

Speed

Victor Couture^{*†}
Gilles Duranton^{*‡}
Matthew A. Turner^{*§}

University of Toronto

This draft: 14 March 2012

ABSTRACT: We investigate the determinants of driving speed in large US cities. We first estimate city level supply functions for speed in an econometric framework where both the supply and demand for driving are explicit. These estimations allow us to calculate a city level index of driving speed and to rank cities by driving speed. Our investigation of the determinants of speed provide the foundations for welfare analysis. Our results suggest that large gains in speed may be possible if slow cities can emulate fast cities and that the deadweight losses from congestion are large.

Key words: roads, vehicle-kilometers traveled, public transport, congestion, travel time.

JEL classification: L91, R41

*Financial support from the Canadian Social Science and Humanities Research Council is gratefully acknowledged.

†Department of Economics, University of Toronto, 150 Saint George Street, Toronto, Ontario M5S 3G7, Canada (e-mail: victor.couture@utoronto.ca).

‡Department of Economics, University of Toronto, 150 Saint George Street, Toronto, Ontario M5S 3G7, Canada (e-mail: gilles.duranton@utoronto.ca; website: <http://individual.utoronto.ca/gilles/default.html>). Also affiliated with the Centre for Economic Policy Research, and the Centre for Economic Performance at the London School of Economics.

§Department of Economics, University of Toronto, 150 Saint George Street, Toronto, Ontario M5S 3G7, Canada (e-mail: mturner@chass.utoronto.ca; website: <http://www.economics.utoronto.ca/mturner/index.htm>).

1. Introduction

We investigate the determinants of driving speed in a sample of large us cities. We proceed in three stages. We first estimate the supply of speed faced by drivers in each of our cities. Our econometric formulation explicitly accounts for the fact that longer trips are faster than shorter trips, and hence that optimizing drivers must behave as ordinary monopsonists facing a downward sloping supply curve. Using the resulting city level supply functions we calculate a city level index of travel speed. This index is analogous to a conventional price index and gives the time cost of a standardized bundle of trips. Our speed index resolves two important problems: it provides a well defined measure of speed for each city and it measures the time cost of travel in a city in way that does not depend on simultaneously determined driver behavior affecting speed.

Our speed index allows us to rank cities by speed of travel. This is of intrinsic interest and provides an alternative to the Texas Transportation Institute's (TTI) widely cited congestion index (Schrank and Lomax, 2009, Schrank, Lomax, and Turner, 2010). However, unlike the TTI index, our index is grounded in economic theory and hence can be more easily interpreted.

We next investigate the relationship between our speed index and inputs into the transportation process, aggregate travel time and roads in particular. Since the supply of roads or aggregate travel time in a city may reflect unobserved determinants of speed, we are careful in this investigation to account for the probable endogeneity of inputs in this estimation.

Finally, we investigate the value of speed using two distinct methodologies. In the first, we consider hypothetical policy interventions which increase speed. This analysis, which is in the spirit of standard analyses of factor productivity, suggests that the value of an increase in speed is measured in the tens or even hundreds of billions of dollars. These results are provocative and provide some information about the value of hypothetical interventions. However, such hypothetical changes are unlikely to occur in equilibrium. In our second investigation of the value of speed, we use our estimates to derive a demand system for driving in our cities. An analysis of this system indicates that the deadweight loss from congestion is on the order of 120 billion dollars per year, larger even than the gains suggested by hypothetical interventions to increase speed. To the extent that much of the benefit from the hypothetical supply improvements are dissipated by increases in equilibrium travel, our results suggest that larger gains are to be had by managing demand for travel, rather than improving supply. Our analysis also suggests that equilibrium

driving may dramatically exceed optimal driving, possibly by a factor of about two.¹

Our results are of interest for several reasons. Two of the principal questions which the literature on transportation economics addresses are ‘what is travel demand?’ and ‘what does the speed-flow curve look like?’² Since the speed-flow curve relates travel time to distance, i.e., time price of travel to quantity, this second question can be restated differently as ‘what is the supply curve for travel?’ Extant empirical investigations ask these questions separately. Therefore, while the study of transportation has a distinguished history in economics, it has almost completely ignored the fact that observed travel behavior results from an equilibrium that depends upon both supply and demand conditions. To our knowledge, our paper is the first to implement an econometric framework which recognizes this problem.³ Our identification strategy exploits the fact that distance differs across different types of trips.

Related to this, we note that the conventional approach to estimating a speed-flow curve involves measuring the speed and number of cars at a particular point on a road and plotting the resulting data. This plot arguably reflects both supply, i.e., that traffic slows down when there are more cars, and demand, i.e., there are fewer cars when traffic is slower. Even if we ignore this issue, the estimation of speed-flow curves is subject to a second important problem. As observed in their classic monograph (Beckman, McGuire, and Winsten, 1956), the interpretation of this curve rests on the assumption that all else is equal. This assumption is problematic. Increasing traffic somewhere almost surely entails changes elsewhere. In particular, measuring what happens at a particular point on a road is likely to reflect driving conditions further along on that road. It is difficult to think of a specific point on a network independently. To avoid this problem, we investigate the ability of an entire city, rather than a particular road segment, to supply speed.⁴ This is also novel.

While the productivity of manufacturing and service sectors is extensively studied, in spite of its size, the transportation sector has received much less attention. Our estimated city level supply functions allow us to investigate the cross-sectional determinants of efficiency in transportation.

¹We focus solely on travel using privately-owned vehicles. We do not provide a detailed cost analysis across various transportation modes in the tradition of Meyer, Kain, and Wohl (1965).

²Much of the work on travel demand follows on the pioneering work of McFadden (1973, 1974). See Train (2009) for a thorough exposition of the state of the art. The first estimation of a speed flow-curve dates back to Greenshields (1935). Important subsequent work include Keeler and Small (1977) and Dewees (1979). See Small and Verhoef (2007) for a review of more recent work.

³In their authoritative book, Small and Verhoef (2007) exposit the supply and demand for travel separately, indeed, in different chapters. The fact that some variables may affect both supply and demand is recognised but only discussed in the context of car purchases.

⁴Bombardini and Trebbi (2011) apply recent results from graph theory to investigate how distortions to urban transportation networks affects their efficiency.

Consistent with the large extant literature investigating productivity in firms (e.g., Syverson, 2011), we find that some cities are dramatically more efficient than others. This suggests that there may be large gains if slow cities can emulate fast cities. Our results also suggest that the supply of transportation exhibits slight decreasing returns to scale: doubling total driving time and the stock of roads in a city leads to about a 5% decrease in speed.

Finally, we note that in 2008 the average driver in our sample spent about 72 minutes driving; the median household devotes 18% of its budget to road travel; in a typical year, the US spends nearly 200 billion dollars on road construction and maintenance; and the value of capital stock associated with road transportation in the US tops 5 trillion dollars (US BTS, 2007 and 2010). By providing better estimates of the fundamental relationships that determine the value and costs of driving, we provide a sounder foundation for the enormous resource allocations affected by transportation policy. More concretely, by estimating the cost of congestion we facilitate the calculation of the net benefits of costly policy responses to traffic congestion.

2. Data

Our data describe aggregate travel behavior in a set of large US cities and the individual driving trips taken by a sample of each city's residents. Our cities are US (Consolidated) Metropolitan Statistical Areas (MSA) drawn to 1999 boundaries. MSAs are census reporting units and are aggregations of counties containing a major urban center and its surrounding region. Our analysis relies heavily on household survey data. To ensure that we observe a sufficiently large number of households in each MSA, in most of our work we consider a sample of the 100 MSAs that are on average most populous over our study period.⁵ For some applications where our household surveys are smaller, we restrict attention to a sample of the 50 largest MSAs.

Data on individual travel behavior comes from the 1995-1996 National Personal Transportation Survey and the 2001-2002 and 2008-2009 National Household Transportation Surveys. In a slight abuse of language, we refer to these surveys as the 1995, 2001 and 2008 NHTS. Each of the NHTS surveys reports household and individual demographics for a nationally representative sample of households. More importantly, the 'travel day file' of each NHTS survey codifies a travel diary kept by every member of each sampled household. For each adult member of participating households we observe the distance, duration, mode, purpose, and start time for each trip taken on a randomly

⁵To be consistent with other data, we rank MSAs by their average population in 1995, 1996, 2001, 2002, 2008, and 2009.

assigned travel day. We eliminate trips entered by non-drivers in order to focus our investigation on the movement of vehicles rather than the movement of people. In our sample of 100 MSAs, the NHTS describes 415,402 trips, 101,771 drivers and 70,826 households in 2008; 168,765 trips, 40,353 drivers and 27,589 households in 2001; and 152,590 trips, 33,879 drivers and 22,604 households in 1995.

We aggregate to describe travel behavior at the MSA level. To estimate MSA vehicle kilometers traveled (VKT) and vehicle time traveled (VTT), we sum the time and distance of each trip over all of an individual's trips. Averaging across drivers, we compute the average distance and time driven by an individual in each MSA. We multiply this individual average by MSA population (from the US Census) to obtain total MSA VKT and VTT.⁶

Our data on MSA road infrastructure are from the 1995, 1996, 2001, 2002, and 2008 Highway Performance and Monitoring System (HPMS) Universe and Sample data. The US federal government administers the HPMS through the Federal Highway Administration in the Department of Transportation. This annual survey, which is used for planning purposes and to apportion federal highway funding, collects data about the entire interstate highway system (HPMS Universe data) and a large sample of other roads in urbanized areas (HPMS Sample data).

The HPMS Universe data describe every segment of interstate highway (IH) and allow us to calculate the number of lane kilometers of IH in each MSA for each NHTS year. To calculate lane kilometers of major urban roads (MRU) in the urbanized parts of an MSA, we sum lane kilometers for the five classes of roads reported in the HPMS Sample data; 'collector', 'minor arterial', 'principal arterial' and 'other highway'. We omit a residual class, 'local roads', that is not systematically reported in the HPMS. To assure that the resulting measures of road infrastructure are comparable to the NHTS surveys, which are collected over two years, we average each of the HPMS variables over the two relevant NHTS sampling years.

Table 1 contains summary statistics for our main variables in the 100 largest MSAs. Means and standard deviations for trip-level variables are reported in Panel A. Trip distance and trip duration increase from 1995 to 2001, from 12.2 to 13.2 km and from 23.7 to 27.8 minutes. Some of the increase in average trip duration is accounted for by a decrease in average trip speed (computed across trips) from 43 to 40 km/h. Average trip duration, distance, and speed are very similar in 2001 and

⁶We delete trips with extreme values for distance, duration, or speed. The exact procedure is detailed in Appendix B.

Table 1: Summary Statistics for the 100 largest MSAs

Variable	1995	2001	2008
Panel A. Trip-level data			
Mean trip distance (km)	12.2 (16.4)	13.2 (17.3)	12.8 (16.8)
Mean trip duration (min)	23.7 (22.5)	27.8 (24.6)	27.4 (23.8)
Mean trip speed (km/h)	43.1 (22.7)	40.0 (22.4)	39.8 (22.8)
Mean trip number (per driver)	4.5 (2.6)	4.2 (2.4)	4.1 (2.3)
Total observed number of trips	152,590	168,765	415,402
Panel B. MSA-level data			
Mean daily VKT ('000,000 km)	60.4 (88.7)	67.7 (96.4)	70.4 (99.5)
Mean daily VTT ('000,000 min)	73.2 (108.6)	89.8 (129.9)	93.9 (135.7)
Mean lane km (IH, '000 km)	2.1 (2.3)	2.2 (2.4)	2.4 (2.4)
Mean lane km (MRU, '000 km)	10.5 (13.5)	11.9 (16.1)	14.4 (18.1)
Mean MSA population ('000)	1,747 (2,673)	1,943 (2,915)	2,095 (3,052)

Notes: Authors' computations. Standard deviations in parantheses. VKT is an estimate of the total vehicle kilometers traveled by privately operated vehicle in an MSA and VTT is an estimate of the total vehicle time traveled by privately operated vehicles in an MSA. IH denotes interstate highways for the entire MSA. MRU denotes major roads within the urbanized area of an MSA.

2008. The average number of trips decreases from 6 in 1995 to 5.5 in 2008. We note that 1995 NHTS survey asks respondents to report the time it took to get to their destination, while the 2001 and 2008 surveys ask respondents to report exact departure and arrival times. This slight difference in wording may partly explain the observed decrease in speed between 1995 and 2001.⁷

Panel B of table 1 reports means and standard deviations for MSA-level aggregates. Average VKT and VTT grow from 1995 to 2008, by 14% for VKT and by 22% for VTT, with much of the increase in VKT accounted for by the sample average MSA population growth rate of 10%. Lanes of interstate highway grow by 14% between 1995 and 2008 while lanes of major urban roads grow by 40%. While MSA boundaries are constant over time, urbanized area boundaries are not, so that some of the growth in lane kilometers of major urban roads reflects the expansion of urbanized areas.

⁷To support this conjecture we note that our results below show that the drop in speed in 2001 is nearly entirely accounted by a small increase of slightly above 1 minute in the fixed cost of each trip. All comparisons between the 1995 NHTS and other years are subject to this caveat.

