and selection into the highest density decile in column 8. The sample is all households driving and non-driving MSA households in columns 1-4, all households in the top density decile in column 5, and driving households in columns 6-8. Standard errors in parentheses. [a], [b], [c]: significant at 1%, 5%, 10%. Demographic controls include a white/asian indicator, log income, log household size, a single indicator, age, age squared, gender, education, and education squared. Geographic controls include average precipitation and its standard deviation and average temperature and its standard deviation. Local socio-economic controls include the share of residents with higher education and its square and log local income.

Table 12: Robustness of selection estimations using local tenure length to measure mobility

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	No high den. location	No high VKT hh	Big Δ & <50	Small Δ & >60	population density	ind. day km as DV	ind. day mn as DV	speed as DV
log 10-km density 1990	-0.070[a]	-0.060[a]	-0.084[a]	-0.054[a]	-0.074[a]	-0.12[a]	-0.018[a]	-0.10[a]
	(0.0039)	(0.0055)	(0.0098)	(0.0088)	(0.0056)	(0.0079)	(0.0043)	(0.0049)
Δ_{90-10} log 10-km density	-0.029	-0.028	0.023	-0.030	-0.067[b]	-0.15[a]	0.00072	-0.15[a]
	(0.030)	(0.033)	(0.063)	(0.45)	(0.026)	(0.041)	(0.031)	(0.023)
Mobility $\times\Delta$ log density	-0.00041	-0.00049	0.012[c]	-0.0048	-0.0053	-0.019[a]	-0.0050	-0.013[a]
	(0.0038)	(0.0037)	(0.0068)	(0.040)	(0.0035)	(0.0047)	(0.0034)	(0.0030)
Mobility	-0.010[a]	-0.0094[a]	-0.0094	-0.024[c]	-0.011[a]	-0.0087[b]	-0.0068[b]	-0.00096
	(0.0038)	(0.0037)	(0.0068)	(0.040)	(0.0035)	(0.0047)	(0.0034)	(0.0030)
F-test 1 p-value	0.00022	0.017	0.21	0.75	0.000092	0	0.000017	0
F-test 2 p-value	0.17	0.31	0.082	0.96	0.76	0.39	0.52	0.042
R^2	0.37	0.34	0.25	0.27	0.37	0.18	0.12	0.14
Observations	74,864	90,662	18,711	19,979	99,875	83,313	85,996	82,849
Number of MSA	275	275	248	252	275	275	275	275

Notes: All regressions include 275 MSA fixed effects. Robust standard errors clustered by MSA in parentheses. [a], [b], [c]: significant at 1%, 5%, 10%. The dependent variable and explanatory variables of interest are in log in all columns. Demographic controls include a white/asian indicator, log income, log household size, a single indicator, age, age squared, gender, education, and education squared. Geographic controls include average precipitation and its standard deviation, and average temperature and its standard deviation. Local socio-economic controls include the share of residents with higher education and its square and log local income. F-test 1 is a joint test of the equality of the coefficients on initial log 10-km density and Δ log 10-km density and of the coefficient on Mobility $\times\Delta$ log density being zero. F-test 2 is a test of the equality of the coefficients on initial log 10-km density and Δ log 10-km density.

Table 13: Selection and mobility using information about the renter/homeowner status of the households

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Period	90 to 10	90 to 10	90 to 10	90 to 10	90 to 10	90 to 10	00 to 10	00 to 10
Household sample	All	All	All	Big Δ	Small Δ	<50	All	All
Initial log 10-km density	-0.081[a]	-0.080[a]	-0.082[a]	-0.085[a]	-0.078[a]	-0.084[a]	-0.080[a]	-0.080[a]
	(0.0053)	(0.0052)	(0.0051)	(0.0067)	(0.0067)	(0.0072)	(0.0052)	(0.0050)
Δ log 10-km density	-0.057[a]	-0.071[a]	-0.074[a]	-0.089[a]	-0.081	-0.083[a]	-0.053[a]	-0.043[b]
	(0.013)	(0.012)	(0.012)	(0.018)	(0.061)	(0.022)	(0.017)	(0.019)
Renter $\times\Delta$ log density	-0.20[a]	0.013	0.040	-0.010	-0.087	0.040	0.033	0.032
	(0.030)	(0.036)	(0.038)	(0.046)	(0.16)	(0.035)	(0.050)	(0.050)
Renter		-0.15[a]	-0.33[a]	-0.14[a]	-0.12[a]	-0.14[a]	-0.15[a]	-0.15[a]
		(0.013)	(0.094)	(0.020)	(0.040)	(0.018)	(0.011)	(0.012)
Renter $\times$ log density			0.015[b]					
			(0.0073)					
Past Δ log 10-km density			-					-0.032
								(0.023)
F-test 1 p-value	0	0.60	0.37	0.95	0.84	0.49	0.10	0.13
F-test 2 p-value	0.026	0.38	0.44	0.80	0.97	0.93	0.078	0.020
R^2	0.37	0.37	0.37	0.36	0.37	0.26	0.37	0.37
Observations	99,875	99,875	99,875	46,942	48,939	39,253	99,875	99,875
Number of MSA	275	275	275	263	267	274	275	275

