The fast, the slow, and the congested: Urban transportation in rich and poor countries

Prottoy A. Akbar[*†] *Aalto University*

Victor Couture[*§] *University of British Columbia*

Gilles Duranton[*‡] *University of Pennsylvania*

Adam Storeygard[*¶] *Tufts University*

August 2023

Abstract

We assemble a new global database on motor vehicle travel speed in over 1,200 large cities in 152 countries. We then estimate comparable city-level indices of travel speed and congestion. Most of the variation in urban travel speed is across countries, not within. National income per capita explains most of this cross-country variation in speed. In rich countries, urban travel is roughly 50% faster than in poor countries. To investigate the link between economic development and mobility, we develop an urban model with endogenous travel, road infrastructure, and land area. The model provides an exact decomposition of how city size, infrastructure, and topography contribute to explaining why urban travel is faster in richer countries. We find that richer countries are faster, mainly because their cities have more major roads and wider land areas. These effects operate by increasing uncongested speed, not by reducing congestion.

Key words: urban transportation, roads, traffic congestion, travel speed determinants, cities
JEL classification: R41, O18

[*]This work is supported by the World Bank, the Social Science and Humanities Research Council, the Transportation Economics in the 21st Century Initiative of the NBER and U.S. Department of Transportation, the Zell Lurie Center for Real Estate at the Wharton School, and the Center for Innovative Data in Economic Research at UBC. We thank Maureen Cropper, Nicole Fortin, Ed Glaeser, Gabriel Kreindler, Thomas Lemieux, Nathan Nunn, Ben Olken, Palak Suri, Nick Tsivanidis, David Weil, and seminar and conference participants for many useful comments, Yizhen Gu and Ben Zhou for a helpful replication of our results using Chinese data, Jessie Handbury for code, Nicolas Gendron-Carrier, Marco Gonzalez-Navarro and Matt Turner for sharing transit data, and Lin Fan, Cindy Feng, Xinyu Ma, Elsie Peng, Serena Xu, William Zheng, and Xianya Zhou for research assistance. The views expressed here are those of the authors and not of any institution with which they may be associated.

[†]Aalto University and Helsinki Graduate School of Economics (email: prottoy.akbar@aalto.fi); https://www.prottoyamanakbar.com/.

[§]Vancouver School of Economics, University of British Columbia (email: victor.couture@ubc.ca); https://www.victorcouture.org/.

[‡]Wharton School, University of Pennsylvania (email: duranton@wharton.upenn.edu); https://real-estate.wharton.upenn.edu/profile/21470/.

[¶]Department of Economics, Tufts University (email: Adam.Storeygard@tufts.edu); https://sites.google.com/site/adamstoreygard/.

1. Introduction

Using data from a popular navigation app, we investigate the determinants of urban travel speed, with a focus on explaining differences between rich and poor countries. Figure 1 shows the average travel speed of motor vehicles throughout the day in the fastest (Flint, United States), slowest (Dhaka, Bangladesh), and most congested (Bogota, Colombia) cities in the world. The figure demonstrates the wide dispersion in travel speed across cities. The fastest city is three times faster than the slowest, and the most congested city experiences travel speeds that are 50% slower during peak hours that last for almost the entire working day.

We find that improved urban mobility is an important and previously undocumented feature of economic development. Most of the wide variation in city travel speed in our data is across countries, not within. A large share of this cross-country variation can be explained by national income per capita. Urban travelers in rich countries experience speeds roughly 50% faster than travelers in poor countries. The top panel of figure 2 illustrates the strong positive relationship between travel speeds and GDP per capita for a large cross section of countries. Cities in the United States, the fastest country in the world, are on average twice as fast as cities in Bangladesh, the slowest country. As shown by the middle panel of figure 2, richer countries are moderately more congested. However, as illustrated in the bottom panel of the same figure, this effect is dominated by faster uncongested speed.

Figure 1: Travel speed throughout the day: The fast, the slow, and the congested

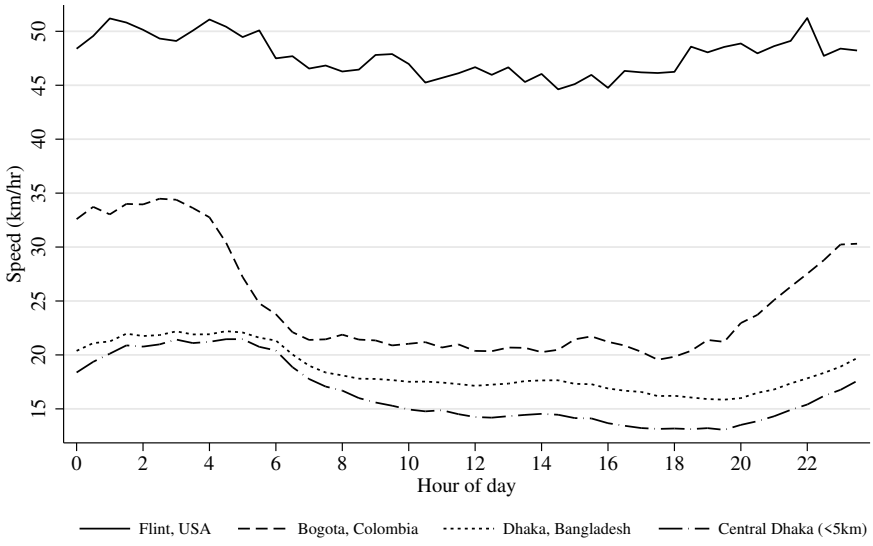

Mean speed for trips with length between 5 and 10 kilometers. Dhaka is the slowest, Flint is the fastest, and Bogota is the most congested city in the world. Central Dhaka refers to trips that take place on average within 5 kilometers of the center of Dhaka. See section 2 for further details.

Figure 2: Speed, congestion, and uncongested speed vs. GDP per capita

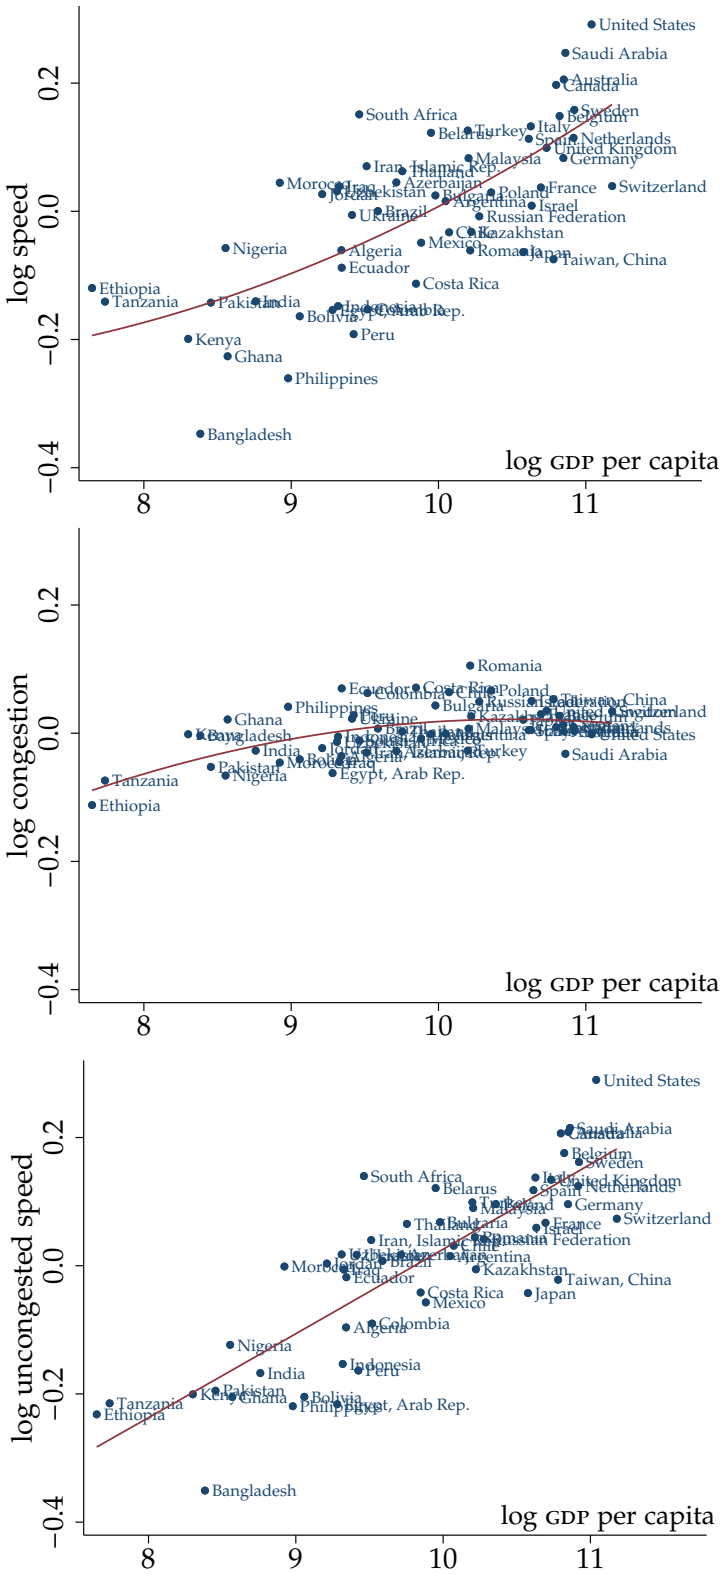

Each point in each scatter is a country. In the top panel, the vertical axis position is the mean of our city speed index, estimated in column 4 of Table 1. The vertical axis in the middle panel is the corresponding congestion index, and in the bottom panel it is the uncongested speed index. The horizontal axis in each scatter is 2018 GDP per capita at PPP in 2017 international dollars. The sample is all 54 countries with GDP data and at least three cities in the baseline sample we use below. The trend line in each panel is a quadratic fit.

To generate these findings and others described below, we produce a new global transportation database, our first contribution. We delineate more than 1,200 large cities in 152 countries that contain 96% of the world's urban population outside China. To produce information on motor vehicle travel speed, we rely on more than half a billion trips simulated on Google Maps. We then estimate city-level speed and congestion indices as in Akbar, Couture, Duranton, and Storeygard (2023) for India. Our database also includes a rich set of city attributes, comparable across countries.[1]

To interpret the link between economic development and mobility, we develop an urban model with endogenous travel speed, road infrastructure, and land area, our second main contribution. Aside from a standard road elasticity, this model introduces two distinct new parameters to capture two important features of the data. A 'spatial scale' elasticity accounts for the increase in speed as a road network of a given length expands into a larger land area. An 'urban crowding' elasticity accounts for lower speeds as a higher population competes with motorized vehicles for road space. In our model, cities in richer countries invest more in fast roads and have lower population density because tax revenues to fund the roads and housing consumption both rise with income. Our model delivers our key estimating equations, a speed and a congestion regression, which we estimate empirically. We validate our findings in three different ways. First, our rich micro-data allow us to estimate the determinants of speed across countries, across cities within a country, and across trips within a city. Second, we verify a number of over-identifying restrictions of our model. Third, we compare the parameters we recover with estimates from the literature.

Our model also delivers an exact decomposition of how the size, infrastructure, and topography of the city contribute to explaining why urban travel is faster in richer countries. This decomposition can be estimated using linear regressions as in Gelbach (2016). The contribution of each city attribute is proportional to the speed elasticity and the income elasticity of that attribute. The mileage of major roads and the land area of the city explain most of the relationship between speed and GDP per capita. That is, urban travel speed is faster in cities with more major roads and larger land areas, and cities in richer countries have more major roads and larger land areas. The correlation between the size of the city population and GDP per capita is relatively small. Therefore, the size of the city population contributes less to explaining the speed difference between rich and poor countries, even though city population is an important determinant of the speed of urban travel.[2] Our

[1] In terms of global datasets of cities with broad coverage, we know only of Angel *et al.*'s Atlas of Urban Expansion (first edition 2012, updated 2016) and Ahlfeldt, Albers, and Behrens's (2022) Prime Locations data. Angel *et al.* (2016) focus mostly on land use over time. Ahlfeldt *et al.* (2022) propose a rich set of city characteristics, but do not measure the mobility outcomes on which we focus. The coverage of these datasets is also much more limited than ours with a stratified sample of only 200 cities for Angel *et al.* (2016) and a sample of 125 large cities for Ahlfeldt *et al.* (2022).

[2] Similarly, cities in richer countries do not have more grid-like networks or a more favorable topography. So while these attributes also affect urban speed, they do not explain any of the speed advantage of rich countries.

model also accounts for the contribution of each city attribute to congestion separately from that to uncongested speed. For instance, more populous cities are slower primarily because urban crowding makes travel slower at all times and, to a lesser extent, because they are more congested.

More generally, our methodology allows us to investigate the speed of urban travel in any city or country. We illustrate this versatility with two additional decompositions. We show that the US is faster than other OECD countries because its cities have lower populations, wider areas, more major roads, and more grid-like road networks. Bangladesh is slower than other poor countries because its cities are crossed by more water bodies, are more populous, and have fewer major roads. Remarkably, the six city attributes in our parsimonious empirical model account for almost all the speed difference between the US and other OECD countries, and between Bangladesh and other poor countries. While infrastructure is most important in explaining why rich countries are faster than poor countries, city population and land area are most important in explaining why the US is so fast, and topography is most important in explaining why Bangladesh is so slow.

Our investigation is important for three reasons. First, urban transportation policy is a major concern worldwide, and our results emphasize several effects relevant for policy.[3] Our main conclusion is that, provided countries continue to develop their road infrastructure and expand physically as they become richer, we expect them to experience improved urban travel speeds. This conclusion stems from the importance of major roads as a determinant of urban travel speed and the strong association between economic development and road construction. Interestingly, these mobility improvements occur because these roads are intrinsically faster, not because the provision of more major roads reduces congestion.[4] In turn, our findings imply that while pricing urban travel congestion in developing countries has other merits, it will not bridge their urban mobility gap with richer countries.

Second, cities exist to reduce the distance between people and facilitate interactions. However, many urban dwellers do not live within walking distance of good jobs, good schools, productive peers, and attractive entertainment options. Instead, they depend on fast and reliable transportation to access such opportunities and benefit from urban interactions. As an example, Borker (2020) finds that women without safe transportation options in Delhi end up attending worse colleges. More generally, the cost of travel is a fundamental quantity in every theory of cities. Fast motorized travel is a necessary condition for the existence of modern metropolitan areas (Heblich, Redding, and Sturm, 2020, Glaeser, 2020). A walking

[3]Until now, knowledge about urban transportation in many parts of the world has been mostly speculative and based on anecdotal or scattered evidence (Gwilliam, 2003). Our main finding that urban mobility improves with economic development, for example, runs against most past priors, while our findings on the importance of city roadway, population, and land area also explain the mobility challenges faced by developing cities.

[4]Since slow travel is costly regardless of its cause, existing cross-city and cross-country indices emphasizing congestion, like TomTom and Inrix, are missing most of the variation in speed. Other advantages of our indices relative to such offerings include studying many more cities by leveraging Google Maps' global reach, and proposing a transparent methodology consistent with an economic model of urban travel.

city on a medieval scale cannot support the complex production networks and diversity of consumption amenities of modern cities. As a result, the kind of motorized travel we study here—which includes private vehicles and public transit like buses that share the main road network—accounts for a vast majority of trips worldwide.

Third, despite the large monetary and time resources devoted to urban transportation by households and governments, little is known about how most cities fare in terms of mobility. Urban policy makers rarely know how their city performs relative to peer cities and why. Even in developed countries, travel surveys are only conducted every five to ten years with limited geographic comparability within each country and none outside. This paper takes a step toward improving the availability of data in transportation. Although our data cannot fully substitute for a traditional travel survey given its lack of information about travelers and their motivation for traveling, it provides detailed information about travel speed and congestion, and their determinants. This data is available at scale, and comparable across cities and countries, and even across different parts of cities.

A large literature focuses on the measurement of vehicular congestion for specific road segments (see Small and Verhoef, 2007, for a review and Anderson and Davis, 2020, or Bou Sleiman, 2022, for examples of more recent work), but fewer studies focus on mobility outcomes for entire cities. For a single city, Akbar and Duranton (2022) for Bogotá and Kreindler (2023) for Bangalore are the main exceptions.[5] For multiple cities, Conwell, Eckert, and Mobarak (2022) measure the distance from city centers that drivers and transit riders can reach in a specified amount of time in the US and various European countries. Otherwise, previous work is limited to single countries such as the US (Couture, Duranton, and Turner, 2018) or India (Akbar *et al.*, 2023). Our investigation of transportation outcomes in rich and poor countries builds on these two contributions. We expand the production function approach to modeling travel speed, introduced in Couture *et al.* (2018), to account for the impact on speed of urban crowding, spatial scale, and unobserved road quality, given their empirical importance. We then integrate this speed production function into an urban model with endogenous land area, roads, and population, which guides our investigation of why richer countries are faster. We implement our approach in 152 countries that contain more than 1,200 cities and represent the full range of incomes, geographies, and urban systems.[6]

[5]Following Ahlfeldt, Redding, Sturm, and Wolf (2015), quantitative spatial models include urban transportation data in a complete model of location choices by firms and residents. These models are often developed to assess the full effect of transportation infrastructure projects (e.g. Tsivanidis, 2019). Our data and estimated elasticities offer better transportation input for these models and highlight important determinants of urban mobility that have so far been largely ignored.

[6]Cross-country work fell out of favor in the 1990s due to identification concerns (Durlauf, Johnson, and Temple, 2005), despite the usefulness of global data in answering questions related to economic development. However, new microdata available worldwide allow researchers to replicate their cross-country analysis within different countries, thus addressing both identification and external validity concerns. In fact, we find remarkably consistent estimates of how city attributes impact urban transportation within different countries. A higher city population, for instance, lowers speed just the same in a poor and a rich country. However, countries vary widely in their endowment of attributes. Hence, countries such as the United States and Bangladesh are fast or slow for different reasons.

2. Creating a global urban transportation database

Investigating the link between economic development and mobility requires a novel urban transportation database. This section describes the creation of that database and our methodology to ensure reliable and comparable data across cities and countries. First, we define city boundaries consistently. We then collect trip-level data on speed, congestion, and route within the boundary of each city. Finally, we assemble a database of city attributes, such as road network characteristics, that can explain the variation in speed and congestion across cities and countries. We rely on sources that are available consistently worldwide, in particular Google Maps (GM) as a source of urban travel information, and OpenStreetMap (OSM) as a source of road network and topographical information.

2.1 City sample and boundaries

Our city universe and definitions come from three sources. The location of the city points is from the World Urbanization Prospects (WUP; United Nations, 2019), which contains all 1,860 cities with a population of at least 300,000 in 2018 according to their estimate. City boundaries are based on two datasets from the Global Human Settlements Layer (GHSL) version 2016A reporting conditions circa 2014-15: (i) the Settlement Model (SMOD; Pesaresi and Freire, 2016), which defines each one-kilometer grid cell as urban or not, and (ii) the 38-meter grid of built area (BUILT; Corbane *et al.*, 2018).

We modify and combine these three datasets in several steps described in Appendix A.1. We initially define urban polygons as contiguous sets of SMOD urban cells. We then restrict them to BUILT cells within these SMOD urban cells. Then we expand them slightly with a 500-meter buffer around these BUILT cells. After limiting attention to urban polygons with a nearby city point, splitting polygons containing more than one WUP city or, more rarely, merging two WUP cities when they are more plausibly conceptualized as one city, we are left with 1,807 cities with explicit boundaries. We further exclude China and South Korea because Google Maps cannot collect data there.[7] Finally, we calculate city population within our boundaries as the sum of population in 100-meter pixels according to WorldPop (Bondarenko, Kerr, Sorichetta, and Tatem, 2020) and drop 130 cities with a population less than half of the initial 300,000 threshold, leaving us with 1,228 cities.

2.2 Trips data

Our trip-level information comes from trips simulated on Google Maps (GM). Within each city, we sample four different types of trip that capture the broad variety of urban travel. We then check the sensitivity of our results to the type of trip studied. Specifically, radial trips

[7] In work in progress, Yizhen Gu and Ben Zhou are replicating many of our results using data from Baidu Maps in China.

are between random locations within 2 kilometers of the city center and points elsewhere in the city. Circumferential trips are from a random origin to a destination 30 degrees clockwise or counterclockwise at the same distance from the city center. Gravity trips are between random origins and destinations, with distance drawn from a truncated Pareto distribution. Amenity trips are from a random origin to an amenity (restaurant, grocery store, train station, etc.) selected by GM based on a combination of proximity and popularity.

Within each city c, we sample a target of $15 \times \sqrt{Population_c}$ trips. We query each trip approximately 30 times (instances) on 30 different days, including 14 at the same time of day to measure reliability. In response to our query of the origin and destination, GM reports the distance and time of the fastest route given current traffic conditions, as well as the time it would take in the absence of traffic.

Between 12 June and 5 November 2019, we successfully collected 582,956,059 instances of 18,967,344 trips spanning the 1,228 cities of our sample. For our three most populous countries, India, the United States, and Indonesia, we collected 66, 57, and 13 million trip instances, in 173, 121, and 29 cities, respectively. Appendix A.2 reports further details on our sample of cities on each continent and descriptive statistics on the trips we collected. Appendix A.3 describes additional data sources for trip instances, such as weather information.

For one instance of each trip, we also track geographic coordinates along the fastest route and combine them with other data sources to compute various trip route characteristics such as composition of road class and elevation changes along the route. We describe the construction of these characteristics of the route of the trip in Appendix A.4.

Trip data quality

GM's speed estimates are based on the location and speed of mobile phones that use the Android operating system, as well as other phones that run Google software, especially GM. Accurate speed measurement thus requires that many drivers provide information, which may be less likely in smaller and poorer cities.

In Akbar *et al.* (2023), we provide evidence of the reliability of GM in Indian cities by comparing simulated travel speeds to actual travel speeds from a mobile app. The accuracy of GM speed data, even in a poorer country like India, is perhaps not surprising in light of the high levels of penetration of smartphones and 2G network coverage in India. We estimate that 99.4% of the 1,228 cities in our largest sample have 2G or better mobile network coverage, and 98.6% of our cities have at least 90% of their population's residences covered. Although city-specific information on cellphones with internet subscriptions is not widely available, 90% of cities in our sample are in countries with more than 36 mobile broadband

subscriptions per 100 residents.[8] This is almost certainly an underestimate among residents of the large cities we study, both because they are richer on average and because they experience better network coverage. More generally, there were 2.5 billion Android phones in the world in 2019, so arguably every traffic jam on the planet has at least one phone sending location-based data to Google.

That said, one may still worry that some cities lack accurate speed information. Indeed, Google's own instruction to GM developers acknowledges that some countries, mostly very small or poor, and territories have no traffic layer (see `https://shorturl.at/xzGI6`). Therefore, we develop a methodology to identify countries with poor traffic data. In cities with real-time traffic information, there should be at least some variation in speed between different instances of the same trip (i.e., between times of day and days of the week). So, our main data quality metric is the share of trips in each city for which the fastest instance of a trip is at least 5% faster than the slowest. We average that metric over all cities in a country to identify countries with the worst data quality. We then eliminate countries with the largest average share of trips that show little speed variation, until we have dropped at most 10% of cities.[9] A country at the tenth percentile cut-off only has 9.6% of trips featuring little speed variation across time of day and day of week, so the vast majority of trips in our sample likely feature real-time traffic information.

Reassuringly, our measure of data quality highly overlaps with Google's own information on the presence of a traffic layer. In particular, dropping countries with the lowest data quality eliminates 35 of the 37 cities in countries that Google identifies as having no traffic layer. We drop the remaining two as a precaution. Eighteen additional cities are in four countries that lack reliable income data (see Appendix A.5). This leaves us with 1,119 cities for our baseline sample.

2.3 Estimating a speed index in each city

We now use the GM trip data to produce speed indices that are comparable across cities. Following Ben-Akiva and Lerman (1985) and consistent with the model we propose below, we consider urban travel as the consumption of trips, each intended to access a destination.

[8]Country-level estimates of Mobile Broadband subscriptions per 100 residents are for 2019 and obtained from the International Telecommunication Union (`https://www.itu.int/en/ITU-D/Statistics/Pages/stat/default.aspx`, last accessed, 13 May 2022)

[9]There are two reasons for dropping countries instead of cites to generate our baseline sample. First, it avoids creating a mechanical correlation between country GDP and city population because cities with poor data quality are predominantly small and poor. Second, GM describes its data quality as varying at the country level rather than city level. A 10% cut-off is arbitrary and likely conservative, so we show robustness of all our main results to keeping our entire sample of cities. We also show robustness to dropping 10% of cities and to dropping 20% of cities, an even more conservative approach that eliminates smaller cities in richer countries like the United States, which almost certainly have real-time traffic data. In Appendix C.1, we also develop an alternative data quality metric based on image recognition applied to the color coding of GM's traffic layer, and obtain very similar results.

The price of travel is in units of time, and other costs, such as fuel, are broadly proportional to time. This idea of urban travel as a set of trips chosen by consumers and priced in units of time naturally lends itself to the kind of price index approach commonly used to study constraints on consumption more broadly (Couture *et al.*, 2018).

Our objective is therefore to obtain a single measure of speed in each city—a speed index—for a standard 'basket' of trips. Although we collected trips using a consistent methodology across cities, we need to condition out systematic sampling variation across cities. For example, we mechanically sample longer trips in larger cities and longer trips tend to be faster. So after obtaining the speed S_i of trip instance i as the ratio of its distance and time from GM, we estimate the speed index in city c as a city fixed effect in a regression of log speed of trip instance i in city c including a set of trip-level controls X_i:

$$\ln S_i = X_i\beta + \mathbb{S}_{c(i)} + \epsilon_i. \tag{1}$$

The resulting fixed effect $\mathbb{S}_c$ is our speed index for each city. Our baseline sample includes all trip instances in our baseline city sample of 1,119 cities.

We consider analogous regressions to estimate two additional indices, for uncongested speed and for congestion. To estimate the city uncongested speed index, $\mathbb{U}_c$, we simply replace the dependent variable $\ln S_i$ in equation (1) with $\ln U_i$, the log of speed in absence of traffic. We define congestion as the ratio of uncongested speed U_i to actual speed S_i: $K_i \equiv U_i/S_i$. We similarly replace the dependent variable $\ln S_i$ in equation (1) by $\ln K_i$ to estimate a city congestion index $\mathbb{K}_c$. Importantly, because the dependent variables in the equations used to estimate the fixed effects for speed, uncongested speed, and congestion are additive ($\ln S_i + \ln K_i = \ln U_i$) and because we estimate the same regression in the three cases, the estimated fixed effects are also additive: $\hat{\mathbb{S}}_c + \hat{\mathbb{K}}_c = \hat{\mathbb{U}}_c$.