Table 2: Mean trip distance in kilometers, by trip purpose, for the 100 largest MSAs

Trip purpose	Frequency (1995-2008)	km 1995	km 2001	km 2008
To/from Work	21.3%	18.2 (19.3)	18.9 (19)	19.4 (19.2)
Work-related business	3.2%	17.5 (21.7)	20.8 (24.4)	20.7 (24.8)
Shopping	23.7%	7.6 (10.8)	8.8 (12.1)	8.5 (11.5)
Other family/personal business	23.1%	9.3 (13.2)	9.8 (14.0)	10.0 (14.1)
School/church	4.5%	11.1 (14.5)	11.0 (13.6)	11.6 (13.2)
Medical/dental	2.8%	13.0 (15.1)	12.9 (14.1)	13.9 (14.9)
Vacation	0.4%	35.5 (42.0)	28.5 (37.3)	23.9 (33.0)
Visit friends/relatives	5.3%	15.0 (20.3)	17.9 (23.9)	17.4 (23.1)
Other social/recreational	15.1%	12.2 (17.0)	12.3 (17.1)	11.5 (15.8)
Other	0.6%	14.6 (19.1)	18.1 (23.8)	26.3 (31.3)


Notes: Authors' computations using all three years of data (pooled together to compute frequencies) by averaging across all trips. Standard deviations in parentheses.

The NHTS data report the purpose of each trip using 10 consistently defined categories such as 'to/from work', 'shopping', or 'medical/dental'. Table 2 shows the mean and standard deviation of distance by trip purpose for the MSAs in our sample. There is significant and persistent variation in average trip distance across trip purposes. Shopping trips are shortest at about 8.5 km on average in 2008. Vacation trips are the longest and average 23.9 km. We note that 'vacation' and 'other' trips occur infrequently in the data and we sometimes exclude them from our analysis.

Figure 1 plots log distance and log (inverse) speed for two groups of trips in Chicago in 2008. The triangles represent commute trips and the circles represent two other groups of trips: school/church and medical/dental. It is clear from the figure that for both groups, speed is higher for longer trips. In fact, this relationship between speed and trip distance is one of the most important features of our data. Consistent with sample averages reported in table 2, for Chicago in 2008 commute trips are about twice as long as the other class of trips described in figure 1 (20.3 km for commutes, 9.0 km for school/church trips, and 12.6 km for medical/dental trips).

In addition to the NHTS and HPMS, we exploit several other sources of MSA level data as explan-

Figure 1: Speed and distance for some Chicago trips in 2008.



Commuter trips are represented by triangles (mean log distance 2.52, plain line). Church, school, medical, and dental trips are represented by circles (mean log distance 1.81, dashed line).

atory variables or as instrumental variables. Specifically, in our investigation of the determinants of MSA driving speed, we consider a number of geographical characteristics of cities (ruggedness, elevation range, cooling and heating degree days, and measures of urban form). We also use variables describing historical transportation networks (1947 interstate highway plan, 1898 railroads, and old exploration routes of the continent dating back to 1528) as instruments for the modern road network. Details about these variables are available in Appendix A and in Duranton and Turner (2011).

3. Speed and the supply of travel

Our unit of observation is a trip made by a particular driver in a particular city. Let i index the set of cities, j index the set of drivers in each city, and k the set of trips by a particular driver in a particular city. Let x_{ijk} denote distance for trip ijk in kilometers, and s_{ijk} the speed of the trip in minutes per kilometer. Actually, s is ‘inverse speed’, minutes per kilometer of trip distance, although to ease exposition we often refer to it as ‘speed’. Let $\tau_{ijk} \in \{1..T\}$ index the possible purposes for trip ijk . Finally, let χ_{ijk}^τ be an indicator variable that is one for trips of type τ and zero otherwise. Since our estimations will rely mainly on within-city variation we often omit the i subscript to increase legibility.

Drivers incur fixed time costs to begin and end a trip. Drivers must also rely more heavily on slow local roads for short trips than for the faster freeways and arterial roads that are available for longer trips. These two facts naturally cause short trips to be slower than short trips. In addition, as figure 1 shows, the relationship between speed and distance is an important feature of our data. Given this, we assume the following speed schedule for trips of different distances in a particular city

$$s_{jk}^s = x_{jk}^{-\gamma} \exp(C + \delta_j + \epsilon_{jk}). \quad (1)$$

This inverse supply curve gives the average speed, i.e., the price in minutes per kilometer, for each kilometer of distance on a particular trip. The parameter δ_j measures drivers' abilities to drive fast on all trips and reflects characteristics such as the driver's skillfulness or the proximity of his or her home to a freeway. The parameter ϵ_{jk} measures drivers' ability to drive fast on a particular trip and reflects events such as stormy weather or road construction. The parameter C is the speed of a trip of unit distance when $\delta_j = 0$ and $\epsilon_{jk} = 0$. Finally, γ is the elasticity of speed with respect to trip distance. Estimating these supply curves, the parameters C and γ in particular, is the goal of our first empirical exercise.

A driver's inverse demand for trip distance x and purpose τ is


$$s_{jk}^d = x_{jk}^{-\beta} \exp(\sum_{\tau=1}^T A^\tau \chi_{jk}^\tau + \eta_j + \mu_{jk}). \quad (2)$$

In equation (2), the 'slope' parameter β ($> \gamma$) is the demand elasticity of speed with respect to distance. This parameter determines the rate of decline in a driver's willingness to pay for a unit of distance as trip distance increases. The 'intercept' parameter A^τ measures the willingness to pay for a trip of type τ and distance one when $\eta_j = 0$ and $\mu_{jk} = 0$. The parameter η_j describes a driver's willingness to give up time for distance. It depends on a driver's innate impatience or value of time. The parameter μ_{jk} reflects trip-specific factors which affect willingness to give up time for distance, for example, how busy a day the driver is having.

Since drivers recognize that their choice of distance affects the speed of travel, they behave as ordinary monopsonists and choose trip distance to satisfy

$$s^d = MC(x) \equiv \frac{d(xs^s)}{dx} = (1 - \gamma)s^s. \quad (3)$$

Figure 2: Supply and demand for trip distance.



That is, the marginal willingness to pay for trip distance equals the marginal cost of trip distance. Note that drivers are still price takers in the sense that they do not recognize that their driving behavior may contribute to congestion in the network and thus shift the whole supply curve.

Figure 3 illustrates this equilibrium. This figure depicts two pairs of marginal cost and demand curves, $(MC_1(x), s_1^d)$ and $(MC_2(x), s_2^d)$ for two different trips. We observe equilibrium pairs of speed and trip distance, points a and c . If we estimate our marginal cost curve on the basis of these points alone, our estimate of the marginal cost curve will look like the thin line connecting these two points. More specifically, we estimate a slope larger than the slope of the inverse-supply and marginal cost curves but smaller than the slope of the demand curve. As figure 3 makes clear, the problem of estimating the marginal cost curve relating trip distance and inverse speed is a textbook example of a simultaneous equations problem. In order to identify our marginal cost curve, we require a source of variation in demand which does not affect supply.

After defining χ_j to be an indicator variable which is one for trips by person j and zero otherwise, using (2) and (1) in (3), and taking logarithms we arrive at the following system of equations,

$$\ln s_{jk} = C + \delta_j \chi_j - \gamma \ln x_{jk} + \epsilon_{jk} \quad (4)$$

$$\ln x_{jk} = D_j \chi_j + \sum_{\tau=1}^{T-1} \tilde{A}^\tau \chi_{jk}^\tau + \zeta_{jk}, \quad (5)$$

where $D_j \equiv \frac{-C}{\beta-\gamma} + \frac{A^T}{\gamma-\beta} + \frac{\eta_j - \delta_j}{\beta-\gamma}$, $\tilde{A}^\tau \equiv \frac{A^T - A^\tau}{\gamma-\beta}$, $\tau \in \{1, \dots, T-1\}$, and $\zeta_{jk} \equiv \frac{\mu_{jk} - \epsilon_{jk}}{\beta-\gamma}$.

Inspection of equations (4) and (5) shows that A^τ , the willingness to pay for a trip of type τ , χ_{jk}^τ

the dummy for the trip being of type τ , and η_j the individual characteristics affecting the demand for trips of driver j , all appear in the distance equation (5), but not in the speed equation (4). It follows that variables measuring these quantities are candidate sources of exogenous variation in demand with which to resolve our simultaneity problem.

In practice, it is hard to think of individual characteristics that affect the demand for trips but not the ability to produce them. For instance, educational attainment affects a driver's opportunity cost of time and hence demand for trip distance. However, education may also affect driving skills and thus the ability to drive at a high speed. This suggests that individual characteristics are unlikely to provide good sources of exogenous variation in demand.

Trip type indicators and measures of mean trip distance by trip type are more defensible sources of identifying variation with which to identify the inverse supply curve described by equation (1). Denote these instruments Z_{jk} . As made clear by the discussion above, valid instruments for trip distance must satisfy two conditions. First, they must predict trip distance conditional on the other controls: $cov(Z_{jk}, x_{jk} | \cdot) \neq 0$ (relevance). We demonstrate that this condition holds below. Second, instruments must be uncorrelated with the error term of equation (4): $cov(Z_{jk}, \epsilon_{jk} | \cdot) = 0$ (exogeneity).

The arguments for the validity of trip type dummies and the validity of mean distance by trip type are different. We begin with trip type dummies. If trip type dummies are orthogonal to ϵ_{jk} then we are not more (or less) likely to observe trips of type τ when such trips are particularly fast. In fact, we suspect that some trips (e.g., 'recreational' trips, to take an example from the data) might be taken with greater propensity when traffic conditions are good, i.e., when ϵ_{jk} is high.

To understand why such a correlation might arise, assume there are only two types of trip: to the gym and to work. Drivers stop going to the gym when there is more than 10 centimeters of snow on the ground and stop going to work when there is more than 30 centimeters of snow (and traffic gets even slower). In this case, trips to the gym will be positively correlated with the error term. This leads to a selection problem which in turn violates the exogeneity condition.

To circumvent this possible problem, we can restrict attention to trips which are not discretionary such as trips 'to and from work' and 'medical/dental' trips.⁸ In addition, we can control for trip characteristics like; month, day of week, and time of day. We are concerned that our instruments

⁸Actually, we only need to restrict attention to trips with the same level of discretion. Arguably, trips 'to and from work' and 'medical/dental' trips have both a low level of discretion. Whether 'shopping' and 'recreational' trips have the same level of 'high' discretion is less obvious.

are correlated with unobserved determinants of speed, the ϵ_{jk} . Adding controls reduces the role of unobserved determinants of speed. Thus, if adding controls does not cause big changes in our estimations then this suggests that our concern that our instruments are correlated with ϵ_{jk} is unfounded.

Turning to the validity of average distances by trip type, we return to the example above. Trips to the gym may be observed on average under better weather (and traffic) conditions than trips to work. This is not an issue as long as this differential selection of trips does not affect the distance of trips to the gym relative to the distance of trips to work. However, suppose that when the weather is worse drivers go to a closer gym and to a closer workplace. If weather conditions affect distances equally for all types of trips, then unobserved weather introduces noise but does not lead to a violation of the exogeneity condition. On the other hand, if more than 10 centimeters of snow causes drivers to choose a closer gym but does not affect the choice of workplace, then mean trip distance by type is correlated with the error term in the speed regressions. More generally, average distances by trip type do not satisfy the exogeneity condition when the unobserved state of traffic differentially affects the distance of different types of trips.

Whether the length of shorter trip types should be more (or less) sensitive to traffic conditions than distance for longer trips types is not obvious. While (on average) distance for short shopping trips may be sensitive to traffic conditions, that of longer recreational trips might be as well. On the other hand, distance for short trips to school may be insensitive to traffic conditions and long trips to work may be similarly insensitive. This said, to avoid a possible correlation between average trip distance and unobserved traffic conditions, we can again restrict attention to trips with low discretion. We can also use extensive controls for trip characteristics, as we do when using trip type dummies as instruments.

To conclude, trip type dummies and mean distance by trip type may fail the exogeneity condition. Unobserved determinants of speed may differentially affect the decision to take particular types of trips. This would create a sample of trips where different types of trips are observed only under systematically different unobserved conditions. On the other hand, unobserved conditions may differentially affect the length of different types of trips, which leads mean trip length by type to be correlated with unobserved determinants of speed. Importantly, if the trip type and average trip distance instruments are not exogenous, they are not exogenous for different reasons. Therefore, if both types of instruments lead to similar estimates this indicates that they are either

both valid or, improbably, that the selection problem for trip type dummies has the same effect on our estimates as the hypothetical simultaneity problem for average distance by trip types.

Apart from concerns about the validity of our instruments, we worry that fast drivers sort into fast cities. As we see in equation (4), the constant term in the speed equation is the sum of the intercept of the inverse-supply curve C (the coefficient of interest), and driver characteristics affecting supply δ_j . Since we only observe drivers driving in one city, C and δ_j cannot be separately identified. Our concern is that fast drivers, those with high δ 's, might systematically choose to locate in fast cities, those with high C 's.

We have two responses to this problem. The first is to consider large areas, the largest us consolidated statistical metropolitan areas (MSAs), as our unit of observation. As long as the problematic sorting of drivers occurs at a smaller scale than our unit of observation, it will not lead to systematic differences between drivers in one MSA and another. Drivers with a desire to drive fast can always locate close to highways in a less densely populated part of nearly any large MSA in the us. Much the same logic is widely used to identify local peer group effects (e.g., Evans, Oates, and Schwab, 1992, for an early example). Our second response to the sorting problem is to parameterize individual effects as a function of observable driver characteristics. In particular, we expect that extensive controls such as age, income, gender, and education are correlated with individual unobservables. Since only the residual ϵ_{ijk} will be confounded with the intercept, we use our controls to reduce this residual as much as possible.