Notes: The dependent variables is log household VKT in all columns. All regressions are estimated with OLS and include 275 MSA fixed effects with demographic controls (a white/asian indicator, log income, log household size, a single indicator, age, age squared, gender, education, and education squared), geographic controls (average precipitation and its standard deviation, and average temperature and its standard deviation) and local socio-economic controls (the share of residents with higher education and its square and log local income). Robust standard errors clustered by MSA in parentheses. [a], [b], [c]: significant at 1%, 5%, 10%. F-test 1 is a joint test of the equality of the coefficients on Initial log 10-km density and Δ log 10-km density and of the coefficient on Mobility $\times\Delta$ log density being zero. F-test 2 is a test of the equality of the coefficients on initial log 10-km density and Δ log 10-km density.

Table 14: Robustness checks for sorting on demographics OLS estimations

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Period	00 to 10	00 to 10	00 to 10	90 to 10	90 to 10	90 to 10	90 to 10	90 to 10
Household sample	All	<50	>60	All	All	Indiv.	All	All
Dependent var.:	an. km	an. km	an. km	stated km	odometer	ind. day km	an. km	an. km
Density:	10 km	10 km	10 km	10 km	10 km	10 km	1 km	NLCD 10 km
Initial log density	-0.082[a]	-0.087[a]	-0.075[a]	-0.12[a]	-0.094[a]	-0.14[a]	-0.053[a]	-0.040[a]
	(0.0053)	(0.0072)	(0.0054)	(0.0075)	(0.0060)	(0.0083)	(0.0037)	(0.0046)
Δ log density	-0.050[a]	-0.049[c]	-0.036	-0.066[a]	-0.077[a]	-0.066[b]	-0.047[a]	-0.039[a]
	(0.017)	(0.028)	(0.035)	(0.024)	(0.024)	(0.029)	(0.0060)	(0.0069)
Past Δ density		-0.037	-0.015					
		(0.028)	(0.039)					
Controls:								
Demographics	Y	Y	Y	Y	Y	Y	Y	Y
Geography	Y	Y	Y	Y	Y	Y	Y	Y
Local socio-econ.	Y	Y	Y	Y	Y	Y	Y	Y
Decade indicators	N	N	N	Y	Y	Y	Y	Y
Decade $\times$ log density	N	N	N	Y	Y	Y	Y	Y
Decade $\times\Delta$ log density	N	N	N	Y	Y	Y	Y	Y
F-test 1 p-value	.	.	.	0.0001	0.027	0	0.0034	0.0033
F-test 2 p-value	0.039	0.10	0.27	0.015	0.40	0.0062	0.22	0.91
R^2	0.37	0.26	0.26	0.42	0.43	0.09	0.37	0.37
Observations	99,875	39,253	40,421	93,602	71,742	121,808	99,874	99,423
Number of MSA	275	274	274	275	275	275	275	275

Notes: All regressions include 275 MSA fixed effects. Robust standard errors clustered by MSA in parentheses. [a], [b], [c]: significant at 1%, 5%, 10%. The dependent variables and explanatory variables of interest are in log in all columns. Demographic controls include a white/asian indicator, log income, log household size, a single indicator, age, age squared, gender, education, and education squared. Geographic controls include average precipitation and its standard deviation, and average temperature and its standard deviation. Local socio-economic controls include the share of residents with higher education and its square and log local income. When decade effects are introduced, households in their 40s are used as reference. F-test 1 is a joint test of the equality of the coefficients on initial log 10-km density and Δ log 10-km density and of the coefficients on decade indicators interacted with Δ log density all being zero. F-test 2 is a test of the equality of the coefficients on initial log 10-km density and Δ log 10-km density.