Table 1 reports the results of the estimation of equation (1), using trip speed as the dependent variable. Column 1 only includes city indicators as explanatory variables and yields an R-squared of 0.32, suggesting large differences in trip speed across cities. Column 2 restricts the sample to weekday trips and adds controls for trip length, the average distance of a trip to the center of its city, indicators for trip type, indicators for time of day departure time in 30-minute periods, and five variables measuring the weather at the time of the query. We first note the large increase in the R-squared to 0.66. Then, unsurprisingly, longer trips and those farther from the center are faster. The magnitude of the trip-length and distance-to-the-center elasticities, together with the fact that we mechanically sample longer trips further from the center within larger cities, suggest that these two controls are important to create comparable speed indices across cities. The estimated trip-type fixed effects are modest and remain so as we add further controls in subsequent columns. This indicates that our speed index does not depend on our exact sampling strategy.

Table 1: Determinants of log trip speed

Dependent variable: log trip speed	(1)	(2)	(3)	(4)	(5)	(6)
log trip length		0.25^a	0.25^a	0.28^a	0.28^a	0.28^a
		(0.0037)	(0.0037)	(0.0037)	(0.0034)	(0.0034)
log distance to center		0.076^a	0.084^a	0.049^a		
		(0.0028)	(0.0028)	(0.0028)		
Gross gradient up				-0.36^a	-0.39^a	-0.40^a
				(0.058)	(0.059)	(0.057)
Gross gradient down				0.30^a	0.27^a	0.29^a
				(0.070)	(0.060)	(0.061)
Gradient missing				-0.017	-0.023^b	-0.018^c
				(0.010)	(0.011)	(0.0099)
Share of primary roads				-0.17^a	-0.17^a	-0.17^a
				(0.0063)	(0.0061)	(0.0060)
Share of secondary roads				-0.19^a	-0.19^a	-0.19^a
				(0.0060)	(0.0057)	(0.0057)
Share of tertiary roads				-0.20^a	-0.20^a	-0.20^a
				(0.0058)	(0.0056)	(0.0055)
Share of residential roads				-0.27^a	-0.27^a	-0.27^a
				(0.0065)	(0.0062)	(0.0062)
Share of other roads				-0.27^a	-0.26^a	-0.26^a
				(0.0079)	(0.0073)	(0.0073)
Share of missing roads				-0.25^a	-0.24^a	-0.24^a
				(0.0062)	(0.0062)	(0.0061)
log intersections				-0.074^a	-0.080^a	-0.081^a
				(0.0031)	(0.0025)	(0.0024)
arsinh turns against traffic				-0.074^a	-0.071^a	-0.070^a
				(0.00089)	(0.00078)	(0.00076)
arsinh density				-0.054^a	-0.050^a	-0.052^a
				(0.0044)	(0.0038)	(0.0042)
Type: circumferential		0.0100^a	0.015^a	0.025^a	0.023^a	0.020^a
		(0.0032)	(0.0031)	(0.0021)	(0.0021)	(0.0020)
Type: gravity		0.020^a	0.024^a	0.026^a	0.025^a	0.024^a
		(0.0022)	(0.0022)	(0.0016)	(0.0015)	(0.0015)
Type: amenity		0.012^a	0.013^a	0.013^a	0.013^a	0.013^a
		(0.0025)	(0.0024)	(0.0018)	(0.0017)	(0.0016)
City effect	Y	Y	Y	Y	Y	Y
Day effect	N	wkday	wkday	wkday	wkday	wkday
Time effect	N	Y	Y	Y	Y	Y
Weather	N	Y	Y	Y	Y	Y
Weight	N	N	Y	Y	Y	Y
City-specific distance to center	N	N	N	N	linear	cubic
Observations	515,019,834	378,535,918	378,009,201	376,459,587	376,459,587	376,459,587
Within R^2	0	0.50	0.48	0.59	0.61	0.62
R^2	0.32	0.66	0.65	0.72	0.74	0.74

Notes: Each column (except column 1) reports parameter estimates for equation (1), regressing log trip speed on log trip length, city, day-of-week, time-of-day (30-minute period), and trip indicators (radial trips are the reference), and five weather variables for trip instances in our main sample of 1,119 cities. Columns 2–6 restrict to weekdays. Column 3–6 weights trips by their time of day according to the share of trips departing at that time of day in the US NHTS and population density within 111 meters of the route. Columns 4–6 add controls for trip route characteristics (motorways are the omitted road class). Columns 5–6 interact log distance to center and city fixed effects. Column 6 further interacts city fixed effects with squared and cubed distance to center. Robust standard errors clustered at the city level are in parentheses. a, b, c: significant at 1%, 5%, 10%.

Column 3 repeats the same specification as column 2, weighting each trip by its time of day, according to the share of trips departing then in the 2017 US National Household Travel Survey (NHTS), and by population density along its route. This weighted regression puts more weight on peak hour trips and less weight on trips in the middle of the night. It also places more weight on trips taken in denser areas where trips are more likely to occur.

Column 4 also includes a number of characteristics of the trip route. To avoid stripping the city fixed effects of these characteristics, we center them by city. The city fixed effects estimated in this regression are our baseline city speed indices.[10] Each coefficient has the expected sign. Moving up is slower, while moving down is faster. Relative to motorways, primary, secondary, and tertiary roads are 16 to 18% slower and other classes of roads are even slower. Trips on routes with more intersections and turns against traffic are slower. Finally, we also estimate an elasticity of speed with respect to nearby population density of -0.054. Columns 5 and 6 allow for the effect of distance to the center to vary across cities, linearly and cubically. Reassuringly, allowing for this heterogeneity does not meaningfully affect any of the other coefficients we estimate relative to column 4.

In Appendix B.1, we report analogs of table 1 for uncongested speed and congestion. In Appendix B.2, we estimate variants of our preferred speed index using alternative samples of trips, such as restricting attention to peak hour travel, and alternative measures of travel speed computed using effective (great-circle) distance. In Appendix B.3, we also verify that the city speed measures obtained in table 1 and in Appendix B.2 are highly correlated.

The results of table 1 estimated in a world sample of cities are close to those in Akbar *et al.* (2023), estimated in a related regression using observations from Indian cities only. That paper also computes speed indices using a wide variety of alternative specifications, and shows that speed indices from equation (1) are not sensitive to restricting trip samples to specific areas of cities, times of day, or types of trip, or to using standard Laspeyre/Paasche price indices, or indices derived from discrete choice microeconomic foundations.

A key concern is that our simulated trips are not representative of the trips people actually take. In the United States, we have access to a nationally representative sample of actual trips from the 2017 NHTS. We use these real trips to benchmark our simulated trips. To do so, we identified 168,874 NHTS trips with origins and destinations within our city boundaries and queried each trip on average 76 times on GM. The details are in Appendix B.4. We do not have routing information for NHTS trips, so we can only estimate the specification of column 2 of table 1. We restricted our attention to 37 cities with more than 1,000 NHTS trips. There is a Spearman rank cross-city correlation of 0.96 between a speed index estimated from these NHTS trips and a speed index estimated using our simulated trip sample. This high

[10] If trip characteristics are appropriately centered and the errors are normally distributed, $\hat{S}_c = \exp\left(\hat{\mathrm{S}}_c + \hat{\phi}_c^2/2\right)$ is a measure of predicted speed for a typical trip in city c where $\hat{\phi}_c$ is the estimator of the standard deviation of the error term ϵ_{ic} for the city. Note that for simplicity, we often directly use the estimated city fixed effects, $\hat{\mathrm{S}}_c$, as an index of speed.

correlation suggests that our trip sampling strategy is representative of actual trips. Overall, our results suggest that speed indices from any broad sample of trips robustly capture speed differences across cities, because slow cities tend to be slow at all times, for all destinations, and all city locations.

Another important concern is that our baseline speed index, which is based on predicted speeds for private vehicles, might not be relevant in contexts where public transit use is widespread. We briefly discuss this idea here and report more details in Appendix B.5. First, we note that even in places with excellent transit, driving in private vehicles remains the dominant mode. In all 13 European countries with harmonized travel surveys reported by Eurostat (2021), at least three fifths of passenger distance traveled was in private vehicles. Second, public transit vehicles are overwhelmingly composed of buses, vans, trams or collective taxis that share the same road network as private vehicles. Only 209 cities in our sample have any kind of transit with its own right-of-way, such as a bus rapid transit (BRT) or subway system and off-road transit modes account for only 5.6% of the viable transit trips we queried (where we define a transit trip as viable when it is faster than walking and requires less than an hour of waiting time prior to departure). When we estimate a driving speed index using only trips whose origin and destination are within relatively short walking distance of on-road transit stops, with no transfers in between, the rank correlation with our baseline speed index is 0.97. Hence, private vehicle travel speeds on routes used by transit are not faster or slower in any way that varies across cities. More generally, given that most transit trips share the road with private vehicles, we expect the speed of both transit and private vehicles to be determined by many of the same factors.

Next, we compare transit and private vehicle speeds more directly. To isolate transit travel time on our trip routes, we exclude walking and waiting times and limit attention to transit trips with a single transit leg on which walking accounts for less than a quarter of total trip distance. We then estimate a transit speed index analogous to our baseline index. This transit speed index has a rank correlation of 0.51 with our baseline speed index. Note that we cannot expect the speed of transit trips to match exactly that of private vehicles. Transit trips are slower, and the difference relative to the speed of private vehicles likely depends on city-specific factors, such as the frequency of stops to pick up and drop off passengers and the payment technology. Our data also often lacks reliable real-time transit information. In most cities, Google Maps reports information based on schedules from transit agencies. In many developing-world cities however, informal transit, which uses the same roadway as private vehicles, represents a large share of transit trips, and is generally not available on Google Maps. For example, there are more than two hundred cities in our sample in which Google Maps returns a viable transit travel time for less than 5% of the trips we query.

Together, this evidence suggests that private vehicles matter everywhere and that city attributes that increase private vehicle speeds are likely to raise transit speeds, both in rich

countries with the most off-road transit and in poor countries with the most informal transit.

2.4 Speed index rankings

Table 2 reports rankings that illustrates the variation across cities in our speed indices. Panel A reports the 20 fastest and 20 slowest cities, while panels B and C report similar rankings for congestion and uncongested speed. The mean value of each index is zero by construction. As an example, a city with a congestion index of 0.05 is about $\exp(0.05) - 1 \approx 5\%$ slower, relative to its uncongested speed, than the average across all cities. The standard deviations of the speed, congestion, and uncongested speed indices are 0.18, 0.05 and 0.17.

The fastest city in the world is Flint, United States, with a speed index of 0.48. The slowest is Dhaka, Bangladesh, with an index of -0.60. This implies a wide speed difference of a factor close to three ($\exp(0.48)/\exp(-0.60)$) between the fastest and slowest cities. Remarkably, 19 of the 20 fastest cities are in the United States, with the exception being Windsor, Canada. Of the 139 US cities in our sample, 86 are among the 100 fastest cities in the world. They are mainly small and midsized cities with a low population density. All the slowest cities are in poor countries, and many are large: Lagos, Nigeria, and Manila, Philippines, accompany Dhaka in the bottom five. Specific countries stand out as being particularly slow. For example, all cities in Bangladesh are among the slowest 15% of world cities.

Turning to our congestion index in panel B, we observe substantially less cross-city variation than for speed. The most congested city is Bogotá, Colombia, with a congestion index of 0.21, and the least congested is Nazret, Ethiopia, with a congestion index of -0.17. Highly congested cities are almost always large. Most of the ten most congested cities are, like Bogotá, in middle income countries (e.g., Moscow, Russia and Mexico City, Mexico), but the next ten are in both rich countries (e.g., New York, United States and London, United Kingdom) and poor countries (e.g., Lagos, Nigeria and Mumbai, India). The least congested cities are generally smaller cities in poor countries, with presumably little motor vehicle traffic.[11]

The fastest uncongested speed ranking in panel C is very similar to the fastest speed ranking in panel A, with an overlap of 15 of 20 cities. Cities with the slowest uncongested speed also overlap with the cities with the lowest overall speed, but the very largest cities, which are also congested, are slightly less prominent in the uncongested speed ranking. Again, specific countries stand out. Nine of the ten cities with the slowest uncongested speeds are in Bangladesh, India, and Nigeria.

[11]The exact ranking of the least congested cities, however, is sensitive to how we clean our data. Our data quality concerns are highest in smaller cities in poor countries, and our data quality metric cannot distinguish a lack of real-time traffic data from a lack of vehicles on the road.

Table 2: Ranking of the 20 fastest, slowest, most congested and least congested cities

Rank	City	Country	Index	City	Country	Index
Panel A:		**Fastest**			**Slowest**	
1	Flint	United States	0.48	Dhaka	Bangladesh	-0.60
2	Greensboro	United States	0.45	Lagos	Nigeria	-0.54
3	Wichita	United States	0.44	Ikorodu	Nigeria	-0.51
4	Tulsa	United States	0.41	Manila	Philippines	-0.49
5	Lansing	United States	0.41	Bhiwandi	India	-0.48
6	Knoxville	United States	0.40	Kolkata	India	-0.47
7	Shreveport	United States	0.39	Arrah	India	-0.44
8	Memphis	United States	0.39	Phnom Penh	Cambodia	-0.43
9	Youngstown	United States	0.39	Mymensingh	Bangladesh	-0.43
10	Toledo	United States	0.38	Aba	Nigeria	-0.43
11	Port St. Lucie	United States	0.38	Bihar Sharif	India	-0.42
12	Rockford	United States	0.38	Chittagong	Bangladesh	-0.42
13	Dayton	United States	0.38	Mumbai	India	-0.41
14	Bakersfield	United States	0.37	Bacoor	Philippines	-0.41
15	Grand Rapids	United States	0.37	Dar es Salaam	Tanzania	-0.40
16	Stockton	United States	0.37	Kumasi	Ghana	-0.40
17	Montgomery	United States	0.36	Mombasa	Kenya	-0.40
18	Springfield, Missouri	United States	0.36	Aizawl	India	-0.39
19	Windsor	Canada	0.36	Bangalore	India	-0.39
20	Oklahoma City	United States	0.36	Shillong	India	-0.39
Panel B:		**Most congested**			**Least congested**	
1	Bogotá	Colombia	0.21	Nazret	Ethiopia	-0.17
2	Krasnodar	Russia	0.19	Dire Dawa	Ethiopia	-0.17
3	Ulaanbaatar	Mongolia	0.18	Songea	Tanzania	-0.17
4	Bucharest	Romania	0.17	Gboko	Nigeria	-0.17
5	Moscow	Russia	0.17	Gondar	Ethiopia	-0.17
6	Bangkok	Thailand	0.17	Potiskum	Nigeria	-0.16
7	Manila	Philippines	0.16	Sikasso	Mali	-0.16
8	Bangalore	India	0.16	Ogbomosho	Nigeria	-0.14
9	Vladivostok	Russia	0.16	El Djelfa	Algeria	-0.14
10	Mexico City	Mexico	0.14	Oyo	Nigeria	-0.14
11	Lagos	Nigeria	0.14	Abakaliki	Nigeria	-0.13
12	London	United Kingdom	0.14	Morogoro	Tanzania	-0.13
13	Mumbai	India	0.14	Gombe	Nigeria	-0.13
14	Yekaterinburg	Russia	0.14	Nasiriyah	Iraq	-0.13
15	Guatemala City	Guatemala	0.14	Ondo	Nigeria	-0.12
16	Panama City	Panama	0.13	Bouake	Côte d'Ivoire	-0.12
17	Nairobi	Kenya	0.13	Safi	Morocco	-0.12
18	New York	United States	0.13	Katsina	Nigeria	-0.12
19	Santo Domingo	Dominican Republic	0.13	Minna	Nigeria	-0.12
20	Delhi	India	0.12	Ijebu-Ode	Nigeria	-0.12
Panel C:		**Fastest uncongested speed**			**Slowest uncongested speed**	
1	Flint	United States	0.43	Ikorodu	Nigeria	-0.52
2	Greensboro	United States	0.41	Dhaka	Bangladesh	-0.49
3	Wichita	United States	0.39	Aba	Nigeria	-0.49
4	Tulsa	United States	0.38	Khulna	Bangladesh	-0.43
5	Knoxville	United States	0.38	Mymensingh	Bangladesh	-0.41
6	Memphis	United States	0.37	Bihar Sharif	India	-0.41
7	Shreveport	United States	0.36	Bhiwandi	India	-0.40
8	Bakersfield	United States	0.36	Lagos	Nigeria	-0.39
9	Lansing	United States	0.36	Kolkata	India	-0.39
10	Hartford	United States	0.36	La Paz	Bolivia	-0.39
11	Grand Rapids	United States	0.35	Port-au-Prince	Haiti	-0.39
12	Youngstown	United States	0.35	Arrah	India	-0.38
13	Windsor	Canada	0.35	Mombasa	Kenya	-0.38
14	Montgomery	United States	0.35	Aizawl	India	-0.38
15	Corpus Christi	United States	0.35	Darbhanga	India	-0.37
16	Toledo	United States	0.35	Dar es Salaam	Tanzania	-0.37
17	Stockton	United States	0.34	Quetta	Pakistan	-0.36
18	Kansas City	United States	0.34	Chittagong	Bangladesh	-0.36
19	Fresno	United States	0.34	Comilla	Bangladesh	-0.35
20	Jacksonville, Florida	United States	0.34	Varanasi	India	-0.35

Notes: Indices in Panel A are estimated city fixed effects from Table 1, column 4. Indices in Panels B and C are fixed effects from the corresponding uncongested speed and congestion regressions reported in Appendix B.1.

Table 3: City attributes summary statistics

		Mean	Min.	25 pctl	Median	75 pctl	Max.
Country Income	GDP per capita (int. US$)	25,213	1,550	6,367	17,225	41,220	110,522
City size	Population (thousand)	1,407	151	325	565	1,206	35,028
	Area (km^2)	359	22	112	187	346	6,025
Road network	Major road length (km)	848	8.4	187	374	788	24,159
	Network griddiness (0-1)	0.17	0.06	0.09	0.12	0.17	0.84
Topography	Water body length (km)	448	0	65	168	388	16,528
	Elevation variance (m^2)	1,694	1	90	451	1,468	64,846

Notes: Sample is 1,119 cities. Income varies at the country level, all other variables vary at the city level. GDP is in PPP 2017 international US$.

These simple rankings of cities showcase several patterns that we explore more systematically in the rest of the paper. In particular, uncongested speed appears to drive more of the speed variation between cities and countries than congestion. Specific countries stand out as particularly slow or particularly fast, so country effects are likely important. Larger cities are slower and more congested, as documented by Couture *et al.* (2018) in the US, and by Akbar *et al.* (2023) in India. Lastly, and perhaps most strikingly, all cities with the fastest speed or uncongested speed are in rich countries, and all the slowest cities are in poor countries. To our knowledge, the gap in urban mobility outcomes between rich and poor countries has never been documented before. The rest of our paper focuses on this investigation.

2.5 City attributes

We now complete our presentation of the data by defining our main city attributes. These city attributes are, at least in theory, potential determinants of the travel speed differences across cities documented above. Table 3 summarizes these variables. Appendix A.5 provides more details on data sources and variable construction, and introduces variables used only in robustness checks.

Income. We measure countries' GDP per capita in 2018 at purchasing power parity in 2017 international dollars, from the Penn World Table (Feenstra and Timmer, 2015, updated 2023). City-level GDP for OECD countries is from OECD (2021). We use separate city-level income data for India (industrial earnings from the Employment and Unemployment Survey of the 2011-12 National Sample Survey, as aggregated to the district level by Li, Rama, Galdo, and Pinto, 2015) and the United States (average household income from the 2015–2019 American Community Survey).

Figure 3: OSM road network of city with highest and lowest network griddiness.

Panel A: Ciudad Obregon, Mexico

Panel B: Charleroi, Belgium

Motorways and primary roads (thinner) are in black. Secondary roads (thicker), tertiary roads, and residential roads (together with service roads and other road types, thinner) are in gray.

City population and area. We calculate city population as the sum of population in 100-meter pixels according to WorldPop (Bondarenko *et al.*, 2020) within the boundaries defined above. We compute city area using the same boundaries.

Infrastructure. Our road network data come from OSM via GeoFabrik in 2020 and processed through OSMnx (Boeing, 2017). We define two main measures of the network. First, we compute the total length of major roads in kilometers. These are roads tagged as motorways, primary, secondary, and tertiary in OSM. Second, we measure the griddiness of the network as the share of edges in a city's road network that are within 2 degrees of the modal edge bearing modulo 90. That is, our griddiness measure captures the share of network edges that conform to the main grid orientation.[12]

The OSM data is crowd-sourced (Nagaraj, 2021), and some cities may have incomplete road networks. To systematically measure the completeness of the network and identify problematic cities, we calculate the share of total distance traveled in our GM trip sample that occurs on roads missing from OSM. To do this, we first map each segment along any GM trip route to the OSM road network. We then classify a GM trip segment as missing in OSM if that segment is more than 30 meters from the nearest OSM network edge. Across all cities in our baseline sample, we find a mean share of missing roads of 15%. However, visual inspection suggests that the missing roads are almost all residential or service roads and

[12]Fuller and Romer (2014), for instance, argue for griddier road networks to improve urban transportation in developing countries.

that the major road network is complete in almost every city in our sample. We nonetheless create a robustness sample omitting the 10% of cities with the highest share of missing OSM roads.[13]

Topography. We measure the total length of water bodies (river centerlines, coastlines, and lakeshores) within each city using information from OSM. We also query each of our trips in walking mode, which allows us to recover elevation at each step and, therefore, measure the variation in elevation within each city.

3. Mobility and economic development

We now investigate the relationship between urban mobility and economic development more formally. We first estimate a regression of speed on only income:

$$\ln \mathbb{S}_c = \ln \Omega_{\text{Base}} + \kappa_{\text{Base}} \ln y_c + \varepsilon_c \,, \tag{2}$$

which, following the decomposition we perform below, we refer to as the 'base' regression.

Table 4 reports results. Each observation is a city, and different columns use different samples of cities. Column 1 shows that for all cities in the world, country-fixed effects alone explain 72% of speed variation. The variation in speed within countries, which has been the focus of previous literature, appears to be much less important. In other words, some countries are slow like Bangladesh, and some are fast, like the United States. In column 2, we replace the country-fixed effect by country GDP per capita. The R-squared of 0.46 suggests that country GDP per capita explains (0.46/0.72 =) 64% of the cross-country variation in travel speed. Increasing GDP by 10% is associated with 1.3% higher speed, which implies that the poorest countries are about twice as slow as the richest ones. Column 3 estimates the same regression, but restricts the sample to cities in the OECD, a group of high income countries. The R-squared of 0.51 shows that even among richer countries, GDP per capita explains most of the cross-country variation in the speed index.

We now turn to investigating whether, within each country, richer cities are also faster. We have city-level GDP per capita for OECD countries, and city level measures of income per capita that we collected separately for India and the United States, the two countries with the most cities in our sample. We first confirm, in column 4, that richer OECD cities are faster, consistent with the result in column 3. However, after adding a country-fixed

[13]This cut-off drops all cities with more than 23% of total distance travel roads missing from OSM. We perform visual comparisons of the OSM network with the GM color-coded network, as well as with aerial and satellite images. Identifying secondary and tertiary roads visually on maps is difficult, but we do not find evidence that motorways, primary, and secondary roads were ever missing, and tertiary roads appear only infrequently missing. Therefore, in robustness checks, we omit tertiary roads from our measure of major roads.

Table 4: Speed vs. GDP per capita

	(1) Speed all	(2) Speed all	(3) Speed OECD	(4) Speed OECD	(5) Speed OECD	(6) Speed US	(7) Speed India
log country GDP (pc)		0.13[a]	0.32[a]				
		(0.023)	(0.049)				
log city GDP (pc)				0.16[a]	-0.13[a]	-0.23[a]	0.019
				(0.037)	(0.023)	(0.052)	(0.028)
country fixed effect	Y	N	N	N	Y	N	N
Observations	1,119	1,119	265	265	265	121	173
R^2	0.72	0.46	0.55	0.19	0.77	0.24	0.00

Notes: Each column reports a parameter estimate from a separate regression of the city speed indices estimated in the Table 1, column 4 specification on GDP per capita and/or city fixed effects. In columns 2 and 3, GDP is country-level GDP at PPP in 2017 international dollars. In columns 3–5, the sample is restricted to cities in OECD countries with a city-level measure of GDP from the OECD, which is the GDP measure used in columns 4–5. In columns 6 and 7, the sample is restricted to the US and India, respectively, and income data are from national sources. Robust standard errors clustered by country are in parentheses in columns 1–5 and robust standard errors in columns 6–7. a, b, c: significant at 1%, 5%, 10%.

effect the sign on city-level GDP per capita switches to negative. In other words, within each country, richer cities are slower. Using data for the United States only, column 6 shows that average household income in each city is strongly and negatively correlated with speed. This is consistent Duranton and Turner's (2012) finding that highways in the United States are disproportionately allocated to declining cities. In fact, the fastest city in the world, Flint, Michigan, is a relatively poor city in a rich country. Finally, column 7 reports result for India only. Unlike in the United States, we find a coefficient on city-level average industrial earnings that is positive, but small and not significant.[14]

We conclude that there is a previously undocumented, first-order relationship between a country's economic development and the speed of vehicular travel in its urban networks. Covering the same unit of distance takes about twice as long in a very poor relative to a very rich country. However, within countries, there is no systematic relationship between city income and speed. In countries like the United States, the political process may allocate road investments to poor or declining places like Flint, while in India richer cities may have better infrastructure.

[14]Akbar *et al.* (2023) estimate related regressions for India. Overall, they document a quadratic relationship between city income and speed, whereby speed initially increases with city income and then decreases as richer cities become more congested.

4. A model of urban travel

Although the relationship between urban mobility and the level of economic development of a country is strong, a higher GDP does not directly increase the speed of travel. To further investigate why richer countries are faster, we propose a model of urban travel, which will guide the next step of our empirical estimation. Starting from a standard production function with unobservable inputs, the model delivers an expression for equilibrium city-level speed in terms of observable city attributes and analogous expressions for speed in the absence of traffic and the congestion factor.