The supply of speed

We now turn to the estimation of travel speed. We start with the estimation of variants of the following equation,

$$\ln s_{ijk} = C_i + Y_j \delta - \gamma_i \ln x_{jk} + T_{jk} \zeta + \epsilon_{ijk}. \quad (6)$$

This equation differs from equation (4) in two regards. First, it includes a vector of trip controls, T_{jk} not present in (4). Such controls are needed since traffic conditions vary by time of the day, day, and week. Second, equation (6) includes an arbitrary vector of individual control variables Y_j . This generalizes equation (4) which restricts attention to individual fixed effects.

We estimate equation (6) using NHTS trips by drivers residing in one of the MSAs in our sample. Table 3 reports our results. Panel A report results based on the 2008 NHTS for trips by drivers in the

100 largest MSAs. Panel B replicates panel A but restricts attention to the 50 largest MSAs. Panels C and D reproduce panel B but are based on the 2001 and 1995 and NHTS. Within each panel and for each estimation, we report the mean values of the MSA intercept C_i and the MSA slope γ_i . For both variables, we report in parentheses the standard deviations of their coefficients within the sample of MSAs at hand. We note that the coefficients reported in table 3 are unorthodox. We estimate equation (6) for each MSA. Table 3 reports the mean and variance of the point estimates of C_i and γ_i across MSAs.

In column 1, we estimate equation (6) without driver or trip controls. The mean value of C_i for the 100 largest MSAs in 2008 appears in the first row of panel A. Its value of 1.407 implies just above 4 minutes for a trip of one kilometer.⁹ This is slightly less than 15 kilometers per hour. The second row of the same column reports the standard error of the mean of the intercepts *across* MSAs. Its value of 0.092 implies an $e^{0.092} \approx 10\%$ difference in speed for a trip of one kilometer. The mean of the standard error *within* MSAs is only 0.031. This suggests that the differences in intercepts across MSAs reflect mostly true differences in speed, not sampling error.

The third row of panel A reports the average of the coefficients for log distance. In column 1, its value of 0.428 implies that speed increases by about $2^{0.428} \approx 35\%$ when trip distance doubles. Alternatively, our coefficients predict a speed of 29 kilometers per hour for a trip of 5 kilometers (instead of just less than 15 for 1 kilometer).¹⁰ The fourth row reports the average standard deviation for these estimates of γ across MSAs. It equals 0.032. Since this is more than twice as large as the mean standard error for γ within MSAs, 0.015, this probably reflects true heterogeneity across MSAs and not just sampling error.

In column 2, we augment the regression of column 1 with several controls for driver characteristics; household income and its square, driver's education and its square, age, and dummies for males, blacks, and workers. We also include trip controls; a quartic in departure hour, and a week-end dummy. In column 3, we include more exhaustive driver and trip controls. For drivers we include 17 dummies for household income, four dummies for education, age, and dummies for males, blacks, hispanics, and workers. For trips, we include 23 dummies for the hour of departure, 11 dummies for the month of departure, and a dummy for trips taken during the week-end.

⁹Since this quantity is an exponential of an average of logs from which we omit the errors, strictly speaking, it is not predicted speed.

¹⁰Of course, we cannot expect this relationship to scale up for extremely long trips. However, 99% of the trips we observe here are less than 83 kilometers and our elasticity still applies within this region.

Table 3: Estimation of inverse-supply curves

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	OLS1	OLS2	OLS3	FE	IV1	IV2	IV3	IV4	IV FE
Panel A. 100 largest MSAs for 2008									
Mean C	1.407 (0.092)	1.338 (0.090)	1.474 (0.092)	1.402 (0.102)	1.281 (0.149)	1.281 (0.143)	1.237 (0.138)	1.214 (0.142)	1.266 (0.760)
Mean γ	0.428 (0.032)	0.426 (0.032)	0.425 (0.032)	0.426 (0.037)	0.360 (0.075)	0.359 (0.072)	0.335 (0.064)	0.346 (0.074)	0.357 (0.402)
Panel B. 50 largest MSAs for 2008									
Mean C	1.407 (0.065)	1.342 (0.067)	1.478 (0.069)	1.399 (0.070)	1.261 (0.098)	1.251 (0.091)	1.222 (0.104)	1.180 (0.087)	1.258 (0.120)
Mean γ	0.424 (0.018)	0.421 (0.017)	0.420 (0.017)	0.419 (0.020)	0.342 (0.046)	0.336 (0.042)	0.320 (0.045)	0.321 (0.042)	0.346 (0.054)
Panel C. 50 largest MSAs for 2001									
Mean C	1.383 (0.072)	1.324 (0.069)	1.454 (0.071)	1.350 (0.067)	1.324 (0.112)	1.320 (0.122)	1.298 (0.118)	1.262 (0.131)	1.244 (0.240)
Mean γ	0.412 (0.022)	0.407 (0.021)	0.406 (0.021)	0.394 (0.022)	0.348 (0.060)	0.345 (0.065)	0.332 (0.065)	0.342 (0.066)	0.341 (0.112)
Panel D. 50 largest MSAs for 1995									
Mean C	1.187 (0.103)	1.171 (0.081)	1.199 (0.081)	1.133 (0.090)	1.121 (0.131)	1.115 (0.136)	1.048 (0.119)	1.070 (0.139)	1.057 (0.163)
Mean γ	0.380 (0.040)	0.375 (0.022)	0.374 (0.023)	0.351 (0.026)	0.340 (0.070)	0.338 (0.072)	0.303 (0.054)	0.326 (0.075)	0.310 (0.079)

Notes: Mean of the coefficients across all cities. Standard deviation of city coefficients in parentheses. OLS estimations in columns 1-4 and IV in columns 5-9.

Controls: No control in column 1. Controls for household income and its square, driver's education and its square, age, dummies for males, blacks, and workers, and a quartic for the time of departure in columns 2 and 5-8. 17 dummies for household income, four dummies for education, age, dummies for males, blacks, hispanics, and workers, 23 dummies for the hour of departure, 11 dummies for the month of departure, and a dummy for trip taken during the week-end in column 3. Driver fixed effects in columns 5 and 9.

Instruments: Mean trip distance for trips of the same purpose in the same MSAs in column 5. Same instrument but computed from the four most similar MSA in term of population in columns 6 and 9. Trip purpose in column 7 (8 categories) and column 8 (2 categories; commutes and other work related trips, shopping, medical and dental, school and church, and personal business).

In columns 2 and 3 of table 3 we constrain the effect of driver and trip characteristics to be the same for all MSAs. This increases the efficiency of our estimations, increases the transparency of the speed indices calculated below, and eases the calculation of these indices. Moreover, since regressions with driver and trip characteristics give similar results to regressions with driver fixed effects, it seems unlikely that this simplifying assumption is important to our estimates of the speed distance relationship.¹¹

We cannot compare estimates of C_i across columns 1, 2, and 3 because of differences in the sets of control variables. We can compare estimates of the distance elasticity of speed, γ , across columns. These estimates are stable.¹² The R^2 s for the different specifications are also stable. The R^2 associated with column 1 when estimating an intercept (C) and a slope (γ) for each city is 56.6%. Adding driver and trip controls in column 2 raises this R^2 slightly, to 57.6%. The more exhaustive controls of column 3 also increase the R^2 slightly, this time to 57.8%. Controls for trip and driver characteristics do not affect our estimation of the distance elasticity.

This does not imply that driver and trip characteristics do not affect speed. They do. While we do not report these coefficients, several are interesting. Women are about 0.5% slower than men. Age is more important. A year of age is associated with 0.3% slower speed. Black drivers drive about 8% slower. Drivers with more education and drivers with higher income are faster, although in both cases the relationship tapers off after a threshold: drivers with a Bachelor degree are about 7% faster than workers with less than high school; drivers from households with annual income around \$60,000 are about 9% faster than drivers from the poorest households.¹³ Our findings on the effect of trip characteristics are unsurprising: week-end trips are about 4% faster than week-day trips; trips departing during the morning peak are about 4% slower than trips in the middle of the night; trips departing during the evening peak are about 10% slower than trips in the middle of the night; there are small differences between months, Winter months are about 1% faster and Fall months about 1% slower than Spring and Summer months.

In column 4, we return to the specification of column 1 and introduce driver fixed effects. The results for this column confirm that driver characteristics do not affect the estimation of

¹¹It would also be of interest to investigate whether the effect of driver and trip characteristics vary across MSAs, e.g., to check if 'peak' hours differ in intensity and duration across MSAs. However, given that our ultimate objective is an understanding of city level determinants of speed, we leave such an investigation for future research.

¹²But this does not imply that the estimated coefficients for the intercept and the distance elasticity are always exactly the same for all cities. We return to this issue below.

¹³This might obviously be related to the state of their vehicles which we do not control for.

our parameters of interest. With driver fixed effects, the mean of both intercepts and slopes are unchanged from column 1.

Column 5 replicates the specification of column 2, but, to instrument for trip distance, uses the mean log distance of other trips with the same purpose in the same MSA. We note that this estimation raises a technical issue. To be consistent with Column 2 we want to estimate a separate IV regression for each MSA.¹⁴ At the same time, we want to constrain the effect of driver and trip characteristics to be the same everywhere. This calls for a two-step approach where the effects of driver and trip characteristics on speed are estimated first from the cross-section of MSAs. We then take these coefficients as given (i.e., treat them as constraints) when estimating a separate TSLS regression for each MSA.

If drivers take longer trips when travel is faster, then OLS estimates of γ are biased upwards. Comparing the IV results in column 5 panel A with the corresponding OLS results in column 2 we see that, as expected, the IV estimates of γ are smaller than the OLS estimates. In column 5, the mean IV mean elasticity of speed with respect to distance is 0.360. The corresponding OLS value from column 2 is 0.426. This 20% difference between the OLS and IV estimates is statistically significant for a majority of cities. These elasticities imply that after controlling for simultaneity in the choice of trip distance and speed, speed increases by only about $2^{0.359} \approx 28\%$ when trip distance doubles as opposed to the almost 35% increase we observe in equilibrium speed.

By inspection of figure 3, finding a lower IV than OLS estimate of C , the intercept of the inverse supply curve, is the logical counterpart of the smaller elasticity (recall that a smaller value of C requires *fewer* minutes per kilometer, and hence a higher speed). Consistent with this, the IV of column 5 results imply a speed of nearly 17 kilometers per hour for a trip of one kilometer. This is about two kilometers per hour faster than is implied by the OLS estimate of column 2.

In column 6, we replicate the IV estimations of column 5, but instrument for trip distance using mean log distance for trips of the same type in the four MSAs with most nearly the same population. The results are close to those of column 5. Mean log distance for trips of the same purpose in MSAs with similar population is a stronger instrument than mean log distance by trip purpose in the same MSA because the instrument is computed using more observations. In most cases we use four MSAs instead of one to calculate the instrument. While weak instruments are an issue for only

¹⁴So that we instrument trip distance in a city by only the instruments for this city instead of the entire set of instruments.

a small number of MSAs in column 5, we prefer the specification of column 6 and we use these regressions to compute our benchmark speed index.

In column 7, we use seven trip purpose dummies as instruments. Despite the different rationales for the validity of the instruments, this yields mean estimates close to those of columns 5 and 6. In column 8, we only use one trip purpose dummy as an instrument. This dummy is defined to be one for commutes and work-related trips and zero otherwise. For other trips, we also restrict our sample to less discretionary trip types; shopping, personal business, school and church, and medical/dental. The slopes we estimate for the inverse-supply curves are close to those from the other IV estimates. While we do not report them here, the standard errors associated with our estimates for each MSA increase when using these more demanding methods of estimation.

Finally, in column 9 we return to the same instruments as in column 6 but apply them to the fixed effect estimation of column 4. This is a very demanding estimation strategy since the effect of distance on speed is identified within driver from the speed differences between long work trips and shorter trips. When estimating the slopes and intercepts for the 100 largest MSAs as we do in panel A, the coefficients are on average the same as estimates for the 50 largest MSAs, but exhibit more dispersion.

Panel B replicates panel A for our sample of the 50 largest MSAs. For more demanding estimations like column 9, the coefficients of interest are more precisely estimated than when we consider the 100 largest MSAs. Even for the specification of column 1, the mean of the standard error on C and on γ is about twice as large for MSAs ranked between 51 and 100 as for the largest 50.

Panels C and D replicate panel B for 2001, and 1995, respectively. The results for 2001 are similar to those for 2008. This is consistent with mean trip distances and speeds being essentially the same in 2001 and 2008. We note nonetheless that the variances of the city intercepts and distance elasticities are larger in 2001 than in 2008. This probably reflects the larger sample drawn by the 2008 NHTS. For 1995, the estimated distance elasticities of speed are close to but smaller than those for 2001 and 2008. The city intercepts are also smaller. This is consistent with the observed reduction in mean speed after 1995.

As a final robustness check, we experiment with alternative functional forms. Figure 1 suggest that for Chicago in 2008 the relationship between log speed and log distance is approximately linear. We see a similar pattern in other cities. To investigate the possibility of non-linearities more formally, we first re-estimate the specification of column 2 table 3 with a common distance effect.

The R^2 of that regression is 0.576. Adding the square and the cube of log distance raises this R^2 slightly to 0.589. While the coefficients associated with these two non-linear terms are statistically significant, they make little economic difference. For a trip of 5 kilometers, the elasticity of speed with respect to distance is 0.490. For a much longer trip of 30 kilometers (at the 90th percentile of trip distance) we find a slightly lower elasticity of 0.420.¹⁵ We prefer to use the linear specifications reported in table 3 because they are more precisely estimated.¹⁶

In sum, three main findings emerge from table 3. First, there is evidence of the simultaneous determination of speed and distance. This will affect our estimates of mean speed for all US MSAs. As the bias may differ across cities, this can also affect our ranking of cities. Second, different IV strategies yield similar results. Third, the inclusion of other trip and driver controls makes little difference to our estimates.