Following Couture *et al.* (2018), we assume that distance traveled is produced using roadway and traveler time, just as regular consumption goods are produced using capital and labor. Over a time interval in city i, a number of travelers H_i and an effective roadway R_i (accounting for both the quality and quantity of roads) produce a quantity D_i of traveler-distance traveled following:

$$D_i = B_i \, R_i^{\phi} H_i^{1-\theta} \, . \tag{3}$$

We refer to ϕ as the roadway elasticity and to θ as the congestion elasticity. In equation (3), travel productivity B_i measures the ability of city i to produce travel distance, conditional on its effective roadway and the number of travelers.

In turn, the travel productivity of a city is determined by how crowded the city is, its spatial scale, and other factors:

$$B_i = {N_i}^{-\mu} \, {Q_i}^{\nu} \, A_i \, , \tag{4}$$

where N_i is population, Q_i is the developed land area of the city, and A_i measures the other determinants of travel productivity, which in our empirical context are its topography. The urban crowding elasticity μ reflects how a higher population reduces travel productivity, holding land area and roads constant. The spatial scale elasticity ν measures how travel productivity improves when the same amount of roadway and population is spread over a larger area. We discuss the empirical content of A_i and the distinction between the crowding and the spatial scale elasticities further below.

Because total distance traveled (per unit of time) D_i is the product of the number of travelers, H_i, and travel speed, S_i, we can rewrite equation (3) as:

$$S_i = A_i \, R_i^{\phi} H_i^{-\theta} N_i^{-\mu} Q_i^{\nu} \, . \tag{5}$$

Therefore, travel speed in city i, S_i, depends on its topography A_i, its effective roadway R_i, the number of travelers H_i, population N_i and land area Q_i. This equation can be thought as an inverse-supply equation for travel where the price of travel (measured by the time cost

or the inverse of speed) depends on the quantity of travelers, city roadway, city population, city area, and topography.[15]

Our specification in equation (5) makes two important distinctions. Crucially, we first distinguish between congestion, θ, which varies throughout the day as a function of the number of travelers, and crowding, μ, which reflect the more permanent effects of having a larger population competing for space with vehicle travelers. These crowding effects include more encroachments by double-parked delivery vehicles, electric scooters, draft animals and walkers, limited visibility as buildings are closer to the roadway at intersections, more time spent at intersections as red lights replace stop signs and roundabouts, more roadwork due to higher usage, etc. Although vehicular congestion has been a major focus of interest in transportation economics (Small and Verhoef, 2007), crowding has received much less attention, if any. Below, we show that crowding is empirically more important than congestion.

Then, we also distinguish between the effects of more roadway captured by the roadway elasticity ϕ and a separate spatial scale elasticity ν. Spatial scale effects capture the possible gains from spreading the same amount of roadway and population on more land, which leads to fewer intersections per unit of distance, greater ease of avoiding natural obstacles, and perhaps a road network that is less complex and easier to manage. Empirically, we show that the spatial scale elasticity of the road network and the road elasticity are similar in size.[16]

We cannot yet empirically estimate the speed equation (5). First, we cannot observe the number of travelers H_i. However, after dropping city indices, the number of travelers H is equal to:

$$H \equiv \frac{N\,t\,d}{S}\,, \tag{6}$$

where t is the endogenous number of trips taken by each resident, and d is the travel distance of these trips, which we take as constant and normalize to unity.[17] In equation (6), speed determines the number of travelers H in a given time interval, itself a determinant of speed

[15]To avoid predicting implausibly large speeds when the number of travelers or population become small, it is reasonable to assume that speed reaches a maximum $S_i = \overline{S}$ when H_i or N_i fall below some threshold. We also note that in our setting, speed is not a choice variable. It is instead given directly by the 'travel conditions' (other travelers and the roadway). We think this assumption is appropriate for dense urban settings where there is much less scope to choose travel speed than on uncongested rural highways.

[16]It may be tempting to write (5) as a function of population density and road density. This formulation however is less general since it constrains $\phi + \nu - \mu = 0$ i.e., constant returns to scaling up city attributes (roads, land, and population) into the speed production function. Instead, we show below that we obtain this equality as an empirical result.

[17]We consider trip distance to be exogenous to avoid too many feedback loops and complicate the derivations further. Empirically, this normalization should be innocuous, because we compute speed indices after conditioning out trip length. For US metropolitan areas, Couture *et al.* (2018) conclude that although trip distance is a fundamental determinant of travel speed, its simultaneous determination with travel speed only appears to be a minor issue.

in equation (5). Intuitively, having more travelers slows down other travelers and, in turn, slower travel further increases the number of travelers and magnifies congestion (Small and Verhoef, 2007). After substituting H into equation (5) using equation (6), we can rewrite speed as:

$$S = A^{\frac{1}{1-\theta}} R^{\frac{\phi}{1-\theta}} t^{-\frac{\theta}{1-\theta}} Q^{\frac{\nu}{1-\theta}} N^{\frac{-\theta-\mu}{1-\theta}} , \tag{7}$$

In equation (7), a city's speed is fully accounted for by its topography, effective roadway, number of trips, land area, and population. However, number of trips is unobserved and we only observe roadway length, not its quality. To address these challenges, we propose a simple theory of travel demand and roadway provision. Assume that resident k in city i enjoys utility

$$V_k = \frac{t_k^{\alpha} \, q_k^{\beta} \, c_k^{1-\alpha-\beta}}{\alpha^{\alpha} \beta^{\beta} (1-\alpha-\beta)^{1-\alpha-\beta}} , \tag{8}$$

where t_k is the number of trips undertaken, q_k is the consumption of housing measured directly in units of land for simplicity, and c_k is a numeraire composite good.[18] The budget of resident k in city i is

$$\frac{1}{S_i} t_k + P_i q_k + c_k = (1-\tau_i) y_i , \tag{9}$$

where the price of a trip of unit length is its time cost, or the inverse of speed.[19] P_i is the price of land, y_i is gross income, and τ_i is the income tax rate paid to fund the roadway. Dropping indices again to ease reading, utility maximization implies that a fixed share of income is spent on each good:

$$\begin{cases} \dfrac{t}{S} = \alpha \, (1-\tau) y , & \text{(10a)} \\ P \, q = \beta \, (1-\tau) y , & \text{(10b)} \\ c = (1-\alpha-\beta) \, (1-\tau) y . & \text{(10c)} \end{cases}$$

We can solve for the equilibrium number of trips by substituting travel demand, using the first condition in equation (10a), into travel supply from equation (7):

$$t = \alpha \, (1-\tau) y \, S = [\alpha \, (1-\tau) y]^{1-\theta} \, \frac{A \, R^{\phi} \, Q^{\nu}}{N^{\theta+\mu}} , \tag{11}$$

[18] Our specification treats trips as a consumption good. We are aware that in reality, urban travel is in large part a derived good, which allows travelers to access work from their residence to earn income as well as access other goods and consume leisure. We sever these important links to avoid obscuring the main insights of our model.

[19] We monetize time directly into the budget constraint to avoid the complications associated with the introduction of a time budget constraint into the consumer problem. We could generalize our specification by specifying that the value of time is proportional to y^{ζ}. Simple algebra shows that in equation (16) below, the coefficients on income would then be $\theta - \zeta$ instead of θ with our specification. The coefficient on population in the same equation would remain θ. When we evaluate this regression, the difference in the coefficient on population and income is small.

where the second equality is obtained after substituting S using equation (7) and simplifying. Distance traveled by each resident increases with income, favorable city topography, more effective roadway, and greater land area, and decreases with population.

We model roadway R as the product of its length L, which we observe, and its quality, which includes hard-to-measure attributes such as roughness of the surface. In turn, we model road quality as a level of technology, for which income is a sufficient statistic, with elasticity σ.[20] Thus, overall roadway provision is given by:

$$R = L\, y^{\sigma}. \tag{12}$$

Combining equations (11) and (12), we write speed as

$$S = [\alpha(1-\tau)]^{-\theta}\, A\, L^{\phi} Q^{\nu} N^{-\theta-\mu}\, y^{\phi\sigma-\theta}. \tag{13}$$

After taking logs, we can implement equation (13) as:

$$\ln S = \ln \Omega_{\text{Full}} + \kappa_{\text{Full}} \ln y + \psi_A \ln A + \psi_N \ln N + \psi_L \ln L + \psi_Q \ln Q\,, \tag{14}$$

where the structural interpretation of the various coefficients in this regression can be read directly from equation (13). In the decomposition below, we will refer to equation (14) as the 'full' regression.

In addition to travel speed, we can also measure speed in absence of traffic, which we model as a situation with a single traveler ($H = 1$). Then, substituting for roads R in equation (5) using equation (12), we obtain the following equation for uncongested speed,

$$U = A L^{\phi} Q^{\nu} N^{-\mu} y^{\phi\sigma}\,, \tag{15}$$

which, after taking logs, leads to a regression equation analogous to (14) for uncongested speed.

The congestion factor K is defined as the ratio of uncongested speed to speed, $K \equiv \frac{U}{S}$. Using equations (13) and (15) and simplifying implies:

$$K = [\alpha(1-\tau) N\, y]^{\theta}\,, \tag{16}$$

which, after taking logs, leads again to a regression analogous to regression (14). Equation (16) makes two important testable predictions. First, the elasticities of congestion with respect to population and income should be the same (this follows from how we specified travel costs). Second, congestion in equation (16) depends only on a subset of the city attributes that affect overall speed in equation (13).

[20] While a full modeling of road quality and its microfoundations is beyond the scope of our paper, we are first thinking about road maintenance, specifically pavement roughness for which there is no available systematic data to our knowledge. Another aspect of road quality is road network management, which includes road network shape, the judicious use of one-way streets, the use of red lights and their appropriate coordination, etc. Some aspects of network shape are observable, while the rest remain unobserved. Finally, there are other speed determinants such as speed limits and vehicle quality that are harder to classify. We experiment with speed limits in our robustness checks but do not have a solution to their obvious simultaneity with observed speeds. Vehicle quality is unlikely to be a binding constraint given the speeds we observe in cities from less developed countries.

Table 5: Determinants of urban travel speed

		Speed index		Uncongested speed		Congestion factor	
		(1) Base model	(2) Full regression	(3) Base model	(4) Full regression	(5) Base model	(6) Full regression
	log country GDP (pc)	0.13^a	0.042^a	0.15^a	0.074^a	0.019^a	0.032^a
		(0.023)	(0.011)	(0.018)	(0.0090)	(0.0059)	(0.0067)
City size	log population		-0.17^a		-0.13^a		0.040^a
			(0.018)		(0.015)		(0.0069)
	log area		0.085^a		0.064^a		-0.021^b
			(0.022)		(0.019)		(0.0087)
Infrastructure	asinh major road length		0.080^a		0.076^a		-0.0040
			(0.012)		(0.013)		(0.0080)
	Network griddiness (0-1)		0.18^a		0.12^a		-0.058^a
			(0.057)		(0.046)		(0.015)
Topography	Elevation variance		-0.0024^b		-0.0011		0.0013^a
			(0.00092)		(0.0010)		(0.00040)
	asinh water body length		-0.071^a		-0.054^a		0.017^b
			(0.019)		(0.019)		(0.0083)
	Observations	1,119	1,119	1,119	1,119	1,119	1,119
	R^2	0.46	0.70	0.63	0.77	0.11	0.43

Notes: The unit of analysis is the city. Columns 1–6 each report parameter estimates of separate OLS regressions on the sample of 1,119 cities with high-quality speed data and country GDP data. Standard errors clustered at the country level in parentheses. a, b, c: significant at 1%, 5%, 10%.

5. Empirical results: Determinants of speed

In this section, we use our global city sample to estimate the equations of our model regarding the determinants of urban travel speed. We then validate our findings using within-country and within-city variation, and we assess the importance of other city attributes not in our model.

5.1 Baseline cross-city results

Table 5 reports results from regressing speed (columns 1–2), uncongested speed (3–4), and congestion (5–6) on the attributes suggested by our model. Columns 1, 3, and 5 include only GDP per capita as an explanatory variable.

Due to the additive properties of our speed indices mentioned above ($\mathbb{S}_c = \mathbb{U}_c - \mathbb{K}_c$), the estimated coefficients are similarly additive. For instance, the coefficient on log GDP per capita in column 1 of 0.13 for speed is equal to the analogous coefficient in column 3 of 0.15 for uncongested speed minus 0.019 in column 5 for the congestion factor. These three regressions underscore that the speed differences between rich and poor countries are

mainly due to the speed vehicles can travel on the roads of their cities. Greater uncongested speed in richer countries greatly dominates increased congestion, which confirms what the scatter plots of figure 2 suggested.

Leveraging again the additive properties of our mobility indices, we can perform a simple variance decomposition. The variance of uncongested speed accounts for 95% of the variance of our speed index across cities, while the variance of the congestion factor only accounts for about 9% of the same variance (and the covariance term for approximately 2%). In Appendix B.6, we show that these results are robust to deviations from the baseline index and the city sample we consider. Even when we focus only on radial peak-hour trips which are subject to more congestion, we still find that the variance of uncongested speed accounts for 91% of the variance of speed across cities.

The specification in columns 2, 4, and 6 additionally include city population, land area, length of major roads, network griddiness, and two measures of topography (elevation variance and water body length). The specification in column 2 for speed estimates equation (13), while columns 4 and 6 estimate equations (15) for uncongested speed and (16) for the congestion factor. We note that the coefficient on country GDP drops from 0.13 in column 1 to 0.042 in column 2. Our set of six city attributes accounts for almost 70% of why richer countries are faster. We also note the relatively high R-squared of 0.70 reported in column 2. Taken together, the seven explanatory variables we consider account for a large share of the variation of speed across cities worldwide.

To interpret the results reported in table 5 through the lens of our model, it is useful to start with the coefficient on population for the congestion factor in column 6. In equation (16), this coefficient provides an estimate of θ, the congestion externality. Its value of 0.040 is modest. Its effects are equally modest: increasing log population by one standard deviation is associated with a reduction in travel speed through increased congestion by slightly less than 4%.

This result is consistent with the variance decomposition reported above, in which congestion plays a minor role. In a related, albeit not fully comparable exercise, Duranton and Puga (2023) use residualized city speed for 180 US metropolitan areas obtained from driver-reported data in the NHTS and from the corresponding simulated trips using Google

Maps. In both cases, they estimate a congestion elasticity of 0.04.[21]

According to the congestion equation (16), the coefficient on income in column 6 also measures the congestion elasticity θ. The point estimate of 0.032 for income is close to that of 0.040 for population and statistically indistinguishable from it. Also consistent with the predictions of the model, we estimate small coefficients on land area and road length. The precisely estimated zero for the coefficient on road length is consistent with the fundamental law of congestion and the evidence provided by Duranton and Turner (2011). The remaining three coefficients on topography and network griddiness are statistically significant in column 6, but they are economically less important and, except for elevation, much smaller than in the speed regression of column 2.

Next, we turn to the estimation of the road elasticity ϕ using equation (13) for speed. This elasticity is given by the coefficient on road length in column 2 of table 5. The estimated value of $\phi = 0.08$ is economically important. A one standard deviation change in our measure of road length corresponds to a nearly 10% change in travel speed. Because the mileage of major roads increases strongly with GDP per capita, this variable plays a major role in our decomposition of why richer countries are faster in the next section. Importantly, this effect of major roads on speed is almost entirely accounted for by increased uncongested speed. Higher road classes allow for faster travel than residential roads and other lower categories of roadway. The road elasticity we estimate is about a third smaller than what Couture *et al.* (2018) find for the US in a much simpler specification that omits many attributes included here. We show below that using within-country and even within-city variation delivers similar estimates for ϕ.

In equation (13) for speed, the coefficient on city land area corresponds to ν, the spatial scale elasticity. Column 2 thus implies $\nu = 0.085$. The coefficient on city land area in equation (15) for uncongested speed is also ν. The estimate of 0.064 in column 4 is close to its value in column 2. This spatial scale effect is also comparable in magnitude to the

[21] Couture *et al.* (2018) also use the NHTS to estimate a congestion elasticity associated with city population in US metropolitan areas. In light of our model, their 0.11 elasticity conflates both congestion θ and crowding μ. This estimate of 0.11 is arguably more meaningfully compared to the population elasticity of 0.17 obtained in column (2) of table 5 or to a value of 0.17 estimated by Akbar *et al.* (2023) in a related regression for Indian cities. Using a different approach, Akbar and Duranton (2022) estimate congestion over an entire city (Bogotá). They find a higher congestion elasticity of about 0.1. Despite a different functional form, the results of Kreindler (2023) in Bangalore are consistent with this value of approximately 0.1. A likely reason for the higher values estimated in Akbar and Duranton (2022) and Kreindler (2023) is that their elasticity is with respect to the number of travelers observed. Instead, our elasticity is measured using population. We expect the number of travelers to be less than proportional to population, as travelers can reschedule their trips or not take them at all as roads become more congested. Then, a large literature estimates the congestion elasticity for specific roads with respect to the number of travelers. This literature usually finds considerably larger elasticities, close to one or above in some cases. See Small and Verhoef (2007) for a full discussion and Anderson and Davis (2020) or Bou Sleiman (2022) for more recent work. An important reason behind this major difference in the value of the estimated congestion elasticity is that road-level estimates do not allow for rerouting (Akbar and Duranton, 2022, Bou Sleiman, 2022) and tend to consider particularly congested roads. In short, and perhaps unsurprisingly, the estimated congestion elasticity gets higher as we allow for more margins of adjustment.

road length effect. A one-standard-deviation increase in log land area is associated with an approximately 8% higher speed. As we show below, this elasticity plays a central role in explaining speed differences between rich and poor countries as well as between the US and other OECD countries. We know of no other estimate of this elasticity in the literature (Ahlfeldt and Pietrostefani, 2019).

Still in equation (13) for speed, the coefficient on city population corresponds to the sum of the congestion externality, θ, and the crowding externality, μ. With an estimated coefficient on population of -0.17 in column 2 and a value of $\theta = 0.040$ obtained from column 6 as discussed above, we find $\mu = 0.13$ for crowding. This is the largest speed elasticity of any of the city attributes that we consider, which underlines the importance of urban crowding as a determinant of speed. A one standard deviation increase in log population is associated with a 16% decrease in travel speed. Interestingly, our estimate of this population elasticity of speed is of the same magnitude as Duranton and Turner's (2018) estimate of 0.11, based on differences in population density within US metropolitan areas.

Using equation (13) one last time, the coefficient on income in column 2 corresponds to $\phi\sigma - \theta$. With the estimated coefficient of 0.042 for income, we recover a value of σ of about 1. This large elasticity of road quality with respect to income implied by the model underscores the importance of unobserved road quality, ranging from road design to road maintenance.

For comparison, we also note the effects of network griddiness and topography variables. A one standard deviation increase in our measure of elevation variance, water bodies, and shape of the road network are associated with a 1.1% decrease, a 3.2% decrease, and a 2.4% increase in speed, respectively. Although these effects are modest, we nonetheless keep in mind that a standard deviation of our measure of water bodies is associated with about the same effect on speed as a standard deviation of population through congestion (see Saiz and Wang, 2023, for an in-depth exploration of water bodies as physical barriers leading to more congestion in the case of Boston).

Overall, our estimations offer new or updated estimates of the key parameters that determine the production of urban transportation: θ, μ, ν, and ϕ. Our findings point to a large role of population through both its crowding and congestion effects. A one standard deviation increase in population is associated with 16% lower speed. More than three quarters of that effect occurs through crowding and less than a quarter through congestion. The other two important drivers of urban travel speed in cities are road length and land area.

The structural interpretation of our results has further interesting implications. In equations (4)–(13), we model distance traveled and subsequently speed using a production function approach, which begs the question of how the production of travel scales. Consider a hypothetical city and increase its population, land area, and road length proportionately. Our results indicate that these changes marginally increase uncongested speed with

elasticity 0.012 while speed, which includes more congestion from more population, is almost constant with elasticity -0.003. These results point at tiny increasing returns in the production of travel with uncongested speed and constant returns to scale once congestion is taken into account.

Next, consider scaling up income. Higher income leads to higher demand for travel and higher road quality. The income elasticity of travel demand is $1 - \theta$, where the unit Cobb-Douglas elasticity is attenuated by increased congestion. A higher income also leads to higher road quality captured by $\sigma\phi$, the product of the income elasticity of road quality σ, and the road elasticity ϕ. So in equilibrium, the elasticity of speed with respect to income in equation (13) is $\sigma\phi - \theta$. The positive elasticity of 0.042 in column 2 of table 5 implies that the effect of unobserved road quality, $\sigma\,\phi$, dominates increased congestion, θ.

Income might also affect speed indirectly through its impact on our six city attributes, in addition to its direct impact through unobserved road quality and congestion. For example, cities in richer countries may have wider areas, which in turn increases travel speed. These indirect effects are important, given how the coefficient on income declines from 0.13 to 0.042 after including city attributes in the model. Therefore, to account for the full effect of country income on urban travel speed, our model must account for how land area, road length, city population, and possibly city topography and road network griddiness vary with income. We extend our model in these directions in section 6, and we decompose the importance of each city attribute in explaining why richer countries are faster. Before doing this, we perform a number of checks to assert the robustness of our findings and assess the role of other city characteristics that we have ignored so far.

5.2 Complementary results

So far, our estimates have been based on global variation in each of our city attributes. Our main identification concerns regard heterogeneity in data quality or model parameters across countries, and possible confounding factors that are missing from our baseline regression. Worse measurement of city roads in poor countries (despite our best efforts to collect comparable data as described in section 2), or stronger effects of major roads on speed are examples of the first set of concerns. Unobserved vehicle quality simultaneously determined with roadway length is an example of the second.[22]

Within-country variation. Our coefficient estimates are remarkably similar when we rely exclusively on within-country variation in table 6. Column 1 replicates our baseline speed

[22]Reverse causation and spatial sorting may also affect identification but we are less concerned. Any causal effect of travel speed on income is arguably of second-order importance (Duranton and Puga, 2023). The possible selection of fast drivers into fast locations is a concern within cities (Duranton and Turner, 2018), but unlikely to matter in our cross-city regressions.

Table 6: Robustness: Within-country impact of city attributes on speed

	(1) Speed all	(2) Speed all	(3) Speed OECD	(4) Speed Poor Countries	(5) Speed US	(6) Speed India
log country GDP (pc)	0.042^a					
	(0.011)					
log population	-0.17^a	-0.15^a	-0.19^a	-0.15^a	-0.17^a	-0.16^a
	(0.018)	(0.011)	(0.027)	(0.011)	(0.025)	(0.016)
log area	0.085^a	0.051^a	0.097^a	0.052^b	0.11^a	0.095^a
	(0.022)	(0.015)	(0.022)	(0.021)	(0.025)	(0.026)
asinh major road length	0.080^a	0.076^a	0.085^a	0.070^a	0.052^b	0.072^a
	(0.012)	(0.0096)	(0.019)	(0.0094)	(0.025)	(0.016)
Network griddiness (0-1)	0.18^a	0.16^a	0.13^a	0.26^a	0.14^a	0.24^b
	(0.057)	(0.029)	(0.0099)	(0.027)	(0.027)	(0.10)
Elevation variance	-0.0024^b	-0.0022^b	-0.0020	-0.0035^b	-0.0012	-0.029^a
	(0.00092)	(0.00092)	(0.0016)	(0.0016)	(0.0019)	(0.0064)
asinh water body length	-0.071^a	-0.040^b	-0.053^a	-0.045	-0.060^a	-0.11^a
	(0.019)	(0.017)	(0.010)	(0.029)	(0.0098)	(0.031)
Country fixed effect	N	Y	Y	Y	N	N
Observations	1,119	1,119	265	425	121	173
R^2	0.70	0.85	0.90	0.62	0.66	0.49
Within (Between) R^2		0.45 (0.55)	0.62 (0.64)	0.41 (0.45)		

Notes: The unit of analysis is the city. Each column reports a version of our full regression for the overall speed index. Column 1 repeats the baseline specification from table 5, column 2. Columns 2–4 report versions of the global, OECD and poor-country specifications with country fixed effects. We define poor countries as having below-median GDP per capita (below $12,337, which corresponds the per capita income of Paraguay.) Columns 6 and 7 are limited to the US and India, respectively. Standard errors clustered at the country level in parentheses. a, b, c: significant at 1%, 5%, 10%.

regression from column 2 of table 5. Column 2 adds a country fixed-effect to that regression. Column 3 and 4 duplicate the specification of column 2 for a sample of cities in mostly rich (OECD) countries and a sample of cities in poor countries. Columns 5 and 6 restrict the same estimation to a single large rich country (US) and a relatively poor country (India), respectively.

The coefficients in every column are uniformly consistent in sign and generally similar in magnitude. In particular, estimates from the United States, a country with few data quality concerns, are very similar to those from India. The stability of coefficients across different contexts is reassuring evidence against parameter heterogeneity, but also omitted variables, because many unobserved factors are unlikely to affect, say, the US and India in a similar way. As an example, unobserved variation in vehicle quality across cities may be a concern in India, but it is almost certainly irrelevant in the United States and within the OECD.

Within-city variation. Our city-level estimates are also consistent with the considerably more granular within-city, trip-level evidence in table 1. City-level variables like the stock of major roads could be correlated with unobserved city attributes that affect speed, but it is harder to think of factors biasing trip-level coefficients given the city fixed effects. Our baseline trip-level regression, column 4 of table 1, includes trip-level analogs of four of our six city attributes: (i) population density along the route for population, (ii) the share of each road class along the route for major road length, (iii) the frequency of intersections and turns against traffic for area, and (iv) gross gradients up and down for the variance of elevation.

We find that trips along highly populated routes are slower. In table 1, the trip-level speed elasticity of population density along a route is -0.054, which is of comparable magnitude as the city-level speed elasticity of population density of -0.085 that can be inferred from table 5.[23] Then, unsurprisingly, trips on major OSM roads (motorways, primary, secondary, and tertiary roads) are faster than trips on non-major (residential and service) roads. After adjusting for the fact that major roads account for only a fraction of total road usage, the trip-level difference in speed between major and non-major roads almost exactly matches the city-level speed elasticity of major roads of 0.08 from table 5.[24] We also find that the frequency of intersections and turns against traffic have a negative impact on trip speed. Finally, trips with a high elevation gradient are slower. Also consistent with our city-level results, we find that all of these effects operate primarily through uncongested trip speed.