4. Speed index

Our objective is to understand the determinants of driving speed in major US cities. Yet, as our model and data make clear, ‘the speed of travel in a city’ is not well defined. Cities do not offer a single speed of travel. They offer a menu of feasible speed and trip distance combinations.¹⁷

Therefore, to describe a city’s ability to supply road transportation we rely on a speed index that is analogous to a standard price index. Loosely, we first calculate the bundle of trips undertaken by an average US driver. Next, using our city level estimates of equation (6), we calculate the time required to complete the average daily bundle of trips in the subject city and in an average city. The ratio of these two times tells us the speed of travel in the subject city in a precise sense. It tells us the time premium required to complete a standard bundle of trips in the subject city relative to an average city.

More formally, suppressing MSA population weights and year indices, the speed (Laspeyres)

¹⁵That the elasticity of speed with respect to distance should eventually decline with trip distance is to be expected since speed is bounded.

¹⁶We also experimented with non-parametric and semi-non-parametric specifications. They generally confirm the finding that a linear specification provides a good first order approximation of the relationship between distance and speed and that the elasticity of speed with respect to distance falls slowly with distance.

¹⁷This obviously makes comparisons across cities difficult. A city with a high C will be fast for short trips but may be slow for long trips if its γ is low. The problem is all the more important since we expect C and γ to be positively correlated, be it only for ‘econometric’ reasons. C and γ are estimated together so that a positive error in the estimation of C should also imply a positive error for γ which enters the same estimating equation with opposite sign.

index for city i is

$$S_i = \frac{\sum_{jk} x_{jk} \exp(C_i - \gamma_i \ln x_{jk})}{\sum_{jk} x_{jk} \exp(C_{US} - \gamma_{US} \ln x_{jk})}. \quad (7)$$

That is, we compute the time that it would take to realize all (weighted) us trip distances in our data at the estimated speed of a given MSA relative to how much time it would take to realize the same trips at the average estimated us speed. We note that, strictly speaking, S_i is an inverse speed index. It measures proportional change in time required to complete a standard bundle of trips in a particular MSA relative to the us average.

Calculation of speed indices

Table 4 reports our preferred speed index for the largest 50 us MSAs in 1995, 2001, and 2008. This index is based on our preferred estimate of equation (3) from column 6 of table 3. In this regression we instrument for trip distance with mean distance for trips of the same type in the four cities most nearly the same size.

For 2008, we find that the time cost of driving in the slowest MSA, Miami, is 27% higher than in the fastest, Grand Rapids. This gap is slightly larger in 2001 and lower in 1995. More generally, among the 10 slowest MSAs, we find the four largest (New York, Los Angeles, Chicago, and Washington), another from the top 10 (Boston), four large cities with a difficult geography (Miami, Seattle, New Orleans, and Pittsburgh) and one city with stringent zoning regulations (Portland).¹⁸

There are some changes in ranking between 1995 and 2008. The rank correlation between the 2008 and 2001 ranking is 0.81 while that between the 2008 and 1995 ranking is 0.62. These correlations are high, particularly given the likelihood that some changes in ranking across years are due to sampling error. However, many of the changes in rank probably also reflect changes in city level fundamentals. For example, in a regression of changes in the speed index from 1995 to 2008 against population changes over the same period and the 1995 value of the same index, we find a coefficient 0.073 (significant at 10%) for the 50 largest MSAs and 0.101 (significant at 5%) for the 100 largest MSAs.

¹⁸Obviously, MSA boundaries also matter. For instance, the metropolitan area of New York contains much more than New York City, including many areas where traffic is fast.

Table 4: Ranking of the 50 largest MSAs, slowest at the top

	2008 Index	2008 Rank	2001 Index	2001 Rank	1995 Index	1995 Rank	Population rank
Miami-Fort Lauderdale, FL	1.13	1	1.14	1	1.09	2	14
Chicago-Gary-Kenosha, IL-IN-WI	1.10	2	1.07	2	1.11	1	3
Seattle-Tacoma-Bremerton, WA	1.06	3	1.06	4	1.02	8	12
Portland-Salem, OR-WA	1.06	4	0.96	19	0.92	28	22
Los Angeles-Riverside-Orange County, CA	1.06	5	1.03	7	1.00	11	2
New York-Northern NJ-Long Isl., NY-NJ-CT-PA	1.05	6	1.05	5	1.07	3	1
New Orleans, LA	1.05	7	1.02	9	0.98	13	37
Pittsburgh, PA	1.04	8	1.02	10	0.98	14	21
Boston-Worcester-Lawrence-Low.-Brock., MA-NH	1.04	9	1.02	11	1.01	9	7
Washington-Baltimore, DC-MD-VA-WV	1.04	10	1.02	8	1.03	6	4
San Francisco-Oakland-San Jose, CA	1.04	11	0.99	13	0.99	12	5
Sacramento-Yolo, CA	1.03	12	0.97	17	0.82	45	24
Houston-Galveston-Brazoria, TX	1.02	13	0.95	21	0.91	31	10
Tampa-St. Petersburg-Clearwater, FL	1.02	14	0.98	15	0.92	26	20
Orlando, FL MSA	1.01	15	0.98	16	0.96	15	26
Philadelphia-Wilmington-Atl. City, PA-NJ-DE-MD	1.01	16	1.04	6	1.03	5	6
Norfolk-Virginia Beach-Newport News, VA-NC	1.01	17	0.89	37	0.93	24	31
Phoenix-Mesa, AZ	1.00	18	0.96	20	0.91	29	13
Las Vegas, NV-AZ	0.99	19	1.06	3	0.87	41	29
Cleveland-Akron, OH	0.98	20	0.89	36	1.05	4	16
Atlanta, GA	0.97	21	0.93	23	0.93	25	11
Detroit-Ann Arbor-Flint, MI	0.97	22	0.93	24	0.95	20	9
Austin-San Marcos, TX	0.97	23	0.93	22	0.94	21	35
San Diego, CA	0.96	24	0.91	31	0.91	30	17
St. Louis, MO-IL	0.96	25	0.92	28	0.83	43	18
Dallas-Fort Worth, TX	0.96	26	0.93	25	0.90	34	8
Denver-Boulder-Greeley, CO	0.96	27	1.01	12	1.01	10	19
Salt Lake City-Ogden, UT	0.95	28	0.92	26	0.94	22	34
San Antonio, TX	0.95	29	0.91	32	0.87	38	28
Indianapolis, IN	0.94	30	0.97	18	0.90	32	30
Jacksonville, FL	0.94	31	0.88	40	1.02	7	44
Hartford, CT	0.93	32	0.87	41	0.81	47	40
Charlotte-Gastonia-Rock Hill, NC-SC	0.92	33	0.88	38	0.87	39	32
West Palm Beach-Boca Raton, FL	0.92	34	0.89	35	0.95	18	43
Richmond-Petersburg, VA	0.91	35	0.84	46	0.86	42	49
Buffalo-Niagara Falls, NY	0.91	36	0.85	44	0.92	27	41
Memphis, TN-AR-MS	0.91	37	0.80	48	0.88	36	42
Minneapolis-St. Paul, MN-WI	0.91	38	0.90	33	0.90	33	15
Columbus, OH MSA	0.91	39	0.86	42	0.95	17	33
Milwaukee-Racine, WI	0.91	40	0.92	27	0.95	19	27
Raleigh-Durham-Chapel Hill, NC	0.90	41	0.84	45	0.82	44	38
Cincinnati-Hamilton, OK-KY-IN	0.89	42	0.99	14	0.96	16	23
Nashville, TN	0.89	43	0.92	29	0.81	46	36
Oklahoma City, OK	0.87	44	0.89	34	0.87	37	45
Rochester, NY	0.86	45	0.88	39	0.87	40	46
Greenville-Spartanburg-Anderson, SC	0.86	46	0.85	43	0.69	50	50
Kansas City, MO-KS	0.85	47	0.77	50	0.93	23	25
Greensboro-Winston-Salem-High Point, NC	0.84	48	0.80	47	0.80	48	39
Louisville, KY-IN	0.83	49	0.92	30	0.78	49	48
Grand Rapids-Muskegon-Holland, MI	0.82	50	0.79	49	0.89	35	47

Notes: Laspeyres index constructed from the estimations reported in column 6 of table 3.

Robustness checks

To assess the robustness of our preferred ranking, we compare it to alternative rankings based on the same data but different aggregation methods or different estimation strategies.

First, the exact construction of our index does not matter. Our speed index is a Laspeyres index. Using a Paasche or a Fisher index instead makes little difference to our ranking of MSAs. Using our preferred estimation strategy and our sample of the 50 largest MSAs, the Spearman rank correlation between our preferred ranking and its Paasche counterpart is above 0.99 for 1995, 2001, and 2008. If we consider the 100 largest MSAs the corresponding correlations are all above 0.97. These findings are not specific to our choice of estimation. We find similarly high correlations for the Paasche and Laspeyres rankings constructed from the output of OLS estimation for column 2 of table 3 (i.e., the OLS estimation that corresponds to our preferred IV).

Next, we compare our preferred ranking, calculated from column 6 of table 3, with alternative rankings calculated from other columns of the same table and with average inverse-speed calculated directly from the data. Starting with the latter, the Spearman rank correlation between our preferred ranking and one obtained based on MSA average inverse speed is 0.68 for 50 MSAs in 2008. The Spearman rank correlations between our preferred ranking and alternative rankings obtained from the OLS estimates of column 1 to 3 of table 3 are between 0.91 and 0.94 for 50 MSAs in 2008. For the fixed-effect estimation of column 4, the correlation is slightly lower at 0.88. For the IV estimations of columns 5, 7, 8, and 9, the correlations are 0.97, 0.90, 0.95, and 0.72, respectively. This last correlation is lower because of the more noisy estimates obtained in column 9 (our most demanding estimation which includes driver fixed effects in an IV regression).¹⁹

For 1995 and 2001, correlations across rankings are slightly lower. For instance, in our sample of the 50 largest MSAs, the Spearman rank correlation between our preferred ranking and the alternative ranking obtained from the OLS estimates of column 2 of table 3 is 0.88 in 2001 and 0.81 in 1995 instead of 0.93 in 2008. For the 100 largest MSAs, the Spearman rank correlation between our preferred ranking and the same alternative based on OLS estimates is 0.85 in 2008. More generally, correlations drop when we use 100 MSAs instead 50.

¹⁹The variance in average normalized inverse-speed across the 50 largest MSAs in 2008 is 0.0924. The variance of a speed index based on the simple OLS estimation of column 1 of table 3 which controls for trip distance is lower at 0.0571. The variance of our preferred speed index is 0.0743. Given our expectation of longer distances when traffic conditions are better, both controlling for distance and then controlling for its endogeneity should reduce differences in observed speed across MSAs. At the same time, more demanding estimation methods introduce greater sampling error. The first effect dominates when going from the raw data to OLS whereas the second dominates when implementing IV estimations.

We draw a number of conclusions from these correlations. First, the relatively low correlation between our preferred ranking and raw measures of speed underscores the importance of controlling for trip distance.²⁰ Second, the high correlations between the indices derived from our various IV estimations suggest that our preferred ranking is not sensitive to the details of our instrumentation strategy provided the relationship between speed and distance is precisely estimated. Third, the relatively high correlations between our preferred ranking and the rankings derived from OLS estimates suggest that controlling econometrically for the simultaneous determination of speed and distance has only a small effect on the final ranking of MSAs.

Comparison with TTI's index

The Texas Transportation Institute produces the best known and most widely reported indices of travel speed. We here investigate the differences between our index and the 2008 and 2009 TTI indices (Schrank and Lomax, 2009, Schrank *et al.*, 2010).

While TTI reports their index annually, their methods changed from 2008 to 2009 (Schrank and Lomax, 2009, Schrank *et al.*, 2010). The 2008 TTI index is based on an estimate of travel speed constructed from data in the HPMS describing road characteristics and traffic levels on interstate highways and other federally funded roads. The 2009 TTI, on the other hand, is based on directly observed travel speeds on highways and major arterial roads for a self selected sample of commercial vehicles and the drivers of privately-owned vehicles. Unlike the TTI index, our index is based on a sample constructed to be broadly representative of all trips in privately-owned vehicles. In addition, the TTI indices do not attempt to control for the simultaneity in the choice of distance and speed.

Aside from differences in the quality of the underlying data and methodology, the TTI indices are based on a different geography than we use (see Appendix A). While neither geography is intrinsically preferable, this difference complicates the comparisons of indices. For the 47 cities reported by TTI also in our sample of the 50 largest MSAs, the Spearman rank correlation of our preferred speed index is 0.67 for the 2008 TTI travel cost index and 0.74 for 2009. For the 71 cities reported by TTI also in our sample of the 100 largest MSAs, those correlations are 0.61 and 0.63 respectively.

²⁰This lower correlation of 0.68 is not specific to our preferred estimation. None of the indices built from the estimations in tables 3 has a correlation with average inverse speed that is above 0.75.

5. Determinants of speed

To investigate why some cities are faster than others, we regress our speed index on a number of its determinants, roads and vehicle travel time in particular. That is, we estimate variants of the following regression,

$$-\ln S_i = \alpha \ln R_i - \theta \ln \text{VTT}_i + X_i\phi + \nu_i. \quad (8)$$

In this equation R_i is a measure of the city's stock of roads, VTT_i is aggregate vehicle travel time for the city, X_i is a set of other city characteristics and ν_i an error term. Note that the dependent variable in this equation is $-\ln S_i$ rather than $\ln S_i$. Therefore, this regression predicts speed or, more precisely, the proportional difference in the speed at which a standard bundle of trips is conducted in a particular MSA relative to the sample average speed for the same standard bundle of trips.