5.3 *Robustness*

Robustness to different city samples, speed indices, and estimation techniques. We assess the effects of alternative criteria to select our city sample in table C.1. We consider various degrees of stringency for data quality. Also, in one regression we exclude all 209 cities that have a BRT or subway system, which do not share the road with private vehicles. We

[23] In equation (14), population and area enter separately in log. One can rewrite that equation to include a coefficient on population density (population/area) and one for population effect through which a larger population decreases speed. The speed elasticity of population density is then mechanically the negative of the coefficient on log area in (14), 0.085. Alternatively, we can try making the trip-level regression and city-level regression more comparable by controlling only for population density in the city-level regression (instead of controlling for area and population separately) and by removing controls for the frequency of intersections and turns in the trip-level regression. In that case, we find coefficients of population density on speed of 0.082 in the trip-level regression and 0.14 in the city-level regression. See tables B.3 and C.5 for regression results.

[24] To relate the trip-level speed advantage of major roads to the city-level speed elasticity of major roads, we first note that major roads account for on average 20% of total road length. Within our trip sample, major roads are 3.3 times more popular as routes than other road types, because Google preferentially routes trips through faster roads. We then estimate a trip-level regression resembling our baseline regression in column 3 of table 1, but where all major road types are aggregated together as in the city-level regression, instead of entering the regression separately as in Table 1 (see table B.3 in Appendix B.) We find that trips on major roads are on average 12% faster than trips on other types of roads. Using these numbers, the predicted city-level speed elasticity of major road is $0.20 \times 3.3 \times 0.12 = 0.079$. This prediction is almost the same as the 0.08 we find in our city-level regression.

find stable coefficients for all the regressions we estimate. These results, together with earlier results on the stability of coefficients across rich, poor, and individual countries, demonstrate that our coefficient estimates are not dependent on a particular set of cities or countries with low-quality data.

We further assess the robustness of our findings to alternative speed indices in table C.2. Relative to our baseline index estimated in column 4 of table 1, these alternative indices range from much simpler indices estimated from specifications including only city indicators to more restrictive indices estimated using only peak-hour trips or only trips within 10 kilometers of the center (to compare cities of fixed size). Again, we find stable coefficients across the regressions we estimate. This shows that our results are not sensitive to the details of how we estimate our speed index.

Our baseline city speed indices, which we use as dependent variables in table 5, are estimates from column 4 of table 1. Their true value is unknown. Because they are estimated more precisely in larger cities for which we use more trips, the errors around these indices are likely heteroscedastic. For our baseline specification in table 5, we report robust standard errors clustered by country. To further explore this problem, table C.3 reports results for alternative estimation techniques using FGLS and WLS. These techniques explicitly account for the measurement error that affects the dependent variable. The results of table C.3 show that this problem is minor. Even for the smallest city, more than 150,000 trips instances are available, so our baseline speed indices are precisely estimated.

Robustness to additional road, city, and country attributes. We first consider additional road variables in table C.4. Including a measure of speed limits increases the R-squared by 3 percentage points and explains about one fourth of the relationship between speed and income unexplained by the six city attributes of our baseline specification. A higher speed limit is associated with higher travel speed, with a large elasticity of 0.26. For instance, Flint, the fastest city in the world, has the second fastest average speed limit at 84 kilometers per hour on average. Although interesting, these results are difficult to interpret because speed limits in part reflect road quality. Reverse causality is also a serious concern.[25] Including other supplementary measures of road network characteristics such as circuitry (actual trip distance divided by straightline distance) makes no noticeable difference to our results. The removal of tertiary roads, for which we are less confident that we have a complete network, from our measure of road length reduces its coefficient by about 15%, as expected given the narrower definition.

We experimented extensively with additional city characteristics. In table C.5, we enrich our baseline specification with indicators for capital cities and for an already large popu-

[25]Speed limits in the US are determined in part using the 85th percentile of observed vehicle speeds on each road segment and the conditions of the road and around the road. See for instance `https://dot.ca.gov/programs/safety-programs/setting-speed-limits` (accessed 02/24/2023) for California.

lation in 1900, with measures of city shape (Harari, 2020), and with estimated population gradient by distance from the city center. These added variables (and others for which we do not report results) are all insignificant and none has a large impact on other coefficients.[26] These results suggest that the six characteristics of cities in our baseline specification are sufficient to capture the relevant features of cities.

We also experimented with adding a broad range of country-level attributes in table C.6. Only three suggestive results emerge. First, cities in countries whose first colonizer was France are 8.4% faster, consistent with evidence from Baruah, Henderson, and Peng (2021) that cities under French colonial rule had integrated city planning and are perhaps endowed with better roads. Then, cities in countries classified as dictatorships are 6.9% faster, consistent with Saiz' (2006) finding that road quality, measured by the share of paved roads, is higher in autocratic countries.[27] Finally, a one standard deviation increase in a country's measure of risk preference from Falk *et al.* (2018) is associated with 2.4% higher speed. While these additional findings are interesting, they are less robust than our core results and do not alter them significantly.

Overall, we conclude from this extensive set of robustness checks that our six-attribute baseline regression is robust and captures most of the important variation in the data.

5.4 Additional results: travel time unreliability

Due to our dense sampling, we can also measure unreliability in travel durations. While we expect that congestion will make travel more unreliable, conditional on mean congestion, unreliability is an additional inconvenience for drivers (Small, Winston, and Yan, 2005, Hall and Savage, 2019). Following Brownstone and Small (2005), we calculate the unreliability of a peak hour trip as the ratio of its 90th and 50th percentile durations across instances on different days at the same time of day (within a 5-minute window). Appendix D reports computational details and results analogous to tables 1 and 5. A city-level unreliability index, calculated analogously to our other city-level indices, has a rank correlation of 0.82 with our congestion index, suggesting that congested cities also tend to be unreliable. The most unreliable city is 21 percent more unreliable than the least.

[26]While historical population is uncorrelated with speed in our global sample, there is a large negative correlation within the United States, where anecdotal comparisons between older transit-based cities like New York and newer car-based cities like Los Angeles often support arguments on the historical roots of low urban travel speeds. Even within the United States however, our six city attributes explains almost all of the unconditional impact of historical population on speed. Using a more comprehensive data source for Europe (Bairoch, 1988), we also find similar but weaker evidence that older European cities are slower (see Appendix C.2).

[27]This also relates to Ades and Glaeser's (1995) argument that political considerations lead to larger central cities in dictatorships. An interaction of capital city and dictatorship is not significant.

6. Decomposing speed differences

This section shows that we can decompose the effect of income on speed into a set of indirect effects, operating through observed variables such as population, area, and road infrastructure and other effects occurring through unobserved channels. This theoretical decomposition corresponds precisely to the empirical decomposition proposed by Gelbach (2016), which we perform to quantitatively account for the importance of various city attributes in explaining why richer countries are faster.

6.1 Accounting for income via other channels

To fully decompose the overall effect of income on speed, we now briefly model the determination of city residential land area, the supply of roadway length, city population, and the link between topography and income.

Since residents consume land directly, the supply of housing is the same as the supply of urban land and is equal to

$$Q = P^{\gamma}\,, \tag{17}$$

where γ is the land supply elasticity. As the price of land increases, so does the supply of urban land. However, the supply of urban land is less than perfectly elastic because expanding cities face a variety of urban costs (Combes, Duranton, and Gobillon, 2019). Following Saiz (2010), the specification in equation (17) has been widely used in the literature.

Equating land demand, aggregated from the second first-order condition in equation (10b), and supply from equation (17) yields an equilibrium where city land area increases less than proportionately with city population and income:

$$Q = Nq = P^{\gamma} = \left[\beta\,(1-\tau) y\,N\right]^{\gamma/(\gamma+1)}\,. \tag{18}$$

In equilibrium, the land area of a city increases with its number of residents and with their demand for land, which, in turn, increases with income per capita. As a check on our model, this equation can be estimated by regressing log land area on log population and log income per capita. When constraining the coefficients for population and income to be the same, we estimate $\frac{\gamma}{\gamma+1} = 0.60$ and thus $\gamma = 1.48$. When we do not constrain the coefficient on log population and log income to be the same, we estimate a larger coefficient for population, 0.74, than for income, 0.35. These values are in line with results from the literature.[28]

[28] In their review of the literature Ahlfeldt and Pietrostefani (2019) recommend using an elasticity of city land area with respect to city population, i.e. $\frac{\gamma}{\gamma+1}$, of 0.53, close to the 0.60 we estimate. Using a very different approach, Saiz (2010) estimates that the population-weighted average elasticity of housing supply in US metropolitan areas with population greater than half a million is 1.75. This implies a value of 0.64 for $\frac{\gamma}{\gamma+1}$, the population elasticity. Again, this is close to what we estimate. More direct estimates of the land supply elasticity in Combes, Duranton, and Gobillon (2021) and Baum-Snow and Han (2021) also yield consistent results.

Roadway length L is financed by the tax τy levied on each of N residents. Its cost increases with the price of land, P. The length of roads is then given by:

$$L \equiv \frac{\tau y N}{P} = \frac{\tau}{\left[\beta\left(1-\tau\right)\right]^{1/(\gamma+1)}} \left(y N\right)^{\gamma/(\gamma+1)} . \tag{19}$$

where the last equality is obtained using equation (18). The length of the roadway increases with population and income, but less than proportionately depending on the elasticity of the supply of housing. In another check on our model, we estimate equation (19), regressing log roadway kilometers on log population and log income to obtain alternative estimates of $\frac{\gamma}{\gamma+1}$. We get a value of 0.88 from the coefficient on log population and 0.76 from the coefficient on log income. These two values are close to each other as predicted by equation (19). When we constrain these two coefficients to be equal, we estimate a value of 0.83 for $\frac{\gamma}{\gamma+1}$, which is of a similar magnitude as the value of 0.60 estimated from city land areas above.

Finally, we make two further reduced-form assumptions to close our decomposition. First we assume that the supply of population in a city is less than fully elastic and given by

$$N = V^{\chi} , \tag{20}$$

where χ is the elasticity of city population with respect to city utility.[29] Second, income and city topography, which determines travel productivity, are also related through

$$A = \Omega_A \, y^{\Gamma_A} . \tag{21}$$

This expression is needed for the decomposition we perform next. In Appendix E, we provide more complex decompositions that do not use this expression and instead treat the topography of cities as exogenous.

6.2 Decomposing the impact of country income on speed

We can now use the model leading to equations (13)-(16) and the additional assumptions contained in equations (17)-(21) to account for the difference in speed between rich and poor countries. To do so, we decompose the impact of income on speed into contributions arising from each of the city attributes we consider and a residual contribution from unobserved road quality and congestion captured through income. The contribution of observed city attributes depends on the product of the speed elasticity of each attribute, derived in equation (13), and the income elasticity of each attribute, which we derive in this subsection.

[29] Underpinning imperfect labor mobility, the literature often assumes rising housing costs (Combes *et al.*, 2019), mobility costs (Kennan and Walker, 2011), or idiosyncratic preferences for each city (e.g. Moretti, 2011, Diamond, 2016). We do not formally model agglomeration effects but they could be embedded into our model by adding a feedback loop where the income of residents would depend on city population.

This approach generates a system of linear equations that we estimate in the next section. They are equivalent to the exact decomposition proposed by Gelbach (2016).[30]

In Appendix E.1 we show how city population, road length, and area can all be written as functions of income and parameters.[31] We note that equation (21) already relates city topography to income directly. After taking logs, we obtain the following set of regressions:

$$\begin{cases} \ln A = \ln \Omega_A + \Gamma_A \ln y \,, & \text{(22a)} \\ \ln N = \ln \Omega_N + \Gamma_N \ln y \,, & \text{(22b)} \\ \ln L = \ln \Omega_L + \Gamma_L \ln y \,, & \text{(22c)} \\ \ln Q = \ln \Omega_Q + \Gamma_Q \ln y \,, & \text{(22d)} \end{cases}$$

where the constants Ω_L, Ω_N, and Ω_Q and the income elasticities Γ_L, Γ_N, and Γ_Q are all (complex) functions of parameters. Following Gelbach (2016), we refer to these equations as the 'auxiliary' regressions.

Finally, returning to the base regression (2) estimated above, $\ln S = \ln \Omega_{\text{Base}} + \kappa_{\text{Base}} \ln y$, we can show that:

$$\kappa_{\text{Base}} = \psi_A \Gamma_A + \psi_N \Gamma_N + \psi_L \Gamma_L + \psi_Q \Gamma_Q + \kappa_{\text{Full}} \,, \tag{23}$$

where κ_{Full} is the income elasticity of speed in regression (13) and the ψs are the elasticities for the other explanatory variables in the same regression. The equality in equation (23) can be verified using equation (13) and substituting for A, N, L, and Q using equations (21) and (E1)–(E3) in Appendix E and taking logs. Empirically, Gelbach (2016) shows that if we estimate equations (14), (22a)–(22d), and (2) by OLS, they also form an exact empirical decomposition. Simply put, we can add "hats" to each element of the decomposition in equation (23).

Intuitively, equation (23) decomposes the total measured effect of income on speed, κ_{Base}, into a sum of contributions through each city attribute and a residual contribution κ_{Full} measured directly through income. As in the decomposition methods familiar from labor economics (e.g. Fortin, Lemieux, and Firpo, 2011), the contribution of a city attribute k (for

[30] For recent prominent uses of the Gelbach (2016) decomposition, see Allcott *et al.* (2019) on the nutrition-income relationship, Bandiera *et al.* (2020) on the effects of a female empowerment program, Stantcheva (2021) on the partisan gap and support for taxation, and Cook *et al.* (2021) on the gender pay gap for Uber drivers.

[31] We treat the tax rate to fund roads as parametric for several reasons. First, we expect this rate to be small given the small share of road infrastructure in national income. The equations in Appendix E.1 show that τ generally appears as $(1-\tau)$ or in terms that we expect to be small. Second, we could easily derive an optimal tax, from the trade-off between more and better roads and a reduction in the consumption of other goods. Given our multiplicative specification for utility in equation (8), the optimal tax rate is a fixed share of income, which depends on the model's parameters. Finally, many countries, including highly decentralized countries like the US and Canada, have a federal road system and federal funding to match. Empirically, a country-specific tax to fund roads gets absorbed into a country fixed effects in our complementary regressions in section 5.2 and we show that this makes little to no difference for the coefficients we estimate for city attributes.

$k = A, N, L, Q$) to the difference in travel speed between richer and poorer countries is the product of the speed elasticity of this attribute (ψ_k) and the income elasticity of this attribute Γ_k. So the importance of, say, land area (Q) in explaining why richer countries are faster depends both on the impact of land area on speed and the extent to which richer countries' cities have larger land areas. Together, $\Gamma_k \psi_k / \kappa_{\text{Base}}$ is the share of the impact of income on speed explained by attribute k. Our model gives the residual (or unexplained) income effect, κ_{Full}, a structural interpretation as $\phi\sigma - \theta$, where $\phi\sigma$ captures the impact on speed of better road technology, while θ captures higher congestion in richer countries.

Finally, we can similarly decompose speed into uncongested speed and congestion, based on the same auxiliary regressions, and on the full regressions for uncongested speed and congestion in equation (15) and (16).[32] With the same set of explanatory variables in the speed, uncongested speed, and congestion regressions, the additive properties of our speed indices ($\mathbb{S}_c = \mathbb{U}_c - \mathbb{K}_c$) imply, for instance, that $\hat{\psi}_N^S = \hat{\psi}_N^U - \hat{\psi}_N^K$. So, as in the model, the estimated contribution of an attribute to explaining the speed-income relationship is the exact sum of its contributions to the uncongested speed-income relationship and to the congestion-income relationship.

6.3 Decomposition results

Table 7 reports results for the auxiliary regressions described by equations (22a)–(22d). These are univariate regressions of city attributes on log income per capita. For instance, the coefficient of -0.10 for log population shows that, in our sample, cities in richer countries have somewhat smaller populations, but the difference is only marginally significant. These auxiliary regressions suggest that cities in richer countries also have substantially wider areas (so population density is lower), much more major roads, somewhat more water bodies, and similar variance in elevation and network griddiness.

Figure 4 shows the decomposition results: the share of the total effect of country GDP per capita on speed explained by each city attribute k, through their impact on uncongested mobility (blue bar; $\Gamma_k \psi_k^U / \kappa_{\text{Base}}$) and congestion (red bar; $\Gamma_k \psi_k^K / \kappa_{\text{Base}}$). For each attribute, these two bars add up to its total contribution in explaining why richer countries are faster (i.e. to $\Gamma_k \psi_k / \kappa_{Base}$ from equation 23).

Two first-order patterns emerge from this decomposition. First, cities in richer countries are faster mostly because they have more major roads and cover more land area. The contribution of roads is more than twice as large as that of land area.[33] Together these

[32] In the model, the coefficients on Q and L in the full congestion regression are zero. Rather than imposing this restriction, we verify it by including all city attributes in the congestion regression.

[33] The contribution of population is not negligible, but we are cautious in its interpretation, because the negative correlation between city population and country GDP is mainly driven by our sample selection.

Table 7: Gelbach auxiliary regressions

	(1) log population	(2) log area	(3) asinh major road length	(4) Network griddiness	(5) Elevation variance	(6) asinh water body length
log country GDP (pc)	-0.10^b	0.28^a	0.67^a	0.033	-0.14	0.14^a
	(0.042)	(0.076)	(0.074)	(0.028)	(0.36)	(0.033)
Observations	1,119	1,119	1,119	1,119	1,119	1,119
R^2	0.01	0.07	0.27	0.05	0.00	0.07

Notes: Robust standard errors clustered by country are in parentheses. a, b, c: significant at 1%, 5%, 10%.

Figure 4: Which city attributes explain why richer countries are faster?

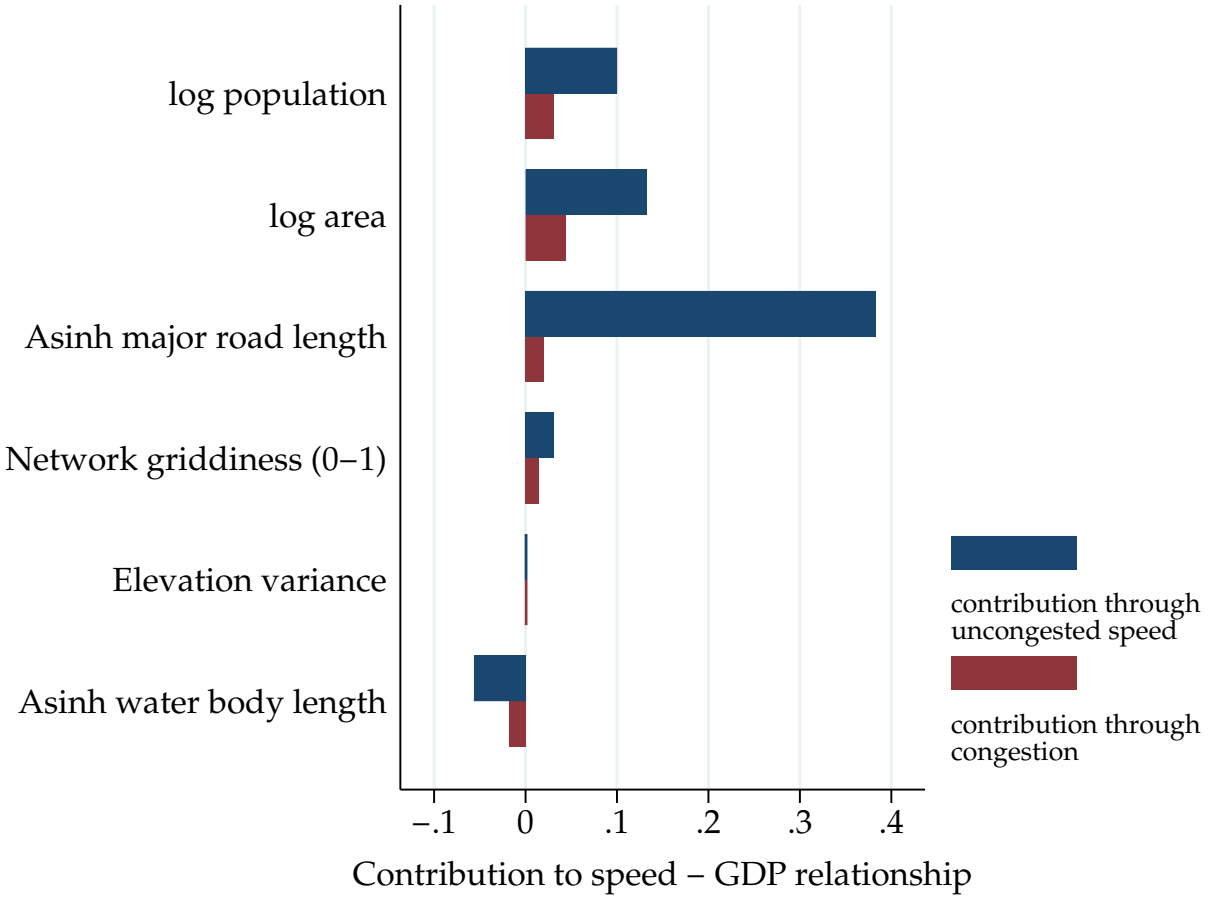

Notes: The blue and red bars are our estimates of $\Gamma_k \psi_k^U / \kappa_{\text{Base}}$ and $\Gamma_k \psi_k^K / \kappa_{\text{Base}}$ from versions of equation (2) specific to uncongested speed and congestion, respectively. They represent the share of the speed-GDP elasticity that is explained by each city characteristic through uncongested speed and congestion, respectively.

two city attributes account for 58% of the effect of country GDP per capita on speed. In other words, infrastructure is more important than city area, and much more important than topography in explaining why richer countries are faster. Second, all city attributes affect speed mainly through their impact on uncongested speed, not through their impact on congestion.

6.4 Why is the United States so fast and Bangladesh so slow?

Our data and decomposition methodology lend themselves to an array of possible investigations, within and across countries. Here, we pursue two such investigations, which focus on the fastest (United States) and slowest (Bangladesh) countries in the world. Specifically, we ask why the US is faster than other rich (OECD) countries, and why Bangladesh is slower than other poor countries with a below-median income per capita.

The regression results for the US vs. OECD decomposition are in panel A of Table 8, and those for the Bangladesh vs. poor countries decomposition are in panel B. The coefficient on the US indicator in the base regression in column 1 implies that the United States is $(\exp(0.24) - 1 =)$ 27% faster than other OECD countries, and the decrease of that coefficient to 0.044 in the full regression of column 2 suggests that our city attributes explain 82% of the speed difference between the US and other OECD countries. A similar computation in panel B suggests that Bangladesh is 20% slower than other poor countries, and our city attributes explain about 75% of that speed difference.

The coefficients on city attributes in both panels are generally close to each other and to those estimated for the whole world in table 5. However, the auxiliary regressions differ widely across the two panels of figure 5, as expected. The auxiliary regressions in panel A suggest that cities in the United States are very different from cities in other OECD countries in ways that make them faster. Compared to other cities in the OECD, US cities are $(\exp(-0.27) - 1 =)$ 24% less populous, cover 72% more area, have 67% more major roads, and have 30% more roads that conform to the road network's main grid orientation. Panel A of figure 5 reports the corresponding decomposition result. The low density of US cities explains most of why they are faster. Major roads matter too, as do griddier road networks. City size variables account for 47% of the speed difference between the US and other OECD countries, infrastructure accounts for 35%, and the share accounted for by topography is negligible.

Similar auxiliary regressions in panel B suggest that cities in Bangladesh are also strikingly different from cities in other poor countries, but on somewhat different attributes, and in ways that make them slower. Relative to cities in other poor countries, cities in Bangladesh are 40% more populous (with similar areas), are crossed by 116% more water

Table 8: Additional Gelbach decompositions

Panel A: United States vs. OECD countries

	Speed index		Uncongested speed		Congestion factor		
	(1) Base	(2) Full	(3) Base	(4) Full	(5) Base	(6) Full	(7) Auxiliary
United States	0.24^a	0.044^b	0.20^a	0.060^a	-0.033^a	0.016^b	
	(0.021)	(0.020)	(0.018)	(0.017)	(0.0054)	(0.0061)	
log population		-0.22^a		-0.16^a		0.060^a	-0.27^a
		(0.026)		(0.022)		(0.0050)	(0.091)
log area		0.100^a		0.084^a		-0.016	0.54^a
		(0.033)		(0.027)		(0.016)	(0.086)
asinh major road length		0.10^a		0.074^b		-0.028^b	0.51^a
		(0.035)		(0.031)		(0.013)	(0.097)
Network griddiness (0-1)		0.12^a		0.069^a		-0.048^a	0.26^a
		(0.022)		(0.020)		(0.0045)	(0.0096)
Elevation variance		-0.0032^b		-0.0039^b		-0.00079	-0.81
		(0.0014)		(0.0018)		(0.00054)	(0.54)
asinh water body length		-0.037^c		-0.029^b		0.0079	0.16^b
		(0.019)		(0.013)		(0.0070)	(0.072)
Observations	265	265	265	265	265	265	265
R^2	0.56	0.84	0.61	0.83	0.18	0.63	

Panel B: Bangladesh vs. poor countries

	Speed index		Uncongested speed		Congestion factor		
	(1) Base	(2) Full	(3) Base	(4) Full	(5) Base	(6) Full	(7) Auxiliary
Bangladesh	-0.22^a	-0.054	-0.20^a	-0.070^a	0.019^b	-0.017	
	(0.016)	(0.033)	(0.014)	(0.017)	(0.0074)	(0.019)	
log population		-0.15^a		-0.13^a		0.021^a	0.34^a
		(0.015)		(0.014)		(0.0066)	(0.095)
log area		0.059^b		0.030		-0.029^a	-0.13
		(0.027)		(0.022)		(0.0100)	(0.10)
asinh major road length		0.079^a		0.100^a		0.021^a	-0.56^a
		(0.013)		(0.013)		(0.0076)	(0.14)
Network griddiness (0-1)		0.13^b		0.060		-0.066^a	-0.014
		(0.057)		(0.052)		(0.022)	(0.010)
Elevation variance		-0.0030^a		-0.0017		0.0013^c	-1.56^b
		(0.00090)		(0.0013)		(0.00071)	(0.65)
asinh water body length		-0.089^a		-0.040^b		0.050^a	0.77^a
		(0.027)		(0.019)		(0.018)	(0.034)
Observations	425	425	425	425	425	425	425
R^2	0.05	0.45	0.06	0.44	0.00	0.37	

Notes: In panel A, the sample is all cities in OECD countries for which we have a GDP measure. In panel B, the sample is all cities in countries with below-median GDP per capita as in table 6. Columns 1–6 each report parameter estimates of separate regressions. Column 7 report coefficients from separate univariate regressions of log country GDP per capita on the attribute in the same row. Standard errors clustered at the country level in parentheses. a, b, c: significant at 1%, 5%, 10%.