Three comments about this estimation are in order. First, since speed multiplied by vehicle travel time equals vehicle kilometers traveled (vKT), equation (4) is equivalent to

$$\ln \text{vKT}_i = \alpha \ln R_i + (1 - \theta) \ln \text{VTT}_i + X_i\phi + \nu_i. \quad (9)$$

Equation (9), and thus equation (8), is the exact counterpart of a standard production function. Vehicle kilometers traveled is our measure of output. Roads and vehicle time traveled are factors of production. α is the share of roads in the production of travel and $1 - \theta$ is the share of vehicle travel time. It follows that $\alpha - \theta$ is also a measure of returns to scale. If $\alpha < \theta$ there are decreasing returns to scale in the production of vKT . The error term ν_i is total factor productivity, the ability of a city to move its residents conditional on its stock of roads and aggregate time spent in cars.

Second, an alternative approach to understanding the role of roads and travel time as determinants of speed would be to include R_i and VTT as controls in equation (6). That is, we could skip the calculation of the speed index altogether and investigate the determinants of speed using individual data. We prefer the approach based on aggregate data, i.e., equation (9) because we think the speed index is of intrinsic interest. With this said, our results are robust to this alternative approach.

Finally, the econometrics problems with estimating production functions are well-known and have received considerable attention (e.g., Akerberg, Benkard, Berry, and Pakes, 2007, Syverson, 2011). Since we observe the true cost of travel, our analysis is not subject to problems associated with unobserved prices. Second, as in the standard case, both factors of production, roads and

Table 5: The determinants of speed, 100 MSAs in 2008

	(0)	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	S_{raw}	S_{OLS1}	S_{OLS2}	S_{OLS3}	S_{FE}	S_{IV1}	S_{IV2}	S_{IV3}	S_{IV4}	S_{IVFE}
log lane	0.087 ^c (0.047)	0.081 ^a (0.019)	0.089 ^a (0.020)	0.090 ^a (0.020)	0.070 ^a (0.019)	0.10 ^a (0.036)	0.11 ^a (0.033)	0.062 ^b (0.029)	0.11 ^a (0.033)	0.22 (0.19)
log VTT	-0.10 ^b (0.045)	-0.10 ^a (0.018)	-0.11 ^a (0.018)	-0.11 ^a (0.019)	-0.091 ^a (0.018)	-0.14 ^a (0.034)	-0.15 ^a (0.031)	-0.099 ^a (0.026)	-0.16 ^a (0.031)	-0.17 (0.18)
R ²	0.11	0.36	0.45	0.45	0.34	0.35	0.41	0.32	0.39	0.01

Notes: OLS regressions with a constant in all columns. Robust standard errors in parentheses. *a*, *b*, *c*: significant at 1%, 5%, 10%. 100 observations per column. In column 0, the dependent variable is average trip speed. In columns 1-9, the dependent variable is minus the log of the index computed from the results of the regressions reported in the corresponding column of table 3.

aggregate travel time, may be endogenously determined. To respond this endogeneity problem we consider a number of instrumental variables estimations and also estimate the dynamic production functions proposed by Levinsohn and Petrin (2003).

Determinants of speed, empirical results

Table 5 reports OLS estimations of equation (8) for our sample of the 100 largest MSAs. In all specifications we include a measure of travel speed on the log of MSA lane kilometers of interstate highways and the log of MSA vehicle travel time. In column 0, the dependent variable is average trip speed taken directly from the data. In columns 1 to 9, the dependent variable is minus the log of the index computed from the results of the regressions reported in the corresponding column of table 3.

Column 6 reports our preferred regression in table 5. In this regression the dependent variable is our preferred speed index calculated from the results of column 6 in table 3. This regression implies an elasticity of travel speed with respect to lane kilometers of road of 0.11. This quantity also corresponds to the share of roads in the production of travel. The coefficient for log VTT implies a negative elasticity of speed with respect to aggregate vehicle travel time of -0.15 . This quantity, $\theta = 0.15$, also implies that the elasticity of vehicle kilometers travelled with respect to travel time is 0.85.

The estimated coefficients for the log of lane kilometers and vehicle travel time are very similar in columns 0-8 of table 5. In all specifications, the coefficient on log lane remains between 0.07 and

0.11 while that on log vehicle travel time remains between -0.09 and -0.16 . Column 9 uses a speed index based on the least precisely estimated speed-distance relationship and yields insignificant estimates.

Appendix C describes extensive robustness tests for the results reported in table 5. We show that the coefficients we find on roads on aggregate travel time and the existence of modest decreasing returns are also robust to our choice of sample, year of data, definition of roads, and measure of aggregate travel time. In this appendix we also tackle the endogeneity of both roads and VTT. In a first exercise, we instrument roads and/or VTT. In a second exercise, we implement the methodology developed by Levinsohn and Petrin (2003). Both exercises largely confirm the results of table 5.

Given the analogy between our estimating equation (8) and a firm level production function, it is interesting to compare our results about the production of travel with what is known about the production of other goods.²¹ We find that the share of roads in the production of travel is about around 0.10, whereas typical estimates regarding the share of capital in conventional sectors tend to be around one third. This suggests that the production of travel is an extremely labour (or time) intensive activity.

We also find evidence of slightly decreasing returns. The literature that estimates the production function for firms contains a variety of results but micro-data estimates are often suggestive of constant returns. Slightly decreasing returns in the production of travel might seem surprising since roads constitute a network and networks are often associated with increasing returns. An explanation may be that urban roads are often organised as 'hub-and-spoke' networks whose hubs are more congested in larger cities. We return to this issue below.

Finally, consistent with extant research on productivity in firms (Syverson, 2011, Fox and Smeets, 2011), we find considerable dispersion in productivity across cities. However, there is much less dispersion in the ability of us cities to produce travel out of roads and vehicle travel time than in the ability of firms to produce output from capital and labour. A city at the 90th percentile in our preferred estimation (column 6 of table 5) produces around 20% more travel from the same inputs than a city at the 10th percentile. Firm level data usually imply that a firm at the 90th percentile of productivity produces 100 to 200% more output from the same inputs than a firm

²¹We do not know of estimates directly comparable to ours in the literature. Combes and Lafourcade (2005) decompose the decline in generalized transportation costs for trucks in France over 1978-1998 and find that changes in the road infrastructure only accounts for 8% of this decline.

Table 6: The determinants of speed, further explanatory variables

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Added:	Emp. central.	Pop. central.	Job/resid. mismatch	E pop. growth	log pop. 1920	s. manuf. emp.	Cooling deg. days	Heating deg. days
log lane (total)	0.098 ^a (0.033)	0.089 ^a (0.032)	0.11 ^a (0.033)	0.093 ^a (0.034)	0.090 ^a (0.033)	0.097 ^a (0.033)	0.10 ^a (0.033)	0.095 ^a (0.033)
log VTT	-0.15 ^a (0.031)	-0.15 ^a (0.030)	-0.15 ^a (0.031)	-0.13 ^a (0.033)	-0.15 ^a (0.031)	-0.14 ^a (0.031)	-0.14 ^a (0.031)	-0.14 ^a (0.032)
Added variable	-0.096 ^c (0.051)	-0.14 ^a (0.054)	-35.1 ^c (20.9)	-0.15 ^a (0.058)	0.015 ^a (0.0050)	0.22 ^a (0.13)	-0.013 ^b (0.0067)	0.0057 ^b (0.0026)
R ²	0.42	0.43	0.42	0.43	0.44	0.43	0.45	0.43

Notes: OLS regressions with a constant in all columns. Robust standard errors in parentheses. *a*, *b*, *c*: significant at 1%, 5%, 10%. 100 observations per column. Dependent variables is $-\log S_{IV2}$ in all columns.

at the 10th percentile.

Other determinants of speed

In table 6 we consider a broader set of determinants of travel speed. In column 1, we introduce a measure of employment centralization, the share of employment within 20 kilometers of the employment weighted centroid of the MSA in 1992.²² In column 2, we use instead the corresponding measure of centralization for population. The results for both columns indicate that more centralized cities are much slower. These findings should be regarded as suggestive; the coefficient for employment centralization is only weakly significant, the R² increases only marginally relative to the benchmark estimation without these additional variables, and the significance of these coefficients often disappears when further controls are added. In column 3, we use a measure of mismatch between employment and residents and find again that a greater mismatch is associated with slower traffic speeds. We have experimented more broadly with ‘urban form’ variables than we report here. Consistent with the results reported in table 6, we find that conditional associations with various measures of density, physical area and employment concentration, occur routinely in our results. Taken together, these findings are strongly suggestive that more compact and centralized cities are slower. This is consistent with earlier results in the literature (e.g., Glaeser and Kahn, 2004).

²²We prefer to use lagged variables to minimize endogeneity problems. These lags typically reduce the significance of the coefficients.

In column 4 of table 6, we turn to a different type of variable, population growth. We find a strong association between slow travel speeds and higher expected growth between 1980 and 2000.²³ A similar but weaker association is found with actual population growth between 1980 and 2008. That population growth should slow traffic is unsurprising in the light of evidence that an MSA's roads adjust slowly to population growth (Duranton and Turner, 2012). In column 5, we replace expected population growth with 1920 population. Given that we also condition for current vehicle time travel which is highly correlated with current population, our negative coefficient implies that cities that have grown less since 1920 are faster. This confirms the finding of column 4. In column 5, we turn to the share of manufacturing in employment in 1983 and find that cities more specialized in manufacturing are much faster. Given that cities with a greater share of manufacturing employment have grown less in population and tend to be more decentralized (Glaeser and Kahn, 2001), the positive sign on manufacturing employment is consistent with our two main findings so far (as well as findings in Duranton, Morrow, and Turner, 2011).

In columns 7 and 8, we introduce two measures of temperature, cooling and heating degree days, and uncover a weak association between slower traffic and more extreme weather conditions. We also experiment with other geographic characteristics of cities such as their elevation range or the ruggedness of their terrain but find nothing. We also found no result for a broad range of socioeconomic characteristics of cities such as their income, education, etc. Finally, we note that introducing all these supplementary explanatory variables has little effect on the coefficients of our two main regressors, roads and vehicle travel time, both of which remain significant and nearly constant in all the columns of table 6.

6. The value of speed

In this section we conduct two distinct exercises to estimate the value of speed. The first describes a series of simple out-of-equilibrium policy experiments. The second conducts a standard partial equilibrium welfare analysis.

²³Again, we use growth predicted by the composition of economic activity of cities in 1980 and the subsequent changes in employment by sector.

Out-of-equilibrium policy experiments

We consider four different counterfactual experiments. Under the first, we increase the speed of 95 of the 100 largest MSAs to the level of the MSA at the 95th percentile of the distribution of speed and leave the five fastest MSAs as they are. Let S_{95} denote the counterfactual speed index associated with this policy. In the second experiment, we increase the efficiency with which MSAs use roads and travel time. More specifically, we take the 95 MSAs with the lowest productivity residual in our preferred estimation of equation (8) and increase their productivity to the level of the MSA at the 95th percentile of the distribution of productivity and leave the five most efficient MSAs as they are. Let S_{TFP95} denote the counterfactual speed index associated with this policy. In the third experiment, we increase road intensity, that is the stock of roads per unit of travel time, in an analogous manner. Let $S_{ROADS95}$ denote the counterfactual speed index associated with this policy. In the fourth experiment we decrease aggregate MSA travel time, holding everything else constant, to the level of the fifth smallest MSA. That is, we ask what would happen if we distributed the drivers in large MSA across several smaller MSAs, holding the provision of roads at the original level. Let $S_{SCALE95}$ denote the counterfactual speed index associated with this policy.

More formally, the last three counterfactual experiments involve computing the following three counterfactual speed indices,

$$-\ln S_{TFP95} = \hat{\alpha} \ln(R_i/\text{VTT}_i) + (\hat{\alpha} - \hat{\theta}) \ln \text{VTT}_i + \hat{v}_i^{95} \quad (10)$$

$$-\ln S_{ROADS95} = \hat{\alpha} \ln(R/\text{VTT})_i^{95} + (\hat{\alpha} - \hat{\theta}) \ln \text{VTT}_i + \hat{v}_i \quad (11)$$

$$-\ln S_{SCALE95} = \hat{\alpha} \ln(R_i/\text{VTT}_i) + (\hat{\alpha} - \hat{\theta}) \ln \text{VTT}_i^{95} + \hat{v}_i. \quad (12)$$

where \hat{v}^{95} is the value of the productivity residual, \hat{v} , at the 95th percentile, $(R/\text{VTT})^{95}$ is the road intensity at the 95th percentile, and VTT^{95} denotes the analogous counterfactual value of VTT .