Figure 5: Which city attributes explain why the United States is so fast and Bangladesh so slow?

Panel A: United States vs. OECD countries

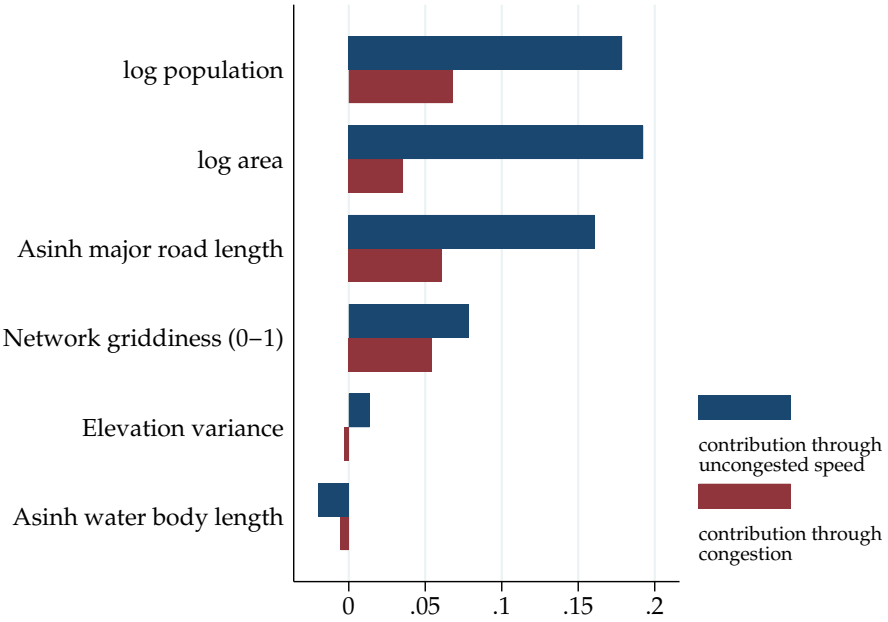

Contribution to speed – United States relationship

Panel B: Bangladesh vs. poor countries

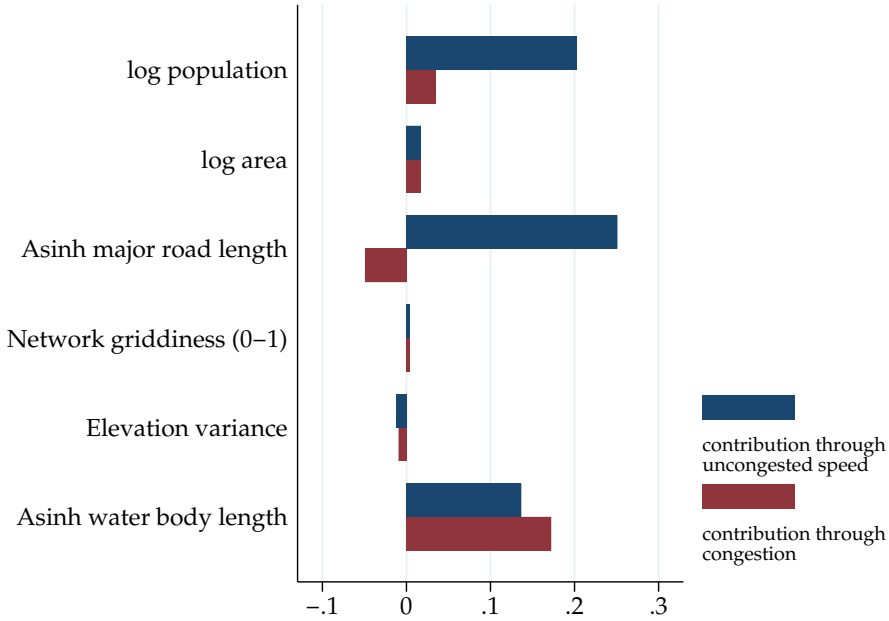

Contribution to speed – Bangladesh relationship

Notes: The blue and red bars are our estimates of $\Gamma_k \psi_k^U / \kappa_{\text{Base}}$ and $\Gamma_k \psi_k^K / \kappa_{\text{Base}}$ from versions of equation (2) specific to uncongested speed and congestion, respectively. They represent the share of the speed-GDP elasticity that is explained by each city characteristic through uncongested speed and congestion, respectively.

bodies, and possess 42% fewer major roads. Panel B of Figure 5 reports that all three of these city attributes play a significant role in making cities in Bangladesh slower. Water bodies (topography) is most important, accounting for 32% of the speed difference between cities in Bangladesh and cities in other poor countries, followed by higher population (23%) and fewer roads (20%).[34]

This exercise demonstrates that our urban model and its parsimonious baseline empirical implementation, with only six city attributes, can account for urban travel speed differences across a wide variety of empirical contexts. Notably, it accounts for almost all the speed difference between the US and other OECD countries, and between Bangladesh and other poor countries.

7. Conclusion

In this paper, we investigate the determinants of urban travel speed. We estimate travel speed and congestion indices for more than 1,200 large cities in 152 countries. We find that most variation in speed occurs between countries, not within them, and that national income per capita can explain a large share of the cross-country variation in speed. As such, we are the first to document faster urban travel speed as a robust feature of economic development.

We then develop an urban model that decomposes this speed-income relationship into the contributions of city attributes that represent their size, infrastructure, and topography. We find that cities in richer countries are faster primarily because they have more major roads, and secondarily because their road networks spread over larger urban areas. Strikingly, these attributes make richer countries faster by raising uncongested speeds, not by decreasing congestion. Through the lens of our model, higher income per capita leads to more tax revenues, housing consumption, and demand for travel. Consequently, cities in richer countries spend more on fast roads and span larger urban areas, while experiencing more travel and thus more road congestion.

Our detailed data and framework can also be used to study specific countries and cities. For example, the six city attributes in our parsimonious baseline model account for almost all the speed difference between the US and other OECD countries, and between Bangladesh and other poor countries. Future research could easily extend our analysis to investigate the effect on speed of many detailed features of road networks beyond those we study here.

These results have several policy implications. While we do not rule out the potential importance of policies directed against congestion in selected large cities, such policies are unlikely to impact the speed gap between rich and poor countries. Building more major roads, even keeping quality constant, would be more impactful, albeit obviously more

[34]We note that cities in Bangladesh also differ from other poor cities by being almost entirely flat, but low elevation variance is not an important determinant of speed.

costly. However, in some contexts, road infrastructure is not the main reason why countries or cities are slow. The ability of Bangladesh to improve urban transportation outcomes appears severely constrained by its challenging topography, large urban populations and limited land area.[35] These results highlight the inherent limitations that hard-to-change city attributes impose on urban transportation. That said, the process of economic development generally improves transportation outcomes, as countries build more roads, expand their urban areas, and develop better road building technology, while incurring only slightly more congestion. We hope that similar analyses, tailored to different case studies, can offer policy makers, in poor and rich countries alike, a set of basic facts about the performance of their urban transportation network.

References

Ades, Alberto F. and Edward L. Glaeser. 1995. Trade and circuses: Explaining urban giants. *Quarterly Journal of Economics* 110(1):195–227.

Ahlfeldt, Gabriel M., Thilo Albers, and Kristian Behrens. 2022. Prime locations. CEPR Discussion Paper No. DP15470.

Ahlfeldt, Gabriel M. and Elisabetta Pietrostefani. 2019. The economic effects of density: A synthesis. *Journal of Urban Economics* 111(0):93–107.

Ahlfeldt, Gabriel M., Stephen J. Redding, Daniel M. Sturm, and Nikolaus Wolf. 2015. The economics of density: Evidence from the Berlin Wall. *Econometrica* 83(6):2127–2189.

Akbar, Prottoy A., Victor Couture, Gilles Duranton, and Adam Storeygard. 2023. Mobility and congestion in urban India. *American Economic Review* 113(4):1083–1111.

Akbar, Prottoy A. and Gilles Duranton. 2022. Measuring congestion in a highly congested city: Bogotá. Processed, University of Pennsylvania.

Allcott, Hunt, Rebecca Diamond, Jean-Pierre Dubé, Jessie Handbury, Ilya Rahkovsky, and Molly Schnell. 2019. Food deserts and the causes of nutritional inequality. *Quarterly Journal of Economics* 134(4):1793–1844.

Anderson, Michael L. and Lucas W. Davis. 2020. An empirical test of hypercongestion in highway bottlenecks. *Journal of Public Economics* 187(104197).

Angel, Shlomo, Alejandro M. Blei, Jason Parent, Patrick Lamson-Hall, Nicolas Galarza-Sanchez, Daniel L. Civco, Rachel Q. Lei, and Kevin Thom. 2016. *Atlas of urban expansion — 2016 edition*. Cambridge (MA): Lincoln Institute of Land Policy.

Bairoch, Paul. 1988. *Cities and economic development: From the dawn of history to the present*. Chicago, Ill.: University of Chicago Press.

[35]Though see Bird, Li, Rahman, Rama, and Venables (2018) for a proposal to expand Dhaka's land area.

Bandiera, Oriana, Niklas Buehren, Robin Burgess, Markus Goldstein, Selim Gulesci, Imran Rasul, and Munshi Sulaiman. 2020. Women's empowerment in action: Evidence from a randomized control trial in Africa. *American Economic Journal: Applied Economics* 12(1):210–259.

Baruah, Neeraj G, J Vernon Henderson, and Cong Peng. 2021. Colonial legacies: shaping african cities. *Journal of Economic Geography* 21(1):29–65.

Baum-Snow, Nathaniel and Lu Han. 2021. The microgeography of housing supply. Manuscript, University of Toronto.

Ben-Akiva, Moshe E. and Steven R. Lerman. 1985. *Discrete Choice Analysis: Theory and Application to Travel Demand*. Cambridge, MA: MIT Press.

Bird, Julia, Yue Li, Hossain Zillur Rahman, Martin Rama, and Anthony J. Venables. 2018. *Toward Great Dhaka: A New Urban Development Paradigm Eastward*. Washington, DC: The World Bank.

Boeing, Geoff. 2017. OSMNX: New methods for acquiring, constructing, analyzing, and visualizing complex street networks. *Computers, Environment and Urban Systems* 65(1):126–139.

Bondarenko, Maksym, David Kerr, Alessandro Sorichetta, and Andrew J. Tatem. 2020. Census/projection-disaggregated gridded population datasets, adjusted to match the corresponding unpd 2020 estimates, for 183 countries in 2020 using built-settlement growth model (BSGM) outputs. WorldPop, University of Southampton.

Borker, Girija. 2020. Safety first: Perceived risk of street harassment and educational choices of women. Processed, World Bank.

Bou Sleiman, Léa. 2022. Displacing congestion: Evidence from Paris. Preprint, CREST - Ecole Polytechnique.

Brownstone, David and Kenneth A. Small. 2005. Valuing time and reliability: Assessing the evidence from road pricing demonstrations. *Transportation Research Part A* 39(4):27–293.

Combes, Pierre-Philippe, Gilles Duranton, and Laurent Gobillon. 2019. The costs of agglomeration: House and land prices in French cities. *Review of Economic Studies* 86(4):1556–1589.

Combes, Pierre-Philippe, Gilles Duranton, and Laurent Gobillon. 2021. The production function for housing: Evidence from France. *Journal of Political Economy* 129(10):2766–2816.

Conwell, Lucas, Fabian Eckert, and Ahmed Mushfiq Mobarak. 2022. More roads or public transit? insights from measuring city-center accessibility. Processed, Yale University.

Cook, Cody, Rebecca Diamond, Jonathan V. Hall, John A. List, and Paul Oyer. 2021. The gender earnings gap in the gig economy: Evidence from over a million rideshare drivers. *Review of Economic Studies* 88(5):2210–2238.

Corbane, Christina, Aneta Florczyk, Martino Pesaresi, Panagiotis Politis, and Vasileios Syrris. 2018. GHS built-up grid, derived from Landsat, multitemporal (1975-1990-2000-2014), R2018A. European Commission, Joint Research Centre (JRC).

Couture, Victor, Gilles Duranton, and Matthew A. Turner. 2018. Speed. *Review of Economics and Statistics* 100(4):725–739.

Diamond, Rebecca. 2016. The determinants and welfare implications of US workers' diverging location choices by skill: 1980-2000. *American Economic Review* 106(3):479–524.

Duranton, Gilles and Diego Puga. 2023. Urban growth and its aggregate implications. *Econometrica* (forthcoming).

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. *Review of Economic Studies* 79(4):1407–1440.

Duranton, Gilles and Matthew A. Turner. 2018. Urban form and driving: Evidence from US cities. *Journal of Urban Economics* 108(0):171–191.

Durlauf, Steven N., Paul A. Johnson, and Jonathan W. Temple. 2005. Growth econometrics. In Philippe Aghion and Steven N. Durlauf (eds.) *Handbook of Economic Growth*, volume 1A. Amsterdam: North-Holland, 555–677.

Eurostat. 2021. Passenger mobility statistics. *Statistics Explained* .

Falk, Armin, Anke Becker, Thomas Dohmen, Benjamin Enke, David Huffman, and Uwe Sunde. 2018. Global evidence on economic preferences. *Quarterly Journal of Economics* 133(4):1645–1692.

Feenstra, Robert Inklaar, Robert C. and Marcel P. Timmer. 2015. The next generation of the Penn World Table. *American Economic Review* 105(10):3150–3182.

Fortin, Nicole, Thomas Lemieux, and Sergio Firpo. 2011. Decomposition methods in economics. In Orley Ashenfelter and David Card (eds.) *Handbook of Labor Economics*, volume 4A. Amsterdam: North-Holland, 1–102.

Fuller, Brandon and Paul Romer. 2014. Urbanization as opportunity. World Bank Policy Research Paper 6874.

Gelbach, Jonah B. 2016. When do covariates matter? And which ones, and how much? *Journal of Labor Economics* 34(2):509–543.

Glaeser, Edward L. 2020. Infrastructure and urban form. Technical report, National Bureau of Economic Research Working Paper #28287.

Gwilliam, Ken. 2003. Urban transport in developing countries. *Transport Reviews* 23(2):197–216.

Hall, Jonathan D. and Ian Savage. 2019. Tolling roads to improve reliability. *Journal of Urban Economics* 113:103187.

Harari, Mariaflavia. 2020. Cities in bad shape: Urban geometry in India. *American Economic Review* 110(8):2377–2421.

Heblich, Stephan, Stephen J. Redding, and Daniel M. Sturm. 2020. The making of the modern metropolis: Evidence from London. *Quarterly Journal of Economics* 135(4):2059–2133.

Kennan, John and James R Walker. 2011. The effect of expected income on individual migration decisions. *Econometrica* 79(1):211–251.

Kreindler, Gabriel. 2023. Peak-hour road congestion pricing: Experimental evidence and equilibrium implications. *Econometrica* Accepted for publication.

Li, Yue, Martin Rama, Virgilio Galdo, and Maria Florencia Pinto. 2015. A spatial database for South Asia. Unpublished manuscript.

Moretti, Enrico. 2011. Local labor markets. In Orley Ashenfelter and David Card (eds.) *Handbook of Labor Economics*, volume 4B. Amsterdam: North-Holland, 1237–1313.

Organisation for Economic Co-operation and Development (OECD). 2021. Metropolitan areas database. `https://stats.oecd.org/Index.aspx?DataSetCode=CITIES`. Accessed: 2021-05-29.

Pesaresi, Martino and Sergio Freire. 2016. GHS-SMOD R2016A - GHS settlement grid, following the REGIO model 2014 in application to GHSL landsat and CIESIN GPW V4-multitemporal (1975-1990-2000-2015). European Commission, Joint Research Centre.

Saiz, Albert. 2006. Dictatorships and highways. *Regional Science and Urban Economics* 36(2):187–206.

Saiz, Albert. 2010. The geographic determinants of housing supply. *Quarterly Journal of Economics* 125(3):1253–1296.

Saiz, Albert and Luyao Wang. 2023. Physical geography and traffic delays: Evidence from a major coastal city. *Environment and Planning B: Urban Analytics and City Science* 50(1):218–243.

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

Small, Kenneth A., Clifford Winston, and Jia Yan. 2005. Uncovering the distribution of motorists' preferences for travel time and reliability. *Econometrica* 73(4):1367–1382.

Stantcheva, Stefanie. 2021. Understanding tax policy: How do people reason? *Quarterly Journal of Economics* 136(4):2309–2369.

Tsivanidis, Nick. 2019. Evaluating the impact of urban transit infrastructure: Evidence from Bogotá's Transmilenio. Processed, University of California Berkeley.

United Nations, Department of Economic and Social Affairs, Population Division. 2019. *World Urbanization Prospects: The 2018 Revision (*ST/ESA/SER.A/420*)*. New York: United Nations.

Appendix A. Data construction

1 City sample and boundaries

Our city universe and definitions come from two sources. City point locations and populations are from the World Urbanization Prospects, 2018 revision (WUP; United Nations, 2019), which contains all 1,860 cities with a (projected) population of at least 300,000 in 2018. City boundaries are based on two datasets from the Global Human Settlements layer version 2016A reporting conditions circa 2014-15: Settlement Model (SMOD; Pesaresi and Freire, 2016), which defines each 1-kilometer grid cell as urban or not, and the 38-meter built-up grid (BUILT). Closely following the approach in Akbar *et al.* (2023), we combine these two datasets in several steps:

1. We split seven "twin" cities that WUP treats as one, but we believe are better conceptualized as two for our purposes. We allocate the total population of the pair to each member based on external population sources. After this, the sample contains 1,867 city points.

2. We alter the location of 184 of these 1,867 cities (including the seven pairs above) as follows. We first compare the coordinates of each city with those reported for the same city (based on a search for name and country) by Google Maps (GM). For any city where the discrepancy is greater than five kilometers, we flag these cities and manually search those cities to determine which is more appropriate. If neither is appropriate, for example, because the GM search returns a smaller place of the same name, we pick another point (train station, historical center) in GM. We further use GM for all cities in the United States.

3. We assign the nearest set of contiguous SMOD cells (SMOD polygon hereafter) to each WUP point.

4. We drop (a) 4 cities that would have had too small a population (well under 300,000) if we carried out an analogous split in step 1 above, (b) 21 cities more than 8.2 kilometers from a SMOD polygon after the spatial join in step 3, and (c) 12 other cities whose SMOD polygon is implausible in some other way. After this, our sample contains 1,830 city points, but fewer associated polygons, because some polygons are the nearest polygon of multiple points.

5. We move WUP points that are outside of but less than 8.2 kilometers from the nearest SMOD polygon to a more appropriate location within the polygon. This does not change the sample but creates a more plausible city center for trip sampling.

6. We merge WUP points that fall within the same SMOD polygon in some cases where we deem it appropriate because they are essentially subcenters of a larger metro area, summing their populations and using the point location of the largest city in the SMOD polygon. 35 points are thus subsumed into larger neighbors, leaving us with 1,795 points.

7. In the remaining cases where multiple WUP points fall within the same SMOD polygon, we split the polygon into smaller polygons, one surrounding each point. Where relevant, we split the polygon at a national border or a large water body (river/strait/bay). In cases where there is no obvious border (international border or large body of water), we split polygons so that the resulting areas assigned to each city are proportional to its population, with an elasticity of 0.57, based on Ahlfeldt and Pietrostefani (2019). This approach results in a few cases where a city has a secondary area that is connected to the primary area only through another city. In these cases we reassign the exclave to the city to which it is adjacent. Thus, 65 polygons containing 176 points are divided into 176 polygons. This leaves us with 1,795 polygons.

8. We drop all cities in China (but not Hong Kong or Macao), because a priori we know that GM estimates do not include traffic there. The resulting set of 1,371 cities is queried in GM.

9. We select all BUILT pixel centroids within each resulting city polygon, and define a 500-meter buffer around them as the final city polygon.

10. We calculate city population as the sum of population in 100-meter pixels according to WorldPop (Bondarenko *et al.*, 2020) within these boundaries. Bondarenko *et al.* (2020) reallocates the population within census tracts to pixels with attributes from satellite data, among other sources, correlated with higher population density. We drop 130 cities with a population less than half of the initial 300,000 threshold. Our exclusion is based on the idea that these are actually small towns assigned large populations due to especially wide administrative boundaries that WUP implicitly adopts.

2 *Trips data*

Target trip sample

We follow the strategy outlined in Appendix A of Akbar *et al.* (2023) to sample trips (origin-destination pairs) between the BUILT pixel centroids within our city polygons and then query these driving trips on Google Maps (GM). In particular, within each city, we target $15\sqrt{Population_c}$ trips, each to be queried approximately 24 times during different times of day and days of the week over several months. That corresponds to roughly 151,000

trip instances in the smallest city in our sample (Kabinda, Congo), 281,000 trip instances in a median-sized city, and 2,204,000 trip instances in our largest city (Tokyo, Japan). Across all 1,241 cities, this sampling strategy is expected to yield 19,236,544 trips and 461,677,056 driving trip instances in total. As we note below, we slightly overshot this target with a median of 26 instances and an average of 31 driving instances per trip.

The time-of-day to query a trip is drawn from a distribution that roughly resembles the distribution of departure times on a typical weekday, as observed in the 2017 US National Household Travel Survey (NHTS). To make sure that we have enough trip instances in every hour of the day, we tweak the distribution to make sure that the quietest hours of the day have no fewer than one-fifth of the trip instances at the busiest hours of the day. To be able to study reliability, we force 14 instances of each trip to be queried at the same time of day on consecutive days for two weeks. For nine other instances of each trip, we let the time of query be drawn independently for each instance (from the same distribution of departure times). We query these nine trip instances half as frequently on weekends relative to weekdays.[36] Finally, for one trip instance (which we query at the same time of day as the reliability sample), we scrape not just the aggregate travel duration and length of each route returned by GM but also the location and content of each step in the directions along the trip's routes (see Appendix A.4 below for details).

In addition to the driving queries, we target four additional instances of each trip: two by "transit" (as denoted by GM) and two by walking. These queries are scheduled to coincide with the query time of day of two of the same trip's driving instances. For transit, we target one trip instance during the day (7 AM - 7 PM) and one trip instance during the night (7 PM - 7 AM).[37] Within these intervals, the query times are randomized such that the overall distribution of query times (within a day) resembles that of the driving trips. As we anticipate systematic differences in transit schedules between weekdays and weekends, our transit queries are restricted to weekdays. The query times on the walking trip instances are drawn from the same distribution of departure times within a day (but not restricted by day of week). On one of the two walking instances of each trip, we force the route to follow the path of one of the driving routes. We do so to collect information on elevation changes (only returned on walking queries to GM) along the driving routes. On all walking and transit queries, we scrape both the aggregate route duration and length, as well as information on the step segments along the route.

[36] We identify the days of the week that constitute the weekend separately for each country in our sample.

[37] If trips have no driving instance during the night (which happens in a few cases), we schedule both transit instances during the day.

Querying trips on GM

The large number of targeted trip instances requires us to be able to make millions of queries on GM per day. We also need to moderate computing power based on varying querying needs at different hours of the day. So, we employ elastic computing and cloud storage services on Amazon Web Services (AWS) for the task. We moderate the frequency of queries at different times of day and days of week by manipulating the number of AWS processors that we have actively querying GM at any point in time. However, this process is imprecise because we can only modify the rate of query in discrete steps.[38]

We do not expect GM to return a route for all our trip queries. Our trip sampling strategy makes it unlikely, but not impossible, to define trips through areas without a viable road network or areas where GM has limited traffic data (see `https://developers.google.com/maps/coverage`).

Query times and days

Our driving queries on GM span 6 months: 12 June to 5 November 2019. The bulk of the queries take place in the earlier months: 37% in June, 19% in July, 16% in August, 15% in September, 12% in October, and 1% in November. Despite our efforts to query some trip instances much less frequently on weekends, the distribution of queries across days of the week is roughly uniform, with the most queries taking place on Thursdays (17.2%) and the fewest queries on Sundays (13.2%).

Our first round of walking queries spans 8 October through 14 December 2019. Our walking queries to GM differ from driving queries in that they also return data on elevation changes along the route. So, to measure upward and downward gradients along our driving routes, in August 2020, we conducted a second round of walking queries where we force the routes to go through points along the corresponding driving route of the trip.

Our transit queries span 11 July through 5 August of 2022. For transit trip instances, we record the departure time on the returned routes since they could differ from the time of query. We are able to distinguish the time spent walking and on transit for each leg of the routes. We can also distinguish between transit modes (such as "Walk", "Bus", "Train", etc.) for each leg of the routes, and compute the time spent waiting between legs.

[38]There are two reasons for this imprecision. First, the number of processors can only be changed discretely, and we do not throttle individual processors' querying speeds (which would be a waste of the computing power we are paying for). Second, our trips span many time zones and each time zone requires a different rate of querying at any given moment. So, in time zones with few processors assigned to them, the variation in query rate is less smooth across different times of the day.

Failed trip queries

We make 583,140,802 driving queries on GM (slightly more than planned) of 18,986,254 trips.[39] We drop all 13 South Korean cities from our sample because less than 1% of the queries there return a route. Outside of South Korea, only 169,336 (0.03%) of the queries do not return any driving route.

Across the remaining 1,228 cities, GM fails to return a route for a total of only 7,002 trips. These failed trip queries are mostly concentrated in a few cities: Macao (73% of trips fail), Sinuiju, North Korea (48% of trips fail), Jerusalem, Israel (13% of trips fail), and Gaza, Palestine (10% of trips fail). Of these, only Jerusalem survive our subsequent data quality checks (described in Section 2.2) to remain in the main sample of cities. In every other city, over 95% of our target trips are queried successfully on GM. In the average city, we successfully query 99.9% of targeted trips (with a coefficient of variation of 0.2%).

Successfully queried trip instances

In the resulting sample of successfully queried trips, the mean and median number of driving instances per trip are 31 and 26, and 99% of the trips have at least 17 instances. Table A.1 reports the counts of trip instances, cities and countries that we manage to successfully query on GM from each continent. Because the number of trips we define in each city scales less than linearly with population, so do the number of trip instances. We also report the counts for the three largest countries in our sample: India, US, and Indonesia.

Each successful trip instance query on GM returns information based on one to five routes. We only retain information on the route with the shortest travel time (under real-time traffic). Panel A of Table A.2 shows the distribution of trip length, duration, and duration in absence of traffic that are reported by GM for our sample of trip instances. The average (median) trip instance involves driving seven (five) kilometers and for 15 (11) minutes. It lasts 13% longer than it would in the absence of traffic. Note that this raw comparison of averages also reflects our oversampling of trip instances during the night and in smaller cities. We also report the distributions of the implied speeds (duration divided by length) and "effective" (haversine) trip lengths. The mean travel speed is 29 kilometers per hour, which is slow relative to the speeds typically observed in US cities but notably faster than typical speeds during peak travel hours in Bogotá or at any hour in Dhaka, as observed in Figure 1 in the main text. Finally, the distance traveled on the average trip is 55% longer than the effective trip length.

Panel B reports analogous distributions for the average trip instance within a city. Mean speeds vary remarkably across cities, with a standard deviation of 6.15, which is more than

[39]Since trips are drawn randomly to be queried, 15,407 (or 0.08% of) trips never got queried during our data collection period.

Table A.1: Counts of countries, cities and trip instances in our sample

	Countries	Cities	Trip Instances
Asia	44	458	242,246,048
Americas	25	337	141,856,368
Africa	43	176	95,374,704
Europe	34	245	83,126,480
Oceania	3	12	20,352,474
India		173	65,694,376
United States		121	57,162,864
Indonesia		29	12,815,559

Notes: The number of countries, cities, and trip instances in each continent for which we have travel information from Google Maps.

Table A.2: Trip summary statistics

			percentile:						
	Mean	St. dev.	1	10	25	50	75	90	99
Panel A: Sample: all trip instances (N=582,956,059)									
Speed	29.07	11.14	12.49	17.48	21.30	26.85	34.20	43.32	67.13
Duration	14.79	13.81	3.20	5.48	7.63	11.42	17.78	27.45	60.33
Duration (no traffic)	13.10	12.21	2.97	5.07	6.98	10.30	15.80	24.02	50.67
Trip length	7.47	10.59	1.33	2.11	3.09	5.05	8.64	15.62	37.99
Effective length	4.82	5.28	1.00	1.31	1.91	3.14	5.21	10.14	25.95
Panel B: Sample: all cities (N=1,228)									
Mean speed	28.76	6.15	17.55	21.85	24.54	27.78	31.89	38.07	44.99
Mean duration	13.21	4.30	7.94	9.33	10.49	12.13	14.79	18.10	26.81
Mean duration (no traffic)	11.78	3.79	7.21	8.52	9.47	10.92	13.08	16.18	24.33
Mean trip length	6.59	2.67	3.23	4.18	4.97	6.20	7.72	9.33	13.92
Mean effective length	4.23	1.41	2.19	2.76	3.22	3.90	4.92	6.06	8.73

Note: Durations are in minutes, lengths in kilometers; and speeds in kilometers per hour.

half the standard deviation of 11.14 at the trip instance level. The slowest cities are more than 2.5 times slower than the fastest ones on average, reinforcing what we already observe in Figure 1.

3 Additional trip data sources

Open Street Map: Road network and topography data

We get information on road networks from Open Street Map (OSM), a crowd-sourced worldwide mapping platform. In May 2020, we used OSMnx (Boeing, 2017) to download the entire

network of driving roads in each of our cities. We then map our GM trip routes onto the OSM road network to measure route characteristics such as the share of the route traversing each road class (e.g. "motorways", "primary roads", "residential roads", etc.) and the number of intersections along the route. In December 2020, we used OSMNX again to download aggregate statistics on each city: intersection count and angle distribution, total road length by type, grid-like road network, elevation of the road, grade and bearing distribution, and the length of the rivers and coastline. We use these measures to construct the city attributes used in our analysis.

Meteostat: Weather data

We obtain hourly weather data (temperature, dew point, humidity, wind speed, atmospheric pressure, and occurrence of weather events such as snow, sleet or thunderstorm) from Meteostat (https://dev.meteostat.net/, last accessed, 9 April 2021). In particular, we query the Meteostat API for weather data from the (up to five) nearest weather stations within 50 km of each city's center (and, where applicable, consider a distance-weighted average of the stations' recordings for each city and hour). We do so for every hour of every day that coincides with the days of our GM trip queries and for all cities that we query. We use these weather variables as controls in our trip instance-level regressions.

We successfully recovered weather data for 1,169 cities. In the rest of the cities, either Meteostat lacks a weather station nearby or any stations were inactive during the six months of our data collection. For 5 cities, the weather readings are always incomplete (e.g., missing temperature). Among the cities with non-missing weather data, the median city has one weather station nearby. The median city-day has weather readings at every hour of the day and the median reading is from a station that is less than six kilometers from the city center. More than 90% of city-days have at least six readings, so we can assign each trip instance the weather reading from the nearest hour of the day for which we have one for the city.

4 Trip route attributes

The construction of many of our regression variables requires the ability to trace the route between each trip's origin and destination. We characterize a trip's route as the sequence of segments defined by the lines connecting the locations of consecutive GPS directions along the route. For instance, the first segment of a route begins at the trip origin and ends where GM provides the first direction (e.g. "turn right"). Every location where GM provides a direction corresponds to the end of the previous step segment and the beginning of a new one. The final segment ends at the destination of the trip. While each instance of a trip could potentially have a different route, due to cost constraints, we only record GPS directions for

one of the driving instances (the most common one) of the trip and treat it as representative of the trip's other instances.

We follow the approach described in Appendix C.A of Akbar *et al.* (2023) to construct various attributes along our trip routes. These include the distance to the city center, upward and downward gradients along the route, the share of the route on each of six road classes, the number of intersections along the route, and the number of turns (in the GM directions). Data on elevation gradients and turns along the route come from our GM queries of walking trips.[40] Data on road class composition and intersections come from OSM. We find that motorways account for 4.5% of the roadway in our cities, primary roads account for 3.6%, secondary roads account for 5.0%, and tertiary roads account for 7.5%. Even though the largest four road classes account for only 20.6% of all urban roadway, they account for 64.5% of the length of our trip routes.

Finally, we use population estimates at 100-meter resolution from WorldPop (Bondarenko *et al.*, 2020) to compute average population density within a 111 meter buffer of each trip's route.

5 *City attributes*

Incomes

Our country-level GDP per capita for 2018 use Expenditure-side real GDP at chained PPPs (in 2017 US dollars) and population from the Penn World Table version 10.01 (Feenstra and Timmer, 2015, updated 2023). These data are missing for seven countries contributing 18 cities in our largest sample of 1,228: Afghanistan, Cuba, Libya, North Korea, Papua New Guinea, Puerto Rico, and Somalia. Thirteen of these 18 cities (all cities in Afghanistan, Cuba, North Korea, and Somalia) are removed from our main estimation sample anyway due to suspected low quality GM data (see section 2.2). We further drop the 13 cities of Venezuela due to poor quality GDP data. The Penn World Table reports a per capita value that would make Venezuela poorer than any other country in the world by a factor of two. The value of Venezuela's total GDP is less than a third of the value of this country's oil exports for that year (OPEC, 2019).

Our city-level GDP per capita for OECD countries were derived from the OECD Metropolitan Database (OECD, 2021), which includes real GDP per capita in USD (constant prices, constant PPP, base year 2015) at the level of Functional Urban Areas (FUA) for the year 2018 (the latest year available at the time of download in June 2021: `https://www.oecd.org/regional/`

[40]Elevation gradients along a trip route are from a walking trip query on Google Maps that was forced to go through way-points that we identified along the trip's driving route. For a small number of trips, we fail to replicate the driving routes via walking, and instead have to use elevation gradients along the fastest unconstrained walking route of the trip. However, we are still missing gradients for roughly 13,000 trips (out of 20 million) where GM is unable to identify even a walking route. We classify gradient along these trips as missing in our trip-level regressions.

regional-statistics/) We also download their shapefiles of the FUA boundaries and compute the area of intersection between each FUA and overlapping cities in our sample. The GDP per capita for a city in our sample is the area-weighted average across all FUAs that it intersects.

We also separately obtain city-level income measures for India and the US. For India, we obtain district-level average daily labor earnings from the World Bank's South Asia Spatial Database (Li *et al.*, 2015) (based on the 2011-12 National Sample Survey) and extrapolate them to our city definitions as described in Appendix C of Akbar *et al.* (2023). For the US, we obtain estimates of the count of households by income bracket and block group from the 2015-2019 American Community Survey. We derive our city-level average income by taking a household-count-weighted average of incomes across all block groups whose geographic centroids lie within our city boundary. We verify that our ACS-based measure of household incomes is highly correlated with our real GDP per capita measure for the US from OECD (2021), with a correlation coefficient of 0.6.

City size

We compute city populations by summing population estimates from WorldPop (Bondarenko *et al.*, 2020) for every 100-meter pixel within our city boundaries.

Topography

The rest of our city attributes come from Open Street Map (OSM) downloaded via GeoFabrik in 2020. We used OSMnx (Boeing, 2017) to determine the entire network of roads and water bodies within the extent of our cities. We compute the total length of water bodies as the sum of the length of river centerlines, streamlines, coastlines and lakeshores within our city boundaries.

Road network

We follow the approach described in Appendix C.B of Akbar *et al.* (2023) to compute our road network characteristics from OSM. In particular, we compute the length of major roads as the sum of the lengths of all roads classified as motorways, primary, secondary, or tertiary. In subsequent robustness exercises in Appendix C.2, we exclude tertiary roads from our measure of road length (to little effect). We measure the griddiness of a city's road network as the share of the network's edges that are within 2 degrees (modulo 90) of the modal edge bearing. This measure is defined more formally in Akbar *et al.* (2023) as the grid-like orientation of the network. We also use our OSM road network to compute additional city attributes such as average block lengths, number of edges at intersections, and speed limits.

Appendix B. Auxiliary analyses

1 Trip-level determinants of uncongested speed and congestion

Table B.1 repeats table 1 using log uncongested speed as dependent variable instead of log speed. Table B.2 repeats again the same set of specifications using log congestion as dependent variable. As noted previously, for each explanatory variable, the coefficient estimated in the speed regression in table 1 is equal to the coefficient estimated in the uncongested speed regression in table B.1 minus the coefficient estimated in the congestion regression in table B.2.

The main result from tables B.1 and B.2 is that the greater speed of trips that are longer or use higher classes of roads more is mainly driven by uncongested speed rather than congestion. Reassuringly, we note nonetheless that in table B.2, lower classes of roads and trips further from the center are subject to less congestion.

2 Trip-level determinants of speed

Table B.3 reports estimates of our baseline trip-level specification using different measures of speed, trip samples, and weighting schemes in columns 1–6. Columns 7 and 8 estimate slightly more parsimonious specifications, relative to our baseline, for use in the calculation of footnote 24. The main feature of table B.3 is that these alternative specifications do not generally affect the sign or magnitude of the coefficients estimated in our baseline specification.

Several secondary results are interesting to note in column 1, which uses effective speed as dependent variable. This column estimates a lower coefficient for trip length, larger speed penalties for lower classes of roads, faster speeds for radial trips and trips farther away from the center, and no effect of nearby population density. This is because longer trips, trips using lower classes of roads, trips closer to the center, trips in denser areas, and non-radial trips are all more circuitous. Below, in table B.4 we show that despite these differences, city-level indices of effective speed look very similar to our baseline speed indices.

3 Correlation among alternative speed indices

Table B.4 reports correlations between our baseline city speed index and several variants. Panel A includes variants from each column of table 1. Column 1, which has no controls, shows the lowest rank correlation at 0.88, emphasizing the relative importance of controlling for trip distance and distance to the city center to ensure that we are considering comparable bundles of trips across cities. The correlations for column 2 and 3 are above 0.99, showing

Table B.1: Determinants of uncongested speed

Dependent variable: log trip uncongested speed	(1)	(2)	(3)	(4)	(5)	(6)
log trip length		0.27^a	0.27^a	0.28^a	0.29^a	0.29^a
		(0.0035)	(0.0034)	(0.0038)	(0.0033)	(0.0033)
log distance to center		0.048^a	0.047^a	0.024^a		
		(0.0027)	(0.0027)	(0.0024)		
Gross gradient up				-0.64^a	-0.65^a	-0.65^a
				(0.058)	(0.054)	(0.052)
Gross gradient down				0.051	0.034	0.048
				(0.069)	(0.060)	(0.060)
Gradient missing				-0.030^b	-0.037^a	-0.033^a
				(0.014)	(0.013)	(0.012)
Share of primary roads				-0.17^a	-0.17^a	-0.17^a
				(0.0063)	(0.0061)	(0.0060)
Share of secondary roads				-0.20^a	-0.20^a	-0.20^a
				(0.0058)	(0.0057)	(0.0056)
Share of tertiary roads				-0.23^a	-0.23^a	-0.23^a
				(0.0056)	(0.0056)	(0.0055)
Share of residential roads				-0.32^a	-0.32^a	-0.32^a
				(0.0068)	(0.0065)	(0.0064)
Share of other roads				-0.32^a	-0.32^a	-0.31^a
				(0.0090)	(0.0087)	(0.0088)
Share of missing roads				-0.27^a	-0.27^a	-0.27^a
				(0.0063)	(0.0064)	(0.0063)
log intersections				-0.063^a	-0.068^a	-0.068^a
				(0.0034)	(0.0025)	(0.0024)
arsinh turns against traffic				-0.085^a	-0.082^a	-0.081^a
				(0.00090)	(0.00080)	(0.00078)
arsinh density				-0.030^a	-0.028^a	-0.031^a
				(0.0038)	(0.0031)	(0.0032)
Type: circumferential		-0.0078^b	-0.0077^b	0.0076^a	0.0067^a	0.0035^c
		(0.0032)	(0.0031)	(0.0020)	(0.0019)	(0.0019)
Type: gravity		0.0058^a	0.0055^a	0.012^a	0.012^a	0.010^a
		(0.0022)	(0.0021)	(0.0015)	(0.0014)	(0.0014)
Type: amenity		0.011^a	0.012^a	0.012^a	0.013^a	0.012^a
		(0.0025)	(0.0023)	(0.0016)	(0.0015)	(0.0015)
City effect	Y	Y	Y	Y	Y	Y
Day effect	N	wkday	wkday	wkday	wkday	wkday
Time effect	N	Y	Y	Y	Y	Y
Weight	N	N	Y	Y	Y	Y
Weather	N	Y	Y	Y	Y	Y
City-specific distance to center	N	N	N	N	linear	cubic
Observations	515,019,834	378,535,918	378,009,201	376,459,587	376,459,587	376,459,587
Within R^2	0	0.53	0.53	0.66	0.68	0.69
R^2	0.34	0.69	0.69	0.78	0.79	0.80

Notes: This table repeats the specifications of table 1 using log uncongested speed instead of log speed as dependent variable. See the notes of table 1 for further information.

Table B.2: Determinants of congestion factor

Dependent variable: log trip congestion	(1)	(2)	(3)	(4)	(5)	(6)
log trip length		0.017^a	0.020^a	0.0070^a	0.0058^a	0.0060^a
		(0.00075)	(0.00087)	(0.0012)	(0.00093)	(0.00092)
log distance to center		-0.029^a	-0.037^a	-0.025^a		
		(0.0011)	(0.0014)	(0.0011)		
Gross gradient up				-0.28^a	-0.26^a	-0.26^a
				(0.033)	(0.030)	(0.029)
Gross gradient down				-0.25^a	-0.24^a	-0.24^a
				(0.030)	(0.028)	(0.028)
Gradient missing				-0.014^a	-0.015^a	-0.015^a
				(0.0051)	(0.0041)	(0.0037)
Share of primary roads				0.0028	0.0016	0.0010
				(0.0022)	(0.0020)	(0.0020)
Share of secondary roads				-0.014^a	-0.014^a	-0.015^a
				(0.0020)	(0.0018)	(0.0018)
Share of tertiary roads				-0.028^a	-0.026^a	-0.026^a
				(0.0021)	(0.0020)	(0.0020)
Share of residential roads				-0.055^a	-0.054^a	-0.054^a
				(0.0023)	(0.0021)	(0.0021)
Share of other roads				-0.054^a	-0.056^a	-0.055^a
				(0.0036)	(0.0035)	(0.0035)
Share of missing roads				-0.025^a	-0.027^a	-0.027^a
				(0.0018)	(0.0018)	(0.0018)
log intersections				0.011^a	0.013^a	0.013^a
				(0.0012)	(0.00091)	(0.00095)
arsinh turns against traffic				-0.011^a	-0.011^a	-0.011^a
				(0.00051)	(0.00045)	(0.00045)
arsinh density				0.024^a	0.022^a	0.021^a
				(0.0019)	(0.0019)	(0.0022)
Type: circumferential		-0.018^a	-0.022^a	-0.017^a	-0.017^a	-0.016^a
		(0.00086)	(0.00098)	(0.00090)	(0.00086)	(0.00085)
Type: gravity		-0.014^a	-0.018^a	-0.015^a	-0.014^a	-0.014^a
		(0.00062)	(0.00073)	(0.00068)	(0.00064)	(0.00063)
Type: amenity		-0.00097	-0.00064	-0.0013^c	-0.00029	-0.00060
		(0.00068)	(0.00079)	(0.00073)	(0.00065)	(0.00062)
City effect	Y	Y	Y	Y	Y	Y
Day effect	N	wkday	wkday	wkday	wkday	wkday
Time effect	N	Y	Y	Y	Y	Y
Weight	N	N	Y	Y	Y	Y
Weather	N	Y	Y	Y	Y	Y
City-specific distance to center	N	N	N	N	linear	cubic
Observations	515,019,834	378,535,918	378,009,201	376,459,587	376,459,587	376,459,587
Within R^2	0	0.33	0.25	0.27	0.29	0.30
R^2	0.08	0.40	0.33	0.35	0.37	0.37

Notes: This table repeats the specifications of table 1 using log congestion instead of log speed as dependent variable. See the notes of table 1 for further information.

Table B.3: Additional trip-level determinants of speed

Dependent variable: log trip speed	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
log trip length	0.17^a	0.27^a	0.35^a	0.30^a	0.28^a	0.28^a	0.31^a	0.20^a
	(0.0060)	(0.0037)	(0.0055)	(0.0058)	(0.0037)	(0.0031)	(0.0044)	(0.0026)
log distance to center	0.12^a	0.055^a	-0.025^a	0.056^a	0.049^a	0.058^a	0.048^a	0.060^a
	(0.0036)	(0.0030)	(0.0038)	(0.0042)	(0.0028)	(0.0025)	(0.0028)	(0.0030)
Gross gradient up	-1.59^a	-0.42^a	0.050	-0.24	-0.35^a	-0.38^a	-0.39^a	-0.58^a
	(0.082)	(0.062)	(0.21)	(0.17)	(0.058)	(0.058)	(0.060)	(0.063)
Gross gradient down	-0.97^a	0.36^a	0.60^a	0.80^a	0.30^a	0.28^a	0.29^a	0.14^b
	(0.079)	(0.073)	(0.20)	(0.16)	(0.070)	(0.079)	(0.071)	(0.071)
Gradient missing	-0.11^a	-0.0090	0.0027	-0.018	-0.017	-0.024^c	-0.0087	-0.0074
	(0.023)	(0.0096)	(0.0090)	(0.020)	(0.011)	(0.013)	(0.012)	(0.012)
Share of primary roads	-0.14^a	-0.16^a	-0.15^a	-0.15^a	-0.17^a	-0.16^a		-0.20^a
	(0.0085)	(0.0058)	(0.0068)	(0.0073)	(0.0063)	(0.0068)		(0.0069)
Share of secondary roads	-0.18^a	-0.17^a	-0.15^a	-0.15^a	-0.19^a	-0.18^a		-0.23^a
	(0.0079)	(0.0056)	(0.0064)	(0.0069)	(0.0060)	(0.0067)		(0.0063)
Share of tertiary roads	-0.22^a	-0.17^a	-0.16^a	-0.14^a	-0.20^a	-0.19^a		-0.26^a
	(0.0081)	(0.0055)	(0.0071)	(0.0073)	(0.0058)	(0.0062)		(0.0061)
Share of residential roads	-0.30^a	-0.24^a	-0.20^a	-0.20^a	-0.27^a	-0.26^a		-0.37^a
	(0.0090)	(0.0062)	(0.0075)	(0.0080)	(0.0065)	(0.0065)		(0.0074)
Share of other roads	-0.34^a	-0.24^a	-0.17^a	-0.21^a	-0.27^a	-0.25^a		-0.33^a
	(0.013)	(0.0073)	(0.011)	(0.011)	(0.0080)	(0.0079)		(0.085)
Share of missing roads	-0.31^a	-0.22^a	-0.17^a	-0.19^a	-0.25^a	-0.24^a		-0.25^a
	(0.013)	(0.0056)	(0.0069)	(0.010)	(0.0062)	(0.0069)		(0.0066)
Share of major roads							0.12^a	
							(0.0031)	
log intersections	-0.020^a	-0.076^a	-0.10^a	-0.075^a	-0.074^a	-0.077^a	-0.083^a	
	(0.0029)	(0.0031)	(0.0037)	(0.0048)	(0.0031)	(0.0027)	(0.0032)	
arsinh turns against traffic	-0.097^a	-0.072^a	-0.049^a	-0.050^a	-0.074^a	-0.072^a	-0.082^a	
	(0.0014)	(0.00089)	(0.0011)	(0.0016)	(0.00089)	(0.00092)	(0.00098)	
arsinh density	0.0062	-0.061^a	-0.074^a	-0.052^a	-0.054^a	-0.059^a	-0.054^a	-0.082^a
	(0.0047)	(0.0048)	(0.0076)	(0.0073)	(0.0044)	(0.0043)	(0.0043)	(0.0047)
Type: circumferential	-0.19^a	0.026^a		0.036^a	0.025^a	0.023^a	0.021^a	0.023^a
	(0.0047)	(0.0020)		(0.0038)	(0.0021)	(0.0018)	(0.0025)	(0.0024)
Type: gravity	-0.14^a	0.029^a		0.035^a	0.026^a	0.022^a	0.025^a	0.026^a
	(0.0039)	(0.0016)		(0.0028)	(0.0016)	(0.0013)	(0.0018)	(0.0018)
Type: amenity	-0.11^a	0.015^a		0.019^a	0.013^a	0.0096^a	0.0075^a	0.016^a
	(0.0039)	(0.0017)		(0.0027)	(0.0018)	(0.0015)	(0.0021)	(0.0019)
City effect	Y	Y	Y	Y	Y	Y	Y	Y
Day effect	wkday	wkday	wkday	wkday	wkday	wkday	wkday	wkday
Time effect	Y	Y	Y	Y	Y	Y	Y	Y
Weight	Y	Y	Y	Y	Y	Y	Y	Y
Weather	Y	Y	Y	Y	Y	Y	Y	Y
Observations	376,459,587	101,504,603	11,316,208	376,497,141	37,647,207	258,774,448	376,459,587	378,009,201
Within R^2	0.30	0.55	0.60	0.64	0.59	0.52	0.57	0.54
R^2	0.47	0.70	0.77	0.73	0.72	0.68	0.71	0.69

Notes: Each column reports parameter estimates for variants of column 4 in Table 1. Column 1 computes speed using effective rather than trip distance. Column 2 uses only peak hour trips (departure time between 7:30 AM and 9 AM or between 4 PM and 7 PM). Column 3 uses only radial trips that going towards the city center during the morning peak and away from the city center during the evening peak. Column 4 weights trips using a congestion weight given by (trip duration/trip duration in absence of traffic)$^{1/\lambda}$ with $\lambda = 0.3$. Column 5 uses a 10% random sample. Column 6 uses only trips within 10 kilometers from the city center. Column 7 replaces all the shares of road type variables with a share of major roads variable (including motorway, primary, secondary, and tertiary roads). Column 8 removes the variables log intersections and arsinh turns against traffic. Robust standard errors clustered at the city level are in parentheses. a, b, c: significant at 1%, 5%, 10%.

Table B.4: Pairwise Spearman rank correlations with our benchmark speed index

Index	Corr.	Std. Dev.	Min	Max
Panel A: Columns from table 1				
(1)	0.878	0.204	-0.686	0.485
(2)	0.997	0.173	-0.612	0.465
(3)	0.997	0.176	-0.625	0.481
(5)	0.960	0.195	-0.512	0.581
(6)	0.925	0.226	-0.558	0.670
Panel B: Variant speed indexes in table B.3				
Effective speed	0.900	0.192	-0.716	0.536
Peak hour	0.987	0.178	-0.588	0.512
Peak hour radial trips	0.907	0.195	-0.589	0.572
Congestion weighting	0.973	0.198	-0.702	0.543
10% random sample	1.000	0.178	-0.597	0.481
Trips within 10km of center	0.988	0.176	-0.523	0.477
Single major roads variable	0.999	0.178	-0.616	0.483
No turn and intersection variables	0.997	0.178	-0.562	0.472

Notes: 1,119 cities. The speed indices used in (5) and (6) in Panel A are adjusted by the city-specific distance to center coefficients evaluated at the sample mean of 7.31 kilometers. For each alternative index, the first column reports the Spearman rank correlation with our benchmark index (table 1 column 4). Other columns report standard summary statistics.

that our speed index is robust to how we weight trips, and to adding other centered trip-level controls. The indices used in columns 5 and 6 of table 1 measure travel speed at the sample mean distance to the center. Their estimates include city-specific distance coefficients (linear in 5, cubic in 6) applied at the sample mean. Although the rank correlations of 0.96 and 0.925 are slightly lower than the previous two, they are still remarkable given the amount of additional structure they impose.

Panel B considers the speed index variants from table B.3. No rank correlation is below 0.90. The only two correlations below 0.97 are with the speed index that substitutes effective speed for actual speed, for which we expect some difference, and with the speed index limited to peak radial trips, which represent about 3% of our sample. For India, Akbar *et al.* (2023) report similarly high correlations, generally above 0.90, between their baseline index and a much larger set of variants.

4 Using trips from the NHTS for the US

Trip sample

We use real trips reported in the 2017 US National Household Travel Survey (NHTS) with confidential geographic identifiers to benchmark our simulated trips. For this exercise, we

ignore NHTS trips that either start or end outside the US, trips that happen via air, rail or water, trips that report a non-positive travel time, and trips with a reported distance less than 100 meters. We drop a small number of trips where the distance between their origin and destination block groups seems inconsistent with the trip's reported distance. In particular, we drop (a) trips longer than 50 kilometers that start and end in the same block group and (b) trips shorter than 30 km between block groups whose centroids are much further apart (at least 100 kilometers more than the trip length). We also drop trips with the highest and lowest 1% of implied travel speeds (reported distance divided by travel time). We are left with 891,441 NHTS trips by 217,023 individuals.