Table 7 describes the value of these alternative policies for our sample of the 100 largest MSAs. Each of the first four columns of the table describes one of the four policies. Each of the three panels describes the effect of our policies in a different year. In the first row of each panel we calculate the change in person weighted travel speed associated with each policy. In the second, we calculate the number of people affected by each experiment. Finally, in the third row of each panel, we calculate the change in aggregate travel time resulting from each experiment, holding distance traveled constant. All calculations in table 7 are based on the estimation of equation (8)

Table 7: Valuing policy experiments for the 100 largest MSAs in 2008

	S_{95}	S_{TFP95}	$S_{Roads95}$	$S_{Scale95}$	S
1995					
Mean Speed (km/h)	63.6	59.9	53.6	57.1	49.9
People affected(millions)	171	171	171	173	
Aggregate Δ VTT (millions of hours)	8,166	5,843	1,377	4,010	
2001					
Mean Speed (km/h)	57.6	52.7	48.9	52.2	48.8
People affected(millions)	190	190	191	192	
Aggregate Δ VTT (millions of hours)	10,350	6,219	2,450	5,791	
2008					
Mean Speed (kph)	57.3	52.3	49.7	51.4	46.5
People affected(millions)	207	205	207	207	
Aggregate Δ VTT (millions of hours)	10,300	5,679	3,073	4,871	

Notes: Authors' computations, using the 100 largest MSAs and regression results from column 6 in Table 5 for 2008, and the equivalent regressions for 2001 and 1995. We report the population of the 95 MSAs below the 95th percentile of each distribution in parentheses.

reported in table 5, column 6. Finally, column 5 of table 7 reports mean speed across the 100 largest MSAs where speed in each MSA is computed from our preferred speed index.

Table 7 suggests a number of conclusions. First, the value of increasing speed is large. Increasing speed to the 95th percentile of the distribution would save about 10.3 billion hours of travel in 2008. Common practice in the transportation economics literature is to value these hours at half the average wage of 23 dollars per hour and consider an average of 0.25 passengers in each car.²⁴ This implies a gain of slightly less than 150 billion dollars.²⁵ The second policy experiment, which involves raising the productivity of cities below the 95th percentile, would imply gains of 82 billion dollars. This suggests that, if slow MSAs can emulate fast ones in their ability to produce travel out of roads and vehicle time, large gains would also occur. Put differently, if we can understand the factors that affect the efficiency with which roads and time are employed, then large savings should be possible. Since using roads and travel time more efficiently is likely to be a lot less costly than building new roads, this suggests that research into the efficiency of transportation networks is important. Finally, increasing road intensity and decreasing aggregate travel time by

²⁴Existing estimates of value of time traveled for commuters generally center around 50% of an individual's hourly wage (Small and Verhoef, 2007). We note that the Texas Transportation Institute uses a much higher number of nearly 24 dollars per hour per person (Schrank *et al.*, 2010).

²⁵The Texas Transportation Institute (Schrank and Lomax, 2009, Schrank *et al.*, 2010) conducts an exercise in the same spirit as this one. Despite its much higher valuation of delay and a broader geographical coverage, it finds smaller numbers; 87 and 115 billion dollars for 2008 and 2009, respectively. The main reason for this difference is that they focus on major roads that measure only a small fraction of travel.

redistributing drivers to smaller MSAs in the third and fourth experiment would imply gains of 44 and 70 billion dollars, respectively.

A simple variance analysis confirms the results of these counterfactual experiments. After re-writing the determinants of speed into road intensity, scale, and productivity as in equations (10) to (12), the variance of $\ln S$ is equal to the sum of the variance of the road intensity term, $\hat{\alpha} \ln(R/\text{VTT})$, the variance of the scale term, $(\hat{\alpha} - \hat{\theta}) \ln \text{VTT}$, the variance of the productivity residual, \hat{v}_i , and twice the covariance between road intensity and scale. In accordance with the counterfactual experiments above, we find that the productivity term accounts for 59% of the variance in log speed whereas the contribution of scale is smaller at 20%. The contribution of road intensity is smaller still at 9%. The remaining 12% are accounted for by the covariance term for scale and road intensity which reflects the fact that larger MSAs have a lower road intensity.

Welfare analysis

We now conduct an alternative analysis of the value of speed in which we explicitly account for the fact that our speed index is the result of an equilibrium process. This analysis is in the spirit of the approach initially proposed by Beckman *et al.* (1956) and Vickrey (1963). We begin with equation (8) and make the substitution $-\ln \Omega_i \equiv \alpha_i \ln R_i + X_i \phi_i + v_i$. Rearranging, and suppressing the i subscript for legibility, we have

$$S = \Omega \text{VTT}^\theta. \quad (13)$$

Note that the left hand side of this expression is speed computed our preferred speed index, not minus the log of that index as in equation (8).

In an equilibrium where road access is unpriced all users will face the same unit cost of travel. Our speed index is precisely this time cost for a standardized bundle of trips. Thus, equation (13) gives the average cost of a kilometer of travel as a function of MSA aggregate travel time. Since vehicle kilometers traveled (the quantity of travel) is equal to vehicle time traveled times speed, i.e., $\text{VKT} \equiv \text{VTT} \times \frac{1}{S}$, we can write $\text{VTT} = S \times \text{VKT}$.²⁶ Substituting in equation and solving for S , we have the following supply equation for VKT :

$$S = AC(\text{VKT}) = \Omega^{\frac{1}{1-\theta}} \text{VKT}^{\frac{\theta}{1-\theta}}. \quad (14)$$

This equation is an aggregate supply curve for automobile travel in MSA i .

²⁶We keep in mind that we are using a speed index that corrects for the endogeneity of the choice of trip distance. The measure of VKT we compute here also accounts for this.

Starting from the average cost curve given in equation (14), simple calculations then imply that the marginal time cost of travel is

$$MC(v_{KT}) = \frac{S}{1-\theta}. \quad (15)$$

The marginal cost for travel reflects both the private value of a marginal kilometer of travel, and also the extent to which this marginal kilometer slows down other drivers. Thus, this marginal cost curve reflects the social cost of time lost to congestion.

In order to complete a partial equilibrium model of travel and speed, it remains only to specify the demand for travel. Write the demand for v_{KT} as $v_{KT} = \Gamma S^{-\sigma}$, where Γ is a constant and σ is the demand elasticity of v_{KT} with respect to the time cost of travel. For our purposes, it is more useful to rearrange this and write it as an inverse demand curve,

$$S = \Gamma^{\frac{1}{\sigma}} v_{KT}^{-\frac{1}{\sigma}}. \quad (16)$$

Equating average cost, equation (14), and demand, equation (16), we can solve for equilibrium v_{KT} ,

$$v_{KT}^{eq} = \left(\Gamma^{1-\theta} \Omega^{-\sigma} \right)^{\frac{1}{1-\theta(1-\sigma)}} \quad (17)$$

Optimal v_{KT} results from equilibrating demand and marginal cost

$$v_{KT}^{opt} = (1-\theta)^{\frac{\sigma(1-\theta)}{1-\theta(1-\sigma)}} v_{KT}^{eq}. \quad (18)$$

Since drivers base their decisions on average rather than marginal cost, equilibrium driving exceeds the optimal. The deadweight loss from this excessive driving is given by,

$$\int_{v_{KT}^{opt}}^{v_{KT}^{eq}} \left[\frac{1}{1-\theta} \Omega^{\frac{1}{1-\theta}} v_{KT}^{\frac{\theta}{1-\theta}} - \Gamma^{\frac{1}{\sigma}} v_{KT}^{-\frac{1}{\sigma}} \right] dv_{KT}. \quad (19)$$

Using equations (16) and (14) to write Ω and Γ in terms of $v_{TT}^{eq} = S^{eq} \times v_{KT}^{eq}$ this expression evaluates to

$$v_{TT}^{eq} \left[\left(1 - (1-\theta)^{\frac{\sigma}{1-\theta(1-\sigma)}} \right) - \frac{\sigma-1}{\sigma} \left(1 - (1-\theta)^{\frac{(\sigma-1)(1-\theta)}{1-\theta(1-\sigma)}} \right) \right]. \quad (20)$$

Normalizing by total travel time we arrive at

$$\Delta = \left[\left(1 - (1-\theta)^{\frac{\sigma}{1-\theta(1-\sigma)}} \right) - \frac{\sigma-1}{\sigma} \left(1 - (1-\theta)^{\frac{(\sigma-1)(1-\theta)}{1-\theta(1-\sigma)}} \right) \right]. \quad (21)$$

This expression gives the deadweight loss from congestion in a particular city as a proportion of the total travel time. It depends on just two parameters, the supply and demand elasticity, θ and σ .

In our preferred regression in column 6 of table 5, we estimate the average supply elasticity $\hat{\theta}$ to be 0.15. We do not estimate the elasticity of the demand for vkt here. However, Duranton and Turner (2011) suggests that the demand for vkt is highly elastic and provides a point estimate of 16.²⁷

Using these values of $\theta = 0.15$ and $\sigma = 16$, Δ in equation (21) is about 0.105. That is, the value of deadweight loss from congestion is equal to about 11% of total travel time. Taking aggregate travel time for 2008 from panel B of table 1, converting from minutes per day to hours per year, and adjusting upwards by 1.25 persons to car, we find 123 million hours per year of loss due to congestion in an average MSA in our sample. This works out to about 60 hours per person per year. If we value this time at 11.5 dollars per hour this is 690 dollars per person per year. In total, for our whole sample of 100 MSAs, the loss is about 12.3 billion hours per year or, again valuing time at 11.5 dollars per hour, 141 billion dollars per year.


The general magnitude of these estimates does not appear to be particularly sensitive to our choices of θ and σ . For example, if, contrary to the findings in Duranton and Turner (2011), we take $\sigma = 1$ so that demand for travel is much less sensitive to its price in time, the deadweight loss from congestion increases from about 123 million hours per year in an average MSA to about 170. If we double the demand elasticity to 32, this quantity decreases to about 98 million hours per year. Keeping σ at our preferred value of 16 and decreasing θ to 0.075 decreases the deadweight loss in an average MSA to about 79 million hour per year. Increasing θ to 0.225 we have a loss of 164 million hours per year. These values of θ are respectively more than two standard errors smaller and larger than our preferred estimate from column 6 table 5.

We can also use our analysis, equation (18) in particular, to calculate the extent to which equilibrium driving exceeds optimal driving. At our preferred values of θ and σ , this calculation indicates that the optimal level of vkt is only about half the equilibrium value. For values of σ close to one, the gap between equilibrium and optimal vkt decreases dramatically to about 15%. The amount of 'overdrive' is less sensitive to other changes in the parameters.

These calculations call for several comments. The first is of course one of caution. Although robust to reasonable variations in our two key parameters, our estimated deadweight loss of about 123 billion dollars annually is conservative in many respects. We only consider the time cost of

²⁷The working paper version of that work (table 6, columns 4 and 8 in Duranton and Turner, 2009) contains an explicit estimation of the demand for vkt which yields $\sigma \approx 16$.

Figure 3: Equilibrium speed and distance.



travel and ignore fuel costs and other car usage costs, e.g., maintenance, pollution and collisions for which estimates are reviewed in Parry, Walls, and Harrington (2007). We also consider only travel by residents and ignore commercial traffic for which the time and fuel costs of congestion are arguably higher.²⁸ Finally, our estimates are based only on the population of the countries 100 largest MSAs. However, these MSAs account for about two thirds of the countries population, and presumably, most of the traffic congestion. With this said, to the extent that congestion occurs outside of these 100 MSAs, our estimates are too low.

Discussion

Figure 3 illustrates our partial equilibrium model of speed and travel. In the figure, the demand curve corresponds to the demand curve given by equation (16), the average cost curve, AC_1 , corresponds to the supply curve given in equation (14), and the marginal cost curve, MC_1 , corresponds to equation (15). In this figure, equilibrium speed and VKT_1^{eq} are determined by the intersection of

²⁸To fix ideas, according to Duranton and Turner (2011), the share of trucks traffic for interstate highways in 228 large US MSAs was 13% in 2003. In the calculations of the Texas Transportation Institute the cost of truck congestion represents slightly more than a quarter of the total cost of congestion. As for fuel losses, they represent only a fraction of the time cost of congestion. To see this, consider an hour lost to drive 10 kilometers. With a car consuming 15 liters per 100 kilometers, this is only 1.5 liters of fuel or about 1.5 dollars at one dollar per liter. This is small relative to a time cost of 11.5 dollars per hour.

supply curve, AC_1 , and the demand curve. Socially optimal speed and vKT^{opt} are determined by the intersection of demand and the marginal cost curve, MC_1 .

The counterfactual experiments performed earlier in this section involve increasing the speed of travel, decreasing minutes per kilometer, holding travel constant at its equilibrium level. In this figure, the value of such an experiment is given by the hatched rectangle whose corners are the points $FCGI$. If we interpret this experiment as shifting down the supply curve in such a way that the cost to provide vKT_1^{eq} is lower by the appropriate amount, then our counterfactual experiment requires an alternative supply curve such as the curve AC_2 in the figure. For every level of vKT , travel is faster along this hypothetical supply curve than along the actual supply curve, AC_1 .

It is obvious from the figure, that the area of the rectangle is likely to overstate the benefits realized from our hypothetical improvement to supply. If we were to improve an MSA's ability to provide travel so that the supply curve dropped from AC_1 to AC_2 , then the new equilibrium level of travel would be vKT_2^{eq} . The increase in surplus associated with this improvement is not the rectangular region $FCGI$, but rather the region whose corners are $FCDE$. Therefore, much of the benefit associated with hypothetical improvements in an MSA's ability to supply travel are dissipated by increased equilibrium demand: we should not expect dramatic shifts in the speed of urban interstate highways at 4 pm to persist in equilibrium.

This means that these counterfactual experiments should be understood as informing us about the sensitivity of the supply of travel to policy interventions to increase network capacity and speed. More specifically, the counterfactual experiments above hint at a number of solutions to improve traffic speed: increase the productivity of cities in their production of travel, increase road intensity, and reduce scale. This last solution involves population dispersion and is thus not feasible. The second solution, road construction, is often proposed in policy discussions. Using data from Duranton and Turner (2012) based on estimates of road construction costs by Ng and Small (2008) and maintenance costs from the US Bureau of Transportation Statistics 2007, increasing the supply of interstate highways by 1% in all of the 100 largest MSAs has an annual cost of 1.45 billion dollars. Even with a low elasticity of the demand for vKT with respect to speed of one, the annual gains are only 874 million dollars. Taking a higher elasticity of 16 as Duranton and Turner

(2009) makes these gain lower still at 269 million dollars.²⁹ While such numbers should obviously be taken with caution, they are indicative of low returns to ‘across-the-board’ road construction at the margin (a conclusion we share with Winston and Langer, 2006). The last, and perhaps, more promising solution suggested by our counterfactual experiments above is an improvement of the management of the road network from ramp metering systems for highway access to the management of intersections, etc. While a detailed analysis of these solutions goes well beyond the scope of our paper, our findings are suggestive of large gains for slow MSAs if they can emulate fast MSAs.

Figure 3 also illustrates our partial equilibrium welfare calculation. In fact, the deadweight loss calculation performed by equation (19) is calculating the area of the shaded region ABC , the deadweight loss from excessive driving. The substance of our empirical results is that this area is large, on the order of 120 billion dollars per year. This is slightly lower than our estimates of the gains associated with raising the speed in MSAs to the level of the fastest ones as represented by the hatched rectangle $FCGI$.

While the counterfactual experiment performed above suggest some supply responses to congestion, the welfare calculation points to another direction: the management of demand. The large deadweight loss estimated above arises because of a wedge between what drivers pay to access the road (the average cost of driving) and the social cost they cause (the marginal cost of driving). Lowering this wedge and the resulting excess driving would lead to large social gains without having to incur extra expenditure on road construction.

Interestingly, figure 3 also points out that we can use our estimations to calculate an optimal congestion tax. This tax is given by the height of the segment AH . With a uniform tax of $|AH|$ per kilometer, we shift up the average cost curve AC_1 so that the intersection of $AC_1(v\kappa\tau_1^{eq}) + |AH|$ is equal to the demand at $v\kappa\tau_1^{eq}$. To calculate this tax, we evaluate the difference between supply and marginal cost at equilibrium $v\kappa\tau$. Letting τ^* denote this optimal tax, we have

$$\tau^* = \frac{1}{1-\theta} \Omega^{1/(1-\theta)} \left[\Gamma^{1-\theta} \Omega^{-\sigma} \right]^{1/(1-\theta(1-\sigma))} \quad (22)$$

²⁹The welfare gain associated with road construction is equal to the reduction in speed caused by road construction times initial $v\tau$ plus the triangle associated with the increase in $v\tau$. Using (8), (13), and the demand for travel time $v\tau = \Gamma S^{1-\sigma}$, we find that the elasticity of (inverse) speed with respect to roads is equal to $-\alpha \times \frac{1/\theta}{1/\theta+\sigma-1}$ and the elasticity of $v\tau$ with respect to roads is $-\alpha \times \frac{1-\sigma}{\sigma}$. Since we consider an increase in interstate highways, we take a value of α of 0.067 which corresponds to the coefficient for this class of roads estimated in column 8 of table 8. For θ , we retain our preferred value of 0.15 as above.

Using equation (14) and (16) we can write Ω and Γ in terms of observed equilibrium vkt and S . This gives,

$$\tau^* = \frac{1}{1-\theta} \left(S^{eq} 1-\theta V^{eq\theta} \right)^{1/(1-\theta)} \left[\left(S^{eq\sigma} V^{eq} \right)^{1-\theta} \left(S^{eq} 1-\theta V^{eq\theta} \right)^{-\sigma} \right]^{1/(1-\theta(1-\sigma))} \quad (23)$$

Evaluating this tax at equilibrium quantities with $\theta = 0.15$ and $\sigma = 16$ gives a tax of about 0.2 minutes per kilometer of travel. If we continue to value time at 11.5\$/hour, this works out to a tax of about four cents per kilometer.

That is, our empirical results indicate that travel in our sample of 100 MSAs is sufficiently congested that even a small constant congestion tax applied to all roads would lead to welfare gains on the order of 100 billion dollars per year. However, we note that our estimations are of course too coarse to provide detailed guidance about congestion taxes but, as just shown, they are useful to establish some magnitudes.³⁰.

7. Conclusion

Road transportation accounts for nearly one fifth of the budget for an average us household and about 2.5 hours per day for an average us driver in 2008. In a typical year, the us spends about 200 billion dollars on road construction and maintenance. Congestion naturally arises in the course of turning these resources into travel. Despite a distinguished history of research in transportation economics, the extant literature has struggled to measure the value of investments in road transportation or the social cost of road congestion.

We make three advances in this agenda. First, we develop an econometric methodology and data to identify city level supply curves for speed. With these supply curves, we construct an index of travel speed for each of the largest us cities. This index number, for the first time, provides a theoretically founded measure of the efficiency with which different locations produce transportation.

Our investigation of the determinants of our speed index suggests that the production of transportation at the city level is subject to slight decreasing returns to scale. This finding provides an empirical basis for the positive relationship between city size and congestion costs that is axiomatic

³⁰Importantly, we consider an average θ of 0.15 but this average is unlikely to be effective for peak hour travel where most of the congestion occurs. With $\theta = 0.6$, optimal 'peak hour tax' is much larger at around a dollar per kilometre. However having a $\theta = 0.6$ for a quarter of travel and $\theta = 0$ the rest of the time yields a total revenue that is not very different from having an average $\theta = 0.15$ that applies all the time

in nearly all extant theoretical models of city size. It also reveals the role of roads, travel time, and unobserved productivity. The presence of economically important variation in the unobserved productivity of cities at producing transportation suggests the possibility of large gains in efficiency if slow cities are able to emulate fast cities. We conduct a rudimentary investigation of city characteristics associated with efficiency. Our findings are unsurprising: dense centralized cities are less efficient. Refining this investigation is an important topic for further research and may provide an empirical basis for the design of cities where travel is provided more efficiently.

Our investigation of the efficiency of travel amounts to the comparison of equilibrium speed with counterfactual, out-of-equilibrium scenarios. In this, it resembles the widely known travel cost indices published by TTI. Fundamentally, these counterfactual cases do not evaluate an economically meaningful definition of congestion. Such a measure of congestion is intrinsically subtle, and requires the calculation of the deadweight loss incurred at equilibrium levels of travel as opposed to optimum. Our estimates allow us a rough, and highly aggregated calculation of the deadweight loss from congestion. This calculation indicates that the losses from congestion are probably even larger than those suggested by counterfactual supply improvements. Since equilibrium responses dissipate much of the benefit from supply improvements, our results suggest that the largest gains in transportation policy are probably to be had by managing demand.

References

- Akerberg, Daniel, C. Lanier Benkard, Steven Berry, and Ariel Pakes. 2007. Econometric tools for analyzing market outcomes. In James J. Heckman and Edward E. Leamer (eds.) *Handbook of Econometrics*, volume 6A. Amsterdam: North-Holland, 4271–4276.
- Bartik, Timothy. 1991. *Who Benefits from State and Local Economic Development Policies?* Kalamazoo (MI): W.E. Upjohn Institute for Employment Research.
- Baum-Snow, Nathaniel. 2007. Did highways cause suburbanization? *Quarterly Journal of Economics* 122(2):775–805.
- Beckman, Martin, C. B. McGuire, and Christopher B. Winsten. 1956. *Studies in the Economics of Transportation*. New Haven, CT: Yale University Press for the Cowles Commission for Research in Economics.
- Bombardini, Matilde and Francesco Trebbi. 2011. Gains from distortions in congested city networks. Processed, University of British Columbia.
- 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):587–633.
- Combes, Pierre-Philippe and Miren Lafourcade. 2005. Transport costs: Measures, determinants, and regional policy implications for France. *Journal of Economic Geography* 5(3):319–349.
- Deweese, Donald N. 1979. Estimating the time costs of highway congestion. *Econometrica* 47(6):1499–1512.
- Duranton, Gilles, Peter M. Morrow, and Matthew A. Turner. 2011. Roads and trade: Evidence from the us. Processed, University of Toronto.
- Duranton, Gilles and Matthew A. Turner. 2009. The fundamental law of road congestion: Evidence from US cities. NBER Working Paper 15376.
- Duranton, Gilles and Matthew A. Turner. 2011. The fundamental law of road congestion: Evidence from US cities. *American Economic Review* 101(6):2616–2652.
- Duranton, Gilles and Matthew A. Turner. 2012. Urban growth and transportation. Forthcoming, *Review of Economic Studies*.
- 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.
- Fox, Jeremy T. and Valérie Smeets. 2011. Does input quality drive measured differences in firm productivity? *International Economic Review* 52(4):961–989.
- Glaeser, Edward L. and Matthew Kahn. 2001. Decentralized employment and the transformation of the American city. *Brookings-Wharton Papers on Urban Affairs* :1–47.
- 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.
- Greenshields, Bruce D. 1935. A study of traffic capacity. *Highway Research Board Proceedings* 14(0):448–477.

- Keeler, Theodore E. and Kenneth A. Small. 1977. Optimal peak-load pricing, investment, and service level on urban expressways. *Journal of Political Economy* 85(1):1–25.
- Levinsohn, James and Amil Petrin. 2003. Estimating production functions using inputs to control for unobservables. *Review of Economic Studies* 70(2):317–342.
- McFadden, Daniel. 1973. Conditional logit analysis of qualitative choice behavior. In Paul Zarembka (ed.) *Frontiers in econometrics*. New York and London: Academic Press, 105–142.
- McFadden, Daniel. 1974. The measurement of urban travel demand. *Journal of Public Economics* 3(4):303–328.
- Meyer, John R., John F. Kain, and Martin Wohl. 1965. *The urban transportation problem*. Cambridge (MA): Harvard University Press.
- Michaels, Guy. 2008. The effect of trade on the demand for skill - Evidence from the Interstate Highway System. *Review of Economics and Statistics* 90(4):683–701.
- Ng, Chen Feng and Kenneth A. Small. 2008. Tradeoffs among free-flow speed, capacity, cost, and environmental footprint in highway designs. Processed, University of California at Irvine.
- Parry, Ian W.H., Margaret Walls, and Winston Harrington. 2007. Automobile externalities and policies. *Journal of Economic Literature* 45(2):1335–1353.
- Schrank, David and Tim Lomax. 2009. Urban mobility report 2009. Technical report, Texas Transportation Institute.
- Schrank, David, Tim Lomax, and Shawn Turner. 2010. Urban mobility report 2010. Technical report, Texas Transportation Institute.
- Small, Kenneth A. and Erik T. Verhoef. 2007. *The Economics of Urban Transportation*. New York (NY): Routledge.
- Syverson, Chad. 2011. What determines productivity. *Journal of Economic Literature* 49(2):326–365.
- Train, Kenneth. 2009. *Discrete Choice Methods with Simulation*. Second edition. Cambridge (MA): Cambridge University Press.
- United States Bureau of Transportation Statistics. 2007. *Transportation Statistics Annual Report 2007*. Washington, DC: US Government printing office.
- United States Bureau of Transportation Statistics. 2010. *Transportation Statistics Annual Report 2010*. Washington, DC: US Government printing office.
- United States Department of Transportation, Federal Highway Administration. 2003a. *Archive Highway Performance Monitoring System (HPMS) Data Item Descriptions 1993 - 1998*. Washington, DC.
- United States Department of Transportation, Federal Highway Administration. 2003b. *Archive Highway Performance Monitoring System (HPMS) Data Item Descriptions 1982-1987*. Washington, DC.
- United States Department of Transportation, Federal Highway Administration. 2005. *HPMS Field Manual*. Washington, DC.
- Vickrey, William S. 1963. Pricing in urban and suburban transport. *American Economic Review (Papers and Proceedings)* 53(2):452–465.

Winston, Clifford and Ashley Langer. 2006. The effect of government highway spending on road users' congestion costs. *Journal of Urban Economics* 60(3):463–483.

Appendix A. Data

Consistent MSA definition: MSAs are defined as aggregations of counties. We use the 1999 definition. The NHTS data contain either a county identifier or an MSA identifier consistent with our MSA definition. Our variables either come from county level data or from maps that we overlay with a map of 1999 counties.

NHTS data (trip-level data): Our trip-level data are from the 1995-1996 National Personal Transportation Survey (NPTS) and its successors, the 2001-2002 and 2008-2009 National Household Transportation Surveys (NHTS). The surveys are sponsored by various agencies at the US Department of Transportation. Detailed documentation on the NHTS sample design and data content is available on the NHTS website (<http://nhts.ornl.gov/>). The data are intended for use by transportation planners in governmental agencies and aim at providing reliable, representative and comprehensive micro-data about the daily travel of Americans.

HPMS data (road infrastructure): To estimate our travel production function, we rely on the Highway Performance Monitoring System (HPMS) data for 1995, 1996, 2001, 2002, and 2008. These data are collected and maintained by the US Federal Highway Administration in cooperation with many sub-national government agencies. Documentation is available in several reports from the Federal Highway Administration (DOT, 2003a, 2003b, and 2005). See also Duranton and Turner (2011).

For the interstate highway system, the HPMS records number of lanes, length, average annual daily traffic (AADT), and county. By construction, road segments do not cross county borders. For segments in urbanized areas, the HPMS also provides an urbanized area code. Since MSAs are county based units, these data allow us to calculate VKT for the urbanized and non urbanized area interstate systems by MSA.