Then we restrict ourselves to urban car trips: trips originating within a metropolitan statistical area (MSA) and where the primary travel mode is car, SUV, van, pickup truck, rental car, or taxi/ridesharing. Finally, we impose one additional data quality restriction. We expect travel times and distances to be strongly correlated, and we confirm so by estimating a quadratic regression of log trip time on log trip distance and fitting the data with an R-squared of 0.72. So we drop the outlier trips whose travel time is not within a 95% confidence interval of our fitted relationship. We are left with 593,929 NHTS trips.

Trip origins and destinations

Our NHTS data identifies the origin and destination of each trip only at the resolution of census block groups. To replicate these trips on Google Maps, we need to define the endpoints of trips more precisely. To do so, we start with the geographic centroids of the reported block groups and adjust the trip endpoints until the straight line distance between them matches the reported trip distance, $RDist$, divided by $\sqrt{2}$ (a good proxy for the ratio between road trip distance and effective trip distance). There are two types of cases to consider.

For trips that start and end in the same block group (around 6%), we only have one centroid. So, we identify the longest possible straight line L within the block group boundary that goes through the centroid, and we define the trip endpoints to be the two points along L that are $\frac{1}{2}\frac{RDist}{\sqrt{2}}$ distance apart from the centroid. Figure B.1 illustrates the process. If this approach places the trip endpoints outside the block group boundaries (because $\frac{RDist}{\sqrt{2}}$ is greater than the length of L), we shift the endpoints to the nearest point along L that intersects the block group boundary.

For trips between different block groups, our approach is similar, but slightly richer. We start by defining a trip between the centroids of the two block groups and determining the travel distance $GDist$ between them along the fastest driving route returned by Google Maps (GM) at the reported departure time. If $GDist$ is similar to $RDist$ (within 500 meters), we assume that the trip between the centroids is an acceptable proxy for the NHTS trip between

Figure B.1: Identifying endpoints of NHTS trips within the same block group

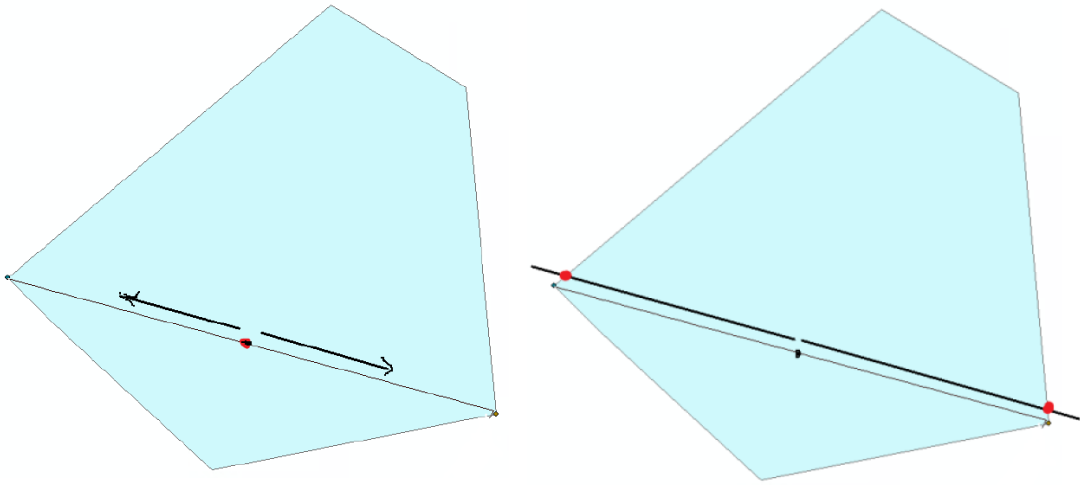

We identify trip endpoints on either side of the block group centroid such that the straight line distance between them is either the reported trip distance divided by $\sqrt{2}$ or the longest possible distance within the block group boundary.

the unobserved origin and destination pair. Otherwise (for 77% of trips), we adjust the trip endpoints along the straight line L connecting the centroids by shifting the endpoints a distance of $\frac{1}{2}\frac{RDist-GDist}{\sqrt{2}}$ away from the centroids. If the GM distance is smaller than the reported distance, we shift the endpoints farther away from each other. If the GM distance is larger than the reported distance, we move the endpoints closer to each other. Figure B.2 illustrates the process. As before, if this approach places a trip endpoint outside of its reported block group, we set the endpoint to the nearest point along L on the block group boundary. To make sure our trips with adjusted endpoints are more reasonable than trips between centroids, we query the adjusted trips on GM and confirm that their GM distances are indeed notably closer to the NHTS-reported distances.

Trip queries on Google Maps

We have 460,425 unique trips (precise origin-destination location pairs) spanning 275 metropolitan areas (MSAs, 1999 definitions).[41] The number of trips per MSA range from 11 (for Pine Bluff, AR) to 43,141 (for Dallas-Fort Worth, TX). We simulate these trips repeatedly on GM to determine driving travel times at random times of day that are drawn from the same distribution of departure times as our main sample of trip queries for the world (as described in Appendix A.2). Once again, we query half as frequently on weekends relative

[41]Note that NHTS trips between the same pair of block groups and similar reported distances may get assigned the same pair of origin-destination location coordinates.

Figure B.2: Identifying endpoints of NHTS trips across block groups

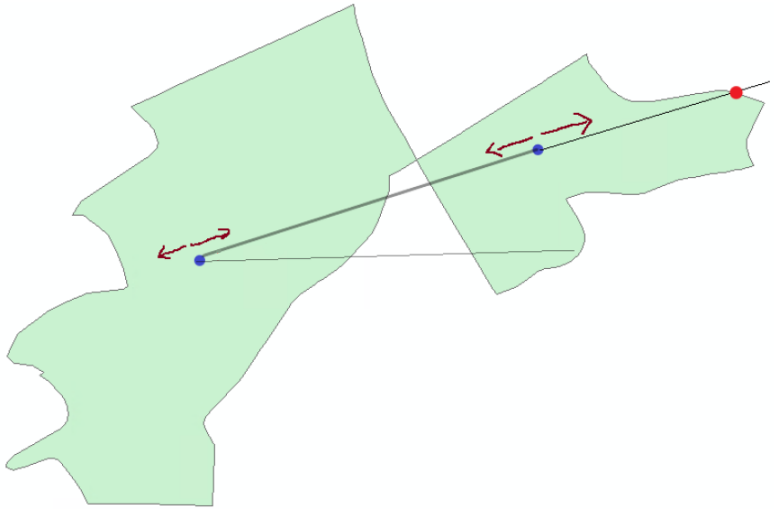

We identify trip endpoints along the straight line connecting the block group centroids to match the reported trip distance.

to weekdays. Between June 19 and October 20, 2019, we completed 33,402,825 queries to GM and managed to get back driving routes for 99.7% of the queried trips. So, the average trip was queried roughly once every other day.

Speed indices with replicated NHTS *trips*

Next, we assign each NHTS trip to a city in our sample only if both the origin and the destination of the trip lie within a 0.2 degree buffer around the BUILT pixels in the city. We ignore trips that either start or end outside our cities' buffer zones. We end up with 168,874 NHTS trips spanning 138 cities. The average city has 1,224 NHTS trips, and the number of trips per city range from 11 to 15,764. The average trip was queried 76 times over four months, and 90% of the trips were queried at least 35 times.

To estimate city-level speed indices, we restrict to cities that have at least 100 or at least 1,000 NHTS trip instances and follow the specification in column 2 of table 1. We try weighting our observations in three different ways: without any weights (just as in column 2 of Table 1), weighting the trip instances by their NHTS trip weights, and weighting the trip instances by the trip's distance to the city center (in five-kilometer bins) where the weights are the proportion of trips in our main sample of queries in the US at each distance bin from the center. Table B.5 reports summary statistics for the estimated speed indices. Their Spearman rank correlation with our benchmark speed index is between 0.96 and 0.97, depending on how we weight NHTS trips, when considering cities with at least 1,000 NHTS

Table B.5: Pairwise Spearman rank correlations for our benchmark speed index with replicated NHTS trips for the US

Index	Corr.	Std. Dev.	Min	Max	Cities
Speed index computed from NHTS trips in cities with:					
$> 1,000$ NHTS trips	0.960	0.094	-0.216	0.136	37
—, trip instances weighted by NHTS weights	0.966	0.103	-0.226	0.0151	37
—, trip instances weighted by trip distance to center	0.970	0.100	-0.264	0.134	37
> 100 NHTS trips	0.918	0.079	-0.246	0.131	98

Notes: The index in each row is compared in the first column to the baseline index, as estimated in column 4 of table 1. The first row uses a similar index computed from replicated data for NHTS trips taking place within the same geography as our baseline index and includes the 37 cities with more than 1,000 NHTS trips. The regression estimating this NHTS speed index has 7,358,839 observations corresponding to as many trip instances of 10,085,671 NHTS trips. In the second row, NHTS trip instances are weighted by their NHTS weight to estimate the NHTS speed index. In the third row, NHTS trip instances are weighted by the proportion of trips in our main sample of queries that are within the same 5 kilometer distance bin from the city center.. The fourth row replicates the first row for all 98 cities with more than 100 NHTS trips. Other columns report standard summary statistics.

trips. This correlation is still 0.92 when we add cities with only between 100 and 1,000 NHTS trips, for which the NHTS is arguably less representative.

5 Mass transit

Data and sample

To estimate comparable transit speed indices, we first queried 83,055,553 transit trip instances from GM in 2022, in order to get information on each 'leg' of each journey: a ride on a single bus or train, or a walk to and/or from a transit stop. We observe the time spent traveling on each leg and its type ("Bus", "Train", "Subway", "Walk", etc.), as well as wait times for transit legs and distances for walking legs. GM relies on General Transit Feed Specification (GTFS) feeds provided by transit agencies. The largest database of GTFS feeds we know of, MobilityData (`https://mobilitydata.org`), has 1,914 feeds from 62 countries. However, the overwhelming majority do not include real-time information on transit vehicle positions. This database implies that 70 or fewer of our cities have such information. Even some of those 70 cities have only partial information, like real-time train but not bus positions.

From those 83,055,553 initially queried trip instances, we first drop 14,544,882 queries that do not return results, and 13,506,691 that return a pure walking trip, leaving 55,003,980 trip instances returning transit directions. We then focus on the 36,346,469 instances of 5,321,629 trips that we define as "viable": faster than walking and requiring less than 1 hour of wait time before departure.

Table B.6: Transit leg summary statistics

	Mean	St. dev.	1	10	25	50	75	90	99
				percentile:					
Panel A: Sample: transit legs with < 25% walking share (N= 330,384)									
Speed	21.59	9.37	10.14	14.49	17.37	20.12	24.11	29.36	48.29
Duration	21.44	13.12	4.57	9.00	12.72	18.28	26.63	38.00	66.12
Trip length	7.35	4.78	1.79	3.10	4.32	6.08	8.99	13.15	23.69

Note: Durations are in minutes, lengths in kilometers; and speeds in kilometers per hour. The sample is transit legs with less than 25% of total transit distance being composed of walking.

Next, we limit to 1,133,569 single-leg viable transit trip instances, consisting of a single transit leg, walking to the origin transit stop and walking from the destination stop. GM reports the total travel distance for the trip instance and the distances on the walking legs. So we derive the travel distance on the transit leg (necessary to compute speed) by subtracting the sum of the walking leg distances from the total distance.

Next, we keep only trips with a road-based transit mode: Trolleybus, Bus, Intercity bus, Minibus, or Share taxi. This leaves 1,032,161 transit legs corresponding to 889,580 unique trips in 691 cities. Each transit trip instance is matched to the driving instance of the same trip that minimizes the difference in departure times between the instances. The average transit trip in this sample has a departure time that is 24.5 minutes from the nearest instance of its matched main sample (private vehicle) trip.

Finally, from the resulting set of transit trips matched to corresponding driving trips, to ensure that the driving and transit trips are truly comparable, we focus on transit trips in which the total walking distance is less than 25% of the total trip distance. This removes cases where the origin and/or destination of the actual transit portion of the transit trip is very different from those of the driving trip. We also drop a small number of trips where the distance between the trip origin and the first transit stop is inconsistent with the reported walking distance to the stop.

Table B.6 reports summary statistics for the resulting sample of 330,384 trip instances in 655 cities. On this sample of driving trips, we first run the specification reported in table 1 column 4. The resulting correlation between this speed index and our baseline index from table 1 column 4 is 0.97. We then estimate a transit speed index analogous to the table 1 column 4 specification (except that we no longer have trip characteristics or weather, and the weights are NHTS weights only instead of weights that combine NHTS weights and population density along the trip). It has a rank correlation of 0.51 with our baseline speed index and 0.53 with the index of transit-like driving trips just described.

Table B.7: Variance Decomposition

Sample	Cities	Uncongested Speed	Congestion factor	Covariance
All	1,119	0.947	0.087	0.017
OECD countries	265	0.666	0.063	-0.135
Poor countries	425	0.713	0.180	-0.053
Top quartile most populous cities	280	0.944	0.086	0.015
Peak hour speed index	1,119	0.931	0.152	0.042
Peak hour radial trips	1,119	0.910	0.174	0.042
Trips within 10 km of center	1,119	0.903	0.104	0.004

Notes: The first four rows report a variance decomposition of our baseline speed index. The first row uses our baseline sample of 1,119 cities. The second row restricts its sample to 265 OECD cities for which we have city GDP data as in table 4. The third row considers only cities that belong to countries in the bottom half of GDP per capita. The fourth row only includes cities in the top quartile of population. The last three rows use our baseline sample of cities but perform a variance decomposition on variants of our baseline speed index. The fifth and sixth row use the speed index computed in columns 2 and 3 of table B.3, respectively. The last row use the index from column 6 of the same table.

6 Variance analysis

Table B.7 reports results for the decomposition of the variance of the speed index, $\mathbb{S}$, into the variance of uncongested speed, $\mathbb{U}$, and the variance of the congestion index, $\mathbb{K}$ using the additive property of these indices: $\mathbb{S} = \mathbb{U} - \mathbb{K}$. Each row reports this decomposition for a different sample of countries, cities, or trips. In most rows, including our baseline sample, we find that the variance of uncongested speed accounts for more than 90% of the variance of speed. The only two exceptions occur in rows 2 and 3 when we consider cities in rich and poor countries separately. Separating cities by income attenuates differences in uncongested speed, which are highly correlated with differences in income, while largely retaining differences in congestion. But even in these cases, the variance of uncongested speed still accounts for at least two thirds of the variance of speed.

Appendix C. Robustness of determinants of speed

This section offers more details on various robustness exercises for our main result in Table 5 described in section 5.3.

1 Robustness to different city samples, speed indices, and estimation techniques

Table C.1 reports results for alternative samples to address data quality concerns. Relative to our baseline specification, table 5, column 2, all the coefficients we estimate remain equally

Table C.1: Robustness to alternative city samples

Dependent variable: log city speed	(1)	(2)	(3)	(4)	(5)	(6)	(7)
log country GDP (pc)	0.042^a	0.027^a	0.046^a	0.052^a	0.042^a	0.040^a	0.040^a
	(0.011)	(0.0094)	(0.0099)	(0.0093)	(0.0096)	(0.011)	(0.012)
log population	-0.17^a	-0.17^a	-0.16^a	-0.16^a	-0.17^a	-0.17^a	-0.17^a
	(0.018)	(0.017)	(0.017)	(0.017)	(0.017)	(0.019)	(0.018)
log area	0.085^a	0.091^a	0.087^a	0.073^a	0.079^a	0.083^a	0.10^a
	(0.022)	(0.022)	(0.022)	(0.024)	(0.022)	(0.025)	(0.022)
asinh major road length	0.080^a	0.081^a	0.081^a	0.090^a	0.089^a	0.082^a	0.078^a
	(0.012)	(0.013)	(0.013)	(0.015)	(0.012)	(0.014)	(0.012)
Network griddiness (0-1)	0.18^a	0.19^a	0.18^a	0.18^a	0.18^a	0.17^a	0.19^a
	(0.057)	(0.059)	(0.055)	(0.054)	(0.056)	(0.059)	(0.059)
Elevation variance	-0.0024^b	-0.0023^b	-0.0025^a	-0.0023^b	-0.0025^a	-0.0023^b	-0.0024^c
	(0.00092)	(0.00090)	(0.00092)	(0.00097)	(0.00092)	(0.00093)	(0.0013)
asinh water body length	-0.071^a	-0.068^a	-0.079^a	-0.077^a	-0.074^a	-0.068^a	-0.096^a
	(0.019)	(0.019)	(0.020)	(0.020)	(0.020)	(0.020)	(0.028)
Observations	1,119	1,197	1,079	966	1,088	1,006	910
R^2	0.70	0.68	0.72	0.74	0.72	0.70	0.70

Notes: Each column displays our baseline specification of table 5, column 2 using speed indices estimated with different samples of cities (but still the same trip level specification as in column 4 of table 1). Column 1 repeats the baseline specification from table 5, column 2. Column 2 includes all cities except those with missing GDP. Columns 3 and 4 drop 10% and 20% of cities with the lowest data quality respectively. Column 5 drops 10% of city with the smallest share of road segments based on color differences. Column 6 drops 10% of cities with the highest share of missing roads on OSM. Column 7 drops all cities with either a BRT or a subway system. Robust standard errors clustered at the country level in parentheses. a, b, c: significant at 1%, 5%, 10%.

Figure C.1: Traffic color comparison for Annaba, Algeria

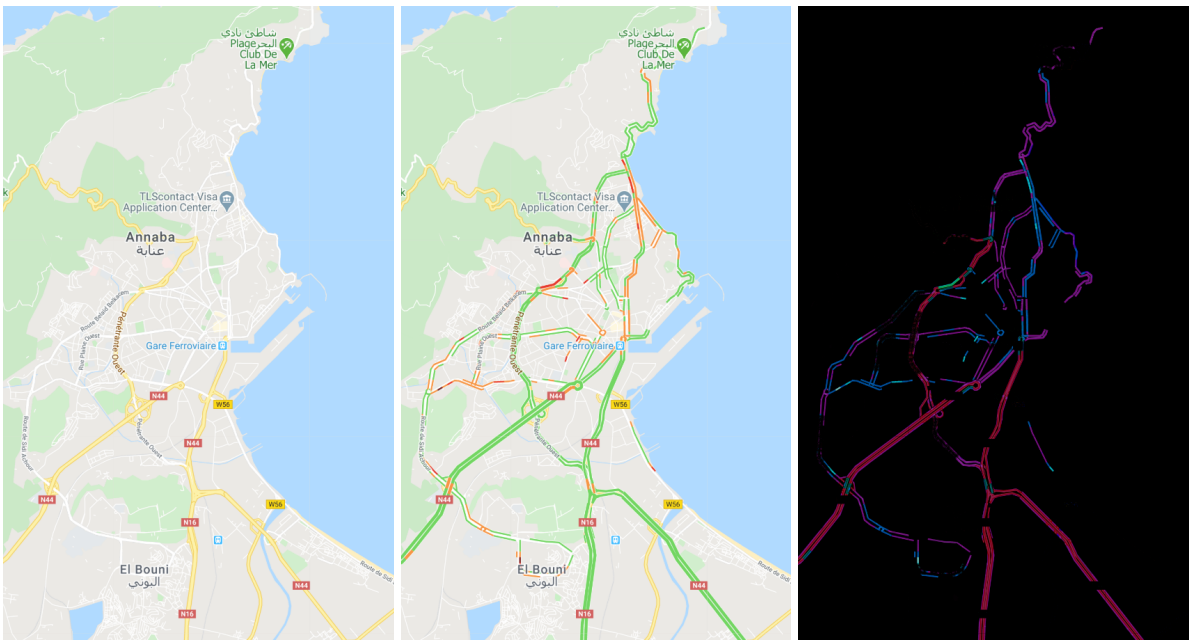

significant and of very similar magnitude, without exception.

Note that column 5 drops 10% of cities with the smallest share of road segments with real-time traffic. We compute this share by measuring whether road segments change color when Google Maps' web browser has the "traffic layer" on. Figure C.1 shows the screenshots that we use for these color comparisons for the city of Annaba in Algeria. We take these screenshots at two different times of the day (10 AM and 3 PM local time), centered at the city center and at the same zoom level for all cities. We use the images without the traffic layer to identify pixels corresponding to road segments (in white and yellow). We use the images with the traffic layer to identify the share of pixels for which Google Maps observes traffic (in green, orange, red, or crimson). Note also that column 7 drops all 209 cities with either a subway (Gonzalez-Navarro and Turner, 2018) or a BRT system as of December 2019 (based on work in progress by Nicolas Gendron-Carrier, Marco Gonzalez-Navarro and Matthew Turner).

Table C.2 reports results using alternative city speed indices as dependent variables for our baseline specification. A full description of these indices is provided in the table's notes. The main feature of table C.2 is the stability of the coefficients relative to our baseline. This feature is discussed further in the main text. Here, we describe a number of secondary results. Column 1 uses a city speed index which essentially averages mean log speed as dependent variable. Not conditioning out trip length and distance to the center when estimating city speed indices leads to a much higher area elasticity of speed at the city level. This is unsurprising given that in larger cities our sampling strategy mechanically draws longer trips and trips further away from the center, which are faster. When using a speed index where congestion features more heavily, such as in column 5 for peak hour trips or in column 7 for trips within 10 kilometers from the center, the income elasticity we estimate is correspondingly lower.

Table C.3 reports results that address the issue that our dependent variable of interest is an index estimated (with error) by the regression described in equation (1) for which the results are reported in table 1, column 4. We first note that heteroscedasticity may not be a major concern, because our baseline estimation of the determinants of city travel speed in table 5 uses 1,119 city effects estimated from a large sample of 376 million observations, and reports standard errors which are already robust and clustered by country. To further address the fact that the dependent variable is estimated, we follow the approach in Combes *et al.* (2019) and reestimate our baseline specification using feasible generalized least squares (FGLS) and weighted least squares (WLS).

The FGLS standard errors do not allow for clustering or robust Huber-White corrections. Hence, the FGLS standard errors in the fourth row of the table are best compared to the unclustered non-robust standard errors reported in the third row of the same table. The two sets of standard errors are virtually indistinguishable. With WLS, we obtain different point

Table C.2: Robustness to alternative speed indexes

Dependent variable: log speed index	(1)	(2)	(3)	(4)	(5)	(6)	(7)
log country GDP (pc)	0.052^a	0.040^a	0.053^a	0.024	0.026^b	0.043^a	0.034^a
	(0.0100)	(0.011)	(0.014)	(0.018)	(0.012)	(0.015)	(0.010)
log population	-0.17^a	-0.17^a	-0.19^a	-0.18^a	-0.18^a	-0.18^a	-0.18^a
	(0.017)	(0.018)	(0.023)	(0.024)	(0.019)	(0.022)	(0.018)
log area	0.18^a	0.071^a	0.070^b	0.078^a	0.082^a	0.065^b	0.10^a
	(0.021)	(0.022)	(0.029)	(0.028)	(0.023)	(0.025)	(0.021)
asinh major road length	0.080^a	0.081^a	0.12^a	0.080^a	0.085^a	0.092^a	0.072^a
	(0.013)	(0.012)	(0.017)	(0.019)	(0.014)	(0.015)	(0.012)
Network griddiness	0.13^b	0.18^a	0.17^b	0.35^a	0.22^a	0.21^a	0.19^a
	(0.061)	(0.056)	(0.069)	(0.050)	(0.059)	(0.068)	(0.059)
Elevation variance	-0.0011	-0.0024^a	-0.0019^c	-0.0037^a	-0.0026^a	-0.0032^a	-0.0023^b
	(0.0010)	(0.00090)	(0.0010)	(0.0011)	(0.00097)	(0.00091)	(0.00095)
asinh water body length	-0.083^a	-0.072^a	-0.086^a	-0.067^a	-0.077^a	-0.077^a	-0.078^a
	(0.025)	(0.020)	(0.022)	(0.014)	(0.020)	(0.020)	(0.017)
Observations	1,119	1,119	1,119	1,119	1,119	1,119	1,119
R^2	0.76	0.70	0.62	0.64	0.67	0.66	0.68

Notes: Each column displays our baseline specification of table 5, column 2 using city speed indices estimated with different trip level regressions. In column 1, the dependent variable is the speed index estimated in table 1, column 1 (a simpler index using city means). Columns 2 uses the speed index from table 1, column 3 (without route characteristics and no weighting). Columns 3 uses the adjusted speed index from table 1, column 6 (using a city-specific speed gradient valued at the sample mean of 7.31 kilometers). Column 4 uses the speed index from table B.3, column 1 (for effective speed). Column 5 uses the speed index from table B.3, column 2 (peak hour trips). Column 6 uses the speed index from table B.3, column 4 (congestion weighting). Column 7 uses the speed index from table B.3, column 6 (trips within 10 kilometers from the center only). Robust standard errors clustered at the country level in parentheses. a, b, c: significant at 1%, 5%, 10%.

Table C.3: Robustness to alternative estimation techniques

Explanatory variable:	log country GDP (pc)	log population	log area	asinh major road length	Network griddiness	Elevation variance	asinh water body length
OLS-FGLS coefficient	0.042	-0.17	0.085	0.080	0.18	-0.0024	-0.071
Robust clustered OLS s.e.	(0.011)^a	(0.018)^a	(0.022)^a	(0.012)^a	(0.057)^a	(0.001)^b	(0.019)^a
OLS s.e.	(0.006)^a	(0.007)^a	(0.009)^a	(0.007)^a	(0.023)^a	(0.001)^a	(0.009)^a
FGLS s.e.	(0.006)^a	(0.007)^a	(0.009)^a	(0.007)^a	(0.023)^a	(0.001)^a	(0.009)^a
WLS coefficient	0.036	-0.18	0.081	0.090	0.19	-0.0020	-0.087
WLS s.e.	(0.006)^a	(0.008)^a	(0.012)^a	(0.009)^a	(0.022)^a	(0.001)^b	(0.012)^a
Robust clustered WLS s.e.	(0.018)^b	(0.025)^a	(0.034)^b	(0.025)^a	(0.054)^a	(0.001)^c	(0.039)^b

Notes: This table reports regression results "horizontally" with each explanatory variable in a different column. The first row of the table duplicates the coefficients of the baseline regression of table 5, column 2. The second row duplicates the (baseline) robust standard errors clustered by country corresponding to this regression, while the third row reports the regular OLS standard errors. The fourth row reports the FGLS standard errors that correspond to the same estimation. The fifth row reports coefficients for the WLS estimation of our baseline specification. The last two rows report the uncorrected standard errors for this estimation and the robust and country-clustered standard errors for the same estimation. Significance levels are reported together with standard errors to avoid ambiguities. a, b, c: significant at 1%, 5%, 10%.

estimates, since smaller cities with less precisely estimated fixed effects are weighted less. Consistent with the results of table C.1 where we vary the sample of cities, the differences in estimated coefficients are small. The standard errors for non-robust and unclustered WLS estimates are also close to their OLS counterparts, albeit marginally larger. The use of clustered and robust standard errors with WLS estimation increases standard errors like it does with OLS. The degree of significance of some coefficients declines but none becomes insignificant. Given our large sample even in the smallest cities, we expect the share of the error in our estimates due to sampling to be low. In this case, WLS is known to be less efficient than OLS (Lewis and Linzer, 2005), hence the larger standard errors. Thus, our preferred comparison is between OLS and FGLS.