Population and employment data: Population data for 1995-1996, 2001-2002, and 2008-2009 is obtained from annual population estimates provided by the US census for these years. These estimates are themselves based on interpolations and extrapolations of population counts made

by the US census. In some regressions we also use population data dating back to 1920, the first census year which allows us to retain our samples of MSAs. We also use employment data from the County Business Patterns to build two variables for our exploration of the determinants of speed. The first is that of manufacturing employment. The second is an exogenous measure of MSA employment growth which interacts the sectoral composition of economic activity in an MSA in 1980 with the national growth of these sectors between 1980 and 2000.

Climate: We use two measures of temperature in US MSAs taken from the data used by Burchfield, Overman, Puga, and Turner (2006).

Urban form: We have available the data underlying Burchfield *et al.* (2006). These data provide fine scale employment data from 1994 zipcode business patterns and from 1990 tract level census data: for every 990 meter \times 990 meter cell in a regular grid covering the whole of the continental US, these data report imputed 1990 population and 1994 employment. These data allow us to calculate the measures of urban form used in the first three columns of table 6; employment centrality, population centrality, and employment-residence mismatch.

The population centrality measure reports the largest share of population that occurs in any ring of radius 20 kilometers whose center lies in the MSA. The employment centrality measure is identical, but is based on employment rather than population.

The mismatch variable provides an aggregate measure of the extent to which people do not live where jobs are (though not of the extent to which people do not live where *their* jobs are). Specifically, let e_i be employment in cell i and p_i be population. Our mismatch measure is

$$\frac{\sum_{i \in \text{MSA}_j} |p_i - \alpha e_i|}{\sum_{i \in \text{MSA}_j} p_i},$$

where α is the inverse share of employment in total population ($\alpha \sum_{i \in \text{MSA}_j} e_i \equiv \sum_{i \in \text{MSA}_j} p_i$). Since we are normalizing by population, this is a per capita measure of mismatch.

Instruments for contemporary interstate highways: Our measures of the 1947 interstate highway plan and the 1898 railroad network are taken from Duranton and Turner (2011) and are documented there. Further discussion of the 1947 highway plan is available in Baum-Snow (2007), Michaels (2008), and Duranton and Turner (2012).

Texas Transportation Institute travel time indices: Until 2009 (using 2008 data), the annual TTI travel time indices were constructed using counts of vehicles on interstate highways and major

urban roads from the average annual daily traffic (AADT) item measured in the HMPS 'Sample' data. More specifically, the TTI methodology subjected the AADT variable to a series of transformations using external information and turned it into a measure of speed. This measure of speed was then 'operationally' adjusted for several factors (e.g., the existence of a ramp metering system to access interstate highways) to obtain a final measure of speed by city. These quantities were then compared to a measure of free flow traffic to infer a travel cost index by city and their travel cost index for the US.

In 2010 (to exploit 2009 data), the TTI paired with INRIX, a leading provider of traffic information, directions, and driver services. Real speed data is now measured using information provided by location devices from vehicles operated by various fleet operators and by the smart phones of voluntary individuals for a subsample of segments of roads covered by the HMPS 'Sample' data.

Importantly, the TTI data cover only urbanized areas. This implies that areas for which the TTI computes its index are smaller than the MSAs that we use. This said, the 'urbanized' part of MSAs will host the vast majority of their residents and jobs. In order to compare our MSA based index with TTI, we merged several TTI urbanized areas to approximate MSAs more closely. To do this, we weighted urbanized areas by their population. The full list of merged TTI urbanized areas is: Washington DC and Baltimore, San Francisco-Oakland and San Jose, Cleveland and Akron, Los Angeles and Riverside-San Bernardino, Denver and Boulder, Greensboro and Winston-Salem, Boston and Worcester. Finally, for Fort Myer-Cape Corral MSA, the only TTI data point is for Cape Corral and for Norfolk-Virginia Beach-Newport News MSA, the only TTI data point is for Virginia Beach. Overall we can match 47 of the 50 largest MSAs and 71 of the 100 largest using data for the 90 urban areas for which the TTI reports an index for 2008.

Appendix B. Computing city aggregate VKT and VTT from the NPTS

As explained in the text, the general idea behind our method for computing MSA totals like VKT and VTT from trip level data is to multiply sample averages by population estimates. With the sample that we use however, this simple multiplication would lead to biased estimates, because of the way in which we removed outliers.

The sample of trips we use consists of all trips by privately owned vehicle entered by a driver. For distance, duration, and speed, we identify the 0.5% of trips with the lowest value and the 0.5% of trips with the largest value and we eliminate all trips by an individual who entered any of these

trips. We remove individuals instead of trips because we want to match population totals using sample averages at the individual level. Including individuals with an incomplete trip schedule would bias our estimated averages downwards. However, before removing the individuals who entered a trip with an extreme value, we compute the share of individuals driving at least one trip by privately owned vehicle. We then multiply this share by MSA population to obtain an estimate of the total number of people, in each MSA, who drive at least one trips by privately owned vehicle. To obtain an estimate of VTT and VKT, we just take the average duration and distance traveled from the clean sample of privately owned vehicle drivers and scale it up by a factor such that the sample size matches our estimate of the population driving at least one trip by privately owned vehicle.

As an example, suppose that in a given MSA, the NHTS contains 5 individuals driving at least one trip by privately owned vehicle, out of 10 individuals in our sample for that MSA. If MSA population is 100, then our estimate of the number of individuals driving at least one trip by privately owned vehicle in this MSA is $(5/10) \times 100 = 50$. If, for the individuals who remain in the sample after cleaning, the average distance driven (distance driven for an individual is the sum of the distance of each trip by privately owned vehicle for which they are the driver) is 20 kilometers, then our estimate of VKT in this MSA is $50 \times 20 = 1000$ vehicle kilometers traveled.

Appendix C. Robustness tests for determinants of speed

Table 8 confirms our findings using different samples of MSAs, years of data, and variables. In column 1 we duplicate our preferred estimation of column 6 of table 5 using only the 50 largest MSAs to find virtually the same results (but a higher R^2 given the greater precision of our index in larger MSAs). In columns 2 and 3, we also duplicate our preferred regression from table 5 but use 100 MSAs for 2001 and 1995, respectively. While the results for 2001 are virtually the same as those for 2008, the coefficient on lanes for 1995 is insignificant and the R^2 is much lower. In columns 4 and 5, we use a log speed index computed from the 2008 and 2009 TTI travel cost indices, respectively, as dependent variables. The results for the 2008 TTI index are close to those of our preferred estimation. With the 2009 TTI index, the coefficient on lanes is small and insignificant. Column 6 substitutes population for VTT finds results similar to the other specifications. In column 7, we use log interstate highway lane kilometers as a measure of roads. Excluding smaller roads leads to a smaller coefficient. Finally in column 8 we control separately for log lane kilometers

Table 8: The determinants of speed, alternative years, samples, and variables

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	S_{IV2} 2008	S_{IV2} 2001	S_{IV2} 1995	TTI_{2008} 2008	TTI_{2009} 2008	S_{IV2} 2008	S_{IV2} 2008	S_{IV2} 2008
log lane (total)	0.11 ^a (0.039)	0.098 ^a (0.031)	0.052 (0.044)	0.12 ^a (0.022)	0.029 (0.023)	0.085 ^a (0.032)		
log lane (interstate highways)							0.067 ^a (0.014)	0.066 ^a (0.014)
log lane (major urban roads)								0.052 ^c (0.029)
log VTT	-0.16 ^a (0.033)	-0.15 ^a (0.030)	-0.096 ^b (0.042)	-0.18 ^a (0.019)	-0.075 ^a (0.019)		-0.11 ^a (0.013)	-0.15 ^a (0.032)
log population						-0.13 ^a (0.031)		
Observations	50	100	100	71	78	100	99	99
R ²	0.54	0.40	0.23	0.71	0.64	0.36	0.47	0.49

Notes: OLS regressions with a constant in all columns. Robust standard errors in parentheses. *a*, *b*, *c*: significant at 1%, 5%, 10%.

of interstate highways and major urban roads. We find a slightly higher coefficient for interstate highways than for major urban roads. Importantly, the sum of these two coefficients is very close to the coefficient for both types of roads considered jointly in our preferred regression.

Table 9 further confirms our findings using alternative estimation techniques. As argued above, vehicle travel time is expected to be determined simultaneously with travel speed. To deal with this problem, in column 1 we instrument for VTT using population. As might be expected, population is a strong predictor of VTT. The estimated coefficient on VTT is -0.14 , only marginally higher than (and statistically undistinguishable from) its OLS counterpart. One may worry that the exclusion restriction associated with this regression is not satisfied since population in 2008 may be correlated with traffic speed through variables other than roads and VTT. This worry may not be as important as it seems. First, Duranton and Turner (2012) document that interstate highways have only a modest effect on urban growth. Second, to explore this question further, in column 2 we duplicate the regression of column 1 but also controls for population growth between 1980 and 2008. The results remain the same. To confirm this finding, in column 3 we use a more exogenous proxy for urban growth, namely expected population growth between 1980 and 2000. This proxy is computed by interacting the sectoral composition of employment in MSAs in 1980 and employment

Table 9: The determinants of speed, 100 MSAs in 2008 IV regressions

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	TSLS	TSLS	TSLS	TSLS	TSLS	TSLS	TSLS	TSLS	TSLS	LEVPET
log lane	0.093 ^a (0.032)	0.085 ^a (0.032)	0.078 ^b (0.034)							0.10 ^b (0.046)
log lane (IH)				0.083 ^a (0.020)	0.087 ^a (0.020)	0.086 ^a (0.021)	0.078 ^a (0.020)	0.082 ^a (0.019)	0.079 ^a (0.020)	
log lane (MRU)					0.056 ^c (0.029)	0.049 (0.030)		0.042 (0.029)	0.033 (0.031)	
log VTT	-0.14 ^a (0.030)	-0.13 ^a (0.030)	-0.12 ^a (0.032)	-0.12 ^a (0.018)	-0.17 ^a (0.035)	-0.16 ^a (0.038)	-0.11 ^a (0.017)	-0.15 ^a (0.033)	-0.13 ^a (0.037)	-0.11 ^a (0.028)
Intrumented										
log VTT	X	X	X				X	X	X	
log lane (IH)				X	X	X	X	X	X	
Instruments										
log pop. 2008	X	X	X				X	X	X	
log highways 1947				X	X	X	X	X	X	
log railroads 1898				X	X	X	X	X	X	
Controls										
$\Delta_{1980-2008}$ log pop.		X								
$E\Delta_{1980-2000}$ log pop.			X			X			X	
log pop. 1950						X			X	
log pop. 1920						X			X	
Observations	100	100	100	95	95	95	95	95	95	300
Overid. p-value	-	-	-	0.26	0.37	0.45	0.26	0.33	0.48	-
First-stage stat.	345	315	269	31.1	31.2	23.7	20.4	20.0	15.7	-

Notes: Regressions with a constant in all columns. Robust standard errors in parentheses. *a, b, c*: significant at 1%, 5%, 10%. Dependent variables is $-\log S_{IV2}$ in all columns.

growth for these sectors between 1980 and 2000.³¹ The results are still very close to the previous estimates.³²

In column 4, we turn to the simultaneous determination of roads and travel speed. Unfortunately, we do not know of any good exogenous predictor of lane kilometers of both interstate highways and major urban roads. However, we can follow Duranton and Turner (2011, 2012) and use the 1947 plan of interstate highways and a map of 1898 railroads to predict contemporaneous lane kilometers of interstate highways in US MSAs. Appendix A provides further details about the

³¹This variable is inspired by Bartik (1991) and further details about its construction can be found in Duranton and Turner (2011).

³²Overall the finding of a slightly less negative effect of travel time on speed in TSLs relative to OLS after controlling for population growth is consistent with higher travel speed having a modest positive effect on population.

construction of these two instruments. As can be seen from the results, these two instruments are good predictors of contemporaneous lane kilometers of interstate highways and pass the appropriate overidentification test. The coefficient on lane kilometers of interstate highways is 0.082. This is slightly above its corresponding OLS estimate of 0.067 in column 7 of table 8. In column 5, we estimate the same regression but add lane kilometers of major urban roads as control. The sum of the two lane coefficients is 0.14, slightly above the sum of the corresponding OLS estimates of 0.12 (in column 8 of table 8). Again, our instruments could be correlated with travel speed through the error term. Cities that received more railroads during the 19th century or cities that were allocated more roads in 1947 may be different in systematic ways. In particular, they were bigger at the time and thus may be spatially organised in a different way and may have a different transportation network relative to more recent cities. In turn, that might affect travel speed. To preclude these correlations, in column 6 we introduce 1920 population (the closest year to 1898 for which we can get population estimates), 1950 population, and expected population growth between 1980 and 2000 as before. Introducing these controls in column 6 changes close to nothing to the estimates of column 6. The marginally significant coefficient on major urban roads is now marginally insignificant.

In columns 7 to 9, we repeat the same strategy as in columns 4 to 6 but now instrument for both lane kilometers of highways and vehicle travel time. The results remain the same as before. Finally in column 10, we use a completely different instrumenting approach which makes use of all three cross-sections of data (and changes in the roadway prior to 1995) and implement the estimation technique for productivity suggested by Levinsohn and Petrin (2003). It is true that, unlike firms, cities do not invest in roads in response to positive demand shocks to maximize profit (be it only because roads are funded mainly by the federal government). Nonetheless we expect road provision to respond to changes in travel conditions. While suggestive, our Levinsohn-Petrin results are very close to our OLS and TSLS estimates.