2 *Robustness to additional road, city, and country attributes*

Table C.4 considers additional attributes of the road network beyond our baseline six-attribute regression. While the main results from this table are discussed in the main text, we provide further details about some of them here. Column 1 adds a measure of average speed limit from OSM. Within each city, we compute a weighted average speed limit for motorways, primary, secondary, and tertiary roads, weighting by their share in total road length across all cities. Speed limit information is often missing from OSM, so our estimation restricts the sample to 704 cities in which all four road types have this information for at least 5% of road segments. In column 1, a 10% increase in speed limit is associated with 2.6% higher urban travel speed, relative to a regression, shown in column 2, that holds constant the sample of 704 cities with speed limit data. Column 3 adds a measure for the average number of lanes, again weighting each road type by its share of total road length in all cities and limiting our sample to 818 cities in which we have lane information for at least 5% of road segments. The coefficient on the number of lanes is insignificant and does not affect other coefficients as shown in column 4, which returns to our baseline specification with the same sample of cities as column 3. Column 5 adds three additional characteristics of the road network from OSM: the average block length, the average number of edges at intersections, and the average circuitry (actual distance divided by straightline distance) of a trip. None of these variables has a significant coefficient, suggesting that our six-attribute regression captures the first-order feature of road networks. Column 6 modifies our major road length variable by removing tertiary roads for which we are less confident about data quality. The coefficient on major roads drops by about 15%, as expected given the narrower definition of major roads. Note that the road elasticity would be smaller still if we kept only motorways. While motorways are fast, they account for a very small share of all roads. Therefore, increasing the share of motorways by 1% has a small impact on the average speed of travel.

Table C.4: Robustness to additional road attributes

Dependent variable: log city speed	(1)	(2)	(3)	(4)	(5)	(6)
log country GDP (pc)	0.034^a	0.046^a	0.043^a	0.044^a	0.038^a	0.044^a
	(0.011)	(0.011)	(0.011)	(0.011)	(0.014)	(0.013)
log population	-0.16^a	-0.18^a	-0.18^a	-0.18^a	-0.17^a	-0.16^a
	(0.018)	(0.020)	(0.018)	(0.018)	(0.017)	(0.017)
log area	0.063^a	0.096^a	0.095^a	0.096^a	0.087^a	0.091^a
	(0.021)	(0.024)	(0.025)	(0.026)	(0.021)	(0.021)
asinh major road length	0.10^a	0.091^a	0.097^a	0.096^a	0.079^a	
	(0.015)	(0.017)	(0.017)	(0.018)	(0.014)	
Network griddiness (0-1)	0.12^a	0.17^a	0.16^a	0.16^a	0.17^a	0.18^a
	(0.044)	(0.058)	(0.057)	(0.059)	(0.062)	(0.057)
Elevation variance	-0.0024^b	-0.0026^b	-0.0024^b	-0.0024^b	-0.0026^a	-0.0023^b
	(0.00091)	(0.0011)	(0.0010)	(0.0010)	(0.00095)	(0.00095)
asinh water body length	-0.059^a	-0.077^a	-0.081^a	-0.082^a	-0.074^a	-0.067^a
	(0.014)	(0.022)	(0.023)	(0.022)	(0.020)	(0.018)
log average speed limit	0.26^a					
	(0.061)					
log average number of lanes			0.021			
			(0.050)			
Average block length					-0.000050	
					(0.00013)	
Average edges at intersections					0.048	
					(0.041)	
Average circuity					0.46	
					(0.37)	
asinh major road length (excl. tertiary)						0.068^a
						(0.011)
Observations	704	704	818	818	1,119	1,119
R^2	0.77	0.74	0.75	0.75	0.71	0.71

Notes: Each column reports a variation of our baseline specification of table 5, column 2. Column 1 adds an average speed limit variable. Column 2 replicates our baseline on the same sample of cities for which the speed limit variable is available. Column 3 adds an average number of lanes variable to our baseline. Column 4 replicates the baseline on the same in which the average lanes variable is available. Column 5 adds variables for average block length, average number of edges at intersections, and average circuitry. Column 6 replaces the major road length variable in our baseline by removing tertiary roads. Robust standard errors clustered at the country level in parentheses. a, b, c: significant at 1%, 5%, 10%.

Table C.5 considers additional city characteristics. The key results are discussed in the main text, and the table's note provides the details of each specification. Here, we provide additional details on some specifications. Column 3 adds an index of city shape from Harari (2020) to our baseline specification. This index captures the compactness of a city, measured as the first principal component of the six main measures of city shape (perimeter, spin, dispersion, range, proximity, and cohesion) in 2010. Column 4 adds a measure of population gradient from the city center. Specifically, for each city we collect the population in each 100-meter pixel from WorldPop (Bondarenko *et al.*, 2020), and winsorize these data at the global

Table C.5: Robustness to additional city attributes

Dependent variable: log city speed	(1)	(2)	(3)	(4)	(5)
log country GDP (pc)	0.072^a	0.043^a	0.043^a	0.042^a	0.043^a
	(0.011)	(0.011)	(0.011)	(0.011)	(0.012)
log population		-0.17^a	-0.16^a	-0.17^a	-0.16^a
		(0.018)	(0.019)	(0.018)	(0.020)
log area		0.079^a	0.079^a	0.085^a	0.051^c
		(0.023)	(0.022)	(0.022)	(0.027)
asinh major road length	0.023^a	0.081^a	0.079^a	0.080^a	0.094^a
	(0.0076)	(0.012)	(0.013)	(0.013)	(0.016)
Network griddiness (0-1)	0.19^a	0.18^a	0.18^a	0.18^a	0.23^a
	(0.052)	(0.056)	(0.055)	(0.057)	(0.058)
Elevation variance	-0.0025^b	-0.0024^a	-0.0024^a	-0.0024^b	-0.0018^c
	(0.00099)	(0.00092)	(0.00092)	(0.00094)	(0.00095)
asinh water body length	-0.11^a	-0.077^a	-0.070^a	-0.071^a	-0.059^a
	(0.021)	(0.026)	(0.019)	(0.019)	(0.022)
log population density	-0.14^a				
	(0.016)				
Harari city shape		0.0073			
		(0.0090)			
asinh population gradient			0.013		
			(0.013)		
Capital city				0.0049	
				(0.014)	
Population in 1900 $> 100,000$					0.0045
					(0.014)
Observations	1,119	1,119	1,119	1,119	610
R^2	0.68	0.70	0.70	0.70	0.73

Notes: Each column reports a variation of our baseline specification of table 5, column 2. Column 1 replaces population and area variables in the baseline specification with a population density variable. Column 2 adds a city shape index from Harari (2020) to our baseline. Column 3 adds a variable for population gradient to our baseline. Column 4 adds a capital city dummy to our baseline. Column 5 adds an indicator variable for whether a city had a population of at least 100,000 in 1900. Robust standard errors clustered at the country level in parentheses. a, b, c: significant at 1%, 5%, 10%.

99th percentile. We then regress the inverse hyperbolic sine (asinh, to address zeroes) of population on asinh of distance to city center and report the resulting regression coefficient. For this purpose, the city center is defined as the pixel with the largest population. We experimented more broadly with variables further describing the internal structure of cities and failed to uncover any robust correlation. Basic city attributes like population, area and, to a lesser extent, our measures of topography, are sufficient to describe the geography of cities for our purpose.

Column 5 adds an indicator variable for whether a city had a population of at least 100,000 in 1900. The data on historical population is from Chandler (1987), and available in 610 cities. We find that having a historically large population has no impact on speed. More

generally, we find no correlation between population in 1900 and current urban travel speed in our world sample.

In a regression not reported here, we restrict attention to the United States, where anecdotal comparisons between old transit-based New York and new car-based Los Angeles often support arguments on the historical roots of slow travel speeds. For this country, we do indeed find a large negative correlation of -0.56 between 1900 population and current urban travel speed. However, almost all of the impact of historical population on current speed is explained by our six contemporaneous city attributes. Within the United States, cities that had historical populations above 100,000 in 1900 are 6.5% slower today. However, after controlling for our six city attributes, these historically large US cities are only 1.2% slower. (After removing current city population from the regression, the negative impact of historical population on speed is still reduced to 2.9%.) In a similar exercise for Europe, we experimented with the more comprehensive historical population data from Bairoch (1988), using the latest available date of 1850. The correlation between population in 1850 and current urban travel speed is -0.33. The elasticity of current city travel speed is -0.025 when estimated with a univariate regression and it falls slightly to -0.020 when controlling for our six main attributes, including contemporaneous city population. So even though there is no worldwide relationship between speed and historical population, we find one for the US and European countries, where historical population data has better coverage and is of higher quality.

Table C.6 considers additional attributes at the country level. Column 1 combines log city population and log area into a log density. Columns 2 and 3 add an indicator variable whether a country is a dictatorship. The dictatorship indicator equals one if the "Revised Combined Polity Score" from Marshall and Gurr (2020) is less than zero as of 2018. Countries classified as dictatorships are 6.9% faster, which is significant at the 1% level. However, there is no significant increase in speed for the capital cities of dictatorial countries. Column 4 adds an indicator variable for whether a country's first colonizer was France. The data on colonial origin is from Mayer and Zignago (2011). Cities in former French colonies are 8.4% faster, which is significant at the 5% level. In regressions not reported here, we experimented with other colonial status but failed to uncover any robust association. There is more suggestive evidence that cities in former Spanish colonies are slower but this effect is hard to separate from a Latin American effect. Column 5 adds a measure of risk preference from the Global Preference Survey (Falk *et al.*, 2018), to capture differences in driver behavior between countries. These risk preferences are elicited from survey participants from 76 countries using a mix of quantitative questions about preferences over lotteries, and qualitative questions about willingness to take risk. We find that a one standard deviation increase in a country's measure of risk preferences increases the speed by 2.4%, which is significant at the 5% level.

Table C.6: Robustness to additional country attributes

Dependent variable: log city speed	(1)	(2)	(3)	(4)
log country GDP (pc)	0.044^a	0.044^a	0.046^a	0.049^a
	(0.012)	(0.011)	(0.011)	(0.011)
log population	-0.17^a	-0.17^a	-0.17^a	-0.16^a
	(0.018)	(0.018)	(0.018)	(0.020)
log area	0.096^a	0.096^a	0.089^a	0.085^a
	(0.021)	(0.021)	(0.021)	(0.021)
Elevation variance	-0.0021^b	-0.0020^b	-0.0025^a	-0.0036^a
	(0.00093)	(0.00093)	(0.00080)	(0.0010)
asinh water body length	-0.067^a	-0.066^a	-0.070^a	-0.069^a
	(0.020)	(0.020)	(0.019)	(0.017)
Network griddiness (0-1)	0.20^a	0.20^a	0.18^a	0.18^a
	(0.054)	(0.054)	(0.055)	(0.052)
asinh major road length	0.073^a	0.073^a	0.077^a	0.071^a
	(0.015)	(0.015)	(0.012)	(0.015)
Dictatorship	0.069^a	0.067^b		
	(0.025)	(0.028)		
Capital city		-0.0061		
		(0.015)		
Capital city × dictatorship		0.013		
		(0.031)		
French colonial origin			0.084^b	
			(0.042)	
Risk preference				0.024^b
				(0.011)
Observations	1,116	1,116	1,119	1,029
R^2	0.72	0.72	0.71	0.73

Notes: Each column reports a variation of our baseline specification of table 5, column 2. Column 1 adds a dictatorship dummy to our baseline. Column 2 adds an interaction of capital city and dictatorship. Column 3 adds a dummy for whether a country's first colonizer was France. Column 4 adds a national level measure of risk preference. Robust standard errors clustered at the country level in parentheses. a, b, c: significant at 1%, 5%, 10%.

Appendix D. Unreliability

Following Brownstone and Small (2005), we calculate the unreliability of a peak-hour trip as the ratio of its 90th and 50th percentile durations across instances on different days at the same time of day (within a 5-minute window). To do this, we drop off-peak and weekend trips and calculate percentiles on trip instance speeds net of city-specific effects for each weekday, to account for, say, systematically faster Thursdays in one city and Mondays in another.

Table D.1 reports trip-level results with unreliability as the dependent variable. This table follows table 1, but does not include weather controls because we want to consider weather-induced variation in travel times as part of unreliability. Several results stand out.

Trip-level unreliability is not well-explained by city fixed effects or other trip attributes. In column 1, city fixed effects explain 10% of the variation and the most saturated specification, in column 6, only explains 16%. While these R-squared values are not directly comparable to those for our main indices, because the unit of analysis is the trip, not the trip instance, they are nonetheless low. As noted in the main text, most route attributes affect unreliability in the same direction as they affect congestion. The main exception is that reliability increases with trip length, consistent with the law of large numbers. In other words, conditional on the segment-specific speed variance, a trip consisting of more segments will revert closer to the mean speed on average.

Table D.2 reports determinants of the city-level unreliability index estimated as a city fixed effect in table D.1, column 4. Table D.2 is analogous to table 5 in the main text. Column 1 first shows that 49% of variation in reliability is cross-country, slightly more than the 45% of variation in congestion that is cross-country. In column 2, unreliability, like congestion, increases with country income. Income explains nearly a fifth of cross-city variation in unreliability, nearly twice what it explains for congestion. In column 3, other city attributes all have effects on unreliability consistent in direction with their effects on congestion. Altogether, they explain unreliability and congestion with similar explanatory power. The precisely estimated zero for the coefficient on road length is analogous to the fundamental law of congestion. However, our city attributes have less impact on the income-unreliability elasticity than on the income-congestion elasticity.

Appendix E. More on the Gelbach decomposition

1 Full derivation of the Gelbach decomposition

Starting from the specification of utility in equation (8), we consecutively eliminate c, t, R, L, V, $\{P, Q, q\}$, and A using equations (10c), (11), (12), (19), (20), (18), and (21), respectively. After some algebra, we obtain an expression for population as a function of income y:

$$
\begin{aligned}
N = \quad & \Omega_N \, y^{\Gamma_N} \,, \qquad \text{(E1)} \\
\text{with:} \quad & \Lambda \equiv \frac{\chi}{(\gamma+1) + \chi[\alpha(\theta+\mu)(\gamma+1) - \alpha\gamma(\phi+\nu) + \beta]} \,, \\
& \Omega_N \equiv \left[\frac{\tau^{\alpha\phi(\gamma+1)} (1-\tau)^{(\gamma+1)(1-\alpha\theta) - \alpha(-\gamma\nu+\phi) - \beta}}{\alpha^{\alpha\theta(\gamma+1)} \, \beta^{\alpha(\phi-\gamma\nu)+\beta}} \right]^{\Lambda} \Omega_A^{\alpha(\gamma+1)\Lambda} \,, \\
& \Gamma_N \equiv \Lambda \left[(1 - \alpha\theta + \alpha\sigma\phi + \alpha\Gamma_1)(\gamma+1) + \alpha\gamma(\nu+\phi) - \beta \right] .
\end{aligned}
$$

Similarly, starting from equation (E1), we substitute for N using equation (19) and then

Table D.1: Determinants of trip unreliability

Dependent variable: unreliability	(1)	(2)	(3)	(4)	(5)	(6)
log trip length		-0.0022[a]	-0.0023[a]	-0.0066[a]	-0.0070[a]	-0.0072[a]
		(0.00044)	(0.00046)	(0.00047)	(0.00044)	(0.00045)
log distance to center		-0.0090[a]	-0.0092[a]	-0.0090[a]		
		(0.00056)	(0.00058)	(0.00055)		
Gross gradient up				-0.10[a]	-0.090[a]	-0.088[a]
				(0.023)	(0.021)	(0.021)
Gross gradient down				-0.049[b]	-0.037[c]	-0.035[c]
				(0.021)	(0.021)	(0.021)
Gradient missing				-0.0010	-0.0019	-0.0030
				(0.0036)	(0.0033)	(0.0029)
Share of primary roads				-0.023[a]	-0.023[a]	-0.023[a]
				(0.0013)	(0.0012)	(0.0012)
Share of secondary roads				-0.037[a]	-0.036[a]	-0.036[a]
				(0.0012)	(0.0012)	(0.0012)
Share of tertiary roads				-0.048[a]	-0.047[a]	-0.047[a]
				(0.0014)	(0.0013)	(0.0013)
Share of residential roads				-0.059[a]	-0.058[a]	-0.057[a]
				(0.0016)	(0.0014)	(0.0014)
Share of other roads				-0.051[a]	-0.051[a]	-0.051[a]
				(0.0016)	(0.0015)	(0.0015)
Share of missing roads				-0.039[a]	-0.040[a]	-0.039[a]
				(0.0013)	(0.0013)	(0.0012)
log intersections				-0.0053[a]	-0.0048[a]	-0.0050[a]
				(0.00047)	(0.00041)	(0.00040)
arsinh turns against traffic				-0.0077[a]	-0.0077[a]	-0.0077[a]
				(0.00021)	(0.00020)	(0.00020)
arsinh density				0.0013[c]	0.00052	-0.00032
				(0.00072)	(0.00073)	(0.00074)
Type: circumferential		-0.0022[a]	-0.0022[a]	0.00054	0.00059	0.00011
		(0.00057)	(0.00057)	(0.00050)	(0.00049)	(0.00048)
Type: gravity		-0.0015[a]	-0.0015[a]	-0.00012	0.000022	-0.00024
		(0.00037)	(0.00038)	(0.00033)	(0.00032)	(0.00031)
Type: amenity		0.0019[a]	0.0020[a]	0.0022[a]	0.0023[a]	0.0020[a]
		(0.00039)	(0.00040)	(0.00034)	(0.00033)	(0.00032)
City effect	Y	Y	Y	Y	Y	Y
Time effect	N	Y	Y	Y	Y	Y
Weight	N	N	Y	Y	Y	Y
City-specific distance to center	N	N	N	N	linear	cubic
Observations	3,958,118	3,958,043	3,953,139	3,935,718	3,935,718	3,935,718
Within R^2	0.00	0.016	0.020	0.048	0.057	0.063
R^2	0.10	0.11	0.12	0.14	0.15	0.16

Notes: Each column (except column 1) reports parameter estimates for the unreliability variant of equation (1), regressing trip unreliability on log trip length, city, time-of-day (30-minute period), and trip indicators (radial trips are the reference), for weekday trip instances in our main sample of 1,119 cities. Columns 3–6 weight trips by their time of day according to the share of trips departing at that time of day in the US NHTS and population density within 111 meters of the route. Columns 4–6 add controls for trip characteristics (motorways are the omitted road class). Columns 5–6 interact log distance to center and city fixed effects. Column 6 further interacts city fixed effects with squared and cubed distance to center. Robust standard errors clustered at the city level are in parentheses. a, b, c: significant at 1%, 5%, 10%.

for P using equation (18) to express road length as a function of income y:

$$L = \quad \Omega_L y^{\Gamma_L}, \tag{E2}$$

Table D.2: Determinants of city level unreliability

		Unreliability index		
		(1) Country fixed effects	(2) Base model	(3) Full regression
	log country GDP (pc)		0.013^a	0.016^a
			(0.0031)	(0.0033)
City size	log population			0.011^a
				(0.0032)
	log area			-0.0014
				(0.0040)
Infrastructure	asinh major road length			-0.00028
				(0.0038)
	Network griddiness (0-1)			-0.050^a
				(0.0062)
Topography	Elevation variance			0.00057^b
				(0.00027)
	asinh water body length			0.0051
				(0.0041)
	Observations	1,119	1,119	1,119
	R^2	0.49	0.18	0.42

Notes: The unit of analysis is the city. Columns 1–3 each report parameter estimates of separate OLS regressions on the sample of 1,119 cities with high-quality speed data and country GDP data. Standard errors clustered at the country level in parentheses. a, b, c: significant at 1%, 5%, 10%.

$$\text{with:} \qquad \Omega_L \equiv \frac{\tau}{[\beta(1-\tau)]^{\frac{1}{\gamma+1}}} \Omega_N^{\frac{\gamma}{\gamma+1}}\,,$$
$$\Gamma_L \equiv \frac{\gamma}{\gamma+1}(1+\Gamma_N)\,.$$

Finally, we relate land area Q to income by substituting for N in equation (18) with its expression in equation (E1):

$$Q = \qquad \Omega_Q A^{\Lambda_Q} y^{\Gamma_Q}, \tag{E3}$$
$$\text{with:} \qquad \Omega_Q \equiv [\beta(1-\tau)]^{\frac{\gamma}{\gamma+1}}\, \Omega_2^{\frac{\gamma}{\gamma+1}}\,,$$
$$\Gamma_Q \equiv \frac{\gamma}{\gamma+1}(1+\Gamma_N) = \Gamma_L\,.$$

2 *Gelbach decompositions with more exogenous variables*

The decomposition proposed above readily generalizes to situations where we consider that both GDP per capita and topography are exogenous. The differences are the following. First, the base regression in equation (2) now includes all the exogenous variables (GDP per capita, the variance of elevation, and the length of water bodies) instead of only GDP per capita in our baseline decomposition. Then, auxiliary regressions corresponding to equations (22b)-

(22d) for the remaining endogenous variables also include all exogenous variables. Finally, there is one decomposition analogous to equation (23) for each endogenous variable, and this decomposition no longer includes topography in its summation.

Recall that the contributions of each endogenous variable are given by $\Gamma_k \psi_k^U / \kappa_{\text{Base}}$ and $\Gamma_k \psi_k^K / \kappa_{\text{Base}}$. In the base regression, the relationship between speed and GDP per capita is slightly stronger, with the coefficient on log GDP per capita at $\kappa_{\text{Base}} = 0.15$ instead of 0.13 in our baseline decomposition. This slightly higher elasticity of speed with respect to income thus leads to slightly lower contributions. The full regressions for uncongested speed and congestion are unchanged. Thus, the coefficients ψ_k^U and ψ_k^K on endogenous characteristics (roads, population, area, and network griddiness) remain the same. The auxiliary regressions now include three explanatory variables instead of one. For roads, the income elasticity Γ_{roads} is lower when including the variance of elevation and the length of water bodies (0.49 instead of 0.67). This slightly reduces the importance of roads in the decomposition. The same happens for land area (income elasticity of 0.09 instead of 0.28). For population, the income elasticity is higher when we include topography (-0.30 instead of -0.10).

Appendix references

Ahlfeldt, Gabriel M. and Elisabetta Pietrostefani. 2019. The economic effects of density: A synthesis. *Journal of Urban Economics* 111(0):93–107.

Akbar, Prottoy A., Victor Couture, Gilles Duranton, and Adam Storeygard. 2023. Mobility and congestion in urban India. *American Economic Review* 113(4):1083–1111.

Bairoch, Paul. 1988. *Cities and economic development: From the dawn of history to the present*. Chicago, Ill.: University of Chicago Press.

Boeing, Geoff. 2017. OSMNX: New methods for acquiring, constructing, analyzing, and visualizing complex street networks. *Computers, Environment and Urban Systems* 65(1):126–139.

Bondarenko, Maksym, David Kerr, Alessandro Sorichetta, and Andrew J. Tatem. 2020. Census/projection-disaggregated gridded population datasets, adjusted to match the corresponding unpd 2020 estimates, for 183 countries in 2020 using built-settlement growth model (BSGM) outputs. WorldPop, University of Southampton.

Brownstone, David and Kenneth A. Small. 2005. Valuing time and reliability: Assessing the evidence from road pricing demonstrations. *Transportation Research Part A* 39(4):27–293.

Chandler, Tertius. 1987. *Four thousand years of urban growth: An historical census*. Lewiston, (NY): Mellen.

Combes, Pierre-Philippe, Gilles Duranton, and Laurent Gobillon. 2019. The costs of agglomeration: House and land prices in French cities. *Review of Economic Studies* 86(4):1556–1589.

Falk, Armin, Anke Becker, Thomas Dohmen, Benjamin Enke, David Huffman, and Uwe Sunde. 2018. Global evidence on economic preferences. *Quarterly Journal of Economics* 133(4):1645–1692.

Feenstra, Robert Inklaar, Robert C. and Marcel P. Timmer. 2015. The next generation of the Penn World Table. *American Economic Review* 105(10):3150–3182.

Gonzalez-Navarro, Marco and Matthew A. Turner. 2018. Subways and urban growth: Evidence from earth. *Journal of Urban Economics* 108:85–106.

Harari, Mariaflavia. 2020. Cities in bad shape: Urban geometry in India. *American Economic Review* 110(8):2377–2421.

Lewis, Jeffrey B. and Drew A. Linzer. 2005. Estimating regression models in which the dependent variable is based on estimates. *Political Analysis* 13(4):345–364.

Li, Yue, Martin Rama, Virgilio Galdo, and Maria Florencia Pinto. 2015. A spatial database for South Asia. Unpublished manuscript.

Marshall, Monty G. and Ted Robert Gurr. 2020. Polity5: Political regime characteristics and transitions, 1800-2018". dataset users' manual. `http://www.systemicpeace.org/inscr/p5manualv2018.pdf` Accessed: 2022-08-05.

Mayer, Thierry and Soledad Zignago. 2011. Notes on cepii's distances measures: The geodist database. Working Papers 2011-25, CEPII. URL `http://www.cepii.fr/CEPII/en/publications/wp/abstract.asp?NoDoc=3877`.

Organisation for Economic Co-operation and Development (OECD). 2021. Metropolitan areas database. `https://stats.oecd.org/Index.aspx?DataSetCode=CITIES`. Accessed: 2021-05-29.

Organization of the Petroleum Exporting Countries (OPEC). 2019. OPEC annual statistical bulletin 2019. Technical report, Organization of the Petroleum Exporting Countries.

Pesaresi, Martino and Sergio Freire. 2016. GHS-SMOD R2016A - GHS settlement grid, following the REGIO model 2014 in application to GHSL landsat and CIESIN GPW V4-multitemporal (1975-1990-2000-2015). European Commission, Joint Research Centre.

United Nations, Department of Economic and Social Affairs, Population Division. 2019. *World Urbanization Prospects: The 2018 Revision (ST/ESA/SER.A/420)*. New York: United Nations.